Learning and the sources of corporate growth

*LEARNING AND THE SOURCES OF CORPORATE GROWTH**

Paul Geroski,
London Business School

Mariana Mazzucato,
Open University

This is an author-produced version of a paper published in ***Industrial and Corporate Change,* Vol. 11 (4). pp. 623-644.** ISSN 1464-3650. This version has been peer-reviewed, but does not include the final publisher proof corrections, published layout, or pagination. The published version is available to subscribers at http://icc.oxfordjournals.org/cgi/content/abstract/11/4/623

*We are obliged to the ESRC for financial support, and to Franco Malerba, Wes Cohen, Giovanni Dosi, Stan Metcalfe, Steve Klepper, and to participants at a Bocconi University a seminar and at the DRUID 2001 Nelson and Winter Conference for helpful comments on an earlier draft and to two anonymous referees**.** The usual disclaimer applies.

Correspondence:
Mazzucato: Economics, Open University, Walton Hall, Milton Keynes, United Kingdom, MK7 6AA, tel. 00-44-1908-654437, fax. 00-44-1908-654488, m.mazzucato@open.ac.uk

ABSTRACT

This paper explores the link between learning and corporate growth by developing different models of learning and showing that they produce observably different models of corporate growth. Using data on the growth of a number of firms in the US Automobile industry during the 20th century, we compare these different models of growth in an effort to identify the major sources of learning which these firms seem to have relied on. Although there are interesting differences between growth processes pre and post the Second War, the basic conclusion that we are drawn to is that learning in this sector is largely unsystematic and opportunistic**.**

JEL Classification: L1 Market Structure, Firm Strategy, and Market Performance, **03** Technological Change.

I. INTRODUCTION

Most people regard knowledge as a key source of competitive advantage and, therefore, of corporate performance. As a consequence, a large literature has grown up over the past decade around the subject of corporate learning. Much of it is concerned with how learning occurs, how knowledge is retained within the firm and how learning affects the strategy and structure of a firm's operations. This literature carries a strong normative presumption that learning is a good thing for a firm to do. The problem is that the evidence that we have on the link between learning and corporate performance is sketchy. This is, of course, no surprise, as neither learning nor the stock of knowledge that it presumably creates is directly observable. Much of the econometric work that has been done to identify the effect of learning on corporate performance comes from work that uses expenditures on R&D or other easily measured inputs as proxies for knowledge accumulation. While this is better than nothing, these are, at best, very limited measures of the stock of knowledge which firms benefit from.

This paper takes a somewhat different approach to the problem of identifying the link between learning and corporate performance. Rather than using one or more imperfect proxies for corporate learning to explain performance, we start from the premise that both the process of learning, and the stock of knowledge that it creates, are wholly unobservable. What is observable, however, are the consequences of learning. What is more – and this is the key to our approach – different types of learning (or, different learning mechanisms) are likely to have different implications for the times series behavior of the observable consequences of learning, namely corporate performance (and, in particular, corporate growth). It follows, then, that one ought to be able to infer something about the nature of the learning which a firm does from observations about it's performance.

To exploit this insight we follow the lead of the R&D literature and posit the existence of a simple and stable relationship between the stock of knowledge possessed by a firm and it's output. We then distinguish five different types of learning, and show that each produces a distinctive time path for output. After reviewing some recent literature on firm learning in Section II, we will outline these models in Section III. Although these different models of learning are not nested, it is possible to next several of them in a more general model. Further, it turns out that they all reduce to a common, baseline specification which emerges from a model in which learning is wholly unsystematic and opportunistic, and we will take this to be our null hypothesis. In Section IV, we apply these five models to firm level data drawn from 85 years of the US automobile industry. We conclude in Section V with a few observations on future directions for research of this type.

II. LEARNING AND CORPORATE PERFORMANCE

The simplest and most familiar story about learning is that which has been built up around "learning by doing" and the learning curve. In this story, firms accumulate experience through production and generate a stock of knowledge which is proportional to their cumulative output. If corporate performance is driven by this stock of knowledge then performance differences between firms should be fairly stable over time (since differences in cumulative output will not change much once firms are very well established). It is usually argued that because learning is essentially an investment in a very specific product and associated production process, firms should pursue learning curve strategies only when consumer tastes or the technological environment is relatively stable. When it's market is turbulent, firms are more likely to gain competitive advantage by pursuing strategies that focus on exploring and creating new product variants (instead of investing even more in existing ones). Amongst other things, this suggests that firms pursuing learning curve strategies are likely to lose market share in periods of turbulence, but will outperform others when their market environment is stable and their customers are price sensitive. [1]

The business literature on core-competencies and the evolutionary literature on innovation has pushed the idea of corporate learning well beyond the notion that learning occurs as a simple "by-product" of doing.[2] The evolutionary variant of this view focuses on the variety of different ways that firms can learn and how this learning is tied to different sources of knowledge and technological capabilities, both of which may be embedded in organizational structures. In this way of thinking, learning evolves over time through the development of specific capabilities and costly investments in "absorptive capacity".[3] Whatever their source, these "capabilities" or "competencies" are widely regarded as being difficult to imitate, and certainly this is true for a competence like absorptive capacity . The fact that the development of a firm's competencies is likely to build on the nature and depth of its prior knowledge means that in this view learning is *path-dependent*: those firms that have developed a significant body of knowledge will be better learners and hence develop more knowledge in the future. This is not dissimilar to the learning by doing story, except that experiential learning is typically narrower than learning through investments in R&D (etc). Further, superior corporate performance based on such competencies is likely to be persistent, and, in fact, it is no surprise to discover, therefore, that this view of learning has developed, in part, to explain persistent inter-firm differences in accounting profitability that have been widely observed.

The other strong, potentially empirically important insight that has emerged from the literature on learning is the observation that learning to different organizational and industry environments, arguing that learning in different environments results in different technological trajectories which then affect future learning patterns:

> "*A first broad property is the diversity of learning modes and sources of knowledge across technologies and across sectors. For example, in some activities knowledge is accumulated primarily via informal mechanisms of learning by doing and learning by interacting with customers, suppliers, etc. In others, it involves much more formalized activities of search (such as those undertaken in R&D labs). In some fields, knowledge is mostly generated internally and specific to particular applications. In others it draws more directly upon academic research and scientific advances. Recent research suggests that this diversity of learning modes may be a major determinant of the diverse patterns of evolution in industrial structures (e.g. in terms of distribution of firm sizes, natality and mortality of firms, corporate diversification*)."[4]

Not only is this observation more or less indisputable, it also yields a classification of different types of learning, including: learning by doing, learning by using, learning from advances in science and technology, learning from inter-industry spillovers, learning by interacting, and learning by searching (this particular list is taken from Malerba, 1992). Depending on which sector that we are examining, some of these are internal to firm's production process or use of products, while others are external to the firm and are related to the development of science or the actions of its competitors. Needless to say, some of these sources of learning generate knowledge that is easy to protect, i.e. knowledge that is easily “appropriable”, and to the extent that this is true, they are likely to lead to long term performance differences between firms.

This literature is both rich and insightful, and it seems to lead to three conclusions: there are many sources of learning (and not just R&D), learning is likely to be sector specific and path dependent, and it is likely to lead to persistent differences in corporate performance. In what follows, we will examine this last conclusion by focusing on learning processes in a specific sector, using a set of models that allow us to trace the observable consequences of many different types of learning. However, linking learning – of whatever variety – to corporate performance is made difficult by the fact that there are numerous ways to measure performance. For many, accounting or stockmarket rates of return are the right choice, if only because they are what firms are supposed to maximize. There are, however, numerous difficulties in measuring rates of return. Corporate growth rates, on the other hand, are not always measured without error, but it is at least clear what is in principle being measured. Further, growth rates of output are typically highly correlated with productivity growth which is, perhaps, the best measure of corporate performance to use in the current context (our date are not rich enough to compute either total factor productivity or, for that matter, accounting rates of return). Finally, insofar as learning is linked to the growth and development of capabilities within the firm, it seems natural to seek – at least in the first instance – to find a reflection of this development in the evolution of that firm’s market position.

