

The Rate of Depreciation of Technological Knowledge: Evidence from Patent Renewal Data

Derek Bosworth and Gregory Jobome¹

ABSTRACT

This paper is critical of studies that assume the rate of depreciation of technological knowledge is exogenously given and constant. It argues that the development of rival inventions and/or the existence of a pool of inventions from which spillovers take place impact directly on the size and value of the stock of knowledge. Patent data seem ideal to test such hypotheses as patents represent a store of R&D knowledge, and the declining value of the exclusive right to use an invention is reflected in the failure to renew patent protection. The empirical model includes not only rival invention and spillover effects, but other variables suggested by the existing literature, such as renewal costs and investment activity (representing new niches for inventions). A new quasi-panel patent data set has been constructed, tracing the survival characteristics of each cohort of patents over the period 1950-75. The data allow the first empirical tests of whether the hazard rate from the patent stock is duration dependent, which we demonstrate is linked to the highly skewed distribution of the value of patents. The long sample period also allows an exploration of whether the influences on the obsolescence of technological knowledge have changed over the post-War period.

1. INTRODUCTION

IN RECENT YEARS, a number of areas of economic research have developed models that require estimates of R&D (or patent) stocks, such as those which use the market value of companies to test the role of intangible assets in determining company performance (recent examples include Hall, 1993 and 2000). In most instances, such stocks are constructed as a perpetual inventory measure, applying a constant rate of depreciation to past R&D expenditures (often assumed to be 15 per cent per annum). Given the shift from 'manufacturing' to 'knowledge-based' production and the anecdotal evidence of the growing significance of intangible assets, it seems important to develop a more rigorous and deeper understanding of the size and rate of depreciation of the stocks of intangible assets (ASB, 1995; Brockington, 1996; Brookings, 1997). The present paper uses patent renewal data² to provide estimates of the hazard rate from cohorts of patents granted in each of the years from 1950 to 1975.³

The hazard rate is used as a measure of the rate of depreciation of the associated technological knowledge, which may be of value both for accounting and R&D decision making purposes. Some care needs to be taken as to what we mean by this. The failure to renew is associated with the fact that it is no longer worthwhile for the patent rights holder to renew the

patent protection. This may be for a variety of reasons, of which the most obvious might be that the technological knowledge has become obsolete. However, it may not be the only reason, for example, the technology might remain 'state of the art', but associated economic conditions may change such that further commercial exploitation of the invention may never again be economically viable. The failure to renew the patent is an indication that the invention was no longer generating monopoly rents for the patent rights holder — this does not necessarily mean that it will not generate returns for society when it becomes a public good. In what follows we will use the terms attrition (referring to patent lapses) and depreciation (referring to the likelihood that the commercial value of patent protection to the patent rights holder has fallen to a value less than the costs of renewal). For simplicity we will assume that the terms patent attrition and commercial depreciation (including the decline in the options value of the patent right) are inter-changeable. Support for this is provided by the positive relationship between the longevity of patents and their commercial value (see, for example, Harhoff *et al.*, 1998).

The traditional view, therefore, is that current knowledge displaces past knowledge in a process of creative destruction, which mirrors that of new products replacing old. While this will clearly be true of particular inventions (i.e. a new pharmaceutical compound is more effective than an existing one), it may not be true of the pool of inventions in total. It has been argued, for example, '...the routine and systematic use of the existing knowledge base ... has given rise to a new economy in which the central organising factor in the process of technology creation is the ability of the system to distribute knowledge so it can be recombined' (OECD, 1994, p. 120). This suggests that it may even be possible for an existing patent to become *more valuable* as new patent disclosure takes place, because the new knowledge suggests new options for the use of the existing knowledge. The present paper explores the extent to which the rate of attrition of patents from the patent stock is determined by the effects of competition from new inventions (i.e. the result of the creative activities of competitors) and the spillovers from the pool of patented knowhow (i.e. arising from the additional profitable options created by the inventive activity of other companies).⁴ We specify a number of variables to represent both 'competitive' and 'spillover' effects, and control for a number of other possible influences (i.e. the costs of renewals, the availability of new investment niches, the economic climate, etc.).

The empirical modelling in the present paper is undertaken using a specially constructed data set of renewal information, which we discuss and illustrate in Section 2.3. In order to simplify our task, the data have been limited to patents granted during the period 1950 to 1975 (with patent renewals running through to 1991), as this leaves the whole of the data largely free from any complications arising from the changes brought about by the *Patent Act, 1977* (such as the move to a 20 year patent life and the introduction of the *European Patent Convention* and the *Patent Cooperation Treaty*). In addition, it means that all patents in the sample have either now lapsed or reached the end of their maximum legal life (which, prior to 1977, was 16 years).

