No. 1862

AN APPLIED ECONOMETRICIAN'S VIEW OF LARGE COMPANY PERFORMANCE

Paul A Geroski

INDUSTRIAL ORGANIZATION

Centre for Economic Policy Research

ISSN 0265-8003

AN APPLIED ECONOMETRICIAN'S VIEW OF LARGE COMPANY PERFORMANCE

Paul A Geroski

Discussion Paper No. 1862
April 1998

Centre for Economic Policy Research
90–98 Goswell Rd
London EC1V 7DB
Tel: (44 171) 878 2900
Fax: (44 171) 878 2999
Email: cepr@cepr.org

This Discussion Paper is issued under the auspices of the Centre's research programme in **Industrial Organization**. Any opinions expressed here are those of the author(s) and not those of the Centre for Economic Policy Research. Research disseminated by CEPR may include views on policy, but the Centre itself takes no institutional policy positions.

The Centre for Economic Policy Research was established in 1983 as a private educational charity, to promote independent analysis and public discussion of open economies and the relations among them. It is pluralist and non-partisan, bringing economic research to bear on the analysis of medium- and long-run policy questions. Institutional (core) finance for the Centre has been provided through major grants from the Economic and Social Research Council, under which an ESRC Resource Centre operates within CEPR; the Esmée Fairbairn Charitable Trust; and the Bank of England. These organizations do not give prior review to the Centre's publications, nor do they necessarily endorse the views expressed therein.

These Discussion Papers often represent preliminary or incomplete work, circulated to encourage discussion and comment. Citation and use of such a paper should take account of its provisional character.

CEPR Discussion Paper No. 1862

April 1998

ABSTRACT

An Applied Econometrician's View of Large Company Performance*

This paper contains a brief survey of recent empirical work on the performance of large companies. It tries to pull together the literature in the form of six stylized facts, illustrating them with data drawn from a single sample. The paper concludes by highlighting the issues which are thrown up for future work. These are: accounting for persistent heterogeneities between firms, accounting for the apparently erratic performance of many firms and, finally, moving away from hypothesis testing driven empirical agendas.

JEL Classification: L11
Keywords: corporate performance, profitability, growth, innovation

Paul A Geroski
Department of Economics
London Business School
Sussex Place
Regent's Park
London NW1 4SA
UK
Tel: (44 171) 262 5050 x 3477
Fax: (44 171) 402 0718
Email: pgeroski@lbs.ac.uk

*This paper is produced as part of a CEPR research programme on *Market Structure and Competition Policy*, supported by a grant from the Commission of the European Communities under its Human Capital and Mobility Programme (no. ERBCHRXCT940653). Many of the ideas and arguments developed in this paper first surfaced during work with some of the author's recent co-authors, particularly Chris Walters, John Van Reenen, Steve Machin, Paul Gregg and Giovanni Urga. The author is also obliged to Steve Martin, Martin Conyon, Geoff Shepherd, Jose Mata, Steve Davies and Karl Aiginger, for helpful comments on an earlier draft. The usual disclaimer applies.

Submitted 27 January 1998

NON-TECHNICAL SUMMARY

This paper presents a guide to a large and growing literature which has examined large company performance using a range of fairly sophisticated econometric techniques applied to panels of data describing the financial and, sometimes, the innovative or productivity growth performance of quoted firms. The variables of interest include company profitability, sales growth, productivity growth and innovativeness.

This literature is interesting to examine for at least two reasons. The first is methodological: most of this work is not, in the first instance, concerned with hypothesis testing. Rather, it is data driven and seeks to identify the interesting features of the data at hand. This produces work which often provides a very sophisticated description of the data but only informal tests of hypotheses thrown up by theory. One of the virtues of this methodology is that it often identifies unexpected relationships that can, and do, lead to major extensions of theory.

The second reason why this literature is interesting is substantive. Many of these studies have thrown up empirical regularities which standard theories of the firm do not predict will be on display. These include the 'facts' that:

- accounting profits display unusual patterns of variation which are statistically incongruent with other measures of performance;
- corporate growth rates are almost random;
- competitive pressures do not induce a strong measure of convergence of performance across firms; and
- most firms innovate only very irregularly.

Overall, the data display two broad features which are likely to feature heavily in future work. The first of these is the repeated observation that corporate performance is heterogenous: even firms in the same industry display idiosyncratic variations in performance, and some of the differences persist for very long periods of time. To understand these differences, it is likely that economists will have to spend more time looking at 'firm competencies'; that is, at the kinds of skills or knowledge which firms possess, and the activities which embody them.

The second broad feature of the data that should stimulate further work is the fact that corporate performance is fairly erratic. Firms seem to grow in very irregular and uneven bursts, possibly powered by the unpredictable arrival of exogenous shocks or endogenously produced innovations. Many economists like to use steady states as a way of interpreting what they observe, and this creates very static habits of thought. Understanding the erratic performance of firms may require the development of models which focus as much on the time of actions, as on their substance. It may also require modelling the process by which external shocks arrive, and how firms respond to them in more detail.

I. INTRODUCTION

Empirical work in Industrial Economics has been transformed over the recent decade or so by the increasing use of new databases. These include very micro data, often recorded at establishment level and taken from the Census, unusual one-off data sets which capture some interesting or policy relevant feature of a particular industry or market, and panels of data on individual firms drawn from a more or less wide variety of industries and observed over long periods of time. Some of these panel data sets focus on small firms, often recording the first few post-entry years of their life, while others concentrate on examining the activities of large, quoted companies. This last group of databases have been used to study corporate performance defined in a number of rather different ways, including profitability, growth, innovativeness, and so on. The goal of this essay is to take stock of what we have learned from this last branch of the new empirical Industrial Economics literature, and to speculate about where we ought to go from here.

One reason for looking at this work is that it challenges many of the presumptions about company performance that have been built up over the years by countless studies of accounting profitability. It turns out that accounting profits display unusual patterns of variation that are sometimes statistically incongruent with the variation displayed by other measures of performance. As a consequence, looking at these other measures of performances can sometimes lead to quite different inferences about what it is that makes firms perform well. A second reason why this work is interesting is that it has used statistical techniques which differ from the cross section techniques commonly used by Industrial Economists a decade or more ago.

Panel data estimation is now almost routine, and the use of times series techniques is gradually increasing. Possibly by coincidence and possibly by consequence, many scholars have also developed new ways of thinking about the empirical work that they do. There is now a fairly respectable school of thought which argues that empirical analysis ought to focus on describing the interesting properties of databases, making links with economic theory only on an opportunistic and largely informal basis. Although rather extreme, this view is a refreshing antidote to the school of thought which tries to turn every regression into a formal test of some hypothesis or other. Working through this literature is a practical way to appreciate the power of this new methodology.

One of the problems of evaluating any body of empirical work is that it is drawn from many different samples using different accounting conventions and covering different time periods in different countries. Although it is valuable to have so much work done in so many circumstances, it is sometimes difficult to know when results are simply an artefact of sample design. To help deal with this problem, I propose to synthesise this literature in the form of a series of stylised facts, and, to help make the exposition reasonably transparent, I will illustrate most of them using a single data set. It is not my intention to formally survey this literature, but I will try to provide enough references to support the claim that what I am talking about is not just an artefact of this particular database. The essay is structured around a list of six stylised facts about company performance which I believe can be extracted from the empirical literature with some confidence. This list is discussed in Section II, and the paper closes in Section III with a few observations on the way forward.

