

Bloom, Nicholas; Reenen, John Van

Working Paper

Real options, patents, productivity and market value: Evidence from a panel of British firms

IFS Working Papers, No. W00/21

Provided in Cooperation with:

The Institute for Fiscal Studies (IFS), London

Suggested Citation: Bloom, Nicholas; Reenen, John Van (2000) : Real options, patents, productivity and market value: Evidence from a panel of British firms, IFS Working Papers, No. W00/21, Institute for Fiscal Studies (IFS), London,
<https://doi.org/10.1920/wp.ifs.2000.0021>

This Version is available at:

<https://hdl.handle.net/10419/90840>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

REAL OPTIONS, PATENTS, PRODUCTIVITY AND MARKET
VALUE: EVIDENCE FROM A PANEL OF BRITISH FIRMS

Nicholas Bloom
John Van Reenen

Real Options, Patents, Productivity and Market Value: Evidence from a panel of British firms.

Nicholas Bloom

*Department of Economics, University College London and Institute for Fiscal
Studies*

John Van Reenen

*Department of Economics, University College London and Institute for Fiscal
Studies*

This Version: 11.11.00

Keywords: Patents; Productivity, market value, real options.

JEL Classification: O33

Acknowledgments

The authors are grateful for comments from Alex Klemm, Stephen Bond, Rachel Griffith, Helen Simpson, David Stout, and Frank Windmeijer. Financial assistance has come from the Leverhulme Trust.

Abstract

Patents citations are a potentially powerful indicator of technological innovation. Analysing the new IFS-Leverhulme database on over 200 major British firms since 1968 we show that patents have an economically and statistically significant impact on firm-level productivity and market value. We also find that while patenting feeds into market values immediately it appears to have a slower effect on productivity. This is potentially because of the need for costly investment in new equipment, training and marketing required to embody patents into new products and processes. This may generate valuable real options because patents provide exclusive rights to develop new innovations, thereby enabling firms to delay their investments. We find that higher market uncertainty, which increases the value of real options, reduces the impact of new patents on productivity. These real options effects have implications for the role of macro and micro stability in the take up of new technologies and productivity growth.

1. Introduction

There is a consensus that technological advance is crucial in the “new economy”. But measuring technology has always been one of the most perplexing problems facing empirical economics. One tradition, epitomised by Solow (1957), is to measure technology as a residual from a production function. The problem is that the residual, no matter how cleverly constructed, is rather like a statistical dust bin - holding a lot of trash as well as a few nuggets of gold. A second tradition, which this paper follows, is to construct observable proxies for technical change. The most popular measure of technology is research and development expenditures (R&D). Unfortunately at the firm level there was no requirement to report R&D expenditures in Britain before 1989, so this hampers the generation of a long time series. Innovation counts have been frequently used in the UK, but the best series for these ended in 1983 (see Pavitt, Robson and Townsend, 1987, Blundell, Griffith and Van Reenen, 1999 and Geroski, 1990).

Counts of patents have also been a popular choice to proxy innovation. And patents themselves contain a wealth of other information (e.g. Lerner, 2000). In particular, the front of a patent details other patents which contributed to the knowledge underlying the new patent. This information can be used in a variety of different ways. We start off with the most obvious use. A patent which is cited many times is more likely to be valuable than a patent which is rarely cited (Griliches, 1990). Other researchers have used patent citations as a “paper flow” to track the way knowledge spills over between organisations and areas (Jaffe, Trajtenberg and Henderson, 1993; Jaffe and Trajtenberg, 1998) and this is a route that we have pursued in complementary work.

We look at the impact of patents on two measures of company performance - productivity and market value. Production functions are more easily interpretable and comparable with other work. Market value is a more forward looking measure, which has attractions for the analysis of an activity whose pay-off may not be for many years in to the future. There is a small literature emerging on the impact

of patent citations on company performance, but all the existing work that we know of is based on U.S. firms (e.g. Hall et al, 2000)¹.

>From our preliminary work with the data it became apparent that while patents have an immediate impact upon market values they take time to affect productivity. One potential explanation is that the new products and processes which are covered by the patents have to be embodied in new capital equipment and training. Firms may also need to undertake expensive marketing and advertising to promote their new products. As such, this will involve extensive sunk cost investments - these capital, training and marketing outlays will be (at least partially) irreversible. But since patents provide firms with the exclusive rights to their new technologies they have the option to wait until making these sunk costs investments. When market conditions are uncertain this will generate valuable real options. Therefore, by giving firms a legally protected right to delay investing, patents provide an excellent test of the importance of real options.

We adapt the developing real options literature to explain the take up of new products and processes covered by patents². The theories developed in this paper predict that higher market uncertainty will lead firms to be more cautious about their investments. We use this theory to then derive empirical predictions on the relationship between patents and uncertainty and empirically test them.

The structure of this paper is as follows. Section 2 describes the database that we have constructed and some of its key features. Section 3 sketches some simple models and the real options extensions that we use to estimate the effects of patenting on company performance. Section 4 details the econometric results and section 5 gives some concluding comments. In short, we find considerable evidence of the importance of technology for firms' productivity and stock market performance. Higher uncertainty, as predicted, reduces this effect of patents on

¹There are, of course, several econometric studies of the impact of patent counts on British firm performance (e.g. Bosworth et al, 2000).

²See, in particular, Dixit and Pindyck (1994), Eberly and Van Mieghern (1997), and Bloom (2000).

productivity but appears to have no significant effect on market value.

2. Data

We combine three principle datasets in constructing the IFS-Leverhulme database. Full details of the matching between the datasets is contained in the Appendix, but we sketch the process here. The first dataset is the Case Western Patent data, the second is the Datastream annual company accounting data, and the third is the Datastream daily share returns data.

To construct the patents data base we used the computerised records of patents granted in the U.S. between 1968 and 1996. This is the largest electronic patent dataset in the world (the European Patent Office records begin only in 1976, and the records are patchy until the mid 1980s). The data is held in Case Western and we received considerable help from their staff in setting up the files. Practically all major patents are taken out in the U.S., so we are screening out many low value patents by following this strategy.

The second and third datasets contain the accounts and share returns of firms listed on the London Stock Exchange. From the population of public firms we selected those whose names began with the letters ‘A’ through ‘L’, which represents a random sample from the whole population. We also added in the top 100 R&D performing firms in the UK that were not already included in this list to maximise the numbers of patents we could collect. Ideally we would have collected information on all firms on the Stock Market, but the resource cost was too great. For all of these 415 firms we used ‘Who Owns Whom’ from 1985 to find the names of all subsidiaries³. We then used these subsidiary names to match to the Case Western Dataset by name.

³There are many problems with only using one year of data to match in the corporate structure. Clearly this changes over time, albeit slowly for most firms. The process of matching is, however, extremely labour-intensive so it was only practical to perform it for one year. In future work we intend to also do the matching for later and earlier years.

