

NBER WORKING PAPER SERIES

PATENTS, CITATIONS AND INNOVATIONS:
TRACING THE LINKS

Manuel Trajtenberg

Working Paper No. 2457

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
December 1987

I am grateful to Tim Bresnahan and Zvi Griliches, and to participants in the productivity workshop of the 1987 NBER Summer Institute, for helpful comments on an earlier draft. This research was supported by funds granted to the Foerder Institute for Economic Research by the NUR MOSHE FUND. The research reported here is part of the NBER's research program in Productivity. Any opinions expressed are those of the author and not those of the National Bureau of Economic Research.

Patents, Citations and Innovations: Tracing the Links

ABSTRACT

The goal is to tackle anew the main problems encountered in using patent data in economic research, namely, the large variance in the value of patents, and the difficulties in matching patents with economic categories. The first is addressed with the aid of patent citations, the second with computerized search techniques for large databases. The proposed solutions are applied to the case of Computed Tomography (CT) Scanners, a pathbreaking innovation in medical technology. The main findings are that patents weighted by citations are highly correlated with the value of innovations, and that important innovations generate further innovative activity (R&D), and hence bring about down-the-line patents.

Manuel Trajtenberg
Department of Economics
Tel-Aviv University
Tel-Aviv 69978
Israel

1. INTRODUCTION

The study of technological change has been hampered all along by the scarcity of appropriate data and, in particular, by the lack of good indicators of innovation having a wide coverage. Patents are the one important exception: these legal documents, going back in some countries two centuries and more, remain the only observable manifestation of inventive activity, covering virtually every field of innovation. However, their use in economic research has been impeded by two formidable problems: first, the enormous variance in the 'importance' or 'value' of patents, and second, the lack of a satisfactory matching between the patent classification system and any set of meaningful economic categories.

The main goal here is to show that one can now bring powerful tools to bear on these problems, and hopefully tackle them successfully. The first can be addressed with the use of patent citations, which appear to be highly informative of the value of innovations. The second with the aid of now widely available computerized search techniques in large databases: those techniques enable the researcher to easily identify and retrieve all the patents issued in narrowly defined 'product classes' of his choice, thus doing away with the need to rely on questionable classification schemes (at least for case studies or small panels).

The proposed methodological solutions are discussed in the context of a particular application, namely, the history of technical advance in Computed Tomography (CT) Scanners, one of the most remarkable innovations in medical technology of recent times. Using the search techniques just mentioned, I gathered the complete set of patents issued in this field, since its advent in

1972, and through the end of 1986 (456 patents in all). With the aid of additional data on CT from a previous study, I analyse the role that patent statistics may come to play as indicators of innovation. In particular, it is shown that patents weighted by a citations-based index can serve as a good proxy for the **magnitude** (or **value**) of innovations occurring over time; on the other hand, simple patent counts are indicative only of the extent of **innovative activity**, as reflected for example in R&D outlays. Pertinent to the recently reawaken demand-pull vs. supply-push debate, I also show that important innovations seem to generate further innovative activity, and hence to bring about down-the-line patents.

In light of those results, it appears that one of the most fertile uses of patent data may be in the study of the **emergence** of new, technological advanced sectors: innovation plays of course a crucial role in the dynamics of those sectors, and moreover, most of the innovation experienced by them over time seems to occur during their initial stages. The advantage of patents and patent citations is that they can be traced right to the very beginning of a sector, whereas conventional industry data usually commence much later, when the thrust of the 'action' may well be over.

One of the first regularities found is that the average number of citations per patent decreases drastically over the years, raising the suspicion that **time** (or **age**) itself might be the main determinant of citations, rather than anything like 'importance'. In Appendix B I put forward a statistical procedure to test for 'age vs. importance', and for truncation biases. Applying them to the case studied here, the objection is strongly rejected, thus legitimizing the use of citations-based indicators.

The paper is organized as follows: section 2 discusses the problems with the use of patents in research, and outlines the proposed solutions. Section 3 lays out the patent data in CT, and section 4 the additional measures of innovation from the previous study. The core of the empirical analysis is carried out in section 5, whereas section 6 makes the case for the existence of supply-push effects, and suggests other extensions. Finally, the results are summarized and interpreted in section 7.