II. SOME STYLISED FACTS

The data which we are going to use as a reference point in what follows consists of a balanced panel 280 large, quoted UK firms observed over the period 1972-1982. It has been drawn from an unbalanced panel of 649 firms who were observed for at least six consecutive years during the period. Attrition bias is one of the major headaches of this kind of empirical work, and comparing calculations made using unbalanced and nested balanced panels is a very easy, informal way to assess its seriousness. For the most part, I will be focusing on seven different measures of corporate performance: accounting profits expressed as a percentage of sales, firm size (measured as the log of sales turnover), the rate of growth of sales turnover, stock market value, the rate of growth of stock market value, the number of "major innovations" produced and the number of patents produced each year by the firms in the sample. The data are presented on four tables: Table I presents a number of descriptive statistics for each performance measure, Table II shows the correlations between them on average over the period, Table III decomposes their total variation into "within" and "between" firm variation using standard analysis of variance techniques, and Table IV presents estimates of some simple autoregressions for each performance measure.

The following six stylized facts outline the most interesting features of this (and related) data set(s):

<u>*Stylised Fact # 1:*</u> ***Accounting profitability tends to understate performance differences between firms, and does not accurately reflect most of the other dimensions of firm performance considered here.***

In 1972, the lowest value of accounting profits on sales in the balanced sample was -.044, while the maximum was .2234. This is a remarkably small range of variation. The range of variation of the growth of sales was 14 times larger than this, while that of the log of firm size was 30 times larger. What is more, the well known problems of measuring accounting profits (including those associated with depreciating capital assets properly) mean that differences in profitability of perhaps as much as .05 or, perhaps, .10 may simply be measurement error. Hence, while there is a clear difference in profitability between the top and bottom of the distribution of firms, it is difficult to be confident that rankings by profitability of any two firms chosen at random away from these tails will reveal substantive (or "significant") differences in performance. Table I tells the same story in a slightly different way. With the exception of the log of sales, the coefficient of variation across firms is much larger for all the other performance measures than it is for profitability. Using this measure of variation, performance differences between firms are nearly 13.5 times more variable when firms are ranked by the growth in their market value than when they are ranked by accounting profitability; performance differences are about 7-7.5 times more variable when measured by innovative activity.

Table I also reveals that accounting profits are (roughly speaking) normally distributed (as are the growth of sales and of market value). This means that most firms are clustered in the middle of the ranking, with a few outliers on the top and the bottom. The log of sales is also normally distributed, meaning that the level of sales is positively skewed. This is also (roughly) true of the other measure of size in the data, market value. Both measures of innovation are also very skewed, and have fat tails (reflecting the innovative output of a small number of very productive firms).[1] These skewed distributions paint quite a different story of performance differences. In the case of firm size, a skewed distribution means that most firms are small, and only a very few are large. In the case of innovative output, most firms produce no innovations or patents, while only a very small number produce most of the total innovative output that we observe. The point is that normal distributions naturally lend themselves to thinking about performance differences as part of a continuum. However, when the distribution of performance differences between firms is skewed, the interesting questions are who is in the tail (which firm is large?, who innovates?) and who is not. Such distributions of performance differences incline one to think discretely, and to look for simple categorical differences between firms.

Profitability is a kind of residual which is left after costs have been deducted from revenues, and it is common to think that it is, for this reason, a natural way to sum up all the difference aspects of the performance of any particular firm. For many people, it is also a natural summary, or even sufficient, statistic of performance because they believe that firms try to maximise profits. Unfortunately, things are not quite this simple. Table II shows the correlations between all seven measures of performance, and it is evident at a glance that accounting profitability is only weakly correlated to any of the other six. In fact, most of the 21 correlations shown on the table are pretty low (the correlations between market value, size and the two measures of innovative output being the exceptions), and it is hard to believe that any one of these measures can reliably be taken as a proxy for any of the others.[2]

The reason why accounting profitability is not a very useful summary measure of corporate performance is that it is statistically incongruent with many of the other performance measures that we are looking at. There are two sources of this incongruence. The first we have already noted, namely the difference in the distribution of performance outcomes across firms: it is very hard to see how variation in a normally distributed variable will help to account for variation in a variable with a pronounced positive skew (or visa versa), unless the two are non-linearly related in a very specific way. Since this is unlikely, it seems hard to believe that the two variables will be causally related, directly or indirectly.

The second type of statistical incongruence between accounting profits and other performance measures is shown on Table III. Using analysis of variance techniques, it is possible to break the total variation in performance across firms over time into two components: "*between*" variation, which reflects differences in firms which prevail on average over a period, and "*within*" variation which reflects variations in the performance of a typical firm over time. Data which display a large amount of "between" variation identify relatively permanent differences between firms which can be explained using cross section analysis, while data which display large amounts of "within" variation suggest that variations in performance over time are the dominant feature of the data. If these differences over time across firms have a common cause, then panel techniques can be used to explain performance differences; if the within variation is idiosyncratic, then the only way forward is to analyse the times series of individual firms one by one. Table III shows that about 60% of the variation in profitability across firms and over time is between firms, a number which I think is rather low in this sample compared to other samples.[3] Apparently permanent differences between firms also seem to be the dominant feature of the data on firm size and innovation. By contrast, the two growth variables display almost no "between" variation, meaning that year by year differences between firms in their rate of growth do not persist for very long.[4] It is hard to understand how one variable which displays mainly "between" variation (like profits or firm size) can hope to explain (or even display a systematic association with) another which displays mainly "within" variation (like growth rates).[5]

Stylised Fact # 2: Corporate growth rates really are very nearly random.

Much of the empirical literature on corporate growth rates has explored the hypothesis that firm size follows a random walk (or, that growth rates are random), comparing it against the alternative that mean reversion induces a convergence in firm sizes over the long run (because larger firms grow more slowly than smaller ones). There are several ways in which this hypothesis can be explored, including regressing firm size in period t, $S_i(t)$, against size in period t-1,

$$(1) \qquad S_i(t) = \alpha_i + \beta S_i(t-1) + \mu_i(t),$$

and testing whether $\beta = 1$ (firm size follows a random walk) or $\beta < 0$ (mean reversion exists).[6] The second row of Table IV shows an estimate of (1) using the whole panel of firms and allowing for firm specific fixed effects. It is clear that $\beta < 0$, meaning that there is a tendency for smaller firms to "catch-up" with their larger rivals, a result which has been reported frequently in the literature.[7] However, the support for this conclusion in the data is actually weaker than it appears at first sight. The fourth row of the table shows an analogous autoregression in growth rates which displays a very small positive and significant estimated value of the co-efficient on the lag dependent variable (which is, nevertheless twice the size of the R^2 for that regression). Read literally, that regression suggests that growth rates between firms diverge, although, given the size of the correlation and the overall fit, it is hardly a prominent feature of the data. It is also inconsistent with the result on mean reversion shown in the second row of the table.[8] All in all, one is hard pressed to see anything other than unsystematic variation in the data on corporate growth rates.[9]

The observation that mean reversion is not an important feature of the corporate growth process shows up in our data in at least two other ways. First, if corporate growth rates are largely random, then the variance of firm sizes will rise over time. This is a clear feature of our data (the variance in log sales rose from 1.51 to 1.61 over the sample period), and, in fact, the variability of growth rates and profitability also rose throughout the period. Second, if firm size follows a random walk, then increments to size will be uncorrelated with each other across firms, while firm size itself will be highly correlated over time across firms. The cross section correlations between firm size in 1973 and size in 1982 across firms is .94, which is exactly what one expects. Differences in growth rates, on the other hand, do not persist at all (the correlation between rates of growth in 1973 and growth in 1974 across firms is -.13). Much the same applies to correlations over time in market value and in the growth in market value.