III. OBSERVABLE IMPLICATIONS OF DIFFERENT LEARNING PROCESSES

We start with the presumption that learning, and the stock of knowledge which results from it, are wholly unobservable. What can be observed, however, is the consequences of learning, and we focus here on output. We posit the existence of a simple and stable association (such as a “knowledge production function”) which links the stock of knowledge possessed by firm i at time t, KN(t), with it’s output, Q(t):

(1) $$Q(t) = A(t)KN(t)^{\alpha}$$

where A(t) summarizes the effects of all other inputs (and anything else) on output rates. Taking logs and first differencing,

(2) $$\Delta\log Q(t) = \Delta\log A(t) + \alpha\ \Delta\log KN(t).$$

Since the rate of growth of KN(t) is, by definition, the rate of learning, it follows immediately that the rate of growth of the firm will vary directly with the rate of learning. For future reference, we use LE(t) to denote the rate of growth of KN; i.e. $LE(t) \equiv \Delta\log KN(t)$. Further, for expositional ease, we will suppose that $\Delta\log A(t) \equiv \varepsilon(t)$, a white noise error process.[5]

Equations (1) and (2) suggests that corporate performance measured by the current rate of growth of the firm is a signal of the current rate of learning. The quality of that signal evidently depends on two things: α, the elasticity of output with respect to knowledge, and the inherent variability in A(t). One simple observation that one can make from this is that the performance of firms in high technology sectors where learning really matters and which are insulated from macroeconomic and other shocks will be more informative about their learning than might be the case in more traditional, cyclically sensitive sectors. A more substantive observation, though, is that cross firm comparisons of learning using this approach are likely to be most informative for a sample of firms with a similar α and a similar variance in A(t); i.e. for firms in the same sector. A third observation is that any inferences about learning that one makes using (1) or (2) are, of course, conditional on how allowance is made for other factors, A(t).

The production function approach apparently underlying Equation (1) appears at first sight to be rather restrictive, or at least mechanistic, and it is important to note that the basic relationship revealed in (1) emerges from almost all output choice models typically used in the theory of the firm literature. A firm that maximizes profits (or even just satisfices) will choose an output rate which depends on it's marginal costs and some parameters of demand, and, since these parameters are affected by learning, it follows that there will be a relationship between the stock of knowledge and output rates.[6] Even if this relationship is very complex (which it is not in most of the commonly used models), (1) can be regarded as a first order approximation to the true relationship between output and those cost or demand parameters affected by learning. It is, therefore, unlikely to paint an terribly inaccurate picture of how learning affects corporate performance (except, perhaps, in being overly simple).

The real problem with (1) or (2) is that neither KN(t) nor the rate of learning, LE(t), are directly observable. However, different types of learning will induce different time paths in KN(t), and so in Q(t). It follows, then, that observing movements in Q(t) over time may cast useful light on the time path of KN(t), and so on the rate of learning. We distinguish five different types of learning process:

(i) unsystematic learning

The simplest story of all about learning is that which says that firms learn things in a wholly unstructured, unsystematic sort of way, opportunistically absorbing whatever their environment (randomly) throws up and, just as likely, forgetting what they have learned previously. In this case,

$$LE(t) = \xi(t), \tag{3}$$

where $\xi(t)$ is an i.i.d. random variable with mean zero and a variance which is constant (or at least so long as the environment which generates the learning opportunities to which the firm passively responds is constant over time). In this case, it follows that

$$\Delta\log Q(t) = \varepsilon(t) + \xi(t) \equiv \nu(t), \tag{4}$$

meaning that firm size follows a random walk. This is, of course, exactly the state of affairs described by Gibrat's Law, and we take it as our null hypothesis in what follows.

(ii) learning by innovation

The basic idea of a learning by innovation story is that virtually all learning can be tied to the appearance of particular product or process innovations.[7] One can think of this in one of two ways: either these innovations embody all the learning that firms actually do, or they act as a signal that intensive (but unobserved) learning has occurred and, amongst other things, that it has produced the particular innovation in question. Either way, the presumption is that relatively little learning occurs between innovations, so that observing the realization of an innovation is tantamount to observing the act of learning. If I(t) is a count of major product or process innovations which are introduced by firm i at time t, then this story might be modeled as

(5) $$LE(t) = \beta I(t) + \xi(t)$$

where $\beta > 0$. It follows, then, that

(6) $$\Delta \log Q(t) = \alpha\beta I(t) + \nu(t),$$

meaning that output follows a random walk with a trend driven by the stochastic arrival of particular innovations.[8]

There are many ways that one can make the specification in (5) richer and possibly more realistic. If major innovations have long lasting effects on performance, or if their short run effect is much smaller than their long run effect, then (6) might be generalized to

(7) $$\Delta \log Q(t) = \alpha\beta(L)I(t) + \nu(t),$$

where $\beta(L)$ is a polynomial in the lag operator L. Similarly, one might generalize (6) by allowing the effects of particular innovations to be firm specific, or to depend on the number of previous innovations that the firm has produced (allowing for some kind of differential ability to use new innovations).

(iii) spillovers

It is, of course, very likely that firms will learn from their rivals, so that learning occurs largely as a result of imitation between firms. There is a well-established tradition of modeling spillovers in the patents and R&D literature, and it is one of two routes that we propose to follow here.[9] To avoid notational clutter, we continue to suppress the i subscript which identifies firm i, but we will use a j subscript to identify variables associated with any or all of i's rivals. If the effects of innovative activity spillover between firms, then (6) becomes

$$(8) \qquad \Delta\log Q(t) = \alpha\beta I(t) + \alpha\beta_j I_j(t) + \nu(t),$$

where β_j measures the effect that rivals innovations, $I_j(t)$, have on the performance of firm i.

There are several problems with (8) as a way of capturing spillovers, but the most serious is that it presumes that all spillovers are associated with observable innovations. It is more than possible that some or all of the entire stock of knowledge, $KN_j(t)$, possessed by firm j might spillover to I in a manner unrelated to the arrival of any particular innovation. If we suppose that this occurs with at least a one period lag, then (1) becomes

$$(9) \qquad \Delta\log LE(t) = \alpha_j \Delta\log KN_j(t-1) + \xi(t),$$

which yields a relationship between the growth rate of firm i in period t and that of its rivals in t-1,

$$(10) \qquad \Delta\log Q(t) = \lambda_j \Delta\log Q_j(t-1) + \nu(t),$$

where $\lambda \equiv \alpha\alpha_j$. Equation (10) shows clearly that a high degree of inter-firm spillovers is likely to lead to a convergence in growth rates and, possibly in the long run, firm size.[10]

(iv) learning by doing

As we noted earlier, the classic source of learning is experiential, and it underlies the famous “learning curve” much beloved of corporate strategists. Although there are many ways to think about “experience”, most accounts focus on cumulative production, $X(t) \equiv \Sigma_t Q(\tau)$, as the main driver of experience and, therefore, of learning,

(11) $$LE(t) = \phi \log X(t) + \xi(t),$$

so that

(12) $$\Delta \log Q(t) = \alpha\phi\log X(t) + \nu(t).$$

We have lagged X(t) in (11) – (12) to avoid spurious correlation.