2. USING PATENT DATA: AN OVERVIEW OF THE MAIN ISSUES

2.a *The Classification Problem*

The use of patents in economic research for any purpose requires, first of all, that we relate them to economic categories of interest, such as products, industries, countries, etc. What makes this task extremely problematic is the enormous size of the Patent File, which consist nowadays of over 4 million US patents, and some 25 million worldwide (about 70,000 new patents are issued annually in the US). There is of course a very elaborate Patent Classification System (PCS) in the US, comprising more than 110,000 patent categories (called subclasses), aggregated into some 400 patent classes, with no well-defined layers of aggregation in-between¹. The problem is that this vast classification scheme has not been designed according to a clear set of criteria (scientific or other), that could be mapped into an

¹ Mirroring the rapid pace of technological advance, the US-PCS is being changed and upgraded continuously, primarily by creating new patent subclasses. Thus, for example, at the time when Schmookler was engaged in his massive classification effort (in the early sixties), there were just about half the number of subclasses that we have today. There is also an international PCS, of similar dimensions and complexity.

economic system². Instead, the PCS has evolved first and foremost according to operational convenience, and with good reason: after all, its main goal is to assist patent examiners in establishing the patentability of new applications, by facilitating the search for related patents. Thus, it is inherently very difficult to create a satisfactory matching between those 110,000 patent sub-classes and any set of economic categories. To wit, the Patent Office constructed in the mid seventies such a 'concordance' between the US-PCS and the Standard Industrial Classification (at the request of the NSF - see OTAF [1975]), but the results have been so far quite dissappointing, primarily because most patent subclasses have been assigned to multiple industries in the SIC.

Aside from various ad-hoc solutions raised in the literature, there have been two major attempts in recent years to deal with the classification problem in a more general way, namely, that of Griliches [1984], and of Scherer[1984]. Griliches and his associates at the N.B.E.R. obviated the problem by using **firms** as the unit of analysis, rather than industry-type categories. The main drawback of that approach is that for many purposes firms are not the most appropriate units to look at, and the grouping of firms by SIC's is a very unsatisfactory solution. Scherer actually classified over 15,000 patents according to industry of origin and industry of use, by examining with a team of experts the contents of each patent. The problem of course is that such a massive and team-dependent enterprise cannot be put

²The Canadian Patent Office does classify patents according to economic categories during the review process, in addition to assigning them to a PCS. Note, however, that this is very different from trying **ex-post** to match patent sub-classes to, say, industries.

forward as an universal research strategy, but has to be regarded as a one-time event. Thus, and even though both projects made big strides in this area, the classification problem remained essentially unsolved.

As already suggested, I propose here a more modest approach to the problem, based upon the availability of powerful techniques for computerized search in large databases. As described in detail in Appendix A, those techniques allow one to identify quite easily all the patents issued in predetermined economic categories, and retrieve them for further analysis. With the aid of these properly classified patents, one can conduct in-depth studies of innovation in single sectors, or comparative studies involving, as of now, a not-too-large number of industries³. Clearly, this does not provide for an all-out solution to the classification problem, but it does open up a promising avenue of research 'in the small', whereas the feasibility of 'universal' solutions of any sort remains dubious.

2.b The problem of a large variance in the value of patents

The second major obstacle in using patents in economic research, resides in the well-known fact that patents vary enormously in their technological and economic 'importance' or 'value'⁴. Thus, the mere counting of patents, at any

³There is nothing in the nature of the approach that precludes the undertaking of more ambitious projects (e.g. large panels). However, and aside from requiring generous research budgets, more experience in searching by industries is needed before larger studies can be contemplated, and eventually conducted on a routine basis.

⁴Thus, for example, Scherer [1965] cites evidence to the effect that the distribution of patent 'values' is highly skewed toward the low end, with a very long and thin tail into the high value side. Moreover, he notes that those Pareto-type distributions might not have finite moments, implying that the 'mean value' of patents is a rather elusive magnitude, and that one should

level, cannot be regarded a priori as a good indicator of the 'amount' of innovation. As with the classification issue, the problem of a large variance is inherent to the patent system as such⁵, and therefore definite solutions can hardly be expected.

An idea that has been often suggested in this context, is to use patent citations as an index of the importance of patents, i.e. to count the number of times that each patent was cited in subsequent patents, and compute with it weighted patent counts (the intention of course is to use those weighted counts as indicators of innovation further down the line). This idea can be traced directly to the widespread use of citations appearing in the scientific literature, in the study of various aspects of the scientific enterprise⁶. However, and contrary to the somewhat arbitrary nature of citations in scientific publications, those in patents are grounded in the Patent Law, and are ultimately decided by a supposedly objective third party, the patent examiner⁷. To quote,

be very cautious in making inferences in that respect from finite samples (see also section 5.a).

