THE JOURNAL OF INDUSTRIAL ECONOMICS 0022-1821
Volume XLVI December 1998 No. 4

HOW TO COUNT PATENTS AND VALUE INTELLECTUAL PROPERTY:THE USES OF PATENT RENEWAL AND APPLICATION DATA*

JEAN O. LANJOUW,† ARIEL PAKES‡ AND JONATHAN PUTNAM§

Patent counts are very imperfect measures of innovative output. This paper discusses how additional data—the number of years a patent is renewed and the number of countries in which protection for the same invention is sought—can be used to improve on counts in studies that require a measure of the extent of innovation. Simple weighting schemes are proposed, which may remove half of the noise in patent counts as a measure of innovative output. We describe models of the patent application and renewal processes whose parameter estimates can be used to assess the value of the proprietary rights created by the patent laws. We illustrate their use with estimates of how the value of patent protection would vary under alternative legal rules and renewal fees and with estimates of the international flows of returns from the patent system. Recent progress in the development of databases has increased the potential for this type of analysis.

I. INTRODUCTION

PATENT DATA have been used both as a source of information on the extent of invention and as a source of information on the value of the protection generated by the patent laws. This paper discusses, in turn, how patent renewal and application data can be used to further our understanding of each of these issues.

Several different measures of the extent of innovation have been used in the literature and each has its strengths and weaknesses. Among the different measures, patents are unique in both the extent of the detail that they contain and in the breadth of their coverage. A patent document provides information on the characteristics of the underlying innovation (for example, its technological area or its citations to related innovations)

* We thank two anonymous referees and the editors for useful comments. Pakes thanks the NSF (grant no. SBR-9512106) for support

† Authors' affiliation: Jean O. Lanjouw, Economics Department, Yale University, 37 Hillhouse Avenue, PO Box 208264, New Haven CT 06520-8264, USA and NBER
email: lanjouw@econ.yale.edu

‡ Ariel Pakes, Economics Department, Yale University, 37 Hillhouse Avenue, PO Box 208264, New Haven CT 06520-8264, USA and NBER
email: ariel@tempest.econ.yale.edu

§ Jonathan Putnam, Economics Department, Yale University, 37 Hillhouse Avenue, PO Box 208264, New Haven CT 06520-8264, USA and NBER
email: jdp@crai.com

and its inventor (both the inventor *per se* and the owner or the assignee of the patent) not available elsewhere. Moreover, patent data are available for all firms and individuals over a very long time period (Sullivan [1994] for example, uses patent renewal data to study 19th century innovation). R&D expenditure data have been the most commonly used alternative indicator of innovative activity. They have the advantage of assigning dollar (and hence, to an economist, easily interpretable) values to the extent of investment in innovation. However, R&D data are more clearly related to inputs into the innovative process than to successful outputs. Their use as an indicator is also hampered by the difficulty of ensuring that the expenditures included under the heading R&D have been recorded in a consistent manner across time and firms and by the fact that these data are, at best, available for a subset of larger firms (and this only in recent periods). Other data that might provide indicators of innovative output—for example, scientific citations databases and innovation surveys--are being developed, but these data are currently very limited in coverage (see Kleinknecht [1996]). It is the unique combination of detail and coverage found in patent data which make them particularly well suited to studies of the efficacy of policies tailored to particular technological areas or specific types of firms, the cross country flows of benefits from the patent system, externalities in the knowledge generation process and many related phenomena.

The problems encountered when using patent data as a measure of innovation in such studies stem primary from the fact that the importance of the innovations protected by individual patents varies widely. As a result, patent counts are very imperfect measures of innovative output. This generates two measurement problems. First, there may be systematic differences in the mean value of the innovations protected by different groups of patents. This makes comparisons of counts a biased measure of differences in the value of the innovations being counted (be it their social or their private value). Second, even among groups with similar *mean* values, the noise in the relationship between patent counts and the value they represent makes it difficult to use counts to study the causes and consequences of variation across groups in the value of innovation.

Recent research has attempted to use additional information from the patent system to refine the patent count measure. Studies using patent renewal data exploit the fact that, in most countries, patentees must pay periodic renewal fees in order to keep their patents in force. Those using patent application data make use of the fact that the same invention may be patented in many countries (producing a patent 'family'). Provided that more valuable inventions generate larger and/or longer-lived patent families, we can use the application and renewal data to attach weights to patents and produce weighted patent count indices that are more precise measures of innovative output than raw patent counts (see Section III).

The second reason for an interest in application and renewal data directly reflects the incentives underlying the application and renewal processes. Patents represent the legal right to exclude others from using an innovation. Thus the private value of a patent is determined by the difference in the returns that would accrue to the innovation with and without patent protection. Since it is this incremental value that determines both application and renewal decisions, application and renewal data contain rather direct information on the value of the proprietary rights created by patent laws and policy, that is, on the *value of patent protection*.

Since patent rights are seldom marketed, application and renewal data are one of the few sources of information on the value of patent protection available. These data can be transformed into a useable form by combining them with models of application and renewal behavior. We can then investigate a host of questions related to the value of patent protection: how it varies with legal institutions, which countries (or firms, or technology groups) gain disproportionately from the patent laws and so on. In other words, the renewal and application data can be used to investigate the efficacy and the implications of a major tool of Intellectual Property Rights (IPR) policy. In addition, the parameter estimates derived from these models are informative about various features of the innovative process, including the nature of the process by which the market for an innovation opens up and the extent to which the returns from an innovation become obsolete over time.

The purpose of this paper is to provide a review of what has been learned from application and renewal data to date and then to consider how such data might be used to further explore important issues surrounding innovation. We begin, in Section II, with a discussion of the alternative frameworks that have been used to map the renewal and application data into estimates of magnitudes of more direct interest. Some of the empirical results currently available are then used to illustrate the potential usefulness of the renewal *cum* application data in studies of the value of patented ideas (Section III) and in studies of the value of patent protection (Section IV). We conclude in Section V with some suggestions of future directions.

II. THE FRAMEWORK OF PATENT RENEWAL AND APPLICATION MODELS

The interest of economists in patent renewal data goes back at least to Nordhaus' thesis [1969]. Pakes and Schankerman [1984] stimulated broader interest in renewal patterns by showing how to use these data to uncover characteristics of the value of patent protection. We begin this section by outlining the framework used in the Pakes and Schankerman study.

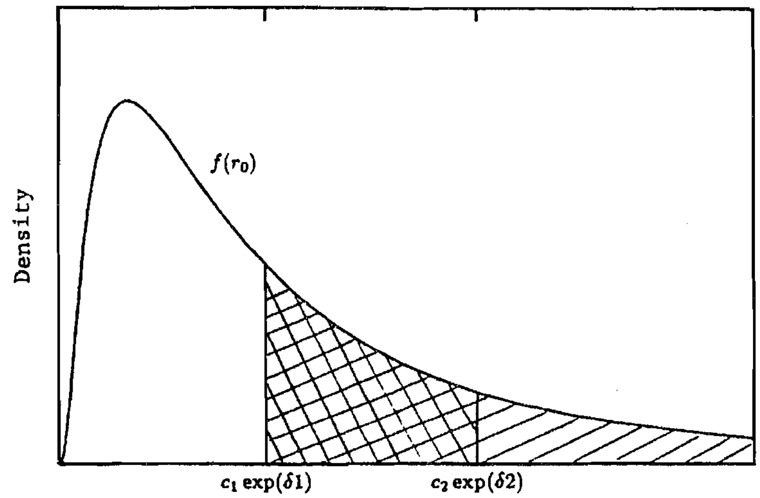

Figure 1
The density of initial returns corresponding to 1 or 2 periods of renewal

The model assumes that an inventor has already filed an application. Each application is endowed with an initial one-period return to patent protection, say r_0, which is assumed to decay deterministically at an annual rate of δ thereafter. Patentees must pay an annual renewal fee to keep their patents in force, and this fee increases with age. A patent owner seeking to maximize the expected discounted value of the (net) returns to patent protection renews his patent at age 'a' if and only if current returns, $r_0 \exp(-\delta a)$, are greater than the current cost of renewal, say c_a.[1] Equivalently, the patent is renewed at age 'a' only if $r_0 > c_a \exp(\delta a)$.