2.1. Patents and Citation Data

The intersection of the two datasets gave us 236 firms who had taken out at least one patent between 1968 and 1996. The total number of patents taken out by this group over the entire period was 59,919, representing about 1% of the 6 million patents taken at the U.S. Patent Office. Table 1 below shows that most of our group of UK firms are involved in only a modest amount of patenting with about half the sample receiving more than 25 patents, while 12 firms received over 1000 patents during the period. This concentration of innovative activity within large firms (the 12 account for 72% of all patents in our data), reflects a similar phenomena in R&D expenditure where the 12 largest enterprises account for about 80% of all R&D expenditure.

	1 or more	10 or more	25 or more	100 or more	250 or more	1000 or more
Firms	236	161	117	75	41	12

Table 2 reports the patenting activity of the twelve largest patenters. This selection of firms reflects the strong performance of the chemicals, pharmaceuticals and the defence engineering sectors in the UK.

Firm	No. Patents
ICI	8422
Shell	7200
Smithkline Beecham	3672
BP	3632
BTR	3432
Lucas Industries	3119
GEC	3054
Hanson	2892
Unilever	2644
Siebe	1876
Rolls-Royce	1575
Glaxo Wellcome	1528

The patents are graphed by their year of application in Figure (5.1). The lesser degree of patenting activity in the latter part of the period reflects truncation bias (on the right) because we collect statistics on patents granted. Since there is a delay between applying for and granting a patent of about two years, this leads to a downward bias towards the end of the period. There is also a truncation on the left of the graph as there may have been patents granted post 1968 which were applied for pre-1968. These caveats apart, there is little discernible trend in the total patents numbers granted to UK firms. There was some decline 1968-1983, a pick-up 1983-1990 and then decline in the 1990s. Interestingly, this broadly mirrors the growth rates in UK productivity.

We also have data on the citations made by any of the other 6 million patents in the main data set to our sample of 59,919 patents. Citations can be taken as an indicator of the technological value of a patent in that those patents which are frequently cited are likely to be more innovative and technologically productive. In Figure (5.2) we plot the histogram of the lag between a patent being taken out and the subsequent citations to that patent. It is clear that citations tend to happen relatively early on in a patent's life when the patent is widely known but technologically still innovative. Interestingly this citation lag still has not completely tailed off even after 20 years. Figure (5.3) plots the histogram of the number of cites per patent from which it is clear that many patents are never cited at all, with a small tail of patents which are frequently cited. The five most cited patents are tabulated in Table 3 below with their patenting topic, the year they were granted and the number of cites made to them over the period 1976 until 1996.

Company	Patent Topic	Grant Year	Cites 1976-96
Shell	Synthetic Resins	1972	221
Grand Metropolitan	Microwave heating package	1980	174
ICI	Herbicide compositions	1977	130
Unilever	Anticalculus composition	1977	97
British Oxygen Corp.	Pharmaceutical Drugs	1975	89

The total number of citations to our patents, dated by the application year of the patent being cited, is plotted in Figure (5.4). Because data on citations is only collected for patents granted after 1976 there is an early downward bias reflecting the fact that for patents granted pre 1976 some of the initial citations data is missing. The discussion in the paragraph above and Figure (5.2) suggests that the loss of these early citations could lead to a serious downward bias for patents taken out pre 1976 since for them this would represent a period of relatively high citation activity. There is also a tail end bias as patents applied for towards the end of the period will only be part of the way through their citations lifecycle, and so will have been cited less often by 1996.

To deal with these biases we use a non-parametric series estimator based on a full Fourier sine and cosine expansion. Following the approach in Hall et al (2000) we assume that the total lifetime number of citations per year is constant throughout our sample. Therefore any observed change in the observed aggregate citation levels is due to time varying levels of truncation bias. We assume that this time varying truncation bias varies smoothly over time according to some piecewise continuous function of time⁴. Our normalization estimator then uses a Fourier expansion to fit a smooth curve to the observed time variation in aggregate citation levels to non-parametrically estimate a truncation bias function.

⁴That our *observed* citation frequency is not smooth over time, even in our sample of almost 60,000 patents, is testament to the extreme skew of the citations data. In data sets such as these which have large second moments the usual weak convergence of the empirical distributions to their underlying distribution is extremely slow (see for example Billingsley, 1986), so that smoothing is usually required.

A Fourier expansion was used because of its ability to approximate to an arbitrary degree of accuracy any piecewise-continuous function (see Churchill and Brown, 1987). We used the first four sine and four cosine terms for an expansion with the base periodicity set at the total time observation period of 30 years⁵. The smoothing property of our estimator can be seen in Figure (5.4) which plots the actual citation frequency and our non-parametric functional estimator. This functional estimator of the time varying citation bias is then inverted to re-weight the citations per patent. This ensures that the normalized citations per patents remain approximately constant over the period.

In calculating a patent based proxy for knowledge stocks it is also more sensible to use a stock measure rather than a flow measure of knowledge as the benefits from a patent are likely to persist into future years. We calculate a set of preferred measures of the stock of patents through the perpetual inventory method so that

$$(\text{Patent Stock})_t = (1 - \delta) \times (\text{Patent Stock})_{t-1} + \text{Patents}_t \quad (2.1)$$

where the knowledge depreciation rate, δ , is set to 30% as in, for example, Griliches (1990). The first year is calculated by assuming a prior steady state growth of patents of 5%. The same perpetual inventory method is used to calculate the citation stock where the flow variable is the citation weighted number of patents. The “5 year cite stock” uses only the first 5 years of citations (after an application) to obtain a citation weighting but without any normalization. Since we select our citation estimation period to run up to 1990 only whilst our citing data runs up to 1996 this means we have 5 years of observations on citations for every patent so that no truncation bias correction will be needed for this 5 year measure. This reduction in the number of years for estimation thus allows us to compare our normalized full citation weighted patent measure and the five year citation measure.

⁵Increasing the length of the base period or using the first three or five terms does not have any significant impact on our results. This is because the first few terms of the Fourier expansion drive the results, as noted for example, by Bertola and Caballero (1994) in a related application.

It is comforting that our three measure of the knowledge stock - the patent stock, the citation weighted patent stock, and the 5 year citation weighted stock - have a strong correlation as demonstrated in Table 4 below. This suggests that whilst each should have its own merit in capturing various aspects of the knowledge stock they proxy a similar measure of the technological innovation stock.

	Pat Stock	Cite Stock	5 Year Cite Stock
Patent Stock	1	0.9871	0.9665
Cite Stock		1	0.9714
5 Year Cite Stock			1

2.2. Firm Level Accounting and Uncertainty Data

The company data is drawn from the Datastream on-line service and represents the accounts of firms listed on the UK stock market. Our initial sample of 415 firms (those whose names began with A-L or were large R&D performers) ⁶ for which we matched patent data was then cleaned for estimation. Cleaning involved ensuring that there are no missing values on sales, capital or employment, deleting firms with less than three consecutive observations, breaking the series for firms whose accounting period fell outside 300 to 400 days due to changes in year end timing, and excluding observations for firms where there are jumps of greater than 150% in any of the key variables (capital, labour, sales). After cleaning we were left with a sample of 404 firms, to which 211 were matched as having patenting subsidiaries (see the Data Appendix for details of this matching process)⁷.