⁵This is so because the importance of a patent - however defined - can hardly be assessed ex-ante, and because it is not the task of patent examiners to make sure that the patents granted are of 'comparable worth'.

⁶Indeed, numerous studies in 'scientometrics' (or 'evaluative bibliometrics') have shown that citation-based indices can serve as good indicators of the 'impact' of scientific contributions, of the 'influence' of scientific journals, etc. Classic works in this field are those of Derek de Solla Price (e.g. Price[1963] and [1975]), and of Cole and Cole [1973]. See also Narin [1976] for an extensive review of scores of studies on the subject.

⁷I am referring here only to the citations appearing on the front page of each patent, under 'References Cited' (Item 56), and not those that may be mentioned in the text of the patent. To the best of my knowledge the latter play only a descriptive role, and do not carry legal weight.

"During the examination procedure the examiner searches the pertinent portion of the patent file. His purpose is to identify any prior disclosures of technology...which might anticipate the claimed invention and preclude the issuance of a patent; which might be similar to the claimed invention and limit the scope of patent protection...; or which, generally, reveal the state of the technology to which the invention is directed" .(OTAF [1976], p. 167).

Thus, and as Campbell and Nieves [1979] argue at length, there is an important legal dimension to patent citations, since they represent a **limitation** on the scope of the property rights established by a patent's claims, that carry weight in court. Moreover, the process of arriving at the final list of references, involving also the applicant and his attorney, apparently does generate the right incentives so as to have all truly relevant patents cited, and **only** those (see Campbell and Nieves [1979], Appendix II). The presumption that citation counts are potentially informative of the 'importance' of patents is thus well-grounded.

In practice, however, patent citations have rarely been used in research (see end of this section), primarily because of two serious obstacles: first, until not long ago it was quite difficult to obtain the necessary data, for reasons similar to those mentioned in the context of the classification problem. Second, in the absence of independent evidence on the **value** of the innovations disclosed in patents, it is virtually impossible to ascertain the merits of a citations-based index or, for that matter, of any alternative indicator of the outcome of innovative activities.

In light of the discussion in the previous section, it is clear that the problem of data availability can be easily overcome nowadays, with the aid of the same search techniques described there: once the relevant set of patents has been identified, the number of citations that each of them received can be

obtained by searching in the 'references cited' field of those same patents (i.e. 'within referencing'), or of all subsequent patents. Better still, some of the patent databases in DIALOG already include citation counts as a standard information item in each patent.

The second problem is much more difficult, since it hinges on our ability to obtain self-justifiable measures of the value of innovations, that could help validate citations-weighted patent counts. The trouble is that the assessment of innovations poses formidable empirical and conceptual difficulties, primarily when it comes to **product** innovations (**process** innovations are better understood, following Griliches[1957] seminal work on hybrid corn). Not surprisingly, few studies of that sort have ever been conducted (notable exceptions are Mansfield et al [1977], and Bresnahan [1986]), and hence the lack of value indicators that could be of help in the present context.

I have addressed those issues extensively in a previous study (Trajtenberg [1983]): using discrete choice models of consumer behavior, I put forward a methodology for measuring the social gains from product innovations, and applied it to CT scanners. The intention is to use those measures here, in order to assess directly the performance of the citations weighting scheme. Aside from the very availability of a well-grounded measure of innovation, the advantage in this case is that both the patent counts and the measures used to validate them refer precisely to the same 'stretch' of innovative activity, i.e. to advances in a carefully circumscribed product-class and time period. Thus, the usual problems that arise when trying to match information belonging to disparate units (as often happens in this context) are altogether absent

here.