The data, organised as a series of cohorts according to the year of patent grant broken down by years of renewal, are used to construct the hazard rate (i.e. the average probability of a patent lapsing, conditional on it having been renewed up to that point). The present study uses OLS techniques to estimate the empirical relationship explaining the hazard rate. In order to avoid problems caused because the hazard rate, δ , takes values between 0 and 1, the dependent variable is specified in logit form, $\log(\delta/(1-\delta))$. The estimated coefficients, which indicate

the effect of a particular variable on the log of the odds of leaving the cohort, however, are reinterpreted to also indicate the percentage and percentage point impact on δ . The model also tests for duration dependence, in other words, whether the conditional hazard rate from each cohort is dependent on the number of years that the patents have been renewed to date. In estimating the model, we control for cohort-specific effects and the effects of the age of the patent for each cohort, and demonstrate that these give a clear indication of the evolution of autonomous depreciation over time. The data are both left and right hand censored, as we explain below, but as the censoring is wholly deterministic, we argue that this should be accounted for by the cohort dummies and individual time dummies that distinguish the age of patents in each cohort.

Evidence is provided from simple exponential decay functions to support the attrition of patent cohorts (average hazard rates) of about 15 per cent per annum, at least for some of the post-War period, which is consistent with the assumption often made in the empirical literature (Hall, 1993). However, we use the data to demonstrate that, consistent with earlier work (Bosworth, 1978), unadjusted measures of the hazard rates do not appear to be exogenously given or constant over time. We return to this assumption after producing new estimates based on econometric techniques and demonstrate that the effects of autonomous depreciation have been increasing significantly over time and that the duration dependence of patent cohorts over the sample period is negative (i.e. the conditional probability of lapse declines with the age of the patent). The results presented below based on renewal data are also consistent with the highly skewed distribution of the commercial value of patents (Scherer, 1965 and 1996).

Section 2 outlines the extension of the traditional market valuation (Tobin's q) models to incorporate the effects of competition and spillovers on the depreciation and obsolescence of intangible assets. It also demonstrates how monopoly-increasing discretionary investments may give rise to a form of competition directly analogous to Schumpeter's 'creative-destruction', within a model that also incorporates spillover effects. Section 3 provides qualitative evidence that the depreciation rates are not constant and are likely to be, at least in part, economically determined. It goes on to report regression results for specifications using economic variables to explain the rate of attrition from the patent stock for a given cohort. Finally, Section 4 provides the main conclusions of this paper.

2. COMPETITION, SPILLOVERS AND STOCKS OF INTANGIBLE ASSETS

2.1 Perpetual inventory measures of knowledge stocks

The discussion of intellectual property in both the accounting and economics literature has focused on some form of perpetual inventory measure of the 'stock' of intangible assets. This approach has been used to construct both the stock of 'R&D knowledge' and the stock of 'patent knowledge'. Examples include a wide range of research, from company case studies (Bosworth and Jobome, 2001) through to the analysis of large-scale firm level panel data sets (Hall, 1993). The expenditures that generate such stocks are generally 'expensed' and, insofar as they have value, they are estimated as a capitalised value of expenses (Donaldson, 1992). However, such values are highly questionable, although we would perhaps not go so far as suggesting, '...capitalised expenses are, as an asset, pure accounting fiction; ... they are truly worthless and should be so treated.' (Donaldson, 1992). In only a limited number of cases has the literature attempted to estimate the rate of depreciation of the stock (Lev and Sougiannis, 1996) and, even here, the rates are assumed to be exogenously determined.

The main literature using these stocks focuses on their effect on enterprise performance. There are two main strands of the literature; one adopts output or total factor productivity as the dependent variable (referred to as the 'knowledge production function' approach — for reviews see Griliches, 1992 and 1995; Mairesse and Sassenou; 1991; Mairesse and Mohen, 1995) and the other uses the market value of the company as the dependent variable (referred to as the market valuation approach — see Hall, (2000) for a review). In the broadest terms, both approaches generally take the view that own-R&D and the pool of R&D outputs from all companies⁶ are both performance-increasing, the latter through spillover effects (Griliches, 1992 and 1995). While the role played by each individual competitor's R&D seems to be central to the competitive process, it has rarely been analysed in the *empirical* literature on firm performance.⁷ This appears to be a crucial omission, as it affects the measure of the 'counterfactual', in other words what the firm's performance would have been if it had not undertaken its R&D.⁸