Table 5 reports summary statistics for this set of 185 patenting firms. Because these are quoted firms the median size is large with sales of about £360 million (in 1985 prices) and a work force of over eight thousand employees. From the last row

⁶ See the data appendix for more details on the selection of this sample.

⁷ This is less than our group of 228 patenters because of both the loss of some firms due to trimming and because of the loss of some years of observations for the remaining firms due to the unavailability or poor quality of data on employment in the early 1970s.

of the table it can also be seen that we generally have a long time series of data on each firm - on average over 20 years for each firm. However, because of the need to deal with the biases discussed above in patent and citation counts we only use the data up to 1993 for raw patent numbers and 1990 for citations. Hence, the average number of observations per firm is 19 and 16 for patents and citation weighted patents, which still represents a relatively long time period per firm. The patent numbers demonstrate the large variation in patenting per firm year with some firms only taking out sporadic patents - as demonstrated by the zero patent observations in Table 3 - and others taking out 409 patents in a single year (ICI in 1974). The total cites number represents the normalized sum of citations for all patents taken out in each firm year.

	median	mean	stan. dev.	min.	max
real capital (£m)	143	744	1,777	1.6	18,514
employment (1000s)	8,398	24,374	42,078	40	312,000
real sales (£m)	362	1,224	2,494	1.15	20,980
real market value (£m)	153	740	1,766	0.29	19,468
patents	3	12.6	34	0	409
total cites	13.7	61.2	157	0	1808
patent stock	10	42.6	113	0	1218
cite stock	49.2	202	507	0	5157
5 year cite stock	26.2	105.9	227	0	2919
uncertainty	1.39	1.47	0.42	0.60	6.6
observations per firm	22	20	7.6	3	29

Notes: Capital, sales and market value are all in 1985 £1,000,000s. 'Patents' is the total number of patents per firm year whilst cites is the normalized total number of cites to a firms patents per year. Uncertainty is the % standard deviation of daily share returns. Sample covers years 1968-96.

In measuring uncertainty we have to capture measure firms uncertainty about future prices, wages rates, exchange rates, technologies, consumer tastes and government policies. In an attempt to capture all factors in one scalar proxy for firm

level uncertainty we use the variance of the firm's daily stock returns⁸, denoted σ_i^2 . In accordance with theories of real options this is a time invariant but firm specific proxy for uncertainty⁹. This measure includes on a daily returns basis the capital gain on the stock, dividend payments, rights issues, and stock dilutions. Such a returns measure provides a forward looking proxy for the volatility of the firm's environment which is implicitly weighted in accordance with the impact of these variables on profits. A stock returns-based measure of uncertainty is also advantageous because the data is accurately reported at a sufficiently high frequency to provide an extremely accurate measure. Our sampling size of 265 recordings per year for the 22 year life of our average firm therefore provides an extremely low sampling variance¹⁰.

3. Models of Patents and Company Performance

We work with a simple Cobb-Douglas production function of the form

$$Q = AG^\alpha N^\beta K^\gamma \quad (3.1)$$

where Q is real sales, G is the knowledge stock, N is number of employees and K is the capital stock and A is an efficiency parameter. Taking logs and introducing subscripts for firm i at time t we have

⁸This measure of uncertainty is also used by other papers in the literature on uncertainty and investment, such as Leahy and Whited (1998).

⁹The real options literature, surveyed in Dixit and Pindyck (1994) for example, always assumes a time invariant uncertainty measure. This drives us to choose a time invariant uncertainty measure so that we can make a close link with the theoretical literature.

¹⁰For example, Andersen and Bollerslev (1998) use high frequency exchange rate data with 288 recordings per period and calculate the implied measurement errors are less than 2.5% of the true volatility.

$$\log Q_{it} = \log A_{it} + \alpha \log G_{it} + \beta \log N_{it} + \gamma \log K_{it} \quad (3.2)$$

We parameterise efficiency, $A_{it} = \exp(\eta_i + \tau_t + v_{it})$, as a function of firm specific fixed effects (η_i), time effects (τ_t) and a random stochastic term (v_{it}). In our empirical application we use patent stocks and citation-weighted patent stocks (PAT) as empirical proxies of G , the knowledge stock.

$$\log Q_{it} = \alpha \log PAT_{it} + \beta \log N_{it} + \gamma \log K_{it} + \eta_i + \tau_t + v_{it} \quad (3.3)$$

We estimate equation (3.3) by within groups (least squares dummy variables) correcting the standard errors for heteroscedasticity.

Market value equations are less well established than production functions. The standard approach pioneered by Griliches (1981) is based on a specification of the form (see also Hall et al., 2000)

$$\log \left(\frac{V}{K} \right)_{it} = \delta \left(\frac{G}{K} \right)_{it} + \tilde{\eta}_i + \tilde{\tau}_t + \tilde{v}_{it} \quad (3.4)$$

where V is the market value of the firm. The left hand side of equation (3.4) is essentially Tobin's average Q. Under perfect competition one would expect this ratio to be equal to unity. The deviation of Tobin's Q from unity is, in this framework, driven by the fact that the firm possesses intangible (G) as well as tangible (K) capital. "New economy" firms with high levels of intangible knowledge capital will have a much higher market value than one would expect if we merely used their fixed capital stock.

3.1. Uncertainty and Real Options

The two models laid out above assume that the knowledge contained in patents can be immediately be used and acted on by firms. Patents, however, represent new products or process innovations whose introduction can involve sizeable investments in additional plant and equipment, hiring and retraining workers, and advertising and marketing. Much of this expenditure will be irreversible - once it is undertaken the initial costs will not be recoverable. Thus, when firms are facing uncertain market conditions then they will possess patent *real options*¹¹. These patent *real options* reflect the value a firm places on its ability to choose the timing of its investment in its patented technologies when this involves sunk costs.

A large theoretical literature has grown up from the seminal papers of Bertola (1988), Pindyck (1988), and Dixit (1989) demonstrating the important role such real options can play in firm's optimal investment strategies, with real options even able to account for more than half a firm's value in uncertain market conditions. As a result real options should play an important role in our approach to modelling investment in innovation. This work emphasizes the role real options play in retarding the response of firms to changing market conditions¹². When market conditions are uncertain firms become reluctant to commit large sums to new investment projects or dismantle old investment projects in case conditions change. This leads to a 'cautionary' investment behaviour. This 'cautionary' ef-

¹¹Apart from partial irreversibility and market uncertainty, the third condition for the existence of real options - that firms can delay their actions - is clearly satisfied in this case where patents give firms the exclusive rights to use their innovations.

¹²See, for example, the early work on threshold behaviour by Abel and Eberly (1996), Dixit and Pindyck (1996), and Bloom (2000).

fect of real options in retarding the response to changing market conditions has been confirmed empirically for physical investment by Guiso and Parigi (1999) in a cross section of Italian firms, and by Bloom, Bond and Van Reenen (2001) in a panel of UK firms.