The curve labeled $f(r_0)$ in Figure 1 shows the density of the initial distribution of returns. Thus patents with initial returns to the right of the vertical line labeled $c_1 \exp(\delta 1)$, or those with returns in the lined area, renew at age 1. The patents which drop out at the second renewal are those with initial returns greater than $c_1 \exp(\delta 1)$ but less than $c_2 \exp(\delta 2)$, or those with returns in the hatched area in the figure.

Assuming a functional form for the initial distribution of returns, Pakes and Schankerman [1984] show that the parameters of this distribution, together with δ, can be estimated by finding the parameter values that make the drop out proportions predicted by the theory 'as close as possible' to those actually observed in the data. The estimates obtained

[1] In this simple deterministic model, if the condition that current returns be greater than current fees is not met in the current period, it will not be met at any age thereafter. This fact delivers the myopic renewal rule given above.

can be used to characterize the distribution of the value of patent protection and its evolution over time (see Section IV).

Pakes [1986] allows a patentee to be uncertain about the sequence of returns that would be earned were the patent to be kept in force. The move to a stochastic model of returns allows for the fact that inventors often apply for patents at an early stage in the innovative process, a stage at which they are still exploring opportunities for earning returns from the use of the information embodied in the patented ideas. In particular, the benefits of protection may increase as the owner learns about the characteristics of the invention and of the market. In this model, because there is a possibility that returns will increase, patentees may find it worthwhile to renew, even if current returns are less than the renewal fee, in order to preserve the option of protection in the future (once a patent lapses, protection is lost forever). Thus, rather than a single decay rate, Pakes [1986] allows there to be a sequence of age specific conditional distributions of returns (the distribution of returns at age '$a+1$' conditional on returns at age 'a'). This additional detail allows us to obtain a deeper understanding of the nature of the innovative process, and therefore of the effects of various policy options.

Lanjouw [1998] estimates a somewhat different stochastic model than that of Pakes [1986] and applies it to more recent and disaggregated data. In addition, she uses a behavioral model that allows for the fact that a patentee must be willing to defend his rights against infringers in order for patent protection to be meaningful. Importantly, this merging of renewal data with available institutional detail allows one to evaluate different legal institutions and IPR policy reforms (see Section IV).

The first attempt to integrate application data into the analysis of the value of patent protection is in Putnam [1996]. He extends Pakes and Schankerman's [1984] analytic framework (which conditions on application) by incorporating the inventor's prior decisions as to whether to apply for patent protection in each country offering such protection. The returns earned as a result of patents in the different countries are allowed to differ both by patent and, for a given patent, by the characteristics of the country. However, mainly for analytic convenience, in this first study Putnam assumes that the returns earned in a given country do not depend on whether the patent is kept in force in a second country. Consequently, inventors apply for a patent in each country where the expected discounted value of net returns (returns minus application and renewal costs) is positive.

Putnam's [1996] study extends the usefulness of patent data in several ways. First, it shows us how to compute estimates of the distribution of the global value of patent protection accruing to inventions. Second, it allows us to study the international flow of returns from patent protection (the returns earned by nationals of one country as a result of the protection afforded their innovations by the patent laws of another country). Third, it provides us with an ability to estimate differences in the cost of application as a function of both the country of origin of the patent and the country of application. Fourth, because all the relevant information about the family size of an invention is available within a few years of first filing, a patent weighting scheme based on applications data can be more timely than one based solely on renewals. Finally, as discussed further in Section III, by combining application data and renewal data we can produce a scheme for weighting patent counts which is more precise than one obtained using renewal data alone.

One troubling aspect of these models is that they rely on assumed functional forms. All evidence points to there being a large number of low-valued patents and a fat 'tail' of higher-valued patents which are filed for in (essentially) all countries and renewed to the statutory limit in those countries. However, the data do not distinguish between different possible realized values of patents in this tail.[2]

The functional forms used in the early studies were chosen partly to allow the estimates to match available information. In studies using a deterministic framework, models were estimated assuming Pareto, Weibull and lognormal distributions for initial returns and the fits obtained by these approaches were compared with each other (Schankerman [1998]; Lanjouw [1992]; Schankerman and Pakes [1986]). In these studies, a lognormal distribution was found to fit the data most closely. Unfortunately, computational constraints in estimating stochastic models, constraints that would have been far less binding if the authors had been using the current generation of computers, limited the robustness analysis that was done on the results from the stochastic models.

Clearly, were we able to obtain exogeneous information on the shape of the tail of the distribution of patent values, we would be able to use the patent renewal and applications data more effectively. In recent work, Harhoff, Scherer and Vopel [1997] have been examining a variety of data in an attempt to determine the value and distribution of innovations in the tail. They have interviewed the owners of German patents that reached the statutory term limit in 1995 to obtain detailed information on the profitability and other characteristics of the underlying innovations, and on the role that patent protection played in helping them to appropriate the returns from those innovations. To the extent that the results of this study are generalizable, the detail it produces on the shape of the tail of the patent value distribution might be just the sort of complimentary information needed in order to increase the precision of measures of

[2] The probability of developing a patent in the tail, together with a probability weighted sum of likely values, influences application and renewal behavior in the early years, and the stochastic models do distil this information from the data.

innovation derived from renewal and application data. As in the studies discussed above, their initial results indicate a highly skewed distribution of values. They fit a Pareto distribution and obtain a value for the Pareto coefficient of less than one. (Formally, at coefficients less than one, projecting the density into regions where we have no data would generate a value distribution with no mean.)

Largely in order to determine the amount of information on the tail of the distribution of patent values available in renewal data, Pakes and Simpson [1989] develop estimation and testing techniques that do not rely on strong functional form assumptions. They develop non-parametric tests of the hypothesis that the returns to patent protection for one group of patents are higher (in a first order stochastic dominance sense) than those for another group. They also show that, with large enough sample sizes and enough variation in renewal fee schedules, patent renewal data are rich enough to identify the entire sequence of conditional distribution functions that determine patent values in the stochastic models. Since in reality there is limited variation in fee schedules, they also show what can be determined from any set of schedules, provided the sample is sufficiently large. While the sample sizes they had available were large enough to make the testing procedures quite useful, due to insufficient data they did not attempt to 'non-parametrically' identify the conditional distributions of returns.

As discussed in the concluding section, the models outlined above may be extended in many directions. However, they already provide enough structure to enable us to illustrate, in the following sections, how patent renewal and applications data might help in improving measures of both the value of patented innovations and the value of patent protection. We begin with the value of patented innovations.

III. WEIGHTED PATENT COUNTS AS A MEASURE OF INNOVATION

As noted above, a measure of the output of the innovative process would help us in analyzing a host of policy and descriptive issues related to the causes and effects of technological change. Simple patent count measures of output have been used extensively (see the review in Griliches [1990]) but because patents protect innovations of widely varying (private or social) value, the patent count measure often runs into difficulties. In particular, if different groups of patents have different mean values, patent count comparisons can be misleading.