2.2 Endogeneity and duration dependence

The discussion of the valuation of patents has thrown some light on the likely influences on patent renewals. A number of studies have utilised patent renewal fees and the attrition of the stock of patents in force to attempt to estimate the total (private) value of the stock of patented knowhow (Pakes and Schankerman, 1978; Pakes, 1986). We illustrate the approach using Pakes (1986), which is perhaps the most sophisticated analysis to date from both theoretical and econometric perspectives. In essence, the value of the patent stock is obtained from information about renewal activity and renewal costs for different ages of patents in a given cohort (in practice, the empirical estimates were based on individual patent data pooled over a period of between 20 and 30 years, depending on the country concerned). The model is based on the hypothesis that patents are granted at an early stage, before the 'uses' of the invention and, hence, the commercial value of the patented information, are fully known.⁹ After patenting, the firm incurs costs of R&D in searching for new, more highly remunerated uses for the invention. Thus, the firm pays the renewal fee in order to maintain the *option* of exploring these avenues; once it fails to pay the renewal fee, the associated knowhow becomes a public good.

The structure of the model allows for three possible outcomes, revealing whether the patent: (i) can never be profitably exploited; (ii) currently does not have a profitable application, but still might; (iii) has a profitable line of use. The patent continues to be renewed when the existing returns that arise from the patented knowledge over the coming period, and the anticipated value of the option from paying the renewal fee (which may give rise to higher future revenues) exceeds the cost of renewal. Thus, the firm's choice of whether to renew or not is based upon an optimal stopping rule (Pakes, 1986). The risk is that in continuing to search for more profitable uses of a given invention, none will be found; alternatively, in giving up the option on the intellectual property, the inventor may miss some new and profitable avenue of use.

The discussion of the Pakes model provides some suggestions about duration dependence (i.e. the conditional probability of a patent being renewed having survived up to that renewal date). Note that, for those inventions that have not found a profitable outlet yet, the expected returns from continuing to search decline with age, given that the most profitable uses tend to be investigated first and the number of years to the maximum life of the patent declines.


This suggests a *positive* duration dependence (i.e. the hazard rate increases with age). However, it is not clear what happens to the proportion of those ‘still hopeful’ (i.e. category ii above) within the remaining cohort; if this rises with age duration dependence could still be negative. The role of competition from subsequent inventions does not appear explicitly in the Pakes’ model. However, Pakes notes that, ‘The second possible outcome is that the absorbing state does not occur, but the experiments do not result in a use for the patented ideas that is more profitable than the current one. In this case current returns decay at the rate $\delta < 1$, as steps forward by other agents in the economy gradually make obsolete the returns from the agent’s own patent..’ (Pakes, 1986, p. 764).

2.3 Attrition from patent cohorts

The main data set concerns information about patent lapses (i.e. where the renewal fee is not paid to the Patent Office) and patent expiry (i.e. where the patent reaches the end of its legal life). The data on lapses and expiry is taken from what was known in earlier years as the *Annual Report of the Comptroller General of Patents, Designs and Trademarks* and more recently as the *Patent Office Annual Report and Accounts*. The data relate to the aggregate patent renewal activity, no systematic data are currently available at lower levels of aggregation in the UK, although this may be possible in the future (Bosworth and Filiou, 2002). The renewal data are set out in each year of the Report as patents renewed for the 4th, 5th, ..., 16th year (and after the 1977 Act, 4th, 5th, ..., 20th year). Renewal fees do not have to be paid until the 4th year and cease to be paid on expiry (i.e. at the 16th or 20th years, depending upon the prevailing Act under which they were granted). Thus, if we observe the 4th year renewals at time t , these relate to patents granted four years earlier, while 5th year renewals relate to patents granted five years previously. In order to trace a sequence of cohorts, the data are manipulated such that we observe all of the subsequent renewal activity for patents granted in a given year. Given that the renewals run from year 4 to year 16 inclusive, this implies that the data are in some sense left and right hand censored, although the censoring is not so complex as that arising from issues of selection and self-selection. We return to the way this is taken into account in Section 3 below.