The incorporation of new products and processes into a firm's production schedule will also be subject to precisely this kind of cautionary effect because of the capital investment, training and marketing required. Hence, firms may be reluctant to exploit these patents in uncertain market conditions when the chances of making an expensive mistakes is high. Since patents provide firms with the exclusive right to use their new innovations they have considerable ability to delay their investments generating potentially substantial real options values.

To incorporate these real options effects we extend the concept of our knowledge stock G into embedded knowledge, G_E and dis-embedded knowledge G_D , where $G = G_E + G_D$. Embedded knowledge represents those product and process innovations which the firm has invested in. Dis-embedded knowledge, however, represents the remaining ideas which the firm has under patent but has not yet committed into actual production.

When conditions are highly uncertain the firm will be more cautious because of the value of the real options associated with embedding new innovations into production. To model this empirically we define $\lambda(\sigma) = \frac{\partial G_E}{\partial G} \frac{G}{G_E}$ to be the elasticity of embedded knowledge to total knowledge, where this ratio depends on the uncertainty of a firms business conditions, σ . This elasticity of embedded to total knowledge, $\lambda(\sigma)$, will be a decreasing function of uncertainty. More uncertain conditions will increase the cautionary effect of real options and lead to a slower

pass through of new innovations into the embedded knowledge stock¹³ - that is $\partial\lambda(\sigma)/\partial\sigma < 0$.

Adapting our earlier production function to incorporate these real options effects and using lower case to denote logs we can therefore write

$$q = a + \alpha g_E + \beta n + \gamma k \quad (3.5)$$

where the embedded knowledge stock, g_E , rather than the total knowledge stock, g , is included in the firm's production function. To evaluate the additional effects of uncertainty on production we take a second order Taylor series expansion of the logged production function in the total knowledge stock, g , and uncertainty σ . This can be used to predict the qualitative effects that uncertainty should play in the adoption of new technologies.

Firstly, holding other factors constant greater levels of total knowledge will lead to a greater level of productivity as some proportion of new innovations will always become embedded in new products and processes¹⁴

$$\begin{aligned} \frac{\partial q}{\partial g} &= \alpha \frac{\partial g_E}{\partial g} = \alpha \lambda(\sigma) \\ &> 0 \end{aligned} \quad (3.6)$$

Secondly, this effect of new innovations on output will be falling in the level of uncertainty because this reduces the rate of incorporation of new ideas into production

¹³This is true both in the short run and the long run because innovation in general, and patents specifically, generally have a falling *private* value over time. Hence, slow incorporation reduces their overall value both in the short run and in the long run.

¹⁴We assume new innovations have an embodied value drawn from some distribution in which the most valuable innovation will always be embodied.

$$\begin{aligned}\frac{\partial^2 q}{\partial g \partial \sigma} &= \alpha \frac{\partial \lambda(\sigma)}{\partial \sigma} \\ &< 0\end{aligned}\tag{3.7}$$

Finally, the direct effect of uncertainty is ambiguous. While uncertainty reduces the speed with which new firms embed new innovations into production there is an ambiguous relationship between the *absolute* level of embedded innovations and uncertainty¹⁵

$$\begin{aligned}\frac{\partial q}{\partial \sigma} &= \alpha \frac{\partial g_E}{\sigma} \\ &\lesssim 0\end{aligned}\tag{3.8}$$

To empirically investigate these effects we include a direct uncertainty term and an uncertainty patenting interaction term. The uncertainty patenting interaction term will then pick up the negative cautionary effect of the term in equation (3.7)

Our estimating equation for the production function will take the form

$$\log Q_{it} = \alpha \log PAT_{it} + \beta \log N_{it} + \gamma \log K_{it} + \psi \sigma_i + \chi [\sigma_i * \log PAT_{it}] + \eta_i + \tau_t + v_{it}\tag{3.9}$$

where the coefficients ψ and χ will pick up the direct and interaction effects of uncertainty. Note that we will not be able to separately identify the linear effect of uncertainty from η_i in the specifications where the latter are treated as fixed effects.

¹⁵While it is likely this sign is negative because greater uncertainty reduces the *rate of* adoption of technology it may not reduce the *absolute level* of adopted technology.

A slower uptake of technologies due to real options should have a much less dramatic effect on the forward looking market value measures. If the adoption of new patented technologies is only delayed by real options effects, then as a forward looking measure, market values should only display a limited response reflecting the limited changes to the total expected discounted cash flow. Therefore, in our empirical equation market value equation shown below, which includes uncertainty interactions, we would expect a lesser and possibly

$$\log \left(\frac{V}{K} \right)_{it} = \delta \left(\frac{G}{K} \right)_{it} + \theta \sigma_i + \zeta \left[\sigma_i * \log \left(\frac{G}{K} \right)_{it} \right] + \tilde{\eta}_i + \tilde{\tau}_t + \tilde{v}_{it} \quad (3.10)$$

insignificant point estimate on the interaction term ζ , while the sign of θ will as before remain ambiguous.

4. Results

Table 6 presents the results of estimating a standard production function on our sample of firms. Column (1) has the OLS estimates of the production function for our complete population of over 2,000 Datastream firms. As expected the coefficients on capital and labour are both positive and significant at conventional levels, and their sum is close to unity (suggesting constant returns in tangible factors). In column (2) we undertake estimation with our preferred within groups estimator which controls for time invariant difference between firms by including firm dummies. Again the coefficients on capital and labour are positive and significant, although slightly smaller than in column (1). Columns (3) compares these within groups results from the whole Datastream sample to our sub-sample of patenters. The higher point estimates on capital and lower point estimates on labour imply that our sample of patenting firms are on average more capital

intensive than lower tech firms (as one would expect). In fact our patenting firms have on average a 20% higher capital to labour ratio than non patenting firms.

Table 6: Basic Production Functions							
<i>Log Real Output</i>	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Firms	All		Patenters				
Log Capital	0.330*** (0.006)	0.289*** (0.010)	0.437*** (0.027)	0.439*** (0.027)	0.468*** (0.031)	0.471*** (0.031)	0.468*** (0.031)
Log Employment	0.650*** (0.006)	0.606*** (0.010)	0.558*** (0.027)	0.554*** (0.027)	0.502*** (0.031)	0.491*** (0.032)	0.502*** (0.031)
Log Patent Stock				0.024*** (0.008)			-0.012 (0.018)
Log Cite Stock					0.030*** (0.007)		0.039*** (0.014)
Log 5 Year Cite Stock						0.031*** (0.007)	
Firm Dummies	no	yes	yes	yes	yes	yes	yes
Time Dummies	yes	yes	yes	yes	yes	yes	yes
Adj. R-Squared	0.901	0.989	0.992	0.992	0.992	0.992	0.992
No. Observations	18,068	18,068	2219	2219	1896	1896	1896
No. Firms	2063	2063	211	211	189	189	189

NOTES: The dependent variable is ‘log real output’. Columns (1) and (2) presents results using our complete Datastream population of all firms, Columns (3) to (7) present the results for our sub-sample of firms with patents. The estimation period covers 1968 until 1993 inclusive for columns (1) to (4), and 1968 until 1990 inclusive for columns (4) to (7) which use the citation data. The *** denotes 1% significance, ** denotes 5% significance and * denotes 10% significance. Standard errors are robust to heteroskedasticity.