As already mentioned, there have been up to now just a handful of studies using patent citations, and all but one (Lieberman [1987]) are outside the realm of economics proper⁸. Lieberman [1987] looked at the impact of patent counts on price changes (as proxies for costs) for a sample of chemical products. He finds that own-patents are positively correlated with prices, and that they become statistically insignificant when weighted by citations. Lieberman provides plausible explanations for these 'negative' results, having to do with the nature of process innovations in those sectors. Carpenter et al [1981] show that 100 'important' patents received more than twice as many citations as a matching sample of randomly chosen patents. They took 'important' to mean patents associated with products that received the IR100 Award of the Journal of Industrial R&D, in 1969 and 1970. Ellis et al [1978] use citation 'networks' to map the technological history of selected fields. Thus, for example, starting from a handful of patents issued in the 1970's in semi-synthetic penicillin, they were able to trace back the key patents in the development of the field. They make use of conventional historical material to validate the 'historiographs' thus constructed. Campbell and Nieves [1979] also put the emphasis on tracing the evolution of a technology, and propose for that purpose a variety of patent-based indicators. In sum, there are some bits and pieces of evidence to the effect that patent citations may be indicative of something like 'importance', but the issue remains wide open.

⁸ Aside from those reviewed here, additional (unpublished) studies putting forward the use of patent citations include Narin and Wolf [1983], and Narin [1983]. However, the intended use of patent data in the later two is for business consulting, rather than academic research.

3. PATENTS IN CT: A FIRST LOOK

Using the search methods described in Appendix A, the complete set of patents granted in Computed Tomography was located and retrieved, from the very start of the field in 1971⁹, and up to the end of 1986, totalling 456 patents¹⁰. In order to appreciate the extent to which the search techniques used here represent a quantum jump in our ability to identify patents in a given field, consider the following facts: The 456 patents in CT are spread over 75 patent sub-classes, the leading one comprising 43% of them, the largest five 69%, and the remaining 31% of the patents being scattered over 70 categories, each of them with no more than 1% of the patents. Thus, had I tried to locate the patents in CT by going over the PCS, I would have probably succeeded in identifying just about 70% of the total. Moreover, even in the sub-classes with the largest numbers of patents in CT, the latter represent only a fraction of the total in those categories (except for the leading sub-class, where 90% of the patents belong to CT); thus, the percentage of patents wrongly selected would have been quite large.

⁹In this case it was very easy to identify the first patent: the origin of Computed Tomography is unequivocally associated with the invention by G. Hounsfield of EMI, England, as described in his US patent # 3778614, applied for in December 1971. Since there were no patents in CT in 1972, I shall treat this first patent as if it had been applied in January 72 rather than in December 71, so as to avoid an unnecessary discontinuity in the data points.

¹⁰The computerized search actually produced 501 patents, but 45 of them were eliminated after a careful examination of their abstracts. Thus, I am certain that all the patents included do belong to CT, but obviously one cannot be equally sure that those constitute all the relevant patents. Still, in this case I am quite confident in that respect as well, since I have been able to cross-check with other sources, including listings of patents from manufacturers of CT scanners.

As is by now standard practice, patents will be dated according to their application, rather than granting date, since the later depends entirely upon the examination procedure at the Patent Office, and hence has nothing to do with the innovation process itself. However, since the availability of patents at the time of the search obviously reflects granting rather than application dates, there is a question as to how well the data cover the period under consideration, particularly the more recent years. In order to answer it we just need to look at the distribution of lags between application and granting¹¹:

Lag (in years)	No. of patents	Cumulative percent
1	105	23.2
2	243	76.3
3	91	96.3
4	11	98.7
5	4	99.6
6	1	99.8
7	1	100.0

Assuming that the distribution is stable, and recalling that the search was conducted in December 1986 (i.e. the set includes all patents granted in CT up to 12-86), I conclude that the data comprise virtually all patents applied for up to (including) 1982, about 96% of the patents applied for in 1983, 76% of those applied in 1984, and a mere 23% of the 1985 patents. Thus, the analysis will be restricted to the period 1972-82, although the citations appearing in the 1983-86 patents will be taken into account.

Now to the data on citations: as mentioned earlier, citation counts can

¹¹The lag is computed as the difference: (year granted - year applied); thus, the second row, for example, means that 76% of the patents applied for in any given year, were granted within the following two calendar years. The distribution is virtually identical if the 1982-86 patents are excluded.

be done in two different ways, namely, counting all citations, or just those appearing in the set of patents belonging to the same field. Each has its own merits, and leads to a different interpretation of the resulting weighted patent counts: in the 'within referencing' case the weighted counts will be associated with the 'value' of the patents for - and in terms of - the specific technological field to which they belong. On the other hand, an all-inclusive index will presumably capture the value 'spilled-over' to other areas as well. Given that the measures of innovation to be used in conjunction with the patent data refer to the gains from advances in CT as such, with no attempt to account for spillovers, the citations data are taken accordingly just from references appearing in patents in CT¹².