Three further observations consistent with this stylized fact are worth making. First, corporate productivity growth rates are also very hard to predict (or account for). Superior productivity growth performance does not persist over time (or even for very long), which means (of course) that differences in the level of productivity between firms persist for relatively long periods of time. Just as aggregate productivity growth is pro-cyclical, so it appears that productivity growth at the level of the firms is highly correlated the rate of growth of output (which is, in turn, noticeably more variable than the rate of growth of inputs).[10]

Second, the stylized fact that corporate growth rates are random is consistent with recent work on the nature of adjustment costs. The most common approach to modelling economic dynamics is to assume that adjustment cost are variable, rising at an increasing rate with the size of the desired change. If this is the case, then firms will have an incentive to spread out their adjustment, acting in anticipation of shocks which are expected to arrive in the future and not fully adjusting to those which have just arrived. The consequence is that data on firm size or growth will be smoothed, and hence predictable. Recent evidence suggests, however, that adjustment costs tend to be fixed and independent of the size of the desired change.[11] In these circumstances, firms will not react to shocks until their effects have accumulated beyond a certain threshold (this is sometimes called an (s,S) response strategy). If adjustment costs truly take this form, then we expect to observe irregular and largely unpredictable responses by firms to the unobserved or unpredictable shocks which they experience. This is basically what we see in the data on the evolution of firm size.

Third and finally, it has often been argued that market shares are much more stable than sales or revenue growth, and might therefore be a more sensible size based performance measure. Those who think this way usually do so because they observe that concentration ratios are fairly stable, and incorrectly infer that if the sum of the leading four or five firm's shares is stable, then this must also be true of the four or five components of that sum. Data on movements in the market shares of particular firms over time rarely sustain this presumption.[12]

<u>Stylised Fact # 3:</u> ***Heterogeneities in performance between firms persist into the long run more or less regardless of how performance is measured. However, performance differences between firms are not constant over time, and many of them widen in recessions.***

The question of how permanent differences in performance between firms are has always been of interest to economists who believe that market forces typically erode excess profits, or lead to a rapid diffusion of new ideas or practices. A large empirical literature has built up testing the proposition that profit differences between firms persist even in the long run using a model very like (1), namely

(2) $\pi_i(t) = \alpha + \beta\pi_i(t) + v_i(t),$

where π is typically one of a number of measures of accounting profitability. The conventional interpretation of the output of these regressions is that the smaller is the estimated value of β, the stronger are the forces which induce convergence; if $\beta = 0$, then "profit differences between firms do not persist".[13] The value of β shown in the first row of Table IV is not untypical of those reported in the literature: it looks like "profit differences between firms persist".

This is, however, just a little bit too simple. Notice that the actual estimated value of β is actually much smaller than it looks. As the final two columns of the table show, the half-life of a shock to profitability is only a year when $\beta = .4894$, and 90% of the shock will be gone in 3.2 years. If this is what the statement "profit differences persist" means, then persistence is hardly likely to be a source of much interest. Nevertheless, profit differences between do firms exist, and, as it happens, they appear to be more or less permanent. Table III shows that most of the variation in profits is between firms, and it turns out that the correlation between firms ranked by profitability is very stable over time (e.g. the cross section correlation in profits across firms in our data between 1973 and 1982 is .44). Both of these observations imply relatively stable, permanent differences in profitability between firms. This conclusion can be also extracted from Table IV if one focuses on estimated values of α rather than β as a way of thinking about persistence. Table IV is a fixed effects estimation, and it turns out that it is impossible to simplify the fixed effects in the profits equation to a single constant, meaning that permanent differences in the profitability of firms exist. The fact that $\alpha_i \neq \alpha$ for all firms i means that "profit differences persist forever", more or less regardless of the value of β.

In fact, this is the conclusion to be drawn from all of the regressions displayed on Table IV. In all cases, the fixed effects in these regression cannot be simplified to a single constant (the decision is marginal in the case of the two growth regressions), meaning that differences between firms are not transitory. For the most part, estimated values of β do not generate very interesting dynamics in any of the performance measures we are examining here (all of the estimated half lives are pretty short). Shocks to sales have an effect which dies away much more slowly than for the other measures of performance (which is what would expect from a series which is practically a random walk), while any systematic dynamics associated with convergence in growth rates are absolutely dwarfed by unsystematic noise. The basic message in this data is that firms are very different, and, although we have not shown it here, this is as true for comparisons between firms in the same industry as it is for firms drawn at random from throughout the economy. Heterogeneity in the short and long run is possibly the most pronounced feature of the data on large company performance.

While differences between firms persist in the long run, they are not constant over time. As we noted earlier the variance of most of the performance measures shown on Table I has risen steadily through the period, meaning that differences seem to have gradually widened over time (at least within our sample period). This increase in variability has, however, occurred rather unevenly over time, and, in fact, much of the increase in variability seems to occur during recessions; that is, there are grounds for believing that performance differences vary counter cyclically.[14] It is hard to know exactly why this occurs, but it is probably because recessions have very selective effects on firm performance: only a relatively small number of firms are very severely affected by recessionary pressures, and some actually prosper.[15]

<u>Stylised Fact # 4:</u> ***Firm size and previous innovative activity are among the characteristics of firms which seem to be associated with systematic differences firm performance.***

It is one thing to observe almost permanent differences in performance between firms, and another thing altogether to account for those differences. Many of the larger correlations shown on Table II seem to involve firm size, whether that is measured by the log of sales or by the market value firms.[16] This is a very commonly made observation, and the inclusion of a firm size variable in many regressions is pretty much routine (particularly in the literature which examines Schumpeterian hypotheses about the innovativeness of large firms). However, firm size also has a second and rather more subtle effect on performance. Many of the estimated parameters in the regressions shown on Table IV seem to vary with firm size, particularly those associated with growth (large firms often have smaller absolute values of β, and display a lower variance in growth over time). What this means is that firm size may have two quite different effects on performance: a ***direct effect*** which shows how much a ceteris paribus increase in size affects performance, and an ***indirect effect*** which arises because the model which describes how the performance of large firms depends various exogenous variables differs from that which describes how small firm performance depends on the same exogenous variables. That is, firm size (and other things) affects performance directly (which is what the usual regression co-efficients measure), and it also affects performance indirectly because it conditions the size of the effects that other things have on performance (i.e. all of the co-efficients in equations vary by size of firm).

Size might be an important source of heterogeneities for any number of reasons. Large firms are often a portfolio of activities whose overall performance is likely to be smoother than specialised firms, and if placing a wider range of bets increases the likelihood of making a big draw, they may also perform better on average. This will be particularly the case if large firms benefit from any kinds of scale economies, pecuniary or otherwise, or learning effects. Many people argue that large firms will be filled with more capable people, may train their employees more thoroughly or may be able to afford the kind of specialist advice which can sometimes make a big difference to performance. Although size is bound to be a reflection of capability in some long run sense, enough large firms under perform relative to smaller rivals to make one just a little hesitant to accept such assertions at face value. It would probably make more sense to develop more direct indicators of "organizational capability" and use them to examine performance differences between firms.[17] One indicator of capability which seems to display systematic associations with performance is innovative activity. It turns out that both profits equations and growth equations tend to differ between innovative and non-innovative firms. Further, firms that do a lot of R&D and produce many innovations are likely to be more flexible or adaptable, and more able to benefit from spillovers, than their rivals.[18]

One final observation about heterogeneities is worth making. Although the data that we are discussing are not organised to show this, it turns out to be the case that the usual Census based definitions of industries do not always seem to be associated with significant differences in performance between firms. In fact, the intra-industry dispersion of performance differences between firms is often almost as large as that across industries. Further, persistence of profits regression or studies of mean reversion in growth rates rarely report any clear tendency for performance to converge within industries, and analyses of the importance of industry factors in generating the overall variability in profitability which we observe typically suggest that industry effects are small.[19] It is possible that these results arise because we are using the wrong definitions of markets in our empirical work, but it is more likely that the diversified activities of most of the firms who are included in these data sets have freed them from the vagaries of competition in particular markets.