The last four columns of Table 6 reports the results from including patents as a proxy for knowledge in the production function. In column (4) we use patent stocks, in column (5) citation weighted patent stocks and in column (6) the “five year ahead” citation weighted patent stock measure. On all the alternative measures, patent stocks are significant at the 0.05 level with an elasticity of about

0.03. This suggests that a doubling of the patents stock would lead to a 3% increase in total factor productivity. In column (7) we include both the patent stock and the citation knowledge stock and find that patents are no longer significant. Thus, citations provide significant information over and above raw patents numbers. This suggests citations could provide a valuable proxy for evaluating knowledge stocks and tracing knowledge flows.

Table 7 reports the results of estimating the impact of patents on firm market value using the conventional average Q specification described in equation (3.4). In column (1) we use the patent stock measure, in column (2) our citation weighted patents stock measure, and in column (3) the “five year ahead” measure and find all three have significant explanatory power at the 5% level. The coefficient in column (2) suggests, for example, that doubling the citation weighted patents stock would increase the value of firms per unit of capital by about 43%. This large estimate of the effect of cited patents on market values captures the market’s expectation of the total discounted rents from patented innovations. These results are larger than those reported for US firms by Hall et al. (2000) where they report coefficients of 0.607 and 0.108 on (patent/capital stock) and (cite patent/capital stock). This appears to be mainly because they chose a 15% rather than a 30% depreciation rate on patents so that their patenting and citation stocks will be approximately twice our size¹⁶. Finally, in column (4) we again compare the predictive power of patents and citation weighted patents and find that citations

¹⁶All our results are robust to using this alternative assumption on the knowledge depreciation rate. For example, if we use a 15% rather than a 30% depreciation rate and re-estimate our market value equations we obtain a coefficient (standard error) of 0.879 (0.327) and 0.246 (0.081) on the (patent stock/capital stock) and (citation stock/capital stock) terms respectively. In our productivity equations we obtain a coefficient (standard error) of 0.035 (0.013) and 0.028 (0.010) on our patent stock and citation stock measures respectively.

provide significant additional information over and above raw patents counts.

Table 7: Market Value with Patents Measures				
$\log(V_{i,t}/K_{i,t-1})$	(1)	(2)	(3)	(4)
Patent Stock/capital	1.620**			-0.352
	(0.537)			(0.828)
Cites Stock/Capital		0.427***		
		(0.147)		
5 Year Cite Stock/Capital			0.519**	0.491**
			(0.221)	(0.243)
Firm Dummies	yes	yes	yes	yes
Time Dummies	yes	yes	yes	yes
No. Observations	2053	1748	1748	1748
No. Firms	205	182	182	182

Notes: The dependent variable is ‘log (market value/lagged capital)’. The estimation period covers 1969 until 1994 inclusive for column (1) and 1969 until 1990 inclusive for columns (2) to (4) which use the citation data which is only available for this shorted period. The *** denotes 1% significance, ** denotes 5% significance and * denotes 10% significance. Standard errors corrected for arbitrary heteroskedasticity.

In Table 8 we conduct some robustness tests on our basic models. In columns (1) and (2) we include both the patent stock and the lagged patent stocks measures. It is the lagged variable which is most informative in predicting productivity, suggesting that patented innovations take some time to enter the production function. In the market value equation, however, the current value of patents per unit capital has the larger coefficient (1.115) and is significant at the 15% level, while the lagged value has a coefficient of (0.661) and is not significant at all¹⁷. This larger point estimate on the current value in the market value equation appears to reflect the forward looking nature of the market value measure. In columns (3) and (4) we lag all our right hand side variables one period to control for the possibility of endogeneity of our current explanatory variables. This does

¹⁷Although individually insignificant, the current and lagged values of they are jointly significant at the 1% level.

not noticeably change our results with significant effects of patents on productivity and market values. We also re-run this specification with all our explanatory variables lagged twice and again find our results looks very similar with a point estimate (standard error) of 0.042 (0.013) on patents in the productivity equation and of 1.01 (0.405) on (patents/capital) in the market value equation. We also look for both structural breaks and a time varying coefficient on our patent measures and find no significant evidence for either, with our main results remaining significant.

Table 8: Robustness Checks				
	(1)	(2)	(3)	(4)
	<i>Real Sales</i>	$\log(V_{i,t}/K_{i,t-1})$	<i>Real Sales</i>	$\log(V_{i,t}/K_{i,t-1})$
Firms				
Log Real Capital	0.468***			
	(0.028)			
<i>Lagged</i> Log Real Capital			0.444***	
			(0.031)	
Log Employment	0.541***			
	(0.028)			
<i>Lagged</i> Log Employment			0.459***	
			(0.030)	
Log Patent Stock	-0.004			
	(0.011)			
<i>Lagged</i> Patents Stock	0.029**		0.036***	
	(0.010)		(0.010)	
Patent Stock/Capital		1.155		
		(0.785)		
<i>Lagged</i> Patent Stock/Capital		0.661		1.362**
		(0.718)		(0.525)
Firm Dummies	yes	yes	yes	yes
Time Dummies	yes	yes	yes	yes
No. Observations	2053	1975	2053	1975
No. Firms	205	182	205	182

Notes: The dependent variable for the columns (1) and (3) is ‘log real sales’ and the dependent variable for columns (2) and (4) is ‘log (market value/capital stock)’ - both are in 1985 prices. The estimation period covers 1968 until 1990 inclusive. The *** denotes 1% significance, ** denotes 5% significance and * denotes 10% significance. Standard errors corrected for arbitrary heteroskedasticity.

We also attempted to address any econometric concerns over the exogeneity of our knowledge stock measures by directly using instrumental variable estimators. However, with a total cross section of around 200 firms this could lead to serious small sample bias for a generalised method of moments type estimator¹⁸. On the other hand the Anderson-Hsiao type estimator can have a poor empirical performance because of the need to first difference the data which removes the levels information. This causes problems for instrumenting highly autoregressive series (like the patents stock) in a first differenced specification (see, for example, Bond and Blundell, 1998). For this reason previous studies using patenting data such as Hall et al. (2000) have focused on using OLS or within groups estimators. Because of these problems with IV estimation we only undertake some exploratory robustness checks here in addition to those in columns (3) and (4) of Table 8, using an estimator which is more appropriate in this long time series setting. This estimator uses the levels information from a within groups estimator but instruments with lagged explanatory variables to deal with any simultaneity problems, and has been used for example, in Mark, McGuire and Papke (2000)

¹⁸The GMM estimator is asymptotic in the cross section on a year by year basis so that it is the yearly number of firms not the total number which is important.

and Bloom, Griffith and Van Reenen (2001)¹⁹.