While the discussion in Pakes (1986) is of relevance to the question of endogeneity, it does not provide any practical evidence. The only evidence that depreciation rates are endogenous comes from the fact that the rates vary across patent cohorts and over time (Bosworth, 1973 and 1978). The hazard rates of patents from the 1950-1975 cohorts (constructed as the number of patents lapsing from the stock of patents that continue in existence up to that renewal date — see equation 6 below), by length of life of patent, are shown at five year intervals in Figure 1. The years of patent grant run from 1950 to 1975, but, for simplicity, only two have been labelled (1960 and 1975). There are clear differences from cohort to cohort, with the hazard rates in the early stages of renewal (i.e. between years 5 and 6 of patent life) varying by up to 5 percentage points. However, despite these differences the figure suggests that the hazard rate constructed from the raw renewal data exhibits positive duration dependence (i.e. the hazard rate from the cohort increases with age). The degree of variation across cohorts suggests that the average hazard rate for each cohort is likely to vary over time. Further confirmation of this can be found in Figure 2, which presents the average annual hazard rates for each cohort in turn. These rise from around 12-14 per cent per annum in the early post-War period to about 15-16

Figure 1: Hazard rates for selected cohort by age of patent



per cent per annum in the first half of the 1960s, after which the rate shows an almost) monotonic decline back to around the value it started from at the beginning of the sample period. This finding gives some support for the average rates of depreciation of R&D stocks, of around 15 per cent, assumed in the literature (see, for example, Hall, 1993).

Figure 2: Average hazard rates by year of grant of the cohort



2.4 IMPLICATIONS OF ENDOGENEITY

In this section, we explore the implications of non-constant, endogenous depreciation rates for the existing literature. For ease of exposition we assume that R&D is the principal activity that determines the magnitude of the firm's intangible assets. The market value of the *i*th company at time *t*, V_{it} , is now written as,

$$V_x = f(K_x, RS_x, RS_x, \sum_{k=1}^q RS_{kx}) \tag{1}$$

where: K is the stock of the firm's tangible assets; RS is the (relevant) stock of R&D knowledge; i denotes the i th company and j the j th (a competitor) firm; spillovers occur from the pool of q firms. The literature suggests that,

$$\frac{\partial V_{it}}{\partial RS_{it}} > 0, \quad \frac{\partial V_{it}}{\partial RS_{jt}} < 0 \quad \text{and} \quad \frac{\partial V_{it}}{\partial [\sum RS_{kt}]} > 0$$

Thus, the contribution of RS_{it} should be measured having controlled for RS_{jt} , and a failure to control for the effects of competitor's R&D may result in an under-estimate of the returns to own-R&D (as own-R&D appears to have less effect than it actually has). This highlights the importance of the counterfactual in estimating the returns to own-R&D. In practice, however, there are at least two problems with this approach. First, as we have already noted, the role of competitor R&D has largely been ignored in the empirical literature. Second, although the equation controls for competitor R&D, the estimated value of i 's stock (i.e. αRS_{jt} , where α is the estimated coefficient) remains unaffected by the research output of firm j . We know of no existing example in the current literature that allows the stocks to be inter-related.

The knowledge production function that broadly corresponds with equation 1 can be written,

$$Y_{it} = g(K_{it}, E_{it}, RS_{it}, RS_{jt}, \sum_{k=1}^q RS_{kt}) \tag{2}$$

where most of the notation carries over, but Y is value added, K is a measure of the tangible capital stock and E is employment. This equation suggests that increases in total factor productivity are improved by the firms own R&D stock and by the overall pool of R&D knowledge,

$\frac{\partial Y_{it}}{\partial RS_{it}} > 0$ and $\frac{\partial Y_{it}}{\partial [\sum RS_{kt}]} > 0$, whereas the effect of competitor R&D is like to have a negative effect, $\frac{\partial Y_{it}}{\partial RS_{jt}} < 0$. Note, however, that where the stocks are pre-constructed on the basis

of assumed rates of depreciation, they have exactly the same problem as those reported in equation 1. In the case of the knowledge production function approach, however, one part of the literature has focused on establishing the rate of depreciation of own-R&D, using some variant of,

$$Y_i = AK_i^\zeta E_i^\phi R_{i,1}^{\rho_1} R_{i,2}^{\rho_2} \dots R_{i,n}^{\rho_n} \tag{3}$$

where, R_{it} denotes research undertaken by firm i at different points in time t ($t=1, \dots, n$). Each unit of R is assumed equally productive in the sense that $\rho = \frac{R}{Y} \frac{\partial Y}{\partial R} = \text{constant}$, but the effectiveness of R&D depreciates with time, (where 1 denotes the most recent year and n the earliest year). Nevertheless, while the pattern of decay of the influence of past R&D can in principle be complex, it is still fixed and exogenously determined outside of the model.