The first two columns of Table 1, graphically displayed in Figure 1, show the basic patent data to be used throughout. Note the smooth, cycle-like path followed by the yearly count of patents: it rises quite fast after 1973, peaks in 1977, and then declines steadily, carrying forward a long and thin tail (presumably extending into the indefinite future). Notice also that the weighting scheme strongly influences the shape of the time distribution, shifting it back towards the earlier period. In fact, the mean of the distribution of simple counts is 70.6 (in number of months elapsed since 1/72, the date of the first patent), whereas that of weighted counts is 54.0. That is, the weighting scheme moves back 'the action' 17 months, centering it around mid 1976, rather than at late 1977. Given the very fast pace at which

¹²In this case it would have not matter much which count was used: in a sample of 30 patents in CT, the correlation between the two counts was of 0.99. Likewise, Campbell and Nieves [1979] report a correlation of 0.73 between what they called 'in-set' and total patents, for some 800 patents in the field of catalytic converters.

TABLE 1
Patents in CT: Counts and Citations, by Year

Year	Patents		Citations		
	Simple counts	Weighted by ^a citations	Average no. per patent	% of patents with: 0	5+
1972	1	73	72.0	0.0	100.0
1973	3	50	15.7	0.0	100.0
1974	21	199	8.5	4.8	76.2
1975	48	242	4.0	12.5	47.9
1976	66	235	2.6	21.2	22.7
1977	115	260	1.3	45.2	11.3
1978	71	126	0.8	54.9	4.2
1979	59	88	0.5	66.1	0.0
1980	26	33 ^b	0.3	84.6	0.0
1981	15	18 ^b	0.2	86.7	0.0
1982	12	13 ^b	0.1	91.7	0.0
1983 ^c	13	14	.	.	.
1984 ^c	6	6	.	.	.
All	456	1357	2.1 ^d	45.1 ^d	16.2 ^d

^a The weighted patent counts are computed as: $\sum_{i=1}^{n_t} (1 + C_i) = n_t + \sum_{i=1}^{n_t} C_i$,


where C_i is the number of citations received by patent i , and n_t is the number of patents in year t (i.e. the simple patent count).

^b These figures are slightly biased downwards (see Appendix B).

^c Partial figures. ^d Averages.

FIGURE 1

Patents in CT: Simple Counts, and Counts Weighed by Citations



□ Simple Patent Count (source: table 1, column 1).

+ Count Weighed by Citations (source: table 1, column 2).

the CT technology evolved, and that the period is just 11 years long, a difference of 1-1/2 years in the means is certainly very significant. Clearly, the reason for it has to be that earlier patents were cited more frequently than later ones. And indeed, Table 1 proves that to be the case: the average number of citations per patent goes down dramatically over time (from 72 to less than 1), the percentage of patents with no citations increases from 0 to 92%, etc.

The crucial question is whether the observed citation frequencies are to be regarded as a 'real' phenomenon, presumably reflecting something like the importance of patents, or just as a statistical artifact, induced primarily by the mere passage of time. Two main concerns arise in this context: first, it could be that older patents are cited more often simply because they have been around longer, i.e. because they have had more opportunities to be cited, since they precede a larger set of patents that could cite them. Second, given that CT is an ongoing technology, it is quite certain that additional patents have been - and will be - issued since the time of the search (12/86). Thus, the data set is necessarily truncated, and that might bias downwards the citation counts of recent patents, since the probability of being cited is likely to decline over time. In that case the observed phenomenon could just be the result of the arbitrariness of the cut-off point.

These are serious *a priori* objections that may arise whenever trying to attach any meaning to citations data, and therefore deserve careful

¹³In fact, the issue of 'age vs. importance' (closely related to de Solla Price's 'immediacy factor') has commanded a great deal of attention in the scientometric literature; however, and to the best of knowledge, so far it has not been addressed with the aid rigorous statistical tests (see for example Line [1970], and Campbell and Nieves [1979]).

scrutiny¹³. In Appendix B I analyse them in detail, and find that neither age nor truncation can account for the observed distribution of citation counts. The issue of age is tackled by constructing an hypothetical 'iso-important' distribution of citations, and testing it against the observed distribution with the aid of a χ^2 test: the null hypothesis that older patents received more citations just because of the passage of time is rejected by a wide margin. As to the effect of truncation, the magnitude of the biases is estimated by extrapolating from the observed distribution of citation lags, and of the application-granting lags. The main finding is that a bias does exist, but the absolute expected number of missing citations to recent patents is so small, that they could not possibly affect the statistical analysis to be performed.