These instrumental variables point estimates are approximately similar to those delivered in tables 6 to 8 but with much larger variances. The point estimates (standard errors) on the patenting stock and citation weighting patenting stocks in the productivity equations are 0.048 (0.026) and 0.065 (0.035) respectively, which are larger than those in the comparable columns (4) and (5) of table 6, but are somewhat imprecisely estimated. For the market value specifications the point estimates (standard errors) on the patenting stock and citation weighted patenting stocks are 2.474 (0.948) and 0.792 (0.293) respectively, which are also somewhat larger than those presented in table 7 columns (1) and (2), although again quite imprecisely estimated. We concluded that, if anything, treating patents as exogenous causes an underestimation of their importance.

Finally, Table 9 reports our results from investigating the effects of uncertainty on the productivity response to patenting. In column 1 the patenting uncertainty interaction term takes the predicted negative sign in our productivity equation, and is significant at the 5% level. The coefficient on the level of uncertainty (σ_i), which is theoretically ambiguous, is also negative but not significant at the 5% level. When we move to the within groups specification in column (2) by including a full set of firm dummies we have to drop this firm specific uncertainty term σ_i

¹⁹To be precise, the data is transformed into orthogonal deviations, where the orthogonal deviations transformation of $y_{i,t}$, labelled $y_{i,t}^*$, has the form (see Arrelano and Bond, 1991)

$$y_{i,t}^* = \left(y_{i,t} - \sum_{j=t+1}^T y_{i,j} \right) \left(\frac{T-t}{T-t+1} \right)^{\frac{1}{2}} \text{ for } t = 1, 2, \dots, T-1$$

The instruments for labour, capital and patents are labour, capital and patents lagged two and three years (also in orthogonal deviations). The instruments for (patents/capital) are also (patents/capital) lagged two and three years (in orthogonal deviations).

since this is collinear with the firm dummies. In column (2) we see that the patenting interaction term is, as before, negative and significant at the 1% level. The size of this interaction coefficient (-0.01) suggests that increasing a firm's uncertainty by one standard deviation (0.42) from the median level of uncertainty (1.39) would reduce the elasticity of productivity with respect to patents from 0.024 to 0.199. Hence, for a one standard deviation increase in uncertainty the patenting effect on productivity falls by about 20%, a moderate but not enormous change.

In column (3) we investigate the levels and interaction effects of uncertainty on the market value in an OLS equation. In line with our theoretical predictions the interaction effect on market values is less significant and of a smaller magnitude than the direct patenting effect. This is because market values are a forward looking measure and so will only incorporate the effects of higher uncertainty to the extent that it impact on total discounted cash flows. The levels effect of uncertainty, however, is highly significant and negative suggesting some direct effects of uncertainty on market values. This could possibly be through a real options effect, which is theoretically ambiguous, or through some other channel such as the greater discount rate associated with more uncertain shares²⁰. In column (4) we include a full set of firm dummies to control for fixed differences between firms. The uncertainty-patenting interaction term remains negative but is now insignificant at conventional levels suggesting only a limited effect of slower patent embodiment on the long term discounted cash flows.

²⁰Strictly speaking the relationship between uncertainty and capital valuation implied by theories such as the Capital Asset Pricing Model (CAPM) or the Consumption CAPM relies on *covariance* (with the market) rather than *variance* concepts. Since covariances and variances are likely to be positively linked, however, this negative statistical relationship between variance and capital valuation is not surprising.

Table 9: Real Options Effects of Uncertainty				
	(1)	(2)	(3)	(4)
	<i>Real Sales</i>	<i>Real Sales</i>	$\log(V_{i,t}/K_{i,t-1})$	$\log(V_{i,t}/K_{i,t-1})$
Firms				
Log Real Capital	0.451***	0.446***		
	(0.015)	(0.020)		
Log Employment	0.517***	0.553***		
	(0.017)	(0.016)		
Log Patent Stock	0.025**	0.038***		
	(0.011)	(0.008)		
$\sigma_i \times$ Log Patent Stock	-0.015**	-0.010***		
	(0.006)	(0.003)		
Patent Stock/Capital			0.913***	1.743***
			(0.338)	(0.447)
$\sigma_i \times$ Patent Stock/Capital			-0.265*	-0.073
			(0.159)	(0.127)
σ_i	-0.036		-0.297***	
	(0.024)		(0.048)	
Firm Dummies	no	yes	no	yes
Time Dummies	yes	yes	yes	yes
No. Observations	2053	2053	2037	2037
No. Firms	211	211	205	205

Notes: The dependent variable for the first two columns is ‘log real sales’ and the dependent variable for the second two columns is ‘log real market value’ - both are in 1985 prices. The estimation period covers 1968 until 1990 inclusive. The *** denotes 1% significance, ** denotes 5% significance and * denotes 10% significance. Standard errors corrected for arbitrary heteroskedasticity.

To account for the possible effects of market-wide bubbles and fads we also calculate a second measure of uncertainty, using the variance of the firm’s daily share returns normalized by the return on the FTSE All-Share index. This measure eliminates common stock market volatility. Results using this normalized measure are almost identical to those reported in table 9, and are available on request from the authors.

5. Conclusions

Patents citations are a potentially powerful indicator of technological innovation. Our analysis of the new IFS-Leverhulme database on over 200 major British firms since 1968 has uncovered some interesting results. First, we show that patents have had an economically and statistically significant impact on firm-level productivity and market value. For example, a doubling of the citation-weighted patent stock increases total factor productivity by 3%. We find that citations are more informative than the simple patent counts that have been used previously in the literature. Secondly, we find that while patenting feeds into market values immediately it appears to have a slower effect on productivity. Thirdly, we find that higher market uncertainty, reduces the impact of new patents on productivity. This is consistent with a simple “real options” effect that has been found to be important in the literature on tangible investment.

There are several future directions to take this stream of research. We have not investigated the technological spillovers that have been a focus of attention in the recent literature. Patent citations are a potentially useful source of information in tracking the flows of knowledge across industries and countries and we intend to use the citations data in combination with R&D to investigate spillovers. A second area of interest is in probing the uncertainty results in more detail. If more uncertain environments reduce the productivity benefits from patents then it is likely that reductions in uncertainty (as is a focus of recent government policy to reduce “boom and bust”) will have effects on firms incentives to innovate. A natural extension of this work is to augment the patenting equations with measures of uncertainty to uncover the importance of volatility in affecting innovation.

DATA APPENDIX

1. Selecting the sample of UK firms

To obtain a manageable sample of firms for the matching stage we took UK firms from UK Datastream with names starting from A to L on which we also have company ownership data from the Leverhulme Company Ownership project²¹. This set of firms was then supplemented by any other UK firms which we believed was likely to be a significant innovator, as proxied by their appearance in the top 100 R&D spenders in the “UK R&D Scoreboard” (DTI, 1997). This resulted in a final sample of 415 firms against which we attempted to match patenting data.

The ISIC breakdown of the selected firms is given in Table A1 below where it can be seen that we have a mix of various sectors, but with a concentration in the traditionally innovative chemicals, pharmaceuticals and engineering sectors.