3. EMPIRICAL EVIDENCE

3.1 Modelling depreciation rates

The main thrust of the present paper is that the time pattern of ρ itself should be determined endogenously, by factors such as rivals' R&D and spillover effects. Consider, for example, the following simple relationships,

$$\rho_2 = \rho_1 - \alpha - \beta \sum_{j=1}^m R_{1j} + \gamma \sum_{k=1}^q RS_1 \tag{4a}$$

$$\rho_3 = \rho_2 - \alpha - \beta \sum_{j=1}^m R_{2j} + \gamma \sum_{k=1}^q RS_2 = \rho_1 - 2\alpha - \beta \sum_{j=1}^m (R_{1j} + R_{2j}) + \gamma \sum_{k=1}^q (RS_1 + RS_2) \tag{4b}$$

$$\rho_n = \rho_1 - n\alpha - \beta \sum_{t=1}^n \sum_{j=1}^m R_{tj} + \gamma \sum_{t=1}^n \sum_{k=1}^q RS_{tk} \tag{4c}$$

where R refers to the R&D expenditure and RS to the R&D stock. There are m competitors whose R&D may (adversely) affect firm i over the n periods that its R&D output would otherwise survive. Spillover effects from the R&D stock occur across a pool of q companies.¹⁰ It is clear from the sequence of equations that ρ represents autonomous (linear) depreciation of the effectiveness of R&D. Endogenising the rate of depreciation of the various stocks in the market value and production functions (such as equations 1 and 2), however, would result in extremely complex functions. However, the general principle is important — that the rate of depreciation of the stock of R&D knowledge is not independent of future competitive and complementary R&D activities and therefore needs to be modelled.

Patent renewal data, however, offer an opportunity to explore the impact of the autonomous depreciation, competition and spillovers on knowledge stocks, at least at an aggregate level. The main advantage of the renewal data is that they offer direct measures of the rate of attrition of the patent stock and, thereby, a proxy for the depreciation of the stock of technological knowledge created in any particular period, corresponding closely to equation 4. Such data allow us to trace the attrition to any patent cohort over time,

$$P_{\tau,t+1} = P_{\tau,t} (1 - \delta_{\tau,t}) \tag{5}$$

where $P_{\tau,t}$ denotes the number of patents, from cohort τ in existence (i.e. surviving) at time t , and $\delta_{\tau,t}$ denotes the proportion lost between t and $t+1$. Therefore, we can write,

$$\frac{P_{\tau,t} - P_{\tau,t+1}}{P_{\tau,t}} = \delta_{\tau,t} \tag{6}$$

where $\delta_{\tau,t}$ can be interpreted as the hazard rate, in other words, the (conditional) probability that a patent from cohort τ will lapse in period t , having survived up to that point in time (Kiefer, 1988; Greene, 1990, pp. 715-727). Values are not constructed for the unobserved beginning or end of period (i.e. 0-4 years — when there is no requirement to pay renewal fees, or from the end of the 16th year — when the patent life has ended).

Note that in the standard perpetual inventory measures of patent stocks, $\delta = \text{constant}$, but, in the present paper, we argue that δ is likely to be a function of a variety of influences suggested by the theory outlined above and by earlier empirical results. Thus, we write the gener-

al form of the empirical specification as,

$$f(\delta_{\tau,t}) = a + b_{\tau}t_{\tau} + c_{\tau}t + dPC_{\tau,t} + ePS_{\tau,t} + f_iX_{\tau,t} + \mu_{\tau,t} \quad (7)$$

The term τ_i denotes a set of cohort dummies — these allow the average autonomous depreciation rate to vary from cohort to cohort and, in addition, the term τ_i denotes a ‘time trend’ estimated for each cohort — this allows for duration dependence within each cohort. We believe that the fact that each cohort effectively has a different constant term and a different slope goes some way to allowing for the left and right hand truncation of the data.

Allowing for duration dependence is an essential feature of the model, enabling the hazard rate to depend upon the length of time the patented knowledge has survived to date. We might observe negative, positive or no duration dependence. The raw hazard rates calculated from the renewal data and at least one earlier piece of empirical research (Bosworth, 1978) suggest positive duration dependence. The outcome, however, is an empirical question, and we demonstrate below that, once we control for other factors, duration dependence is negative (i.e. the hazard rate falls with the length of period the ‘invention’ has survived up to that point). We show, however, that this result is still consistent with the highly skewed nature of returns to R&D, whereby a very large percentage of inventions are worth little and a very small proportion are commercially extremely valuable (Scherer, 1965 and 1996).