This is illustrated rather dramatically by Schankerman and Pakes [1986]. Previous studies had shown the patent/R&D ratio declining rapidly over time in most Western countries. This, in turn, had created a concern that we had entered into a period of 'technological exhaustion' in which the potential for further productivity growth was small (see Evenson

[1984] and for a more recent review, Kortum [1997]). Using the Pakes and Schankerman [1984] model, but allowing the parameters of that model to vary with both the patent cohort (the filing year) and the country, Schankerman and Pakes [1986] compare aggregate *patent count* indices to their estimated *patent value* indices for each of the UK, France and Germany for the period 1955 to 1975. Table I is taken from their study. On the basis of this table they conclude that '...one cannot draw inferences on changes in the value of cohorts of patents during this period from changes in the quantity of patents, for there have been large (and largely offsetting) changes in the "quality" (or mean values) of patents...'. Since their estimates come from renewal data, Schankerman and Pakes [1986] refer to the value of patent protection rather than the underlying value of the patented ideas, but at least *a priori* one would think that the two are closely related.

Pakes and Simpson [1989] draw similar conclusions after applying their non-parametric testing procedures to Finnish and Norwegian aggregate patent renewal data. It is notable that the inverse correlation between quantity and quality across cohorts of patents was seen even in Finland where, unlike in the other countries, patent counts increased over time. Schankerman [1998], using data from France disaggregated by technology group for the period 1969–81, also finds that decreases in patent counts were partially offset by increases in the average value of patents in his data.[3]

In addition to changes over time in the mean value of patents, it is also possible to discern differences in their value in other dimensions. In the Norwegian and Finnish data, Pakes and Simpson [1989] find that, conditional on the cohort and nationality of the patentee, patents from different ISIC industry categories have different value distributions. They derive a rough ordering across industry groups: pharmaceuticals, lumber, wood and paper, as well as machinery and chemicals dominate; followed by electronics, fabricated metals and stone, clay and glass; then heavy industries and finally a 'low-tech' grouping. Conditional on cohort and ISIC categories, they do not find (significant) differences in the value distribution across different patentee nationalities. However, in similar

[3] There are several possible reasons for the average quality of the patents in a cohort to be inversely related to their quantity. The simplest is suggested by the estimates of the shape of the patent value distribution. All of the studies discussed above find that most patents are of very little value. Thus, any policy change which, say, lowers the value of patenting in general (an increase in application costs, a change in the mood of the courts against patentees, a more restrictive anti-trust stance, for instance) could cause large changes in the number of low-valued patents applied for and, consequently, an increase in the mean value among those which remain. That is, many types of policy variations could induce a negative intercohort correlation between average patent values and patent counts. It is not necessary, of course, that the relationship be offsetting—results presented in Lanjouw [1992] suggest that, in the case of German computer patents, large increases in numbers were accompanied by increasing average values.

TABLE I
INDICES OF QUANTITY, QUALITY AND TOTAL VALUE PER SCIENTIST AND ENGINEER FOR COHORTS OF PATENTS AT SELECTED YEARS, 1965-1975[a]

	United Kingdom			France			Germany		
Year	P/SE	$\bar{V}$	V/SE	P/SE	$\bar{V}$	V/SE	P/SE	$\bar{V}$	V/SE
1965	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00
1968	1.03	1.06	1.09	0.89	1.15	1.02	0.88	1.49	1.31
1972	0.69	1.18	0.81	0.70	1.40	0.98	0.64	1.39	0.89
1975	0.58	1.63	0.95	0.56	1.25	0.70	0.53	1.96	1.04

Note:
[a] P, $\bar{V}$ and V denote the number of patent applications, the mean value of patent rights, and the total value of patent rights, respectively. SE refers to the number of scientists and engineers.
Source: Schankerman and Pakes [1986 Table 6].

tests on French and German data performed by Schankerman [1998] and Lanjouw [1992], equality is rejected in both nationality and technology dimensions (although with stronger evidence of differences in the technology dimension).

Even when comparing groups of patents with similar mean values, the large variance in patent values generates a degree of noise in patent count measures which makes them extremely difficult to use in studies of innovation. There are varying estimates of the fraction of the total variance in patent values captured by the patent count measure and all indications are that the quality of the patent counts measure depends on the type of data. For example, Griliches [1990] estimates that the variance in *numbers* of patents across firms is just 7% of the variance in the *value* of patents across firms. The fraction of the variance in patent values captured by differences in patent counts is likely to be even lower in the within firm across time dimension (see Pakes and Griliches [1980]), but higher when we aggregate up to inter-industry differences in patent counts (see Lach [1995]). Independently of the data under consideration, however, there is little doubt that the variability in patent values significantly reduces the efficacy of patent counts as a measure of invention.

Indeed one of the longest lasting debates in the history of economic measurement has been whether the noise and the biases in patent count measures can be made small enough to make patent counts useful measures of innovative output in economic studies (see for example, the papers of Kuznets and Sanders and the comments of Schmookler in Nelson [1962]). We reiterate here that these problems with the patent count measures are particularly unfortunate since the best alternative, R&D expenditure data, is not comprehensively available. Moreover, the R&D data that are available are not broken down by technology group and contain neither the detail on ownership nor on relationships to other inventions, found in patent data.

Both renewal and application data can be used to develop a weighted patent count measure which mitigates the problems in the standard patent count measure. The idea is straightforward. Rather than simply counting the number of patents, we partition them into groups (say J of them) by the age at which the patent was allowed to lapse (at which the renewal fee was not paid) and/or by the set of countries in which patent applications were filed. We then construct a patent value index, say VI, as

$$VI = \sum_{j=1}^{J} w_j N_j, \tag{1}$$

where N_j is the number of patents in group 'j' and w_j is the weight associated with that group. To construct this index we need the set of weights, $\{w_j\}$.

There are a number of alternative ways to determine these weights. One is to regress a measure of the value (private or social) of innovation on the N_j and let the data estimate the $\{w_j\}$. The choice of dependent variable would determine the interpretation of the weights. Profit and firm value would produce alternative indices of the contribution of innovation to private returns, while a more aggregate (industry or economy-wide) productivity variable would be more likely to produce measures which are closely linked to the social value of innovation (see Griliches [1979] and Jones and Williams [1996] for discussions). The relationship between the private and the social value of innovation is at the core of most innovation policy, so a comparison of the weights obtained from measures of private and social value and of how the patterns of weights differ across time, industry or ownership category will be of major interest. Alternatively, we could develop weighting schemes that measure the externalities from innovation more directly (as in Jaffe's [1986] construction of spillover measures).

The econometric problems likely to be encountered in using regression techniques to obtain precise estimates of the weights needed are similar to those which arise in distributed lag estimation.[4] However, the situation here seems more sympathetic to the econometrician than the typical distributed lag situation. First, there is, at least potentially, a wealth of patent data. Furthermore, if we are willing to assume some connection between the value of patent protection and the value of patented ideas,

[4] Two come immediately to mind. The first is that the within group error in the patent count measure of patent values will produce an error in variables problem in estimation. We note that in cases where R&D data are also available they are a potential instrument. The second is the fact that the $\{N_j\}$ sequences are likely to be highly correlated across observations (especially in studies where a major source of the variation in the data is across time). This is likely to produce a precision problem similar to the problem we often find in distributed lag estimation.

there is substantial prior information on the structure of the weights that ought to be of significant help in estimating them. We discuss this latter possibility in more detail now.