Table A1: Industry Breakdown of Patenting Firms

ISIC	Industry	No. observations
3100	Food and Beverages	216
3200	Textiles and Apparel	55
3400	Paper and Paper Products	49
3500	Chemicals & Pharmaceuticals	409
3600	Non-Metallic Minerals	143
3700	Basic Metal Industries	93
3800	Engineering & Metal Products	793
3900	Other Manufacturing	111
4000	Electricity, Gas and Water	10
5000	Construction	20
7200	Transport and Storage	11
689	Other Services	267

2. Matching Up the Patent Data

²¹See, for example, Bond and Chennells (2000), for more details.

The main patent data set contains information on over 6 million patents granted between 1968 and 1996 by the United States Patent Office. This information includes indicators on the name of the inventor, the location at which the patent was taken out and the patents which the patent cites itself. These patents also have an assignee code which is an indicator matching up the patent with the organisation which registered it, and there are about 140,000 different patenting assignees. Since firms register patents under a variety of different assignee names, usually the parent name of the firm or one of their subsidiaries, it is not possible to directly match up these assignees with the ultimate parent. For example, Glaxo PLC.. has 354 patents listed directly under the assignee "GLAXO GROUP LIMITED", 196 patents listed under the assignee "GLAXO LABORATORIES LIMITED", and a further 80 patents listed under an assortment of other assignee names such as "GLAXO INC" and "GLAXO CANADA". Whilst the linkage between the Glaxo parent group and its subsidiaries are obvious in this case due to the common parent firms name in many other cases this link is less clear. For example, one of the largest patenting firms in our group is "BTR PLC" whose three largest patenting assignees are called "DUNLOP LTD", "STEWART-WARNER CORPORATION" and "BTR INDUSTRIES", where only the third assignee would show up under a direct computerised name search. Because of the inadequacy of direct computerised name matching we had to undertake a careful two stage process to try and match up all the assignees of our ultimate parent list of UK quoted firms to our list of 140,000 patenting assignees. This was carried as follows:

Stage 1

We selected the larger assignees - deemed to be those with 10 or more patents - which numbered about 12,000 out of the initial list of all 140,000 assignees and accounted for 5.2 million (or 87%) of all of the 6 million patents and undertook a manual match against these. This was an extremely lengthy procedure which involved individually looking up each of these 12,000 assignees by turn in "Who Owns Whom 1985" to check whether its ultimate parent was a UK company. If the ultimate parent was a UK company we then name checked this against our sample of 415 quoted UK companies to see whether it was owned by a firm in this group. If the assignees was owned by one of our 415 quoted UK companies we then typed in the appropriate Dscore matching details.

Stage 2

For the remaining 128,000 assignees which had registered less than 10 patents (and accounted for only 0.8 million patents) we had to rely on direct computerised name matching by searching on key string words in the ultimate parent name and then cross checking with "Who Owns Whom" to ensure this was a correct match. For example, string searching on Glaxo welcome revealed two additional patenting assignees with less than 10 patents which would have been missed in the first stage of matching- "GLAXO CANADA, INC" with 2 patents and "GLAXO OPERATIONS UK LIMITED" with 3 patents. Whilst this procedure is less desirable (we proportionally matched only a third as many assignees by computer compared to by hand) it was the only feasible matching procedure. Since only 13% of the total patents stock number is contained in this longer list of small firms the degree of omitted patenting should anyway not be great.

When this matching process was complete we aggregated the patenting infor-

mation assignees up to their common parent company to yield firm level patent statistics.

3. Extending the employment data series

There are two main problems with the UK employment series. First, employment information was only recorded for a sub-sample of Datastream firms prior to 1982. For those firms where data was missing we matched in the information from the EXSTAT database of company accounts which contains employment data back to 1972. Secondly, prior to 1982 UK companies were required to report their total UK labour force, but from 1982 onwards have been required to report their global labour force, leading to a break in the series. Some firms report the global employment series all the way through, but many report only UK employment before 1982. Since we have a long time series of patenting data pre 1982 we have tried to bridge across this gap rather than drop the earlier data. To bridge this gap we have use two methods. For firms who report both UK and global labour force figures post 1982 we use the post-1982 UK/global ratio to extrapolate the pre-1982 global labour force based on the reported UK labour force. For the remaining firms we assume that the growth in total labour force between 1981 and 1982 is equal to the average of the growth rates two years before and two years after this split. This enables us to generate a forecasted 1982 UK labour force figure, from which we can obtain an estimated UK/global labour force split and extrapolate the global labour force pre-1982 using the reported UK labour force. Since the mean (median) share of global employment based in the UK in 1982 is 0.87 based on those firms which report both figures²² the errors arising

²²Because the more internationalised firms are more likely to continue reporting this UK/global employment split this may actually underestimate the share of global employment

from trying to extrapolate the global labour force from the UK labor force exercise should not be great. However, to check the robustness of our results to this we re-estimated all our regressions including a dummy variable for extrapolated labour data and found no significant change in the estimated results.

based in the UK, so that its true average is probably greater than 0.87.

Bibliography

Abel, Andrew. and Eberly, Janice. (1996) "Optimal Investment with Costly Reversibility", *Review of Economic Studies*,. vol.63m pp. 581-593.

Andersen, Torben. and Bollerslev, Tim. (1998) "Answering the Skeptics: Yes, Standard Volatility Models do Provide Accurate Forecasts", *International Economic Review* 39:4.

Arellano, Manuel, and Bond, Steven, (1991) "Some Tests of Specification for Panel Data: Monte Carlo Evidence and an Application to Employment Equations", *Review of Economic Studies*, 58, pp277-297.

Bertola, Giuseppe. (1988), "Irreversible Investment", Massachusetts Institute of Technology doctoral thesis, re-printed in *Research in Economics* (1998), vol. 52, pp. 3-37.

Bertola, Giuseppe. and Caballero, Ricardo. (1994) "Irreversibility and Aggregate Investment" *Review of Economic Studies* 61: 223-246.

Bloom, Nicholas, (2000), "The Real Options Effects of Uncertainty on Investment and Labour Demand", IFS working paper no. w00/15.

Bloom, Nicholas, Griffith, Rachel, Van Reenen, John (2001), "Do R&D Tax Credits Work? Evidence from an International Panel of Countries 1979-1998", forthcoming *Journal of Public Economics*.

Bloom, Nicholas, Bond, Steven, Van Reenen, John (2001), "The Dynamics of Investment under Uncertainty", IFS mimeo.

Billingsley, Patrick (1968), "Probability and Measure", 2nd edition, New York: John Wiley & Sons.

Blundell, Richard, and Bond, Steve, (1998). "Initial Conditions and Moment

Restrictions in Dynamic Panel Data Models”, *Journal of Econometrics* 87, 115-143.

Blundell, Richard, Rachel Griffith, John Van Reenen (1999) ‘Market Structure and Innovation: Evidence from British Manufacturing Firms’ *Review of Economic Studies* 66(3), 228, 529-554.

Bond, Steve, Harhoff, Deitmar and Van Reenen, J. (1999) ‘Investment, R&D and Financial Constraints in Britain and Germany’ IFS Working Paper 99/5.

Bond, Steve and Chennells, Lucy (2000), ”Institutional Share Ownership and Dividend Behaviour”, IFS Mimeo.