The equation also includes ‘competition’ and ‘spillover’ effects, as well as the influence of other variables. Thus, $PC_{\tau,t}$ represents a measure of the number of patents in competition with cohort τ patents still in existence at time t . Likewise, $PS_{\tau,t}$ denotes a measure of the pool of patents from which cohort τ might benefit if still in existence at time t . The literature suggests that a variety of other variables, X , may be important in explaining renewal activity, including the cost of renewals. The renewal cost variable underpins the literature on the commercial value of patented knowledge (Pakes and Schankerman, 1978; Pakes, 1986). We return to the precise definition and measurement of all of these variables below.

3.2 SAMPLE PERIOD

While, in principle, it is possible to undertake econometric analysis using equation 7 for a longer sample period, the present paper focuses on the period 1950 to 1975, primarily because of the effects of the Second World War on patenting activity prior to this period and the effects of the major changes to the patenting system brought about by the *Patent Act, 1977*. Note, however, that, while the grant data only run to 1975, the renewal data still run through to 1991. The period from 1950 to 1977 covers a period governed by the *Patent Act, 1949*, which was essentially a domestic system (although it followed many years of international harmonisation of domestic laws), under which both domestic and foreign inventors applied for protection within the UK. Patents were awarded to the first applicant and protection lasted for a maximum of 16 years. The procedure involved application, examination and grant, with a period for objection by other ‘inventors’. It normally took a considerable time to process applications, resulting in around a three year lag between application and grant.

The Patent Act, 1977 represented a significant departure in both the law and the administrative procedures of the Patent Office. From the viewpoint of the present study, the most important aspect is that the system became considerably more complex at this point.

Applications continued to be processed for some years under the 1949 Act, although the vast bulk had gone through the system before the end of 1981. The first grants under the 1977 Act appeared in 1978 and, by 1979, exceeded those under the 1949 Act. The key complicating factor, however, was the move towards an international patent system (as opposed to a system based around the international harmonisation of domestic laws), with the introduction of the *European Patent Convention* and the *Patent Cooperation Treaty*.

A further key change was in the procedure leading to grant, with the new sequence involving: application, first publication, request for examination, examination and second publication (grant). This change is particularly relevant as it was intended, in part, to help alleviate some of the back-log in unprocessed applications, resulting in a speedier passage through the system. Objections could be lodged at the stage of first publication (unlike the US system in which publication still does not occur until after grant). In practice, the introduction of the 'international route' did much to relieve the growing pressure on the UK domestic system.¹¹ Finally, as we have already noted, the *Patent Act, 1977* changed the maximum life of a patent from 16 to 20 years.

At the present time, therefore, in order to test the model, we have restricted the sample to the years 1950 to 1975, such that none of the data are affected by the major changes in 1977. This has the further advantage that all patents taken out in our sample period have either lapsed or reached their legal maximum patent life (i.e. by 1991).

3.3 DEFINITION OF THE VARIABLES

3.3.1 Dependent variable

Following our discussion of the attrition rate, δ (see equation 6), we now describe the form used for the dependent variable in the empirical estimates. Given that $\delta_{\tau,t}$, which denotes the attrition from patent cohort τ at time t , has an upper value of unity, we follow the standard practice and transform this variable into logit form, writing the dependent variable,

$$\delta' = \log \left[\frac{\delta}{1-\delta} \right] \tag{8}$$

This has the property that $(\delta/(1-\delta)) \rightarrow 0$ as $\delta \rightarrow 0$ and $(\delta/(1-\delta)) \rightarrow \infty$ as $\delta \rightarrow 1$. Thus, $\log []$ can be interpreted as the log of the odds of leaving the sample during the period in question. Below,

we interpret $\frac{\partial \delta'}{\partial x}$ as the percentage change in the odds of leaving the cohort

$$\frac{\partial \delta'}{\partial x} = \frac{1}{\left[\frac{\delta}{1-\delta} \right]} * \frac{\partial \left[\frac{\delta}{1-\delta} \right]}{\partial x} \tag{9}$$

It is fairly easy to demonstrate the relationship between this and the probability of leaving the cohort. Rearranging equation 9 yields,

$$\frac{1}{\delta} \frac{\partial \delta}{\partial x} = (1-\delta) \frac{\partial \delta'}{\partial x} \tag{10}$$


which is the percentage change in the hazard rate caused by a unit change in x . In addition, the percentage change in the hazard rate caused by a one per cent change in x can be obtained by multiplying through equation 10 by the mean of x .