(v) learning using internal resources

None of the learning mechanisms thus refer to any limit or constraint on the ability of firms to learn. In fact, most scholars believe that there are constraints on the ability of a firm to learn, and that these depend on a set of capabilities usually referred to as “absorptive capacity”. In addition, there may also be constraints on the speed with which firms accumulate knowledge (analogous to Penrose effects). Both of these observations effectively mean that the rate of growth of the stock of knowledge of the firm is likely to depend on its level and perhaps also on the recent increase in that stock. Further, it is generally believed that there may be increasing returns to knowledge; i.e. that knowledge gained today facilitates the acquisition of further knowledge tomorrow. This too will create a link between the stock of knowledge maintained by a firm today and tomorrow. Either way, it seems reasonable to believe that learning might depend on

(13) $$LE(t) = \delta\log KN(t\text{-}1) + \theta LE(t\text{-}1) + \xi(t),$$

which means that :

(14) $\Delta \log Q(t) = \rho \log Q(t-1) + \psi \Delta \log Q(t-1) + \mu(t)$,

where $\mu(t) = \varepsilon(t) + \alpha\theta\varepsilon(t-1) + \alpha\delta \log A(t-1)$, $\rho = \alpha^2\delta$ and $\varphi = \alpha\theta$.[11]

Equation (14) is a specification which is familiar from a large empirical literature on the growth of firms.[12] The most common version of (14) used in that literature sets $\varphi = 0$ and generates estimates of ρ, interpreting it as a measure of the degree of "*reversion to the mean*" (or, in more modern parlance, "*convergence*"). Equation (14) gives one way of thinking about how reversion to the mean occurs. The parameter δ reflects the effects of absorptive capacity. If $\delta > 0$ (which, of course, implies that $\rho > 0$), then learning is easier the larger is the current stock of knowledge possessed by the firm; i.e. increasing returns to knowledge accumulation prevails. In this situation, firms do not converge to a common size (or size distribution) in the long run. If, on the other hand, $\delta < 0$ (i.e. $\rho < 0$), then diminishing returns prevail (and, indeed, knowledge will gradually depreciate over time). In this case, we will observe reversion to the mean and convergence. When $\delta > 0$ knowledge and, therefore, size differences between firms become magnified over time; when $\delta < 0$, firms eventually converge in size. $\theta < 0$ (or, equivalently, $\psi < 0$) indicates the existence of diminishing returns to growth, and may, therefore, reflect limitations on the ability of firms to absorb knowledge over time and/or to turn that knowledge into increased growth. If, on the other hand, $\theta > 0$ (or, equivalently, $\psi > 0$), then firms will display sustained periods of high (or low) growth.[13]

in short

Using the basic framework, (2), we have outlined five empirical models of corporate growth based on five different learning mechanisms. These five can, in principle, be distinguished using a minimum of observable variables (we need only two): unsystematic learning (equation (4)) induces a random walk in firm size; learning by innovation (equations (6) or (7)) creates a correlation between current period growth rate of particular firms and the current and lagged innovations which they produce; learning by spillovers (equations (8) and (10)) create correlations between the growth of particular firms and either the innovations produced or the growth of their rivals; learning by doing (equation(12)) creates a correlation between current period growth and cumulative output; and learning from internal resources (equation(15)) creates a correlation between the growth of a particular firm and it's size lagged and/or previous growth.

The five models that we have discussed produce empirically observable differences in the times series behavior of firm output growth, and all of them reduce to model of unsystematic learning. However, the other four are not nested amongst themselves, which makes comparing them difficult. It is, however, possible to combine (10) and (14) into a slightly more general model:

$$\Delta \log Q(t) = \alpha_0 + \alpha_1 \log Q(t-1) + \alpha_2 \Delta \log Q(t-1) + \alpha_3 \log Q_j(t-1) + \alpha_4 \Delta \log Q_j(t-1) + \mu(t), \tag{15}$$

where, as before, the subscript "j" denotes rivals (in fact, we shall break these up into two groups: the big three car makers (GM, Ford and Chrysler) and the rest, in what follows). When $\alpha_3 = \alpha_4 = 0$, (15) reduces to that of learning via internal resources (equation (15)); when $\alpha_1 = \alpha_2 = \alpha_3 = 0$, (15) reduces to the output spillovers model (equation (10)); and when $\alpha_2 = \alpha_3 = \alpha_4 = 0$, it reduces to unsystematic learning (equation (4)) which is our null hypothesis

IV. PATTERNS OF CORPORATE GROWTH IN THE US AUTOMOBILE INDUSTRY

Our goal in what follows is to use these several different models of learning to try to identify the major sources of corporate growth for firms. Since it seems reasonable to believe that learning occurs in different ways in different sectors, we have drawn our sample of firms from a single industry – the US Automobile industry. Further, as learning is typically taken to lead to persistent performance differences between firms, we have collected as long a times series of data as possible to enable us to measure just how long persistent performance differences exist. Finally, as it seems plain that different environments (possibly including different stages of the industry life cycle) may lead to different types of learning or lead to different consequences of learning on corporate performance, we have tried to collect data across at least two apparently different competitive regimes for our sample: the pre-War period (1910-1941) and the post-War period (1949-1998).

the data

Our data covers the period 1910-1998 for the US automobile industry. The firms (and time periods) included in our sample are shown on Table I. Although quite a number of US and foreign owned firms have operated in this industry over the years, we limit our investigation to those for which enough times series data was available to make sensible computations (needless to say, they are also the ones with the most significant market shares). There are only two observables in the models outlined in Section III: the output rate of each firm in each year and the number of innovations they produced in each year.[14]

Figure I displays total industry output from 1910-1998, while Figure II shows the annual growth rate of total industry output over the period (omitting the War years). Figure III plots the total number of innovations (product and process) produced in each year (up to 1981, when our innovation data ends). Two observations seem to be worth making about this data. First, it seems clear that there are two rather different periods which can be distinguished in the data: the pre-War (1910-1941) and the post-War (1949-1998) periods.[15] Indeed, the (second) War generates a “natural” break in the data, since car production in the US ceased for those several years. This division separates an earlier period in which the industry faced quite a lot of turbulence in technology and demand from a later one where both tastes and technology were, arguably, more stable. The pre-War period saw the establishment of the industry, while the post-War period saw it rise to dominate the US economy and then mature. Growth rates in the early period were higher but much more erratic than they were in the later period, when growth was (relatively) steady and sustained. Data for foreign firms is only available for the post-War period since the first significant entry of foreign firms began in 1965.[16]

Figure III illustrates the second observation that we wish to make, namely that innovative activity gradually declined over the whole history of the industry. This is so whether we simply count innovations, or try to asses their influence – in fact, hardly any post-War innovations received a weight above 5 in the transilience scale devised by Abernathy et al. (1983), while many received 6s and 7s in the pre-War period. There is a (surprising but clear) negative correlation between innovative activity and market size in the data, a (not very surprising) positive correlation between mean market growth and innovative activity and a (surprising) positive correlation between the variability of growth and innovative activity.[17]

Table II shows the mean and the standard deviation of the growth rates for all the domestic and foreign firms individually, and for all firms taken together. It also shows the *p*-values from the Shapiro-Wilks test for normality. As we saw on Figure II, the market as a whole grew more slowly after the War than before, but the interesting observation which springs out from Table II is that all of the individual firms in our sample (except Chrysler) experienced a higher average growth rates in the pre-War period than post-War. This relatively rapid post-War growth for individual firms in a market that was growing more slowly came at the expense of the many smaller firms who populated the market in the pre-War period and then exited after the War. The pre-War period is also characterized by a higher variance in growth rates for both those individual firms who operated in both periods, and for the market as a whole. Some of the foreign owned firms in our sample have average post-War growth rates that are similar to, or higher than, those recorded by the Big 3 firms in the pre-War period (Honda, Toyota Mitsubishi); VW and Mazda, however, had negative average growth rates.[18] The distribution of growth rates is, on the whole, normal (the exceptions are Packard, Studebaker, Honda and Toyota, and their departure from normality is not too egregious).