A starting point is to assume that the average value of the ideas embodied in the patents in a particular group is proportional to the value of patent protection in that group. If we then assume that the distribution of the value of patent protection is estimated correctly by our renewal and/or application model, we can obtain the weights needed for Equation (1) up to a factor of proportionality from working out the implications of our parameter estimates. As yet there are no results available where both the choice of countries in which to apply for a patent and the subsequent renewal decisions are analyzed simultaneously. Consequently we review estimates of weights for groupings determined by the length of time renewed in particular countries and for groupings determined by the countries in which a given patent was applied for separately.[5]

Table II summarizes this information. Columns 1, 2 and 3 are taken from Lanjouw [1998] and are calculated for the 1975 cohort of German computer and pharmaceutical patents (the paper gives estimates for four technology groups based on thirty cohorts of data). Columns 4 and 5 are computed from Pakes [1986] and are based on all patents applied for between 1951 and 1979 in France and on all patents granted that were applied for between 1952 and 1972 in Germany. The weights in all but the second column sum to one (we come back to Column 2 below).

Lanjouw's estimates imply that the average computer patent lapsing at age four is worth three times as much as one lapsing at age three. All of Lanjouw's estimates show patent weights growing approximately linearly in lapse age until the group of patents that survive to the maximum age of 20 years. The latter group are about 50% more valuable, on average, than those that survive 19 years. Pakes' estimates show somewhat lower weights for patents allowed to lapse in the early ages, especially the weights based on the French data. Thereafter they grow at a fairly steady rate, again until the group of patents that renew until the statutory limit. These are estimated to be 100% more valuable than the patents that lapse the year before the statutory limit. Note that all columns imply that, were we to attempt to estimate the weights econometrically, a three or four parameter family would suffice.

Most of the differences between columns are explainable. For example, the smaller weights for patents that lapse early on in France might have

[5] A more complete database would allow a two-way classification of patents into countries of application and the age at which the patent was allowed to lapse in each country. Since inventions with patents in more countries are typically renewed longer in each country (Putnam [1991]), weights for a two-way classification cannot be built up from the 'marginals' of the two one-way classifications.

TABLE II
ESTIMATED MEAN VALUE WEIGHTS BY AGE OF LAPSE AND BY SIZE OF PATENT FAMILY[a]

Age of Lapse	Lapse Age Weights						Family Size Weights	
	Computers		Pharmaceuticals	France	Germany[b]		Number of Countries	Family Size Weights
	(1)	(2)	(3)	(4)	(5)			(6)
3	.004	.017	.003	.000	.000		2	.000
4	.012	.023	.016	.000	.000		3	.002
5	.018	.028	.022	.001	.001		4	.003
6	.022	.032	.028	.001	.003		5	.004
7	.026	.037	.029	.002	.005		6	.006
8	.027	.042	.035	.003	.008		7	.008
9	.032	.046	.040	.005	.012		8	.011
10	.033	.052	.036	.007	.017		9	.016
11	.035	.058	.039	.010	.024		10	.020
12	.038	.065	.045	.014	.034		11	.026
13	.046	.072	.043	.019	.046		12	.034
14	.049	.080	.049	.029	.056		13	.044
15	.056	.089	.049	.043	.075		14	.066
16	.061	.100	.057	.058	.100		15	.092
17	.072	.112	.070	.078	.125		16	.137
18	.084	.126	.080	.100	.149		17	.177
19	.103	.136	.110	.123			18	.354
20	.111	.159	.101	.152	–			
No lapse	.169		.146	.357	.344			

Notes:
[a] Columns 1–5 give weights for patents that lapse in the indicated age (Column 2 is the truncation weight for patents renewed *through* the indicated age) or that are filed in the indicated number of countries (Column 6).
[b] In the period investigated by Pakes the statutory maximum length of German patents was 18 years; during the period investigated by Lanjouw the maximum changed from 18 to 20 years.
Sources: Columns 1–3 are estimates for Germany derived from Lanjouw [1998]. Columns 4–5 are derived from Pakes [1986]. Column 6 is derived from Putnam [1996].

been expected from the fact that the French data contain all patent applications while the German data contain only those patents which were eventually granted. The differences across columns in the weight given to patents renewed until the statutory term limit is a bit more troubling. Patents which are renewed until the statutory term limit are an 'open ended' group. That is, since there never is an observed renewal fee that induces these patents to drop out, there is a sense in which we never have an upper bound to their returns. (See the non-parametric analysis in Pakes and Simpson [1989].) The model determines the value of these patents from a combination of the fact that the possibility of a patent becoming high valued has an effect on initial renewal decisions and from functional form assumptions. Thus we might expect this weight to be more model dependent than the others—which is just what we find.

The weights in Columns 1, 3, 4 and 5 apply directly to cohorts of patents which are older than the statutory term limit. If we also want to weight the patents in younger cohorts, cohorts that are truncated from the right (that is, for which we have not yet seen the entire sequence of lapse proportions), we will need weights for combined groups ('truncation weights'). These are based on the mean value of patents which are renewed up through a given age. Column 2 provides an example of the needed figures. Comparing Columns 1 and 2 we see that the average value of all computer patents renewed through age three is just over four times the size of the mean value of those lapsing in age three (0.017 versus 0.004). If a cohort were only ten years old, we could assign the weights from Column 1 to inventions whose lapse we have actually observed and the age ten weight of 0.052 from Column 2 to all patents still in force in age ten.

Column 6 presents corresponding estimates from Putnam [1996] for patent families. These show the relative mean value of patent families of different sizes.[6] For example, the average patent family with applications filed in four countries was worth about one and a half times as much as one with applications in three countries. These weights grow approximately log-linearly, again excepting the weight for 18-country families. These are the largest families in his data and are estimated to be approximately twice as valuable, on average, as a 17-country family.

We close this section by noting that the complementary data needed to fully exploit better indices of the value of innovation are largely available. For example, the concordance in Kortum and Putnam [1997] enables the mapping of patents from the technology classification used by the patent examiners to a standard industrial classification (both by the industry of use and by the industry of origin of the patent) and, therefore, the construction of patent value indices by industry. Together with census and survey of manufacturing data available for the standard classification of industries, this allows one to study innovation at the industry level of aggregation. Also, the ownership information in patent documents allows the matching of our patent value indices to input data by the source of funding (private, government funded and by agencies within the government funded category) and by the location of the R&D activity (universities, national research labs, private firms, joint ventures and so on). This makes it possible to examine the efficacy of different institutional arrangements for the production of innovation. By combining weighted patent value indices with the citation information which is also included in the patent documents (see Trajtenberg [1990]; Jaffe, Henderson, and Trajtenberg [1993]), values could be attached to the spillovers that lie at

[6] Putnam [1996] also gives a fuller set of weights which depend not only on the number of patents in patent families but also on the country composition of the patent applications. We have chosen this aggregation to simplify the presentation.

the heart of many of the public policy issues associated with R&D policy. Patent documents also indicate the number of International Patent Classification subclasses to which a patent is assigned. With patent value indices it is possible to examine whether this measure of the breadth of an innovation is indeed related to its private or social value (Lerner [1994]).

We now move on to a consideration of the use of the patent renewal and application data in studies of the value of patent protection and IPR policy.

IV. THE VALUE OF PATENT PROTECTION AND IPR POLICY

As discussed in the introduction, the estimation of patent renewal and applications models gives us some of the few pieces of information available on the value of protection itself, as distinct from the value of the underlying innovation. We now consider what insights have been gained about the nature of the learning and obsolescence processes. We then discuss what has been learned about the distribution of the value of patent protection and its evolution over the lifespan of a group of patents. The section closes with a discussion of how the results from these models can be used to understand the incentives generated by patent systems and how they can be used to analyze the effects of intellectual property policy reforms.