4. ESTIMATES OF THE VALUE OF INNOVATIONS IN COMPUTED TOMOGRAPHY

As already mentioned, the intention is to relate patent counts weighted by citations, to the independent measures of innovation from Trajtenberg [1983], so as to test whether the former could serve as indicators of the latter. Following is a brief description of these measures.

The first problem encountered in trying to quantify product innovations, is that there is nowhere to be found a natural - let alone unique - way of doing so in meaningful economic terms. In my previous study I chose to associate the 'amount' of innovation occurring in a certain product class at time t , with the social gains stemming from the technical advances embedded in the set of products offered at t . That is, given a monetized 'social surplus' function $W(\cdot)$, and the sets of products S_t and S_{t-1} offered in two successive

periods, the amount of innovation is defined simply as $\Delta W_t = W(S_t) - W(S_{t-1})$ ¹⁴. Thus, the question "how much innovation is there" is taken to mean how much do consumers benefit from the underlying technical improvements, using as a yardstick their preferences for the attributes of the products.

The methodology used to compute the values ΔW_t consists in estimating discrete choice models of demand (more specifically, the multinomial logit model), and integrating the ensuing probabilistic demand functions so as to obtain the measures of consumer surplus:

$$(1) \quad W_t = \ln \left\{ \sum_i^{m_t} \exp \{ \phi(z_{it}, p_{it}) \} / \lambda_t \right\}$$

where z is the vector of attributes, p price, m the number of alternatives (different brands) in the choice set, λ the estimate of the marginal utility of income, and $\phi(.)$ the relevant portion of the estimated indirect utility function. Lastly, the differences ΔW_t are computed from the yearly values of (1), for every pair of adjacent years. In short, ΔW is just like a compensating (or equivalent) variation measure, except that it is designed for the assessment of product changes in discrete choice sets, rather than of price changes in a given set of products. Thus, as time goes by new brands are added to the set S , existing products get better (in the sense of having 'more' of some of the z 's), older models drop-out, and so forth. All those changes are encompassed in what is referred to as 'product innovation', and

¹⁴In practice ΔW is confined to changes in consumer surplus, since net aggregate profits were nil during the period studied. As to S , it comprises the prices and main performance characteristics of the products in the set.

it is their compound value for the consumer that ΔW is measuring.

Noting that ΔW refers to the incremental gains accruing to the representative consumer, I compute also the total gains associated with the innovations at t , as follows,

$$(2) \quad TW_t = \Delta W_t \left[n_t + K \left(\sum_{r=0}^t \Delta W_r \right) \int_{t+1}^{\infty} f(r) e^{-r(r-t-1)} dr \right] = \Delta W_t (n_t + n_f)$$

where n_t is the number of consumers buying at t , $K(.)$ is the ceiling of the diffusion process (that shifts-up as a consequence of successive innovations), $f(.)$ the diffusion path, and r the interest rate. Thus, TW simply multiplies ΔW_t by the number of current beneficiaries from the innovations at t , n_t , plus the discounted number of those that will benefit from those same innovations in the future, n_f , the later being assessed on the basis of the observed diffusion process.

I gathered for the earlier study a comprehensive data set on CT scanners, including the prices and attributes of all scanners marketed in the US since the inception of the technology in 1973 and up to 1982, details of all sales to hospitals and clinics (i.e. who bought which scanner and when), R&D expenditures of firms on CT, etc. Applying the methodology just sketched to these data, yearly estimates of ΔW and of TW were obtained (see table 2): those are the figures that will be related to patent counts. Note that the two measures follow a very similar pattern over time: they are very large - and rising - at first, and then decline dramatically, carrying a long tail into the future. Thus, and eventhough the 'action' in terms of R&D, entry, etc.,

TABLE 2
Measures of Innovation and other Data on CT Scanners^a

Year	ΔW^b	TW^b	$R\&D^b$	# of firms	# of new brands	# of new adopters
1973	2.99	638	20.6 ^c	3	1	16
1974	8.71	6926	22.6	8	1	74
1975	1.51	1503	59.7	12	4	216
1976	4.78	5959	96.1	13	11	317
1977	0.94	997	79.7	14	14	328
1978	0.12	79	64.3	11	6	211
1979	0.14	73	56.1	9	5	209
1980	0.07	30	46.4	8	2	177
1981	0.18	79	37.9	8	3	101
1982	0.20	87	37.9	8	8	.