Stylised Fact # 5: Most firms are, at best, irregular and erratic innovators.

Many people like to assess corporate performance using some measure of "innovativeness". For the most part, measures of R&D intensity are used for this purpose. R&D activity is not only heavily skewed within industries, but it is typically very stable over time (i.e. "between" firm variations in R&D intensity dominate "within" firm variations over time). Aside from convenience, this reliance on R&D measures has little to recommend it. When we talk about "innovativeness" as being desirable in a firm, we typically mean that the firm produces lots of useful product and process innovations, something can happen whether or not a firm allocates a large percentage of its revenues to a cost category labelled " R&D expenditures". In the data we have been using here, there are two measures of innovative "output", and they paint a picture of innovativeness which is slightly different from that which emerges from studies using R&D expenditures.[20]

Table IV and Figure IV tell much the same story: firms that patent tend to do so reasonably regularly, while those who launch a major innovation do so much less regularly. One way to think of these correlations is in terms of "history dependence"; that is, by asking whether the fact that a firm innovates today increases the chance that it will do so again tomorrow. This might occur if there were any kind of learning or dynamic increasing returns to scale associated with innovative activity.

Alternately, innovative activity might be driven by exogenous variations in "technological opportunity" which reward lucky firms with a seemly endless stream of innovation outputs, while others are forced to mine less fertile areas of technological knowledge. The data suggest that patents are more history dependent than major innovations, meaning that knowledge of which firms patented in period t provides useful information for predicting which firms will patent again in t+1, but that the same information is less useful for predicting who will produce a major innovation in the future.[21] This observation probably will not surprise those who believe that patenting is basically an input measure linked to R&D spending, or that firms with large R&D labs create strong incentives for active scientists to patent regularly.

However, Table IV and Figure IV do somewhat exaggerate the predictability of innovative activity, since the correlations on both tables are both driven by the most salient feature of the data on innovative activity, namely that some firms innovate (from time to time) while others never innovate. One gets quite a different view when one tracks the innovative activity of particular firms over time, looking at the length of spells in which they regularly patent and innovate. Most of the firms who produce an innovation or a patent do so every once in a while, and very few do so in successive years; i.e. few firms record innovation spells lasting longer than a year (a few outliers produce multiple outputs year in and year out). The data show almost no sign of learning or increasing returns to innovative activity over time, except when very high thresholds of activity have been reached.[22]

Needless to say, if innovative activity is as erratic as the data suggests, then it is probably not to hard to understand why corporate growth rates are so unpredictable. Firms evidentially grow in bursts, many of which may be associated with innovations, major and minor. It is also not hard to understand why a firm which innovates and grows erratically is likely to see its market value rise in an unpredictable fashion over time. What is more of a mystery is why one does not see any real reflection of this erratic activity in accounting measures of profitability.

__Stylised Fact #6__: ***Attrition bias does not seem to be a major problem, and inferences made about firm performance using data on survivors is often robust to the inclusion of data on non-survivors.***

Virtually all of the calculations discussed in this paper have been done for two samples: a balanced sample of 280 firms observed throughout the whole period, and an unbalanced sample of 649 firms observed for at least a substantial part of the period. An inspection of the tables reveals that very few of our inferences depend on which sample we use. It is clear that measures of variability are understated when applied to survivors alone, but mean performance levels of profitability and other factors are not all that much higher in survivor only samples.[23] This is, of course, only a very informal way to test for the presence of attrition bias, but the fact is that most of the formal tests reported in the literature come to very much the same conclusion. The regressions shown on Table IV, for example, are remarkably insensitive to the omission or inclusion of non-survivors.

What this means is one of two things: either the random noise in measures of performance like growth or innovativeness is so large that it that it makes the effects of attribution bias very hard to detect, or that the criteria for survival is orthogonal to many of the performance measures which typically interest us. This observation raises the interesting question of just what it is which determines the survivability of large firms. In principle, firms face two types of selection pressures: normal product market selection and selection in the capital market (which allocates funds to successful firms, and sanctions the disappearance of others through mergers and take-overs). In practice, capital market criteria are probably more important for the population of firms that we are looking at, and they seem to be based much more on size than profitability or growth or innovativeness.[24] This means that measured correlations between size and other dimensions of company performance may be more sensitive to sample composition than correlations involving other variables of interest.

III. SO, WHERE DO WE GO FROM HERE?

This brief tour of the econometric literature on firm performance is just a snapshot of a body of work which will almost certainly grow and develop. As it does, my guess is that active scholars will end up thinking about three issues: heterogeneities between firms, the erratic performance of firms and the question of empirical methodology. I will close with a few observations on each of these issues.

Heterogeneities between firms

Heterogeneities between firms are easily the most salient feature of the data that we have been discussing, and it seems clear that if we are going to understand the determinants of company performance, then we are going to have to track down the sources of this heterogeneity. Most of the theoretical work currently being done by economists on the theory of the firm is about contracts or the effect that certain types of transactions costs have on firm structure. Interesting as this literature is, it is hard to argue that it is anything other than completely orthogonal to the kind of empirical work that we have been looking at here. There is, however, a large and growing management literature on "firm competences" which may help structure the next step forward.[25] This literature purports to identify the deep seated sources of differences between firms by looking at certain types of activities which firms undertake. Stimulating as it is, the competences literature has yet to progress much beyond tautology. It is much easier to invent competences which might explain what we observe than it is to independently try to measure competences and then try to correlate them with observed performance differences between firms. Taking this work forward empirically seems like a worthy challenge.

We are also going to have to think carefully about how we model heterogeneities between firms. As we noted earlier, many exogeneous variable of interest have both ***direct*** and ***indirect*** effects on performance. Direct affects are those which measure the relationship between an exogeneous variable x and performance, π, given the value of other exogeneous variables z. Indirect effects arise whenever the effect that z has on π depends on x, meaning that changes in x change the partial relation between z and π. We usually measure direct effects by including the relevant exogeneous variable x into a linear regression. Indirect effects are more complicated, since they imply that the slope of that regression varies with x. In the limit, this means that each firm that has a distinct value of x should have it's own regression, or, more practically, that groups of firms with similar values of x should not be mixed with other groups whose value of x differs markedly in a regression. We have been used to developing empirical models which can be applied to large numbers of firms, but the message from the literature that we have reviewed here may be that we ought to start thinking more about models which apply only to specific groups of firms, and how they might differ from models which apply to other groups of firms.

This observation also tells us something about the kinds of data that we will need to generate if we are going to make any progress in understanding heterogeneities between firms. It is clear that cross section work is completely unsuitable for exploring this kind of problem. Panel data enables one to correct for fixed or permanent differences between firms very easily (using fixed effects), but panel models usually force slope co-efficients to be the same across firms. Further, measurement error and the well known bias which arises in short times series sometimes makes panel data estimates of slope parameters unreliable and hard to interpret (particularly those associated with lagged dependent variables). To deal with hetereogeneity in both direct and indirect effects, then, we are likely to need to turn to longitudinal data. This means using our regressions to help understand the history of particular firms, and developing empirical models which enable us to compare the histories of different firms in terms of a set of estimable parameters.

The erratic performance of firms

Most economic theory is pretty static in character, and focuses on long run equilibria or steady states. This tradition puts us into the habit of looking for almost permanent differences in performance between firms, differences which we hope can be associated with variations in the basic features of costs and demand. Whenever we look at firm performance through the lens provided by accounting profits data, we see these presumptions reflected back to us: profitability differences between firms seem to be almost permanent, and the only interesting features of the data are these cross section differences. The short run dynamics of profitability seem muted and not very interesting.