We have recomputed virtually all of the work discussed thus far (and that to be discussed anon) using relative growth rates (i.e. the growth rate of firm i less the average industry growth rate). This has the advantage of normalizing for a number of common industry shocks, and it is also consistent with much of the dynamic modeling done by evolutionary economists. As it happens, relative growths display the same basic patterns as are shown on Table II, except that relative growth rates are somewhat less likely to be normally distributed than (absolute) growth rates.

We also computed the correlations between the growth rates of the different firms What is interesting about these correlations is that, with the exception of those within the Big 3, they are rather small. Ford, GM and Chrysler all seem to expand and contract pretty much together, but other US and foreign firms displayed much more idiosyncratic patterns of growth. That is, correlations in the growth rates of most of the firms who operate in what is apparently the same market are rather low. This is rather surprising since one would normally expect firms in the same industry to be subject to much the same cost and demand shocks. One might argue that this is evidence that competition within the Big 3 is much higher than it is between members of the Big 3 and the rest or within the rest, or one might argue that they operate in different market segments or belong to different strategic groups within the broader automobile market.

Finally, we examined the correlations between current and recent past growth rates for each firm in the sample. Although these correlations are slightly higher in the post-War period than in the pre-War period, the simple fact is that they are all very small (Honda and, to a lesser degree, Toyota and Nissan are exceptions to this observation). The obvious conclusion is that is that high (or, for that matter, low) growth rates simply do not persist for long over time; that is, that firms do not, on the whole, enjoy long periods of sustained success (or failure). This, of course, implies, that performance differences (as measured by growth rates) between firms are unlikely to persist for long either. This, in turn, suggests that size differences will persist; i.e. that firm sizes are not converging to a common size or distribution of sizes. Autocorrelations of relative growth rates are stronger than those of (absolute) growth rates, but still relatively small. This lack of persistence in growth rates over time is, of course, consistent with the null hypothesis that learning is largely unsystematic.

regression results

We now report the results from the regression results for different learning models reviewed in Section III. We start with equation (15), then report the results from a regression of the learning by doing model and a regression for the learning by innovation model. We test the robustness of each model by adding three further variables (not reported in the tables to save space): a time trend, the growth rate of GDP, and the growth rate of the average industry stock price (the last two were lagged once, differenced, and in logs). In each case we report whether the results are qualitatively different when relative growth rates are used instead of absolute growth rates. All variables, except the innovation variables, are measured in logs.

The results in Table III indicate that in the pre-war period, the only independent variable that significantly affects each of the firm's growth (except that of Chrysler) is the lagged output variable (α_1), and it is negative. The growth of Chrysler and Packard is also significantly affected by the growth of the other Big 3 firms. It is interesting that the rate of growth of firm output (internal resources) is always insignificant. In the post-war period, lagged output is no longer significant for the Big 3 firms—only for American, Honda, Toyota, Mazda and Mitsubishi, where it is negative. Lagged output is significant only for American and Mitsubishi. The growth rate of the Big 3's output is significant in the case of Studebaker and Mazda. In both periods the coefficient on lagged output (α_1) is negative, confirming that growth is not explosive (reversion to the mean). In all cases, it is possible to reduce (15) either to equation (14) with $\psi = 0$, or, more frequently, to the null of unsystematic learning, equation (4). Further, the R^2s using (15) are not large, meaning that despite using a relatively full model (and much fuller than is typically used in the growth of firms literature) there is a very large amount of unsystematic movement in corporate growth rates.[19]

These results do not change much when *relative* growth rates are used in equation (15) instead of absolute growth rates. In the pre-war period, it is again lagged output that is the most significant variable, but this time GM is the exception instead of Chrysler (i.e. all firm's growth rates are significantly affected by their past output except for Chrysler). Also, whereas the growth rate of past output (internal resources) was insignificant in the case of absolute growth rates (in both periods), in the post-War period the relative growth rates of Studebaker, American Motors and Mazda are significantly affected by this variable. To check for the robustness of the results, we added to Equation (15) the following variables: the growth rate of GDP, a time trend, and the growth rate of the average industry stock price. The growth rate of GDP and the growth rate of the industry stock price proved to be significant (at the 1% level), and while in the pre-War period they did not alter the significance of lagged output (for GM and Ford), they did render lagged output insignificant in the post-War period (again for GM and Ford). This is true for both the case of absolute and relative growth rates.

In sum, the results using (15) suggest that the main driver of corporate growth is unsystematic learning. Although there are traces of capability formation conditioning growth, these traces are not strong and, in any case, point to a pattern of diminishing returns. However, all of the drivers of learning that we have assessed thus far are latent, and we have tried to capture their effects by lagged output, lagged growth and rivals lagged output or growth. Before we decide not to reject the null, we need to explore the other two learning mechanisms discussed in Section III. They both have a corresponding observable in our data (which is why they do not next easily into (15)): for learning by doing, it is cumulative output at time t-1, X(t-1), while for the innovations mechanism it is the innovations data developed by Abernathy et al. (1983), simple or weighted.

It turns out that all the regressions of the form of (12) that we ran produced insignificant correlations between growth and cumulative output, and the various statistical models that we estimated could all be reduced to the null, equation (4). It is, therefore, hard to see any clear basis for the learning by doing model in our data[20]. Perhaps this is because the persistent advantages that a firm builds up by travelling down a learning curve may not be persistent if it forgets things quickly (Benkard, 2000).

The other learning mechanism involves learning through innovations, whether they be the firm's own innovations (equation (6)) or those of it's rivals (equation (8)). Since it no doubt takes time for the new innovations that a firm introduces to affect its growth, we introduced the different innovation variables with a 5 year lag (and we looked at the effects that different lags might have on the results). We also used several alternative measures of innovation, in each case first adding product and process innovations together and then separating them. In each case, the innovation variables were included on their own and then along with the lagged output variable. All the innovation variables (except some of the dummies) were weighted using the weighting scheme devised by Abernathy et al. (1983) discussed above. The different measures of innovation that we used were: the number of own innovations (product plus process) in the last 5 years, the total number of rivals' innovations in the last 5 years (i.e. total innovations minus own innovations); the total industry innovations in the last 5 years; the squared industry innovations in the last 5 years; a dummy variable (not weighted) indicating whether the firm innovated in the last 5 years; a dummy weighted by the transilience scale; a dummy variable indicating whether rivals innovated or not in the last 5 years; and ranked dummies (equal to number of innovations if the firm innovated, and 0 if firm did not innovate). Each of these innovation variables were entered into equation (4) first separately and then together.

The results indicate that innovation does not affect firm growth very much. In fact, the only time that it does seem to have a significant positive effect is on the relative growth of GM and Chrysler in the *pre-War* period when the independent variable is the number of own innovations (product plus product) in the last 5 years. However, the significance of GM's innovations disappeared when the various robustness checks were implemented (the growth rate of GDP reduced the significance of GM's innovations). With absolute growth rates, Chrysler was the only firm whose growth rate seemed to be significantly affected by this innovation variable—in fact, Chrysler was clearly the firm for which the innovation variables most significantly affected firm growth. We also allowed the coefficient on innovation to vary by introducing interaction terms: a time trend multiplied by the number of innovations, the growth of GDP multiplied by the number of innovations and the number of rival innovations multiplied by the number of innovations. For both the case of absolute and relative growth rates, the only interaction term that was significant (in both periods) was the growth of GDP times the number of innovations. In the case of Ford and Chrysler, this particular interaction term was significant at the 1% level. The coefficient on lagged output (when included) was significant and negative in all the innovation regressions.