3.3.2 Independent variables

Renewals and Renewal Fees

A further important feature of the institutional system is that patents only continue in existence after their fourth year of life, on payment of renewal fees. The existing literature suggests that renewal fees have an effect on patent survival (see, for example, Pakes, 1986). Figure 3 shows renewal fees for year 5, 10 and 16, deflated by the producer price index. The figure demonstrates the way in which renewal fees increase with the age of the patent and over time. During the first half of the period, renewal fees were kept at constant nominal rates for a number of years, before being increased to allow for the effects of inflation. This meant that the real costs of patent renewals tended to fall during periods when nominal prices were pegged (and more so the higher the rate of inflation), before jumping upwards when they were adjusted (again followed by a period of fall). In the latter part of the period, fees were adjusted on a more continuous basis to keep in line with (or ahead of) the rate of inflation — this can be seen quite clearly in the relatively ‘smooth’ series after 1980, as shown in Figure 3.

Figure 3: Trends in real patent renewal fees



The other feature of the data on renewal fees is that the relative fees charged for renewal at different lengths of life of the patent has changed over time. Figure 3 again shows this by comparing the real cost of renewals in the 5th, 10th and 16th years. We can demonstrate this quite simply by taking the ratio of the 16th/5th and 10th/5th in 1950 and comparing these, with the same ratios in 1991. In 1950, the ratios are: 2.01 (10th/5th) and 3.23 (16th/5th); in 1991, the ratios are 1.51 (10th/5th) and 2.77 (16th/5th). In fact this is an interesting result from a policy perspective, as there are grounds for the Patent Office to subsidise applications (i.e. for reasons of disclosure) and to penalise longer-lived patents (to minimise the social costs arising from the patent monopoly).

From the point of view of the present paper, however, it demonstrates that there is considerable variation in real renewal fees, both over time and across different ages of patents from a given cohort. The variables included are the real cost of renewals (Cost1) and the correspon-

ding variable squared (Cost2) — see Table 1.

Competition

A priori we anticipate that existing patents will have to fight off competition from new inventions. We represent this competition by the cumulative sum of subsequent patents following the grant of the patent cohort in question. Subsequent patents build upon the knowledge disclosed in earlier patents and, hence, in principle, ought to be associated with superior inventions. However, invention is a risky business both from technical and commercial viewpoints and, while, on balance, subsequent inventions should displace earlier inventions, there is no guarantee that this will always be the case. Indeed, we argued in the introduction that subsequent patent disclosures may reveal new information that indicates new options for existing patented knowledge. Thus, we use a measure of subsequent patents that continue in force up to the point that the renewal decision is made about the existing patent (Comp1). Our assumption is that Comp1 should have a competitive effect increasing the rate of attrition from the cohort in question.

We also include Comp2, which is a measure of patents granted subsequent to the cohort in question, but which are allowed to lapse prior to the renewal decision in question. These were originally expected to be of a higher quality than earlier patents and, therefore, given they have not been renewed, this is an indication that they failed to find a profitable outlet. Thus, the lapse of subsequent patents associated with ‘failed inventions’ is not going to raise or lower the attrition rate from the cohort in question because of any competitive effect, but may raise it insofar as it acts as a signal to some of the earlier and potentially inferior inventions that no profitable outlet is likely to emerge.

Spillovers

As in other studies of this type, it is not immediately clear how to define the pool from which spillovers occur. In the absence of other information, the pool is generally defined as the total stock of R&D or patent knowledge. In the case of the present paper, given that we are arguing that patent disclosures may form the source of such spillovers, the corresponding variable to the one that appears in the R&D literature appears to be the cumulative sum of all patents granted up to that point in time. Given that patents prior to the beginning of the sample period can be assimilated within the constant term, we proxy the pool by the cumulative sum of patents over the post-War period up to the date of the renewal in question.

Quality of competition, quality of the cohort and duration dependence

A number of measures of the quality of the patents being renewed are constructed. The most obvious is that based upon the assumption that applications through the domestic system by UK residents are likely to contain a higher proportion of more trivial ideas than those from abroad. Ideally, however, we would like this information about the breakdown of the subsequent patents in force, but, in order to do this we would require information about the foreign *versus* domestic composition of renewals, which is not available in the published statistics. Nevertheless, we can include proxies for this variable for each cohort by estimating the proportion of domestic to foreign patents at the time of grant (Qual1). Clearly, however, this variable is constant across the life of any particular cohort and only varies across cohorts. Given the other variables described below, it appears a fairly weak candidate.

The form of the data, which appears as a time series of cohorts, lends itself naturally to the inclusion of cohort dummies, τ_t . These take the form of a dummy variable for each year, $\tau_t=1$ for cohort t in every year that cohort is present and $\tau_t=0$ otherwise (note that we exclude the first potential dummy for each sample used for estimation, which then acts as the base group). The coefficient on this variable can be interpreted as the autonomous rate of depreciation for each cohort, where cohorts with higher rates of decay are associated with a lower average quality of inventive output.