The interesting question is whether all of this just says that we see only what we are looking for. If, by chance, we had grown up thinking that the only interesting measure of firm performance was growth, we would never think in this way. Looking at the data on growth rates would make us focus on the kind of short run, unanticipated internal and external shocks which seem to be the dominant driver of growth rates. Since the interesting feature of the data on firm growth is the idiosyncratic variation over time in the performance of particular firms, we would never think too seriously about doing cross section work on performance differences between firms, or look for the determinants of steady state or optimum firm sizes. Instead, we would undertake longitudinal econometric or qualitative case studies of particular firms, and try to develop for theories of growth which generate history dependent paths of evolution.

What is interesting about these speculations is that they are consistent with the intuition that many of us have that it is rarely possible to talk about a particular firm in a manner independent of time, or to identify clearly and unambiguously what it is good at and what it thinks that it wants to do. Firms change their management teams fairly regularly, many of them actively buy and sell subsidiaries at a strikingly high rate (particularly those in the US and UK), some innovate or start new businesses from time to time, almost all of them change their strategies at both corporate and at particular business unit levels frequently, and, periodically, most also try to change their culture or corporate style. Further, these changes are fuelled by a steady undercurrent of struggle between stakeholders inside and outside the firm. Firms are political organizations ruled by coalitions, and this makes them slow and awkward in responding to external shocks, and sometimes unstable. It is only when an unusually charismatic and capable individual (or well integrated management team) dominates an organization for a relatively long period of time that one can see a persistently clear corporate identity. All of this is to say that the inner lives of many firms is often rather turbulent, and much more so than appears on the balance sheet. Maybe this is something which we should think more systematically about.

Methodology

The work that we have discussed in the essay is identifiable by both its subject matter and its style. The literature is not, on the whole, hypothesis testing driven. Rather, most of this work has involved the application of relatively sophisticated econometric techniques to interesting data sets with the idea of uncovering the interesting properties of the data. Those scholars who do this kind of work and are well versed in economic theory are often surprised by features of the data which, in a more formally staged test, would form the basis of a persuasive rejection some widely held hypothesis or (less grandly) commonly shared presumption. However, for the most part the process of uncovering surprises in data sets is less much formal than hypothesis testing, and this is reflected in the careful and often very cautious conclusions which are drawn from this kind of work.

On the whole, I am inclined to think that this kind of methodology is a significant improvement on previous practice. Empirical work which is driven by a hypothesis testing agenda seems to me to suffer from (at least) two problems, and both of them stem from weaknesses in the kind of theory which economists generate. First, economic theory is rarely expressed in a testable form. Most of the models which theorists produce are very simple, and empirical tests of the predictions generated by these models are also bedevilled by the need to control for the many, many factors which are absent in the model. Since it is often extremely difficult to control for outside factors in a persuasive way, most of the tests which are reported in the literature are conditional on quite a number of untested assumptions. They are, therefore, much weaker than they appear and often lead to more arguments than conclusions.

Second, theorists often talk about things like "*technological opportunity*", "*knowledge stock*", "*market power*" and "*efficient levels of ...*", "*elasticities of ...*" or "*expected values of ...*" virtually everything, as if these things were directly observable.[26] They are not of course, and this means that most tests start with the problem of trying to observe things which are the subject of the test (or at least central to it). One approach to this problem asserts that there exist one or more proxies for every observable that we might be interested in. The proponents of this view are typically driven by a hypothesis testing agenda, and they often very creative in the ways that they construct proxies and interpret their data. However, they sometimes succumb to wishful thinking, and, in any case, often find it difficult to dispute alternative interpretations of their data. The alternative approach asserts that observables should be interpreted literally (for what they are) and not creatively (for what they might be), and that unobservables are, unfortunately, just that. This alternative philosophy dovetails neatly with data driven empirical agendas (such as those in much of the work discussed above). The work that results is often little more than a sophisticated statistical descriptions of what we observe, but at least it is possible to test the accuracy of such descriptions.

The real power of data driven methodologies is that they have the potential to drive theory in new and productive directions. The conventional division of labour between theorists and empiricists is that the latter are supposed to follow the former, identifying the wrong turnings which have been made through carefully carried out hypothesis tests. The great problem with this is that it leads to theorizing in a fact free vacuum, particularly when (as is the case in economics) the ability of empiricists to mount persuasive tests of theoretical proposition is weak. Data driven empirical methodologies have the potential of taking more of a lead. Such work identifies interesting facts and makes fairly precise characterizations of puzzling phenomena, and these observations ought to provide an agenda for those theorists who are interested in understanding the world as we see it. I certainly hope that some of them will find the six stylized facts discussed here stimulating.

NOTES

[1] Cohen and Klepper, 1992, report a similar skew in the R&D intensity of 2491 line of business units within three digit US industries over the period 1974-77. The number of these firms that did no R&D varied between 5% and 50% across industries (see also Bound et al, 1984). I believe that advertising expenditures have a similar pattern within industries, being typically much more concentrated than sales. This impression is derived from various case studies (including some of those reported by Sutton, 1991) plus the data used in Geroski and Toker, 1996.

[2] These weak correlations are also a feature of the large literature which regresses measures of profitability on a number of these other variables and on variables like market share; see Machin and Van Reenen, 1993, Haskel and Scaramozino, 1997, and others for recent work and Schmalensee, 1989, for a survey of earlier work of this type. Dickerson et al, 1997, examines the effects of changes in market structure induced by mergers on profitability (and contains references to the earlier literature), while Lichtenberg, 1992, looks at the effects on productivity. Similarly, R&D intensity and measures of innovative output tend to have weak but positive correlations to profitability (see Geroski et al, 1993, amongst others). A much earlier literature which looked at the relationship between firm size and accounting profitability has slipped from sight, mainly because of concerns that the correlations might reflect differences in accounting practices between large and small firms. For work, linking R&D to other measures of firm performance, see Pakes, 1985, Lach and Shankerman, 1989, and others.

[3] See, for example, Geroski and Machin, 1993, whose (unbalanced) panel of 539 UK firms observed over the period 1972-83 displayed 93% between variation.

[4] For example, Geroski, Urga and Walters, 1997, tracked the growth in net assets of 77 firms over the period 1955-1985, and observed that differences in growth between any two firms chosen at random persisted for about two years on average. No spells of superior relative growth performance lasted longer than 7 years, and the sign of the difference between the growth rates of any two firms chosen at random changed 7 times on average over the period (most pairs of firms had well over 10 sign changes).

[5] The data shown on Table II are hard to reconcile with the hypothesis that managers trade profits off for growth (for surveys of earlier work testing this hypothesis, see Mueller, 1987, or Hay and Morris, 1991, pp. 356-67). However, managers whose compensation is more closely tied to size than profits (which seems typical: see Gregg et al, 1993, Conyon et al, 1995, Main et al, 1996, and references cited therein), may expand their firm too much and generate the kind of negative correlation between size and profitability displayed on Table II.

[6] A somewhat more common version of this test involves running the regression $G_i(t) = \alpha_i + \gamma S_i(t-1) + \varepsilon_i(t)$, where $G_i(t)$ is the rate of growth of firm i in year t, and testing whether $\gamma = 0$ (firm size follows a random walk) or $\gamma < 0$ (mean reversion).

[7] For recent work, see Dunne and Hughes, 1994, Hart and Oulton, 1996, and references cited therein. These simple convergence regressions are not altogether satisfactory, and it is not clear how much one can learn about long run tendencies (i.e. convergence) from cross section or (short) panel data. For work using long times series data and a times series based test of convergence, see Geroski, Urga and Walters, 1997, who find no evidence supporting the convergence hypothesis.