To check for simultaneity in the relationship between output growth and innovation, we also looked at the effect of growth on innovation, i.e. innovation as the dependent variable. We did this only for the Big 3 firms since it is for only those firms that we have innovation data. The weighted number of own innovations in the last 5 years (the only innovation variable that was significant above) was regressed on the lagged output, the rate of growth of lagged output, the lagged output of the Big 3, the lagged output of the other firms, the difference lagged output of the Big 3, and the lag of the rivals innovation. The only variable that came out significant was the lagged rivals innovation, indicating that although there may be no innovation spillovers on output growth, there are innovation spillovers on innovation. Why these do not transmit into output growth is a question we cannot answer here. These results held for both absolute and relative growth rates and in both the pre-War and post-War period.

V. CONCLUSION

If learning is an important source of competitive advantage, then firms who accumulate skills and knowledge appropriate to their environment will outperform those who do not. More interesting and possibly more useful is a second insight, namely that the way that firms learn ought to be visible in how they perform. In this paper, we have used this second insight to pursue the empirical implications of five different methods of learning on corporate growth performance in the US automobile industry in the pre and post-War periods. The bottom line is this: although one can detect traces of most of the different learning mechanisms in the data, it is hard to find any systematic evidence which supports any hypothesis other than that which asserts that learning is unsystematic and random. There are some signs that lagged output rates are correlated with growth rates, an observation which our models help us interpret as a consequence of capability accumulation (with diminishing returns), but the overwhelming impression that one gets from the data is that growth rates are unsystematic. If learning drives growth, then it must follow that most learning is opportunistic and unsystematic, and does not, on the whole, lead to sustained differences in the stock of knowledge between firms.

Much of what we have said here differs from the conventional wisdom about corporate performance because of the particular performance measure that we have been looking at. As is well known, data on accounting profitability are highly autocorrelated over time, and suggest that it is possible for firms to be persistently successful for many successive years. Further, the data shows that performance differences between firms persist for long periods of time. Accounting profits are, however, not the only measure of firm performance that one might legitimately examine, and measures like sales growth, productivity growth and (slightly less clearly) shareholder rates of return are much more variable over time and across firms. It is not clear whether there is one single best way to measure firm performance, but it does seem clear that relying only on the properties of accounting profitability to make inferences about performance patterns is liable to be highly misleading. This said, it is possible to be slightly uneasy about using annual growth rates to assess performance, and it is possible that their relatively high volatility might contain an unusually large number of measurement errors or reflect the effects of noisy events (like strikes). The important point, though, is this: there is a real choice to be made between using accounting returns or growth rates (annual or smoothing by averaging over time), and this is because their statistical properties are deeply incongruent.

Our conclusions about learning in the US Automobile industry are, in a sense, not all that surprising. There is a long and rich literature which suggests that the growth rates recorded by many firms in many sectors are, roughly speaking, random (or negatively correlated to lagged output rates), and that is basically what we have observed. The various models that we have set out in this paper have enabled us to interpret this result in terms of different learning processes, but it is, of course, possible to think about it in other ways.[21] One way or the other, the interesting observation is future growth performance is very hard to predict from current or recent past performance, or, for that matter, from average industry growth performance. The implications of this observation for how we think about corporate performance is, we think, profound, and goes well beyond the US Automobile industry. It is commonplace to think that firms differ from each other, but these empirical results say something much stronger than this: not only are performance differences not constant over time, but they vary in unpredictable ways over time. The common practice of accounting for differences in performance between firms by invoking an unobservable asset (like the stock of knowledge or some set of core capabilities) is fine, but these results suggest that one cannot suppose that these invisible assets are durable (or that their effects on performance are systematic and stable), or that it will be easy to predict how well today’s assets will fare tomorrow. Nothing in this life is permanent, and this seems to apply particularly clearly in the case of the sources of corporate growth.