Learning and Depreciation

Parameter estimates derived from the renewal models are informative about specific features of the innovative process, such as the speed of learning, exploitation and obsolescence of innovations. One of the early interests in estimating models of the renewal decision was to obtain a measure of the rate of obsolescence of the private returns to innovations. This rate of obsolescence is needed to weight R&D investments in the construction of knowledge 'stocks' (analogous to physical capital stocks) in the analysis of the private returns to R&D investments and then again to convert estimates of gross rates of return to rates of return net of depreciation costs. The early deterministic models of renewals found rates of obsolescence in the private value of patents that were considerably higher (on the order of 25%) than the depreciation rates typically used in the construction of physical capital stocks. This has marked implications for comparisons of estimates of the private and social rates of return to innovative activity. (See Griliches [1979] or Pakes and Schankerman [1984] for details on these points.)

As noted in Section II, the later, stochastic, renewal models allow for greater variation in the pattern of returns over time. In these models, returns may depreciate over the life of a patent, but they may also increase, particularly in the initial years, as the innovation is developed and the patentee learns about the market. From aggregate data, Pakes [1986] finds that most learning is over by the fifth age of protection, and few patents yield higher returns after that point. With disaggregated data and a model which allows for zero returns in the early ages, Lanjouw [1998] obtains a similar result, with the bulk of learning completed after four years and most learning being over after seven years. In the learning period, many patentees discover that their patents are worthless, while others learn that substantial returns can be earned from their patents. After this period, returns decay more or less at a common rate.

The Magnitude and Distribution of Value

Currently, renewal model estimates of patent values are available for Germany, France, and the UK at an aggregate level (Schankerman and Pakes [1986]; Pakes [1986]; Sullivan [1994]). For France, Germany and India they are also available disaggregated by the nationality of the patentee and by the type of technology (Schankerman [1998]; Duguet and Iung [1997]; Lanjouw [1998]; Fikkert and Luthria [1996]).

Estimating a deterministic model, Schankerman and Pakes [1986] find that the average value of a patent from cohort 1970 *applied for* in France and the UK was quite low at 11,250 DM (1975 Deutchmarks). In Germany the average value of a patent *granted* was about 27,300 DM. (just over a third of the patents applied for in Germany were granted). In all countries the distribution of values was very skewed. One percent of applications in France and the UK had values in excess of 112,500 DM, while in Germany one percent of granted patents had values above 193,000 DM. Disaggregated value estimates for the 1975 German cohort are in Table III (from Lanjouw [1998]). Consistent with the results noted above, they indicate quite a wide variation across technologies in the average value represented by a single patent (from 17,500 to 49,700 1975 DM) and a considerable skew in the distribution of value. Similar results were obtained for France by Schankerman [1998].

The applications model estimated by Putnam [1996], described above in Section II, gives us the first quantitative measure of the total value of the global patent rights obtained on an innovation (in contrast to the value obtained as a result of the protection afforded by the patent laws in a given country). This is particularly interesting because it allows us, for the first time, to measure the flows of the value of patent rights across borders and therefore to answer the question of the extent to which different countries benefit disproportionately from the existence of the international patent system.

TABLE III
VALUE DISTRIBUTIONS BY TECHNOLOGY
GERMANY—COHORT 1975

Computers			Pharmaceuticals			Textiles			Engines		
Percentile	Value [a,b]	Cumulative % of Total	Percentile	Value	Cumulative % of Total	Percentile	Value	Cumulative % of Total	Percentile	Value	Cumulative % of Total
50%	13.0	9.7%	50%	10.6	4.6%	50%	8.5	7.6%	50%	33.2	15.4%
	(1.3)			(1.4)			(2.0)			(3.3)	
75	32.0	32.2	75	35.0	23.6	75	23.3	28.9	75	66.4	39.4
	(2.9)			(3.3)			(4.9)			(6.1)	
90	59.7	59.8	90	77.7	51.9	90	46.3	56.9	90	115.0	65.3
	(5.2)			(6.6)			(9.6)			(10.2)	
95	83.3	74.6	95	115.5	68.9	95	66.1	72.3	95	153.7	78.5
	(7.1)			(9.4)			(13.2)			(13.0)	
99	143.8	92.3	99	210.1	90.3	99	119.3	91.6	99	249.0	93.7
	(11.8)			(16.7)			(23.0)			(20.5)	
99.9	245.6	98.9	99.9	332.4	98.6	99.9	191.0	98.8	99.9	433.0	99.0
	(19.9)			(26.2)			(34.4)			(35.0)	
Mean Value		23.5 DM		27.4 DM			17.5 DM			49.7 DM	
		(2.1)		(2.6)			(3.7)			(4.6)	
# of Granted Patents		1,172		1,251			1,069			1,594	

Notes:
[a] Values are net of annual renewal and administration fees, as well as 476 DM for the application, examination and publication costs (assuming examination at age 7 and publication at age 9). Calculations use 15000 simulation draws. All values are in thousand 1975 Deutschmarks.
[b] Estimated standard errors for the value percentile (*vper*) estimates due to error in parameter estimates $\hat{\omega}_N$ are calculated using a Taylor approximation: $vper(\hat{\omega}_N) \approx vper(\hat{\omega}_0) + \Gamma(\hat{\omega}_0)'(\hat{\omega}_N - \omega_0)$. The unknown gradient matrices $\Gamma(\omega_0)$ are approximated with central finite difference gradients calculated at the point $\hat{\omega}_N$. Thus $\hat{\mathrm{Var}}(vper(\hat{\omega}_N)) = \hat{\Gamma}(\hat{\omega}_N)'(\hat{\mathrm{Var}}(\hat{\omega}_N))\hat{\Gamma}(\hat{\omega}_N)$. Similarly for the estimated standard errors of the mean values.
Source: Lanjouw [1998] Table 4].

TABLE IV
MEAN VALUES OF PATENTS GRANTED AND HELD
SHARE OF R&D REPRESENTED BY INTERNATIONAL PATENT HOLDINGS[a]

TOP 5 COUNTRIES
1974 PATENT APPLICATION COHORT

Country	Mean value Granted	Relative to GDP[b]	Mean value Held	Share of National R&D
US	449 DM	0.12	231 DM	14.3
Japan	164	0.13	242	21.4
Germany	277	0.26	185	34.0
France	201	0.24	164	19.6
UK	172	0.31	180	32.3

Notes:
[a] Values are expressed in thousands of 1975 DM
[b] (Mean value of patent) * 1,000,000/GDP.
Source: Putnam [1996] and *IMF International Financial Statistics* [1977].

Table IV presents estimates of the mean value of patents granted and held in each of the top five OECD countries. (Value is again calculated as the present value of annual returns to patent protection, net of application and renewal fees, with renewal decisions made optimally.) The data are restricted to include only inventions for which patent applications were made in more than one country. In 1974 this was about 28% of all inventions reaching patent offices in the OECD. The selection of only internationally-filed inventions eliminates the lowest valued patents from the analysis. As a result, the figures in Table IV are not directly comparable to those given above for renewal models (as the latter include inventions protected in a single country).

Column 1 of Table IV shows the average value of a patent granted in each country. Column 2 indicates how the value of protection relates to the size of the market. This varies substantially. For example, while Germany and Japan have economies of similar size, there is a large disparity between the expected value of protection in the two countries. The average internationally-filed patent granted by Japan appears to have been worth less than in other countries (about 164,000 1975 DM, while the corresponding value in Germany is 277,000 DM). Although some of this could be due to compositional differences (for example, a higher share of chemical and pharmaceutical inventions in Germany), Putnam's results suggest that inventors perceive Japanese patent protection to be worth less, holding the 'quality' of the invention constant, than a German patent.