[8] Market value is sometimes thought of as a measure of firm size, but it does not display the same times series properties as the log of sales in our sample. The third row of Table IV suggests a process in which firms with large market value get even larger relative to their smaller rivals over time. This may be an artefact of the small sample, or of the particularly simple times series model used on Table IV. The growth of market value equation is not consistent with the market value equation (it suggests a process by which growth rates are equalised, meaning that persistent but stable differences in market valuation between firms will persist over time).

[9] There is a noticeable difference between the literature on small and/or new firms (where support for convergence is much stronger - see Evans, 1987, Mata et al, 1995, Dunne et al, 1989, and many others) and that on large firms. This makes sense, since it is much more plausible to imagine that the growth of small or new firms is driven by the kind of transitional dynamics embodied in equation (1) than it is for large, well established firms.

[10] See Geroski, Small and Walters, 1997, Nickell et al, 1992, Nickell, 1996, Nickell et al, 1997, Griliches and Mairesse, 1983, and others.

[11] A classic account of adjustment costs which focuses on managerial limits to growth and suggests proportional adjustment costs is Penrose, 1959. Recent work based on Census data includes Caballero et al, 1995 and 1997; Geroski and Gregg, 1977, and Davis et al, 1996, also observe that many responses to recessions are "lumpy" and not obviously consistent with a proportional costs model of adjustment. See also Hamermesh and Pfann, 1994, for a survey of a wide range of work on adjustment costs.

[12] Davies and Geroski, 1977, observed the market shares of the leading five firms in UK three digit industries in 1979 and 1986, and found that "within" variations in shares accounted for 86% of total variation. Market concentration, on the other hand, was relatively stable; see also Baldwin, 1995, Cable, 1997, and the survey by Caves, 1997.

[13] See Geroski and Jacquemin, 1987, Mueller, 1986 and 1990, and many others.

[14] This observation is not inconsistent with Schumpeter's vision of "gales of creative destruction" which he thought might be unleashed during recessions, and also with more recent pit-stop theories of recession which argue that the opportunity cost of change will be lower during downturns (e.g. see Hall, 1991, Caballero and Hammour, 1994, Cooper and Haltiwanger, 1993, and others). The evidence seems to suggest that most of the restructuring which occurs in recessions in involves shedding jobs, closing plants and redeploying workers; mergers (which tend to be pro-cyclical) are as important a driver of major corporate restructuring as recessions are (see Geroski and Gregg, 1997).

[15] Geroski and Gregg, 1997, report that the quintile of firms shedding the most jobs account for 94% of job losses in the 1990 recession, while the 10% of firms who experienced the biggest fall in profits lost more than the largest 2100 UK firms did as a group (the top 40% of firms actually increased their profitability in the recession). See also Davis et al, 1996, for work on the distribution of job creation and job destruction using US establishment data taken from the Census which reaches similar conclusions.

[16] The correlation between the production of major innovations and size shown on the Table is also reported in Blundell et al, 1995, but it is probably due to sample selection The innovations data set which is the source of the major innovations data includes not just large quoted firms, but also smaller firms who, in fact, account for a disproportionately large share of the total quantity of innovations produced; see Pavitt et al, 1987.

[17] For work in this area, see Henderson and Cockburn, 1994, Levinthal and Myatt, 1994, Roberts and Dowling, 1997, Roberts, 1997, Helfat, 1997, and others.

[18] Such firms are often believed to have superior "absorptive capabilities"; see Cohen and Levinthal, 1989 and 1990, Malerba, 1992, and others. For work on direct and indirect effects on performance associated with innovative activity, see Geroski et al, 1993, and Geroski and Machin, 1993.

[19] Schmalensee, 1985, reported very large effects, but subsequent work suggests that industry factors are relatively unimportant; see Rumelt, 1991, and others. Cubbin and Geroski, 1987, looks at the same question using a dynamic model of profitability drawn from the persistence of profitability literature, and also find weak industry effects.

[20] For work on R&D, see Bound et al, 1984, Harhoff, 1994, Hall and Mairesse, 1995, and others. It is possible to argue that patents are more of a measure of research inputs than outputs, largely because they record the acquisition of technical knowledge which is a necessary but hardly a sufficient input into the production of a successful product or process innovation. For survey discussions of these issues, see Griliches, 1990, Pavitt and Patel, 1995, and others.

[21] In fact, not only can past patents and major innovations be (somewhat) helpful in predicting the future values of either variable, but patents seem to lead major innovations (that is, patents seem to "Granger cause" innovations); for evidence at the level of the firm, see Geroski, Van Reenan and Walters, 1997a, and for evidence at aggregate level, see Geroski and Walters, 1995.

[22] For work on the persistence of innovative activity, see Geroski, Van Reenen and Walters, 1997b, Malerba and Orsenigo, 1995, Malerba et al, 1996, Cefis, 1996, and others.

[23] Geroski and Gregg, 1997, compared the performance of a sample of firms that survived the 1990-1992 UK recession with a (smaller) sample of firms who failed during the recession, and observed a fairly robust continuity in characteristics: firms that failed has lower values of the characteristics of firms (e.g. profitability, liquidity, and so on) who were badly affected by the recession, and they, in turn, had lower values of these characteristics than those firms who were not seriously affected by the recession. Hence, an examination of the question: "which firms are badly affected by the recession?" produced much the same answer regardless of whether failed firms were included in the sample being analysed or not.

[24] See Singh, 1971, Mueller, 1987 or Scherer and Ross, 1990, pp. 164-174.

[25] See Peteraf, 1993, and Teece et al, 1997, for overviews; for popular discussions, see Hamel and Prahalad, 1994, and Collis and Montgomery, 1995. Some of the empirical work stimulated by this literature was cited in footnote 17 above.

[26] It is not, for example, entirely obvious what "company performance" is and how it ought to be measured (see Nickell, 1995, for some observations on this), and it is certainly the case no one believes that the seven measures which we have examined here are equalivalent or perfectly substitutable (indeed, many will not regard firm size as a performance measure at all).

REFERENCES

Baldwin, J. (1995) The Dynamics of Industrial Competition, Cambridge University Press, Cambridge.

Blundell, R., R. Griffith and J. Van Reenen (1995) "Dynamic Count Data Models of Technological Innovation", Economic Journal, 105, 333-344.

Bound, J. et al (1984) "Who Does R&D and Who Patents?" in Griliches, Z. (ed), R&D, Patents and Productivity, University of Chicago Press, Chicago.

Caballero, R., E. Engel and J. Haltiwanger (1997) "Aggregate Employment Dynamics: Building from Microeconomic Evidence", American Economic Review, 87, 115-137.

Caballero, R., E. Engel and J. Haltiwanger (1995) "Plant Level Adjustment and Aggregate Investment Dynamics", Brookings Papers on Economic Activity, 2, 1-54.

Caballero, R. and M. Hammour (1994) "The Cleansing Effects of Recession", American Economic Review, 84, 1350-68.

Cable, J. (1997) "Market Share Behavior and Mobility", Review of Economics and Statistics, 79, 136-141.

Caves, R. (1997) "Industrial Organization and New Findings on the Turnover and Mobility of Firms", forthcoming, Journal of Economic Literature.

Cefis, E. (1996) "Is There Any Persistence in Innovative Activities?", mimeo, European University Institute.

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

Cohen, W. and D. Levinthal (1990) "Absorptive Capacity: A New Perspective on Learning and Innovation", Administrative Science Quarterly, 35, 128-152.

Cohen, W. and S. Klepper (1992) "The Anatomy of Industry R&D Distributions", American Economic Review, 82, 773-799.

Collis, D. and C. Montgomery (1995) "Competing on Resources: Strategy in the 1990s", Harvard Business Review, July-August, 118-128.