Figure I: Total Industry Output

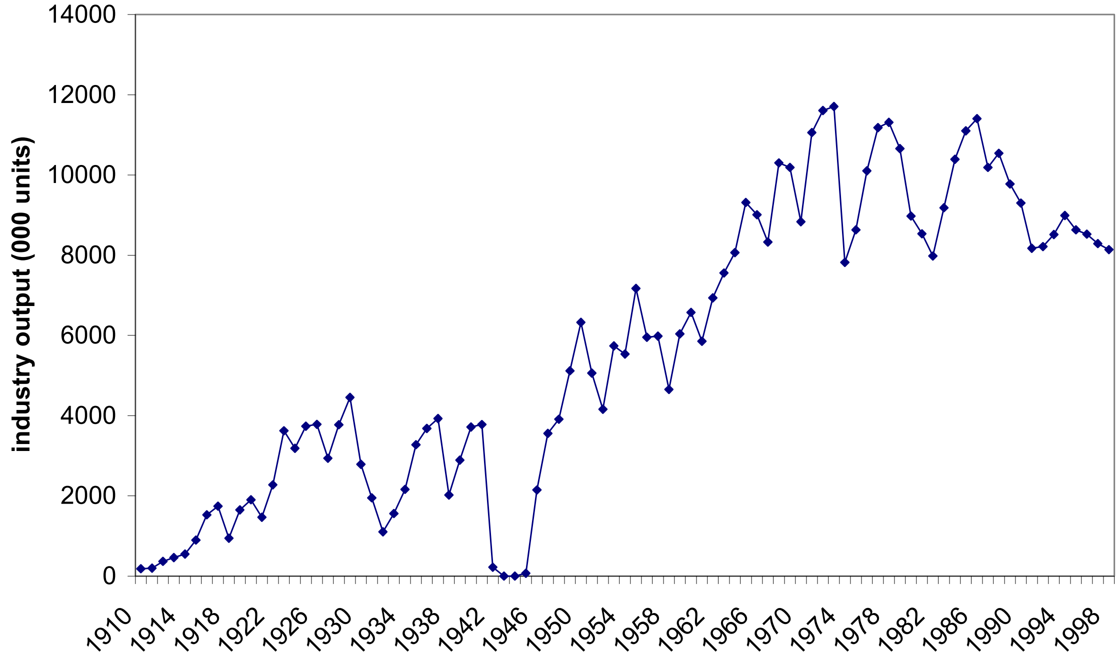

Figure II: The Rate of Growth of Industry Output

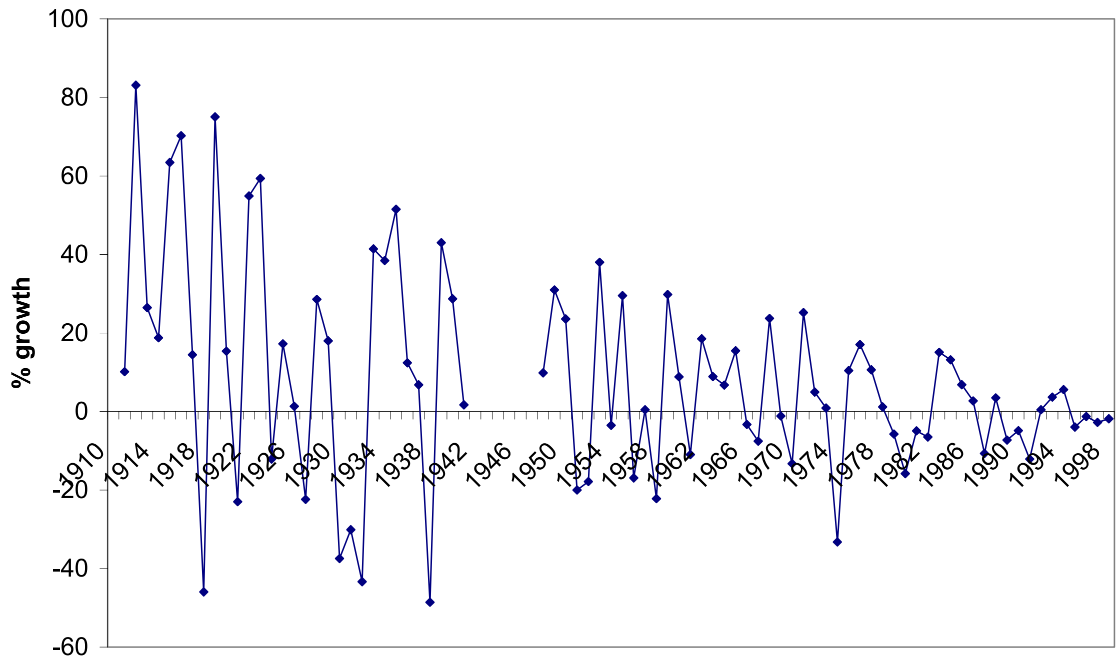

Figure III : Product and Process Innovations*

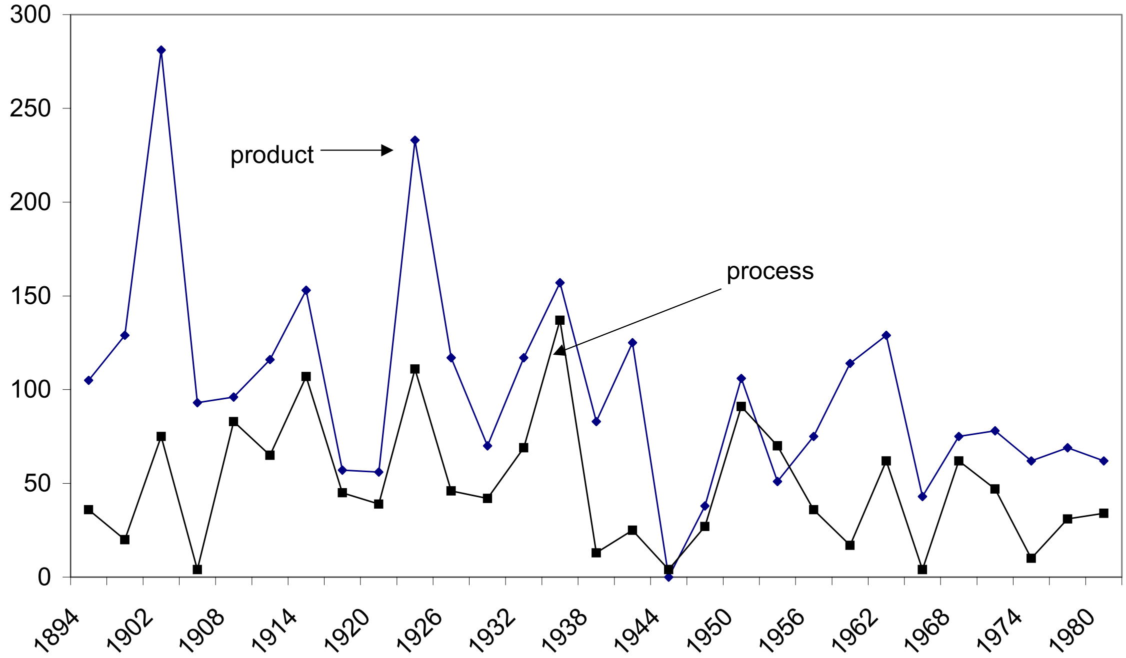

* 3 year moving average, weighted using "transilience scale" in Abernathy et al. (1983)

Table I: The Firms and Sample Periods

Big 3 firms:	
GM	1910-1998
Ford	1910-1998
Chrysler	1925-1998
Other US owned firms:	
Packard*	1910-1941
American	1946-1985
Studebaker	1946-1966
Foreign owned firms:	
VW	1965-1998
Honda	1971-1998
Nissan	1965-1998
Toyota	1996-1998
Mazda	1985-1998
Mitsubishi	1985-1998

* There is no output data for Packard for the period 1926 – 1934

Table II: Descriptive Statistics

Absolute Corporate Growth Rates

	Mean pre	Mean pst	St.Dv pre	St. Dv pst	Norm pre	Norm pst
GM	0.1244	0.0109	0.3857	0.1853	0.51407	0.5109
Ford	0.0994	0.0172	0.536	0.1946	0.2402	0.2268
Chrysler	0.1271	-0.0002	0.3411	0.4415	0.2597	0
Packard	0.194		0.5899		**0.0102**	
Stud		-0.2632		0.7078		**0.0345**
AMC		0.0045		0.249		0.3052
VW		-0.011		0.2518		0.5751
Honda		0.1643		0.2321		**0.0002**
Toyota		0.1209		0.2289		**0.0001**
Mazda		-0.0095		0.1253		0.7457
Mitsubishi		0.0838		0.1896		0.1249
TOTAL	0.1312	0.0169	0.3947	0.1484	0.2276	0.9259

Relative Corporate Growth Rates

	Mean pre	Mean pst	St. Dev. pre	St. Dev. pst	Normal pre	Normal pst
GM	0.0212	-0.0041	0.2370	0.0765	**0.0059**	**0.0018**
Ford	0.0238	0.0004	0.4536	0.1127	**0.0154**	0.7295
Chrysler	0.1264	-0.0149	0.1449	0.1499	0.3555	0.4718
Packard	0.0016		0.4047		0.5388	
Studebake	-0.0421	-0.2878	0.2316	0.6334	0.2793	**0.0006**
AMC	0.1130	-0.0251	0.3154	0.2832	0.3249	0.2729
VW		-0.0070		0.2413		0.9226
Honda		0.1757		0.2265		**0.0007**
Nissan		0.1188		0.2850		**0.0005**
Toyota		0.1242		0.2302		**0.0003**
Mazda		0.0144		0.1097		0.0801
Mitsubishi		0.1077		0.2153		**0.0272**

Cols. 1-4: Mean (pre and post-War) and standard deviation (pre and post-War) of firm growth rates.
Cols. 5-6: Shapiro-Wilk test for normality on growth rates (**bold** values indicate that can reject the null hypothesis that the variables are normally distributed).

Table III: Regression Results for Equation (15)