Column 3 of Table IV provides the average value of a patent held by citizens of each country. Japan is again an outlier, this time holding patents whose value is higher, on average, than those held by citizens of other major countries. This feature of Japanese patenting is associated with two factors: (1) a relatively large share of Japanese inventions are filed abroad (particularly given its distance from other countries) and (2) a disproportionately large fraction of their patents are filed in the US, a country which, with its large market, grants the most valuable patents (see Column 1).

Table V presents a subset of the patent 'trade' results from Putnam [1996]. For each of the three patent granting countries under consideration, the table presents the following five columns of data: the numbers of patents granted by the country to citizens of other specified 'source' countries; the percentages these numbers represent of the total patents granted by the country to foreigners; the estimated mean value of the associated patent rights; the total value of the patent rights granted to citizens of the source countries and lastly the share these represent of the total value of patents granted by the country to foreigners. Each row of the table represents a different source country. Thus, the table is a 3×3 matrix of 5-column blocks; the i, j^{th} block of the matrix gives information about patents originating in country i and patented in country j. Because the model applies only to patents filed internationally, the diagonal blocks of the matrix are omitted.

The table shows, for example, that among inventions whose first filing occurred in 1974, Japanese inventors obtained 5,239 patents in the United States, while US inventors obtained 3,836 patents in Japan. The mean value of a Japanese-origin patent in the US was about 325,000 1975 DM, while that of a US-origin patent in Japan was about 148,000 DM. The total value of patents granted by the US to Japan was estimated to be about 1,703 million DM, while Japan granted about 569 million DM worth of patents to US inventors. Thus the US 'trade deficit' with Japan amounted to over 1,134 million DM in 1974 (which represents almost 90% of its total deficit with the developed world). Note that in contrast to the changes in measures over time which were discussed in Section III, in this case international differences in counts seem to *underestimate* differences in value.

IPR Policy

Recall that one of the main motivations for maintaining an intellectual property rights system is to increase the extent to which inventors can capture the returns to investments in R&D. The ratio of the total private value of patent protection to related R&D expenditures is one measure of the 'implicit subsidy' created by a patent system. This can be compared to other incentive policies such as R&D tax breaks or direct government funding of research.

Using results from disaggregated models of returns in France and Germany, respectively, and using trade data to determine the portion of

TABLE V
THE INCIDENCE AND VALUE OF INTERCOUNTRY PATENTING[a]

1974 PATENT APPLICATION COHORT

	Granting Country														
	US					Japan					Germany				
Source Country	No. of patents	% of total	Mean value	Total value	% of total	No. of patents	% of total	Mean value	Total value	% of total	No. of patents	% of total	Mean value	Total value	% of total
US			–			3836	41.0	148	569	37.5	9658	33.8	239	2307	39.8
JP	5239	27.5	325	1703	18.0			–			4415	15.5	201	892	15.4
DE	4838	25.5	637	3076	32.5	2275	24.3	164	373	24.6			–		

Note:
[a] These data are based on patents filed internationally, that is, in at least one country outside the home country. Mean values are expressed in thousands and total values in millions of 1975 DM.
Source: Putnam [1996].

world industry R&D expenditure which is related to patent protection in those countries, Schankerman [1998] and Lanjouw [1998] find subsidy rates in the order of 10–15% for most technology areas. Similar rates are found for aggregate data by Pakes [1986] and Schankerman and Pakes [1986].

Putnam [1996] calculates the ratio of the value of *international* patent rights held worldwide by inventors in a country to the *total* R&D expenditures by that country. Table IV, Column 4, shows the implicit subsidy rates he finds for five countries. They lie between 14 and 34 percent. Both Germany and the UK appear to benefit significantly more than the other countries from patenting abroad; they received an 'implicit subsidy' amounting to about 33% of their R&D. The figure for France and Japan was about 20%. The stock of international patent rights held by US inventors, despite being the largest stock worldwide, was just 15% of the value of US R&D expenditure. Note that the subsidy rate that accounts for the international value of patent rights is higher than that based on the value of patent rights in a single country.

The value of the legal right to the exclusive use of an innovation and hence the incentives created by a patent system depend on features of the system which can, and do, vary across countries and time. For example, statutory patent terms range from zero years (excluded products, such as pharmaceuticals in many countries) to 20 years (EU; soon to be joined by all members of the WTO). Application and renewal fees have also varied widely—the latter from zero (the US until recently) to over 3,300 DM (for the twentieth year of protection in Germany). By affecting the duration of patents, these variables not only help determine the private incentives to invest in research, but they also affect the social costs of allowing patent monopolies. To date, fees and term limits have been set in an *ad hoc* fashion, usually to cover patent office costs. However, there is a potential for these features to be designed so as to create a given level of expected private returns at a lower social cost. (For theoretical discussions see Nordhaus [1969] for an early investigation of optimal statutory limits, and Cornelli and Schankerman [1996], who consider the choice of renewal fee schedules.) Estimates from structural renewal or applications models can be used to obtain empirical measures of one side of this tradeoff: how changes in the fee schedule and the statutory term influence the private value of patent rights and hence the incentives to invest in research.

Enforcement, and the costs of enforcement, also have a bearing on the benefits of patent protection. The lack of enforcement figures highly in IPR negotiations with developing country governments, while the high costs of prosecuting infringements is receiving a great deal of attention elsewhere. The US has seen sweeping changes to its IPR law and practice since the early 1980s, including the establishment of a new federal court of appeals to hear patent cases. Renewal data can be used to measure how sensitive the returns to patenting are to various changes in the law or legal policy. Lanjouw [1996] embeds an infringement/litigation model into a renewal model to take into account the fact that patentees must be willing to enforce their rights for protection to be meaningful. While in the absence of data related to filed court cases this approach relies heavily on structural assumptions, it provides the only measure to date of the 'hidden' effects of legal policy changes on the number of patents and their duration—effects which are over and above changes in direct legal costs. Moreover, legal data are increasingly available and have recently been matched to patent data. (See Lanjouw and Schankerman [1997] for an analysis of the characteristics of litigated patents.)

Tables VI and VII present estimates of the effects of specified changes in fees, statutory term limits and of legal policy on the benefits of patent protection. In each table the distribution of the value of protection for the 1975 cohort of German computer patents is simulated several times. In the first case, the simulation is based on the parameter estimates from a renewal model and the actual characteristics of the IPR system in Germany—with legal fees paid by the losing party (the English Rule, ER) and a statutory term limit of 20 years. The first column in each table shows the mean and percentiles of this value distribution, net of application and renewal fees. Subsequent columns show simulated value distributions after a policy change. (For details on the construction of these estimates and further simulated policy experiments, see Lanjouw [1996].) Since we have not yet estimated a model which analyzes the application and subsequent renewal decisions jointly, the results in these tables *condition* on the number of applications in a particular cohort and then ask how the net returns to protection for those patents (returns minus renewal fees for the years the patents are in force) change under alternative institutional arrangements. That is, a more complete model would allow the quantity and quality of applications to also vary with the institutional arrangement, but such a model has yet to be estimated.