Conyon, M., P. Gregg and S. Machin (1995) "Taking Care of Business: Executive Compensation in the UK", Economic Journal, 105, 704-714.

Cooper, R. and J. Haltiwanger (1993) "The Aggregate Implications of Machine Replacement: Theory and Evidence", American Economic Review, 83, 360-382.

Cubbin, J. and P. Geroski (1987) "The Convergence of Profits in the Long Run: Inter-firm and Inter-industry Comparisons", Journal of Industrial Economics, 35, 427-442.

Davies, S. and P. Geroski (1997) "Changes in Concentration, Turbulence and the Dynamics of Market Shares", forthcoming, Review of Economics and Statistics.

Davis, S. J. Haltiwanger and S. Schuh (1996) Job Creation and Job Destruction, MIT Press, Cambridge Mass.

Dickerson, A., H. Gibson and E. Tsakalotos (1997) "The Impact of Acquisitions on Company Performance: Evidence From a Large Panel of UK Firms", Oxford Economic Papers, 49, 344-361.

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

Dunne, T., M. Roberts and L. Samuelson (1989) "The Growth and Failure of US Manufacturing Plants", Quarterly Journal of Economics, 104, 671-688.

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

Geroski, P., S. Machin and J. Van Reenen (1993) "The Profitability of Innovating Firms", Rand Journal of Economics, 24, 198-211.

Geroski, P. and S. Machin (1993) "Innovation, profitability and Growth Over the Business Cycle", Empirica, 20, 35-50.

Geroski, P. and C. Walters (1995) "Innovative Activity over the Business Cycle", Economic Journal, 105, 916-928.

Geroski, P. and A. Jacquemin (1988) "The Persistence of Profits: A European Comparison", Economic Journal, 98, 375-390.

Geroski, P. and P. Gregg (1997) Coping with Recession, Cambridge University Press, Cambridge.

Geroski, P. and S. Toker (1996) "The Turnover of Market Leaders in UK Manufacturing Industry, 1979-1986", International Journal of Industrial Organisation, 14, 141-158.

Geroski, P., G. Urga and C. Walters (1997) "Are Size Differences Between Firms Permanent or Transitory?", mimeo, London Business School.

Geroski, P., J. Van Reenen and C. Walters (1997a) "Innovations, Patents and Cash Flow", mimeo, London Business School.

Geroski, P., J. Van Reenen and C. Walters (1997b) "How Persistently Do Firms Innovate?", Research Policy, 26, 33-48.

Geroski, P., S. Machin and C. Walters (1997) "Corporate Growth and Profitability", Journal of Industrial Economics, 45, 171-190.

Geroski, P. I. Small and C. Walters (1997) "External Economies, Technology Spillovers and Corporate Productivity Growth", mimeo, London Business School.

Gregg, P., S. Machin and S. Szymanski (1993) "The Disappearing Relationship Betwen Directors Pay and Corporate Performance", British Journal of Industrial Relations, 31, 1-10.

Griliches, Z. (1990) "Patent Statistics as Economic Indicators", Journal of Economic Literature, 28, 1661-1707.

Griliches, Z. and J. Mairesse (1983) "Comparing Productivity Growth: An Exploration of French and US Industrial and Firm Data", European Economic Review, 21, 89-119.

Hamel, G. and C.K. Prahalad (1994) Competing for the Future, Harvard Business School Press, Boston, MA.

Hall, B. and J. Mairesse (1995) "Exploring the Relationship between R&D and Productivity in French manufacturing Firms", Journal of Econometrics, 65, 263-293.

Hall, R. (1991) Booms and Recessions in a Noisy Economy, Yale University Press, New Haven.

Hamermesh, D. and G. Pfann (1994) "Adjustment Costs in Factor Demand", Journal of Economic Literature, 34, 1264-1292.

Harhoff, D. (1994) "R&D and Productivity in German Manufacturing Firms", mimeo, University of Mannheim.

Hart, P. and N. Oulton (1996) "Growth and Size of Firms", Economic Journal, 106, 1242-1252.

Haskel, J. and P. Scaramozzino (1997) "Do Other Firms Matter in Oligopolies?", Journal of Industrial Economics, 45, 27-46.

Hay, D. and D. Morris (1991) Industrial Economics and Organization, Oxford University Press, Oxford.

Helfat, C. (1997) "Know-How and Asset Complementarity and Dynamic Capability Accumulation: The Case of R&D", Strategic Management Journal, 18, 339-360.

Henderson, R. and I. Cockburn (1994) "Measuring Competence? Exploring Firm Effects in Pharmaceutical Research", Strategic Management Journal, 15, 63-84.

Lach, S. and M. Shankerman (1989) "Dynamics of R&D and Investment in the Scientific Sector", Journal of Political Economy, 97, 880-904.

Levinthal, D. and J. Myatt (1994) "Co-evolution of Capabilities and Industry: The Evolution of Mutual Fund Processing", Strategic Management Journal, 15, 45-62.

Lichtenberg, F. (1992) Corporate Takeovers and Productivity, MIT Press, Cambridge Ma.

Machin, S. and J. Van Reenen (1993) "Profit Margins and The Business Cycle", Journal of Industrial Economics, 41, 29-50.

Main, B., A. Bruce and T. Buck (1996) "Total Board Renumeration and Company Performance", Economic Journal, 106, 1627-1644.

Malerba, F., L. Orsenigo and P. Peretto (1996) "Persistence of Innovative Activities, Sectoral Patterns of Innovation, and International Technological Specialization", forthcoming, International Journal of Industrial Organization.

Malerba, F. and L. Orsenigo (1995) "Schumpeterian Patterns of Innovation", Cambridge Journal of Economics, 19, 47-65.

Malerba, F. (1992) "Learning by Firms and Incremental Technical Change", Economic Journal, 102, 845-859.

Mata, J., P. Portugal, and P. Guimaraes (1995) "The Survival of New Plants", International Journal of Industrial Organization, 13, 459-481.

Mueller, D. (1986) Profits in the Long Run, Cambridge University Press, Cambridge.

Mueller, D. (1987) The Corporation: Growth, Diversification and Mergers, Harwood, London.

Mueller. D. (ed) (1990) The Dynamics of Company Profits, Cambridge University Press, Cambridge.

Nickell, S., S. Wadhwani and M. Wall (1992) "Productivity Growth in UK Companies: 1975-1986", European Economic Review, 36, 1055-1091.

Nickell S., (1995) The Performance of Companies, Basil Blackwell, Oxford.

Nickell, S. (1996) "Competition and Corporate Performance", Journal of Political Economy, 104, 724-746.

Nickell, S., D. Nicolitsas and N. Dryden (1997) "What Makes Firms Perform Well?", European Economic Review, 41, 783-796.

Pakes, A. (1985) "On Patents, R&D and the Stock Market Rate of Return", Journal of Political Economy, 93, 390-409.

Patel, P. and K. Pavitt (1995) "Patterns of Technological Activity: Their Measurement and Interpretation", in P. Stoneman (ed) Handbook of the Economics of Innovation and Technological Change, Basil Blackwell, Oxford.

Pavitt, K., M. Robson and J. Townsend (1987) "The Size Distribution of Innovating Firms", Journal of Industrial Economics, 35, 297-316.

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

Peteraf, M. (1993) "The Cornerstones of Competitive Advantage: A Resource Based View", Strategic Management Journal, 14, 179-191.

Roberts, P. (1997) "Innovation and Firm Level Persistent Profitability", mimeo, Australian Graduate School of Management.

Rumelt, R. (1991) "How Much Does Industry Matter?", Strategic Management Journal, 12, 167-185.