	pre-war					**post-war**										
	GM	FORD	CHRS	PACK	STUD	GM	FORD	CHRS	STUD	AMER	VW	HONDA	NISSAN	TOYOTA	MAZDA	MITS
Q (t-1)	-0.259	-0.555	-0.114	-1.430	-0.627	-0.178	-0.281	-0.322	0.213	-0.297	-0.365	-0.131	-0.144	-0.174	-0.865	-0.592
	(2.22)*	(2.58)*	(0.56)	(3.05)*	(2.58)*	(1.08)	(1.81)	(1.99)	(0.82)	(2.95)**	(1.91)	(2.58)*	(1.94)	(2.23)*	(4.21)**	(2.15)
D Q (t-1)	-0.146	0.112	-0.078	0.471	0.390	-0.152	-0.002	-0.331	-0.681	0.508	0.220	-0.049	0.015	0.287	0.455	1.073
	(0.71)	(0.54)	(0.10)	(1.55)	(1.46)	(0.74)	(0.01)	(1.82)	(1.40)	(3.06)**	(1.09)	(0.23)	(0.09)	(1.59)	(2.14)	(2.69)*
Q big3 (t-1)	0.230	-0.063	-1.187	1.397	0.631	-0.053	0.080	0.105	-1.330	0.157	0.068	-0.026	-0.115	-0.210	-0.796	-1.110
	(0.98)	(0.68)	(3.04)*	(2.90)*	(2.06)	(0.30)	(0.42)	(0.78)	(2.43)*	(1.00)	(0.16)	(0.14)	(0.44)	(1.28)	(3.70)*	(1.21)
Q other (t-1)	0.059	0.387	0.364	-0.778	-0.480	-0.015	-0.015	-0.032	-0.235	-0.072	-0.244	-0.343	-0.150	0.006	-3.791	4.275
	(0.26)	(1.50)	(1.38)	(2.01)	(1.75)	(0.40)	(0.40)	(0.78)	(0.83)	(1.82)	(1.99)	(1.40)	(0.70)	(0.04)	(3.87)*	(1.67)
D Q big3 (t-1)	0.034	-0.145	0.764	0.035	-0.482	0.148	-0.011	0.352	-0.336	-0.392	0.297	-0.035	0.215	0.000	-0.807	3.499
	(0.17)	(0.58)	(1.08)	(0.09)	(1.49)	(0.64)	(0.04)	(1.81)	(0.52)	(2.10)*	(0.64)	(0.14)	(0.71)	(0.00)	(2.69)*	(3.14)*
D Q other (t-1)	0.028	-0.043	-0.438	0.006	-0.234	-0.041	-0.029	-0.062	0.321	0.002	0.111	0.024	-0.194	-0.142	1.920	-5.442
	(0.12)	(0.16)	(0.96)	(0.02)	(1.18)	(0.32)	(0.21)	(0.45)	(0.50)	(0.04)	(0.24)	(0.08)	(0.37)	(0.46)	(2.23)	(2.23)
Constant	-0.261	3.950	14.548	2.978	2.967	3.698	3.056	3.265	20.586	2.131	6.871	7.117	5.854	5.553	78.507	-38.326
	(0.24)	(2.81)**	(3.01)*	(1.40)	(1.94)	(2.74)**	(2.13)*	(2.07)*	(2.21)*	(1.19)	(1.20)	(1.38)	(0.94)	(1.42)	(4.18)**	(1.10)
Observations	31	31	15	14	22	51	51	51	19	38	32	26	32	31	12	12
R-squared	0.23	0.41	0.64	0.76	0.49	0.19	0.19	0.28	0.65	0.41	0.27	0.69	0.47	0.71	0.91	0.84
F-test	0.3427	0.0315	0.1282	0.0566	0.0769	0.1478	0.1361	0.0192	0.0263	0.0089	0.2070	0.005	0.0089	0.0000	0.0177	0.0649
Ramsey	0.4573	0.1274	0.0659	0.3808	0.3327	0.6709	0.8004	0.6936	0.2062	0.3826	0.3253	0.0020	0.2580	0.0032	0.7892	0.7336
Likelihood	-9.902	-15.854	2.731	2.267	-1.223	19.826	17.516	16.98	-9.658	9.829	3.8518	20.247	12.875	26.792	22.340	13.986

Absolute value of t-statistics in parentheses* significant at 5% level; ** significant at 1% level
all variables in logs
Q=annual units produced
D= difference
big 3= GM, Ford, Chrysler
other=non big 3 firms

REFERENCES

Abernathy W.J. and K. Wayne (1974), "Limits to the Learning Curve," *Harvard Business Review, 52: 109-120.*

Abernathy, W.J., K. Clark and A. Kantrow, (1983) *Industrial Renaissance: Producing a Competitive Future for America,* Basic Books, New York

Aoki, M. (1986), "Horizontal vs. Vertical Information Structure of the Firm," *American Economic Review,* Vol. 76: 971-83.

Benkard, L. C. (2000), "Learning and Forgetting: The Dynamics of Aircraft Production," *American Economic Review*, 90 (4) 1034-1054.

Cohen, W. and D. Levinthal (1989), "Innovation and Learning: The Two Faces of R&D, Implications for the Analysis of R&D Investment," *Economic Journal,* 99: 569-596.

Cohen W. and D. Levinthal (1989) "Absorptive Capacity: A New Perspective on Learning and Innovation," *Administrative Science Quarterly*, 35: 128-152.

Coriat, N. and G. Dosi (1998), "Learning How to Govern and Learning how to Solve Problems: On the Co-Evolution of Competences, Conflicts, and Organizational Routines," in A. D. Chandler, Jr., P. Hagstrom and O. Solvell (eds) The Dynamic Firm: The Role of Technology, Strategy, Organization, and Regions, Oxford University Press, Oxford.

Dosi, G. and L. Marengo, (1994) "Some Elements of an Evolutionary Theory of Organizational Competences", in R. England (ed) Evolutionary Concepts in Contemporary Economics ,), University of Michigan Press, Ann Arbor, Michigan.

Dosi, G., O. Marsili, L. Orsenigo and R. Salvatore (1995), "Learning, Market Selection and the Evolution of Industrial Structures," *Small Business Economics*, 7: 411-436.

Dunn, P. and A. Hughes (1994) "Age, Size, Growth and Survival:UK Companies in the Late 1980's", *Journal of Industrial Economics*, 42, 115 – 140.

Evans, D. (1987) "Tests of Alternative Theories of Firm Growth", *Journal of Political Economy*, 95, 657 – 674.

Geroski, P. (1995) "Do Spillovers Undermine the Incentive to Innovate?", in S. Dowrick (ed) Economic Approaches to Innovation, Edward Elgar, Aldershot.

Geroski, P. (2000) "The Growth of Firms in Theory and Practice", in N. Foss and V. Mahnke (eds) , Competence, Governance and Entrepreneurship, Oxford University Press, Oxford.

Geroski, P. and S. Machin (1993), “Innovation, Profitability and Growth over the Business Cycle,” *Empirica*, 20: 35-50.

Geroski, P., S. Machin and C. Walters (1997) “Corporate Growth and Profitability”, *Journal of Industrial Economics*, 45, 171 – 189.

Geroski, P. and M. Mazzucato (2000), “Advertising and the Evolution of Market Structure in the US Car Industry”, mimeo, London Business School.

Geroski, P., and M. Mazzucato (2001) “Modelling the Dynamics of Industry Populations”, forthcoming, *International Journal of Industrial Organization*.

Geroski, P., I. Small and C. Walters (2000) “Agglomeration Economies, Technology Spillovers and Corporate Productivity Growth”, in R. Barrell, G. Mason and M. O’Mahoney (eds) Productivity, Innovation and Economic Performance, Cambridge University Press, Cambridge.

Geroski, P., S. Lazarova, G. Urga and C. Walters (2001) “Are Differences in Firm Size Transitory or Permanent?”, mimeo, London Business School.

Griliches, Z. (1979) “Issues in Assessing the Contribution of R&D to Productivity Growth”, *Bell Journal of Economics*, 10, 92 – 115.

Griliches, Z. (1991) “The Search for R&D Spillovers”, mimeo, NBER, Cambridge Mass.

Griliches, Z. (ed) (1984) R&D, Patents and Productivity, University of Chicago Press, Chicago.

Hart, P. and N. Outlton (1996) “Growth and Size of Firms”, *Economic Journal*, 106, 1242 – 1252.

Henderson, R. and K. Clark (1990), “Architectural Innovation: The reconfiguration of existing product technologies and the failure of established firms,” *Administrative Science Quarterly,* 35: 9-30.

Hunker, J.A. (1983), Structural Change in the US Automobile Industry, Lexington Books, Lexington, MA.

Kawahara, A. (1998), The Origin of Competitive Strength, Springer, New York.

Malerba, F. (1992), “Learning by Firms and Incremental Technical Change,” *The Economic Journal,* July 1992, p. 8545-855

March, J.G. (1991), “Exploration and Exploitation in Organizational Learning,” *Organization Science*, 2: 71-87.