Table VI shows how patent value and the average revenue collected by the patent office change with the alterations in renewal fee schedule. The reform being considered is indicated in bold type at the top of each column. As an example, we examine, in Column 3 of the table, the effect of switching to a schedule, denoted CS, which may correspond more closely to an optimal design (based on the analysis in Cornelli and Schankerman [1996]).[7]

[7] In their model, firms are heterogenous in R&D productivity and, as a result, social welfare can be enhanced by a fee schedule which encourages relatively more R&D to be done by the high productivity firms. This result is reached in their model when fee schedules are low in the first years and increase substantially more rapidly than existing schedules in the latter years.

TABLE VI
SIMULATED VALUE DISTRIBUTIONS FOR COHORT 1975
COMPUTERS (1,172 PATENTS)

Simulation	(1)	(2)	(3)	(4)
Variables – (Bold are changes from Base level)				
If – fee shifting	ER	ER	ER	ER
T – statutory limit	20	20	20	20
c – renewal fees	Base	**(0.9)*Base**	**CS**[a]	**zero**
Value Distributions[b]				
Mean	24,329 DM	24,568 DM	25,577 DM	27,948 DM
	(2,147)			
Percentile:				
50% (median)	13,922 DM	14,092 DM	15,868 DM	17,733 DM
	(1,295)			
75	32,686	32,949	33,602	37,207
	(2,881)			
99.9	239,868	240,699	235,983	247,799
	(19,944)			
Average Fees Paid	1,606 DM	1,511 DM	539 DM	0 DM
Decrease in Average Fees	–	95 DM	1,067 DM	1,606 DM
Increase in Value Gross of Fees	–	144	181	2,013
Total Increase in Net Value	–	239	1,248	3,619
Percent of Base Value	–	101.1%	105.1%	114.9%
Elasticity	–	−0.10	–	–

Notes:
[a] CS refers to renewal fees of zero for ages 1–10, increasing annually by 500 DM thereafter.
[b] Values, in 1975 DM, are net of annual renewal and administration fees, as well as application, examination and publication costs (assuming examination at age 7 and publication at age 9). Calculations use 10000 simulation draws. Estimated standard errors for the value (*vper*) estimates calculated using a Taylor approximation: $vper(\hat{\omega}_N) \approx vper(\omega_0) + \Gamma(\omega_0)'(\hat{\omega}_N - \omega_0)$. Gradient matrices $\Gamma(\omega_0)$ are approximated with central finite difference gradients calculated at $\hat{\omega}_N$.
Source: Lanjouw [1996, Table 3].

TABLE VII
SIMULATED VALUE DISTRIBUTIONS FOR COHORT 1975
COMPUTERS (1,172 PATENTS)

Simulation	(Base)[a]	(2)	(3)	(4)	(5)	(6)	(7)
Variables – (Bold are changes from Base level)							
If – fee shifting	ER	**AR**	**AR**	**AR**	**AR**	ER	ER
T – statutory limit	20	20	**17**	**17**	**17**	20	20
w – win probability	.91	.91	.91	**.94**	.91	**.94**	.91
LF – legal fees	Base	Base	Base	Base	**1.25*Base**	Base	**1.25*Base**
Value Distributions[b]							
Mean	24,329 DM	20,503 DM	15,011 DM	15,326 DM	13,338 DM	25,292 DM	24,210 DM
Percentile							
50% (median)	13,922 DM	5,980 DM	5,096 DM	5,096 DM	5,096 DM	15,039 DM	13,802 DM
75	32,686	27,682	11,296	11,399	10,778	33,192	32,651
99.9	239,868	237,037	232,285	232,285	227,687	239,868	239,868
Percent of Base Value	100%	84%	62%	63%	55%	104%	100%
Elasticity	–	–	–	0.64	−0.45	1.20	−0.02

Notes:
[a] 'Base' refers to parameter values estimated from the renewal model and other data. The other cases are selected changes.
[b] Values, in 1975 DM, are net of annual renewal and administration fees, as well as application, examination and publication costs (assuming examination at age 7 and publication at age 9).
Calculations use 10000 simulation draws.
Source: Lanjouw [1996].

The benefit to patentees of a decrease in renewal fees is twofold: they pay less for each year of protection and, because it is less costly, they also tend to take advantage of more years of protection. The first row in the bottom section of Table VI indicates the size of the first benefit, a benefit which is, of course, a direct cost of the reform in terms of lost revenue to the patent office. The second row gives the size of the second effect—the increase in patent value which is due to protection being sought for more years. Together, these figures give the total change in the average net value of patent protection that would result from the indicated policy reform.

Column 2 shows that decreasing the renewal fee schedule by ten percent has a very moderate impact on both patent office revenue and on the incremental benefit obtained by patentees. Changing to the alternative fee structure, CS, would cause a dramatic fall in average payments to the patent office—in the order of two-thirds, with, at the same time, a relatively small additional benefit to patentees (181 DM on average). Finally, the simulation in Column 4 demonstrates the effect of abolishing renewal fees altogether. Absent fees, the average value of a patent increases by 3,619 DM, or 15 percent, with a net increase in average value of 2,013 DM.

In Table VII each of Columns 2 to 7 demonstrate the substantial impact that legal policy reforms can have on the value of patent protection (with the set of policy changes again indicated in bold at the top of the columns). The first row in the bottom section of the table gives the mean value of protection after the indicated policy change as a percentage of the pre-reform mean value, while the second row gives the elasticity of response to the policy. For example, changing from a system where the loser pays the legal fees for both parties (the English Rule, ER) to one where litigants pay their own legal fees (the American Rule, AR) lowers the average value of a computer patent by 16 percent. Moving to the shorter term available in the US until recently takes another 22 percent off of the mean value. To put this in perspective, given the 21,515 patents in cohort 1975, the figures in the first columns of Table VII imply that moving from the current German system to one more closely resembling the American system would lead to a fall of about 200 million DM in the total value of protection received by inventors as a result of patents granted in Germany in a single year.

Increases in legal fees can also have a substantial impact on value—under the American system the elasticity of response is -0.45 (compare Columns 3 and 5). The level of legal fees makes less of a difference under the English system, with an elasticity of only -0.02 (compare Columns 1 and 7). On the other hand, patent value under the English system is very responsive to changes in the patentee's probability of winning at trial (compare Columns 1 and 6). Note that these changes in value are due solely to patentees renewing for fewer years—any increase in actual legal expenditures would have an additional, more direct, impact on net value.

V. CONCLUDING COMMENTS

The potential of patent renewal and applications data has only just begun to be exploited. However, these data have already yielded some useful insights.

The patent value estimates derived from a number of the models discussed above indicate that weighting patent counts, based on their years of renewal or family size, can go a long way toward removing the noise in simple patent counts. The estimates also suggest that the model used to estimate the appropriate weights can be restricted to a linear or log-linear 'slope' coefficient (at least after the first renewal year or two) with a free weight for the patents in the tail of the distribution. Parameter estimates from the models of the decision to renew have yielded estimates of the depreciation rate of innovation, suggesting that it is considerably more rapid than that typically assumed for physical capital. Similarly, the stochastic renewal models give some idea of the speed of learning during the development of new innovations. All estimates of the value of patent protection indicate a skewed distribution of returns. The level of average returns implies a noticeable (although not overwhelming) subsidy to R&D in single-country models, with a distinctly higher subsidy rate when global patent protection is considered. Finally, simulations suggest that policy reforms can have a significant impact on the distribution of returns.

There are many directions in which this type of analysis could be developed further. The most straightforward is to combine renewal and family information with other economic data. Then model parameters can be allowed to differ with various aspects of the economic environment which contribute to determining the value of innovations and the value of protecting them. Progress is being made in this direction by Duguet and Iung [1997] who have combined renewal data for French European Patents with a wide variety of firm-level and macro data for that country. For example, they use responses from an innovation survey to construct sectoral measures of the rate of imitation of new products and of the extent to which innovation is demand driven or pushed by technological factors. They then estimate a renewal model in which they allow the value of protection and the speed of depreciation to vary with these sectoral characteristics.