Bosworth, Stoneman and Greenhalgh (2000) “Patents, trademarks and market value: Evidence from British firms” University of Warwick, mimeo.

Churchill, Ruel, and Brown, John (1987), ”Fourier Series and Boundary Value Problems”, 4th edition, New York: McGraw Hill.

Dixit, Avinash. (1989). ”Entry and Exit Decisions under Uncertainty”. *Journal of Political Economy* 97 (June): 620-638.

Dixit, Avinash K., and Pindyck, Robert S. (1994) ”Investment Under Uncertainty”, Princeton University Press, Princeton, New Jersey.

Eberly, Janice. C., and Van Mieghem, Jan. A. (1997), ”Multi-Factor Dynamic Investment under Uncertainty”, *Journal of Economic Theory*, vol. 75, pp. 345-387.

Geroski, Paul (1990) ”Innovation, Technological Opportunity and Market Structure” *Oxford Economic Papers*, 42, 586-602.

Griliches, Zvi (1981) “Market Value, R&D and Patents” *Economics Letters*, vol 7 (2), 183-187.

Griliches, Zvi, (1990), "Patent Statistics as Economic Indicators: A Survey", *Journal of Economic Literature*, pp. 1661-1707, vol. XXVIII

Guiso, Luigi. and Parigi, Giuseppe. (1999) "Investment and Demand Uncertainty", *Quarterly Journal of Economics*, February, pp.185-227.

Hall, Bronwyn, Jaffe, Adam and Trajtenberg, Manuel (2000) 'Market Value and Patent Citations: A First Look' NBER Working Paper 7741.

Henderson, Rebecca, Jaffe, Adam and Trajtenberg, Manuel (1993), "Geographic Localization of Knowledge Spillovers as Evidenced by Patent Citations" *Quarterly Journal of Economics*, Vol. CViii, pp. 577.

Jaffe, Adam and Trajtenberg, Manuel (1998) 'International Knowledge Flows: Evidence from patent citations' NBER Working Paper No. 6507.

Leahy, John V. and Whited, Toni. M. (1998) "The Effects of Uncertainty on Investment: Some Stylized Facts", *Journal of Money Credit and Banking*, vol. 28, pp. 64-83.

Lerner, (2000) "50 Years of Patent Office Practice" National Bureau of Economic Research Working Paper 7477.

Pavitt, K. Robson, N. and Townsend, J. (1987). "The size distribution of Innovating Firms in the UK: 1945-1983", *Journal of Industrial Economics*, vol 35 297-316.

Pindyck, Robert (1988) "Irreversible Investment, Capacity Choice, and the Valuation of the Firm" *American Economic Review* 79 (December): 969-985.

Solow, Robert. (1957), "Technical Change and the Aggregate Production Function", *Review of Economics and Statistics*, vol. 39, pp. 257-263.

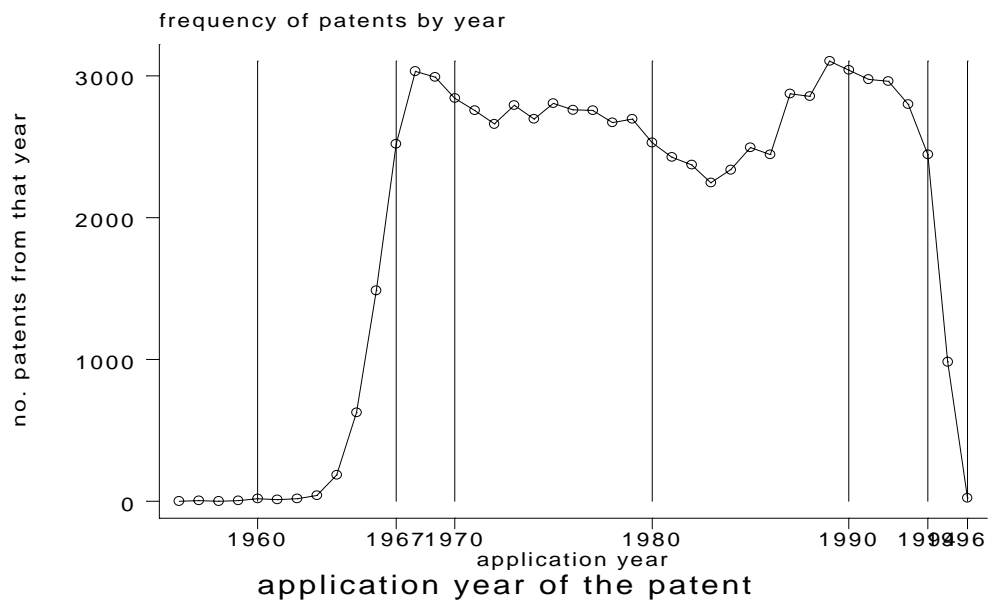


Figure 5.1: Patents per Year

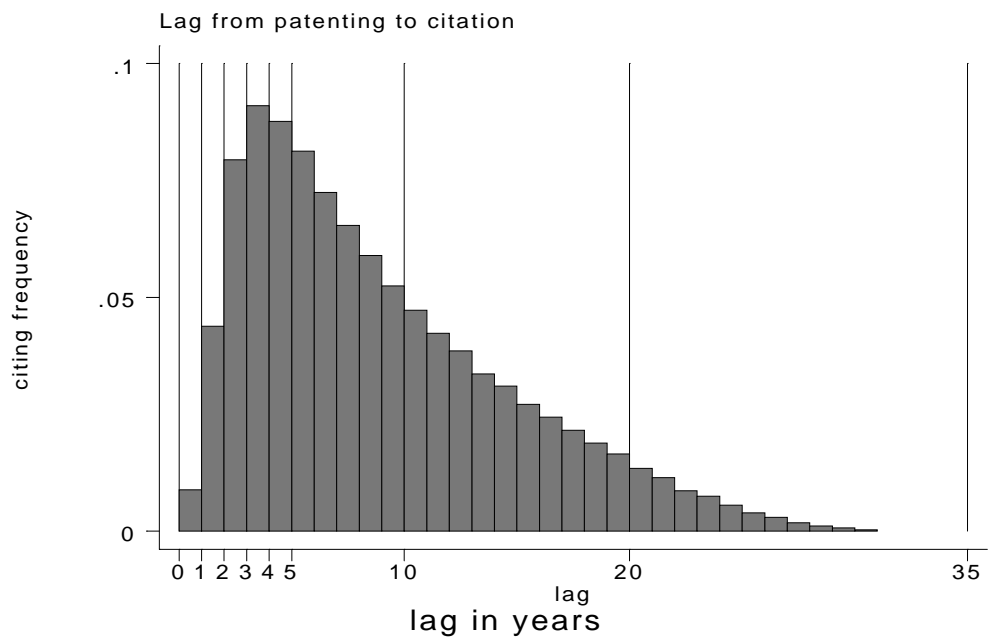


Figure 5.2: Lags Between Patenting and Citing




Figure 5.3: Distribution of Cites Per Patent




Figure 5.4: Actual and Normalizing Average Citations per Year