Nadiri, I. (1993) “Innovation and technical Spillovers”, mimeo, NBER, Cambridge Mass.

Nelson, R. R., and S.G. Winter (1982), An Evolutionary Theory of Economic Change, Cambridge, MA: Harvard University Press.

Pavitt, K. (1984), “Sectoral Patterns of Technological Change: Towards a Taxonomy and a Theory,” *Research Policy,* 3, pp. 343-73.

Penrose, E. (1959), The Theory of the Growth of the Firm, Oxford, Basil Blackwell

Prahalad, C.K. and G. Hamel (1990), “The Core Competence of the Corporation,” *Harvard Business Review,* May-June 1990, p. 79-91

Spence, M. (1981), “The Learning Curve and Competition,” *Bell Journal of Economics*, 12: 49-70.

Sutton, J. (1997) “Gibrat’s Legacy”, Journal of Economic Literature, 35, 40 – 59.

Teece, D. (2000) Managing Intellectual Capital, Oxford University Press, Oxford.

Teece, D., G. Pisano and A. Shuen (1997), “Dynamic Capabilities and Strategic Management,” *Strategic Management Journal*, 18(7): 509-533.

Tushman, M. and P. Anderson (1986) “Technological Discontinuities and Organizational Environments,” *Administrative Science Quarterly*, 31: 439-465

Yelle, L.E. (1979), “The Learning Curve: Historical Review and Comprehensive Survey”, *Decision Sciences,* 10, 302 - 328.

White, L. (1971), The Automobile Industry, Harvard University Press, Cambridge, MA.

NOTES

[1] For a survey of early empirical work on the learning curve, see Yelle 1979; Spence, 1981, explores the relation between learning and market structure, while Abernathy and Wayne, 1974, pursue some of the strategy implications.

[2] This is an enormous literature, one that arguably began with Penrose, 1959 and Nelson and Winter, 1982; for more recent work, see Prahalad and Hamel, 1990; Teece et al. 1997; Coriat and Dosi, 1998; Dosi and Marengo, 1994, Teece, 2000 and (many) others.

[3] See Cohen and Levinthal, 1989, and others. Absorptive capacity is not just another competence (like marketing skills or the ability to manage geographically dispersed project teams): it is likely to be a result or by-product, of the firm's active engagement in learning activities: it is both a cause and a consequence of learning.

[4] Dosi et al, 1995. For a sectoral classification of patterns of technical change, Pavitt, 1984; see also Malerba, 1992, Dosi and Marengo, 1994, and others for further work in this vein.

[5] The simplest, least data demanding assumption to make about the evolution of A(t) over time is that it is driven by a large number of small, idiosyncratic cost and demand shocks. Under well known conditions, the effects of all of these shocks can be accurately described by a simple white noise process.

[6] For example, if demand is $P = \varphi_0 - \varphi_1 Q$ and costs are $C = \varphi_2 Q$, then the optimal output choice of the firm will be $Q = (\varphi_0 - \varphi_2)/2\varphi_1$. Hence, if costs are reduced or demand is increased by some sort of learning that lowers φ_2 or increases φ_0, the rate of growth of output will be positively related to the rate of learning.

[7] It would certainly be interesting to explore the relationship between learning and investments made to develop product or process innovations; i.e. R&D. However, we have been unable to compile an accurate series on R&D for the firms in our sample.

[8] This specification is very similar to that one often encounters in the literature which assesses the effects of patents and major innovations on corporate performance; see, for example, the papers in Griliches, 1985; Geroski et al, 2000,cites some of the more recent literature.

[9] For surveys of empirical work on spillovers, see Griliches 1979, 1991, Geroski, 1995, Nadiri, 1993, and many others.

[10] Previous work using corporate growth Equations like (10) have interpreted the co-efficient λ_j as a measure of agglomeration economies; see, for example, Geroski et al, 2000. Equation (10) is also very similar to the Lotka-Volterra Equation which is used by ecologists to describe the rate of growth of two populations which inhabit the same niche; see Roughgarden, 1996, Chapter 21.

[11] Much the same end might be achieved by supposing that R&D investments are driven by firm size (say, because firms adopt the decision rule of allocating x% of their revenues to R&D), and that (lagged) R&D drives growth. Since R&D is not measured (it is unavailable in our data), substituting it out yields a relationship between growth and lagged size.

[12] For work on corporate growth rates in this tradition, see Dunn and Hughes, 1994, Evans, 1987, Hart and Oulton, 1996, Geroski et al, 2001 and many others.

[13] Equation (14) leads to conclusions that are similar to those discussed in Geroski, 2000, which shows that a set of unobserved but durable competencies will induce a moving average in growth rates over time. In that model, one expects $\theta > 0$, reflecting the positive and persistent effects that such competencies have on corporate performance over time.

[14] The data is as follows: OUTPUT: Individual firm output (number of autos produced, or units) and total industry units from 1900-1999 are from annual editions of *Wards Automotive Yearbooks* (first editions, reporting data starting in 1904, are published in 1924). Although firm-level units were collected for only 8 domestic firms and 5 foreign firms (the first foreign firms entered in 1965), the *total* industry sales include the units shipped by all existing firms—e.g. in 1909 that includes the output of 271 firms. Units produced follow the same general qualitative dynamic as that of net sales in dollars but is preferred due to its greater precision (sales figures are affected by idiosyncratic accounting items). INNOVATION: Firm and industry level innovation figures were taken from a chronological list of product and process innovations (listed separately) in the auto industry from 1893 to 1981 (Abernathy et al., 1987). The authors devised a weighting scheme, a "transilience scale," to evaluate each innovation in terms of its overall impact on the production process: 1's represent those innovations that had little or no impact on the production process and 7's those innovations that were very disruptive for the production process. Although the innovation data only goes to 1981, this allows us to broadly compare 40 years in the pre-War period (during which most of the radical innovations occurred) with 40 years in the post-War period.

[15] We have experimented somewhat with altering the precise dates which define the beginning and end of these periods, and it seems to make little difference quantitatively (and no difference whatsoever qualitatively) to the results which are reported in what follows.

[16] Entry by foreign firms eventually led to a noticeable decline in the incumbents (i.e. major US owned automobile manufacturer's) market shares and a sharp rise in industry advertising: see Geroski and Mazzucato, 2000, and, more generally on the history of the industry, Hunker, 1983, Kawahara, 1998, and White, 1971, amongst many others.

[17] There is also a clear negative correlation between market size and concentration (the market became more concentrated as it grew), not least because of the extensive entry which occurred early in the pre-War period: see Geroski and Mazzucato, 2001.

[18] Mazda grew steadily until the early 1990s and then suffered a major fall in production; Volkswagon's early penetration into the market was reversed in the late 1960's, and their recovery (partial as it is) only dated from the mid-1990's. Honda, Nissan, Toyota and Mitsubishi as recorded fairly steady (and, in some cases, spectacular) market penetration over the post-War period.

[19] We also assessed the null more directly by computing Dickey-Fuller unit root tests (with and without time trends) on the data on firm size. In the pre-war period, it was impossible to reject the hypothesis that there is a unit root for all firms, and this was true for most of the foreign owned firms in the post-war period. These results are not quite consistent with those reported in the text using (15), although both set of results incline one towards accepting the null.

[20] We find that cumulative output variable is highly correlated with the lagged output variable, making cumulative output appear significant when lagged output is included as an independent variable in (12)—but this is of course a false illusion.

[21] There are, of course, other ways to interpret this result: see, for example, Geroski et al, 1997, who argue that a random walk in firm size is consistent with a model in which firms hold rational expectations about the future and choose an output trajectory subject to adjustment costs. Sutton, 1997, examines several other possible explanations.