While this dummy provides some information about the average quality of the cohort, it does not control for the marginal quality of the remaining patents on which the renewal decision is being made. In other words, the average quality of patents in ‘younger’ age categories is likely to be lower than the corresponding quality amongst the ‘older’ categories. In the present study, we use the cumulative sum of patents lost from the cohort up to the renewal date in question as a proxy for the quality of that cohort (Qual2). We would expect this to be a powerful variable insofar as the earlier renewal behaviour for the cohort reflects the propensity to renew at the margin.

Table 1 Summary: Notation and Definition of Variables

<i>Notation</i>	<i>Name in Table 2</i>	<i>Definition</i>
$\delta_{\tau,t}$	$\frac{P_{\tau,t} - P_{\tau,t+1}}{P_{\tau,t}}$	hazard rate
$f(\delta_{\tau,t})$	$\log\left[\frac{\delta}{1-\delta}\right]$	dependent variable: log of the (conditional) odds of leaving the cohort
τ_t	Present but not shown in Table 2	average quality of the cohort: year dummy
t_t	Time1	duration dependence: time trend for each cohort
$PC_{\tau,t}$	Comp1 Comp2	cumulative sum of subsequent surviving patents cumulative sum of subsequent patents lost
$PS_{\tau,t}$	Spill1	spillovers from the pool of all patents prior to renewal
$X_{\tau,t}$	Cost1 Cost2 Qual1 Qual2	renewal fees (real) renewal fees squared (real) quality of cohort: ratio foreign to domestic patent grants quality of cohort: cumulative sum of renewals for that cohort up to the point of renewal
	Econ1 Econ2 Inv1 Inv2	GDP level (real) GDP growth rate (real) investment level (real) investment growth rate (real)

Finally, we estimate the 'duration dependence' of each cohort using a series of time trends τ_t , one for each cohort t . This provides information about whether the conditional probability of the patent lapsing increases, remains constant or decreases with the length of life of the patent. In effect, the duration dependence coefficient is estimated within the larger sample from the 11 renewal categories available for each cohort. Given the inclusion of Qual2, whose cumulative nature is likely to have some degree of trend effect for each cohort, this will provide a fairly exacting test of the role played by the duration dependence variable. Similar to the cohort dummy variable described above, the time trends are represented by an overall (average) time trend (reported in the table of results) and a series of cohort time trends, omitting the first cohort (reported in Figure 4).

Economic activity

The precise timing of the termination of patent life within the maximum period allowed may be dependent on the prevailing economic conditions. In particular, we hypothesize that the hazard rate is likely to be lower during boom periods and higher during recession, other things being equal. Thus, we use the level and the rate of growth of GDP (Econ1 and Econ2 respectively) as indicators of the economic climate (both in real terms). However, since the work of Schmookler, patenting has always been linked to investment activity, although this relationship was developed in terms of the effects of investment on the incentive to invent, rather than the effect on renewal activity (see especially Schmookler, 1966). In order to cover this possibility, we also include series reflecting the level and rate of change in investment in plant and machinery (Inv1 and Inv2, both in real terms).

3.4 ECONOMETRIC RESULTS

This section presents the econometric results explaining the attrition from the series of cohorts from 1950 to 1975. It is important to point out that the period covered is a very long one, and while the arguments outlined above about the sign of the variables appear robust, it seems important to check for major shifts in parameter values by sub-dividing the period. It should be stated from the outset that we are expecting to find changes in parameter estimates over this period. The War-time research, for example, gave rise to a number of radical new areas of invention that were exploited in the early post-war period (i.e. jet engines, radar, etc.). In addition, the opening up of national markets led to significantly higher international competition (i.e. competition from foreign inventions) in the latter part of the sample period. In addition, the increasing availability of electronic patent data bases, with high speed search engines revolutionised the ability to utilise the disclosed information. Investigation of a number of sub-periods shows that the coefficient estimates shift fairly gradually, but continually with time. In what follows, therefore, we provide results for the sub-periods of 1950 to 1959 (beginning of period) and 1966 to 1975 (end of period), as well as for the period 1950 to 1975 as a whole.

The main results are set out in Table 2. Note that the first column for each of the three periods shows the impact of a unit change in the dependent variable on the log of the odds of the patent not being renewed (and the second column gives the associated t -statistic). Given that the log of the odds is not always easy to understand, the third column shows the percentage impact on the hazard rate of a one per cent change in the independent variable. Finally, given that the attrition is itself measured as a rate (see equation 6), the fourth column gives the per-