Scherer, M. and T. Ross (1990) Industrial Market Structure and Economic Performance, 3rd rd, Houghton Mufflin, Boston.

Schmalensee, R. (1989) "Inter-industry Studies of Structure and Performance", in Schmalensee, R. and R. Willig (eds) Handbook of Industrial Economics, North Holland, Amsterdam.

Schmalensee, R. (1985) "Do Markets Differ Much?", American Economic Review, 75, 341-351.

Singh, A. (1971) Take-overs: Their Relevence to the Stock Market and the Theory of the Firm, Cambridge University Press, Cambridge.

Sutton, J. (1991) Sunk Costs and Market Structure, MIT Press, Cambridge Mass.

Sutton, J. (1997) "Gibrat's Legacy", Journal of Economic Literature, 35, 40-59.

Teece, D., G. Pisano and A. Shven (1997) "Dynamic Capabilities and Strategic Management", Strategic Management Journal, 18, 509-533.

TABLE I: Simple Statistical Descriptions of Company Performance*

	mean	standard deviation	co-efficient of variation	skewness	kurtosis
accounting profits/sales	.0455 (.0402)	.0357 (.0396)	.7846 (.9851)	.0727 (-.3163)	6.465 (9.215)
log of sales	3.989 (3.230)	1.596 (1.628)	.4001 (.504)	.3527 (.5647)	2.721 (3.092)
rate of growth of sales	.1342 (.1233)	.1680 (.1808)	1.252 (1.466)	-.9500 (-.3831)	25.47 (18.21)
market value	62.64 (36.06)	219.7 (160.9)	3.507 (4.462)	12.87 (17.15)	245.9 (443.7)
rate of growth of market value	.0503 (.0849)	.5308 (.5116)	10.55 (6.026)	.0611 (.0241)	3.690 (3.963)
number of "major" innovations produced	.1260 (.072)	.7061 (.5209)	5.604 (7.235)	11.73 (15.41)	192.6 (338.0)
number of patents produced	2.508 (1.404)	14.69 (10.59)	5.857 (7.543)	14.46 (19.81)	245.5 (466.4)

* for a balanced panel of 280 large, quoted UK firms over the period 1972-1982 (one year is lost through differencing for the two growth rates); corresponding figures for an unbalanced panel of 649 firms observed over the same period are given in brackets. See the Appendix for precise definitions of the variables used.

TABLE II: Correlations Across Different Measures of Performance*

	profits/ sales	log sales	growth sales	market value	growth market value	innova-tions	patents
accounting profits/sales	1.0						
log of sales	-.1760 (-.0293)	1.0					
rate of growth of sales	.2462 (.2980)	.0263 (.0802)	1.0				
market value	.0979 (.0939)	.4455 (.4183)	.0133 (.0226)	1.0			
rate of growth of market value	.1905 (.2164)	.0776 (.0043)	.0364 (.1015)	.0864 (.0597)	1.0		
number of "major" innovations produced	.0438 (.0439)	.2694 (.2528)	.0155 (.0200)	.6003 (.5944)	.0294 (.0198)	1.0	
number of patents produced	.0390 (.0395)	.2974 (.2766)	.0183 (.0194)	.5723 (.5775)	.0107 (.0011)	.4125 (.4186)	1.0

* for a balanced panel of 280 firms over the period 1973-1982; corresponding figures for the unbalanced panel of 649 firms are given in brackets.

TABLE III: Analysis of Variance*

	% "within" variation	% "between" variation
accounting profits/sales	40%	60%
log of sales	9%	91%
rate of growth of sales	89%	11%
market value	34%	66%
rate of growth of market value	96%	4%
number of "major" innovations produced	35%	65%
number of patents produced	7%	93%

* for a balanced panel of 280 firms, 1972-1982 (one year is lost for the two growth rates through differencing).

TABLE IV: Simple Autoregressions*

	β	t-statistic	R^2	implied half life	implied 10% life
accounting profits/sales	.4894 (.4874)	26.89 (34.34)	.5937 (.5558)	1	3.2
log of sales	·.8560 (.8294)	141.1 (156.7)	.989 (.988)	4.5	15
market value	1.239 (1.243)	102.3 (141.1)	.924 (.926)	-	-
rate of growth of sales	.0417 (.0240)	2.031 (1.56)	.0220 (.0246)	0.2	0.7
rate of growth of market value	-.1598 (-.1606)	7.861 (10.89)	.015 (.009)	-	-
number of "major" innovations produced	.0982 (.0879)	5.168 (6.479)	.460 (.440)	0.3	1
number of patents produced	.5031 (.4911)	29.42 (39.82)	.927 (.926)	1	3.4

* these are regressions of the form: $x_i(t) = \alpha_i + \beta x_i(t\text{-}1)$ for each of the performance variables, $x_i(t)$, taken in turn. The column labelled "t-statistic" reports (absolute) values of the standard test of the null that $\beta = 0$. Estimates for the unbalanced panel of 649 firms are given in brackets below those for the balanced panel of 280 firms. To compute the implied half life, we set $\alpha = 0$ and solve for the value of T which satisfies $x(T)/x(0) = 0.5 = \beta^T$; the "10% life" satisfies $0.1 = \beta^T$. These are computed only for the balanced panel estimates of β.

DATA APPENDIX

Both the balanced and unbalanced samples were taken from the DATASTREAM on-line firm accounts service, the Science Policy Research Unit's (SPRU) innovations database and the SPRU patents database. DATASTREAM financial data covers the population of firms listed on the London Stock Exchange from 1972 onwards. Firms with fewer than six continuous time series observations of our principal variables were dropped, as were firms whose principal operating industry (defined by sales) was outside manufacturing. Firms who were involved in large scale merger activity were also excluded from the samples.

The SPRU innovations database consists of over 4370 major innovations, defined *as "the successful commercial introduction of new or improved products, processes and materials introduced in Britain between 1945 and 1982"*; Pavitt, Robson and Townsend (1987) describe the data more fully. Innovations are enumerated by *innovating unit* - the parent or subsidiary relationship (if applicable) of each innovation holding firm has been ascertained using annual editions of Dun and Bradstreet's *Who Owns Whom* for the UK and the Republic of Ireland. An aggregate yearly innovation count has then been generated at the parent company (DATASTREAM) level for firms which are not "independent". When matched to firms accounts we have captured about one third of all SPRU innovations in the period 1972-1982. The remaining innovations accrue mainly to smaller "independent" firms who are not listed on the London Stock Exchange.

The SPRU patents database consists of an aggregate annual count of the number of patents awarded to UK registered firms by the US Patents and Trademarks Office in the period 1969-1988. There are over 7450 UK firms in receipt of US patents in this period. The parent or subsidiary relationship (if applicable) of each patent holding firm has been ascertained using annual editions of Dun and Bradstreet's *Who Owns Whom* for the UK and the Republic of Ireland. An aggregate yearly patent count has then been generated at the parent company (DATASTREAM) level for firms which are not "independent". When matched to firms accounts we have captured about 30% of all SPRU patents in the period 1972-1982. The remaining patents accrue mainly to smaller "independent" firms who are not listed on the London Stock Exchange.

The net result is a balanced panel of 280 firms observed continuously over the period 1972-82, and an unbalanced panel of 649 firms observed for at least six consecutive years during that period.

Profits are defined pre-tax, sales are in normal values and the market value series is taken from the London Share Price Database (LSPD), and it is defined as the product of the firm's outstanding shares and the price of these shares on December 31st. This date was chosen to obtain the best possible match between the dating of the firm's innovations and patents and its market value. An average price for the three months preceding the firm's accounting year was also constructed from monthly observations (last trading day) of share prices to iron out atypical fluctuations but this had little effect on results.

This data is essentially that used in Geroski, Machin and Walters, 1997.