These models can be extended to consider policy interactions further. We have already discussed some initial results regarding legal policy reforms. Current theoretical work is making progress in modelling the litigation and settlement process. These models can contribute to specifying renewal and application decision rules which reflect the legal environment. Just as the value of patent rights is determined by policies that effect their enforceability, the returns to patent protection may be strongly affected by other policies. For example, the position taken by governments regarding the enforcement of anti-trust regulation in situations involving patent monopolies has varied over time and across countries (see Gallini and Trebilcock [1995] for a discussion). Incorporating information about the type and timing of these changes in a renewal/application model would allow us to obtain a measure of the size of their effect on the incentives generated by the patent system.

The databases relevant to this type of analysis are improving very rapidly in quality and size as countries computerize their patent offices and court systems. This should allow the estimation of more finely differentiated models. In addition, the growth in our computational abilities has been nothing short of spectacular. This opens up further potential avenues to extend the current work. For example, models could be developed which recognize the possible interdependence between the returns that an inventor receives from patent protection in one country and those that he receives from protection in another, an interdependence that would be reflected in both his application and renewal decisions in both countries. The increase in both data and computational power should also enable us to make significant progress in non-parametric analyses of patent value distributions and in exploring the robustness of results to modelling assumptions.

ACCEPTED MARCH 1998

REFERENCES

Cornelli, F. and Schankerman, M., 1996, 'Optimal Patent Renewals', The Economics of Industry Group Discussion Paper no. EI/13. STICERD. The London School of Economics.

Duguet, E. and Iung, N., 1997, 'R&D Investment, Patent Life and Patent Value: An Econometric Analysis at the Firm Level', Document de travail G 9705. Institut National de la Statistique et des Etudes Economique.

Evenson, R., 1984, 'International Invention: Implications for Technology Market Analysis' in Z. Griliches (ed.) *R&D, Patents and Productivity*, Chicago, University of Chicago Press.

Fikkert, B. and Luthria, M., 1996, 'Estimates of the Private Value of Patent Protection in India', Mimeo. University of Maryland.

Gallini, N. and Trebilcock, M., 1995, 'Intellectual Property Rights and Competition Policy', Mimeo. University of Toronto.

Griliches, Z., 1979, 'Issues in Assessing the Contribution of Research and Development to Productivity Growth', *Bell Journal of Economics*, Vol. 10, no. 1 (Spring).

Griliches, Z., 1990, 'Patent Statistics as Economic Indicators: A Survey', *Journal of Economic Literature*, Vol. XXVIII, pp. 1661–1707.

Harhoff, D., Scherer, F. M. and Vopel, K., 1997, 'Exploring the Tail of Patent Value Distributions.' Mimeo. Harvard University.

Jaffe, A., 1986, 'Technological Opportunity and Spillovers of R&D: Evidence from Firms' Patents, Profits, and Market Value', *American Economic Review*, Vol. 76, no. 5, pp. 984–1002.

Jaffe, A., Henderson, R. and Trajtenberg, M., 1993, 'Geographic Localization of Knowledge Spillovers As Evidenced By Patent Citations', *Quarterly Journal of Economics*, Vol. 108, pp. 577–598.

Jones, C. and Williams, J., 1996, 'Too Much of a Good Thing? The Economics of Investment In R&D', Mimeo. Stanford University.

Kleinknecht, A., 1996, *Determinants of Innovation: The Message from New Indicators*, New York, McMillan.

Kortum, S., 1997, 'Research, Patenting, and Technological Change', *Econometrica*, Vol. 65, no. 6, pp. 1389–1420.

Kortum, S. and Putnam, J., 1997, 'Assigning Patents to Industries: Tests of the Yale Technology Concordance', *Economic Systems Research*, Vol. 9, no. 2, pp. 161–175.

Kuznets, S., 1962, 'Inventive Activity: Problems of Definition and Measurement', pp. 19–53, in R. R. Nelson (ed.) *The Rate and Direction of Inventive Activity*, Universities NBER Conference Series No. 13, Princeton University Press.

Lach, S., 1995, 'Patents and Productivity Growth at the Industry Level: A First Look', *Economics Letters*, Vol. 49, pp. 101–108.

Lanjouw, J. O., 1998, 'Patent Protection in the Shadow of Infringement: Simulation Estimations of Patent Value', *The Review of Economic Studies*, (forthcoming).

Lanjouw, J. O., 1996, 'Beyond Lawyers' Fees: Measuring the Indirect Effects of US Intellectual Property Reforms', Mimeo. Yale University.

Lanjouw, J. O., 1992, *The Private Value of Patent Rights: A Dynamic Programming and Game Theoretic Analysis of West German Patent Renewal Data, 1953–1988*, Ph.D thesis. London School of Economics.

Lanjouw, J. O. and Schankerman, M., 1997, 'Stylized Facts of Patent Litigation: Value, Scope and Ownership', NBER Working Paper no. 6297.

Lerner, J., 1994, 'Patenting in the Shadow of Competitors', *Journal of Law and Economics*, Vol. 38, no. 2, pp. 463–496.

Nordhaus, W. D., 1969, *Invention, Growth, and Welfare: A Theoretical Treatment of Technological Change*, Cambridge, M.I.T. Press.

Pakes, A., 1986, 'Patents as Options: Some Estimates of the Value of Holding European Patent Stocks', *Econometrica*, Vol. 54, pp. 755–784.

Pakes, A. and Griliches, Z., 1980, 'Patents and R&D at the Firm Level: A First Report', *Economics Letters*, Vol. 5, pp. 377–381.

Pakes, A. and Schankerman, M., 1984, 'The Rate of Obsolescence of Patents, Research Gestation Lags, and the Private Rate of Return to Research Resources', in Z. Griliches (ed.) *R&D, Patents and Productivity*, NBER Conference Series. Chicago, The University of Chicago Press.

Pakes, A. and Simpson, M., 1989, 'Patent Renewal Data', *Brookings Papers: Microeconomics*, pp. 331–410.

Putnam, J., 1991, 'What Can We Really Learn from the (International) Patent System?' Mimeo. Yale University.

Putnam, J., 1996, *The Value of International Patent Protection*, Ph.D. Thesis. Yale University.

Sanders, B., 1962, 'Some Difficulties in Measuring Inventive Activity', pp. 53–92, in R. R. Nelson (ed.) *The Rate and Direction of Inventive Activity*, Universities NBER Conference Series No. 13, Princeton University Press.

Schankerman, M., 1998, 'How Valuable is Patent Protection: Estimates by Technology Fields', *The RAND Journal of Economics*, (forthcoming).

Schankerman, M. and Pakes, A., 1986, 'Estimates of the Value of Patent Rights in European Countries during the Post-1950 Period', *The Economic Journal*, Vol. 96, pp. 1052–1076.

Schmookler, J., 1962, 'Comment', in R. R. Nelson (ed.) *The Rate and Direction of Inventive Activity*, Universities NBER Conference Series No. 13, Princeton University Press.

Sullivan, R., 1994, 'Estimates of the Value of Patent Rights in Great Britain and Ireland, 1852–1876', *Economica*, Vol. 61, pp. 37–58.

Trajtenberg, M., 1990, 'A Penny for Your Quotes: Patent Citations and the Value of Innovation', *The Rand Journal of Economics*, Vol. 21, no. 1.