^a Source: Trajtenberg [1983].

^b In millions of constant 1982 \$.

^c This figure refers to total R&D expenditures from 1968 through 1973.

peaks later-on, the bulk of the gains from innovation are generated during the first few years. This time profile seems to be typical of the emergence of entirely new products, and may be attributed to an initial phase of sharply increasing returns both in the 'production' of innovations, and in the utility derived from them, followed promptly by the setting-in of diminishing returns in both dimensions.

5. PATENTS AS INDICATORS OF INNOVATION: THE STATISTICAL EVIDENCE

The main hypotheses to be put forward are (note that they refer to a given technological field - or industry - as it evolves over time):

- H1: Patent counts weighted by citations are good indicators of the **magnitude**, or **value**, of innovations, but simple counts are not.
- H2: Simple patent counts are good indicators of the extent of **innovative activity** taking place in a field, and are therefore related to R&D expenditures.

These hypotheses are examined in the simplest possible way, namely, using pairwise correlations between the different variables, for two reasons: first, since the point of departure is that ΔW , TW , $R\&D$, etc. accurately represent the phenomena of interest, the only remaining question is whether patents (which are at best just an indirect manifestation of the same phenomena), closely follow the path of those variables over time. Any conventional measure of association along the time dimension will do, certainly the Pearson coefficient. Second, and more pragmatically, there are just ten observations in all (1973-82), and that is obviously too little to estimate anything but correlations.

5.a The First Hypothesis

In order to proceed with the statistical analysis, the hypotheses have to be formulated in more detail, i.e. one needs to define more precisely the scope of the patents counts, and the time structure of their links with the measures of innovation. With regard to scope, the question is whether the counts should include all patents in CT, as done until now, or just those granted to firms that were active in the market for CT scanners (the latter accounted for 66% of all patents, and for 80% of all citations). Since ΔW and TW were computed on the basis of the CT scanners actually marketed in the US, we would expect these measures to be more highly correlated with the patents granted to manufacturers of CT. That would not be so only if the appropriability of the patents issued to the other assignees had been extremely low, i.e. only if the manufacturers of CT scanners benefitted from the innovations done by other inventors as much as they did from their own.

The issue of timing (i.e. of the lag structure) is much more complex, primarily because it is not clear what is precisely the information conveyed by the application dates of patents: innovation is obviously not a one-shot event, but a continuous time process, starting from the formulation of a novel idea, and ending with the introduction in the market of the product embedding that idea (and hence with ΔW). Sometime along the way the inventor files for a patent application, presumably before the development stage is completed; but, how much earlier, and what determines the length of the lag? Eventhough the answers are far from clear, one can at least identify two of the main forces at work: the first has to do with the stringency of the application

requirements set by the Patent Office, the second with the technological characteristics and the competitive structure of the industry. Thus, the more stringent the requirements are, and the more intense the technological rivalry in the field is, the shorter will be the lag between patent counts and ΔW . Beyond that, however, no general prior can be put forward regarding the expected lag: in principle the patent might be filed just before the innovation reaches the market, or much before; notice also that there is no reason for the lag to be constant over time.

There is, however, further information on dates that may shed light on the lag structure: in addition to the application date in the US, many patents have a 'foreign priority application date', and some make reference to earlier 'Related US Application Data',¹⁵. To quote in relation to the former,

"It is common to seek patent protection on a single invention in several countries...International multiple patenting has been facilitated by a treaty which permits applications filed in a foreign country within a year of filing in the home country to be accorded the home country filing date. However, the treaty requires that...the initially filed patent application must be identified by country, serial number and filing date, and that the 'priority' of this filing be claimed." (OTAF [1977], p.17).

Thus, for example, according to OTAF [1977] 30% of all patents granted in the US in 1975 contain foreign priority data. In the case of patents in CT, 56% of them mention a foreign application date, with an average lag between it and the application date in the US of fourteen months (see table 3)¹⁶. Such a

¹⁵ Sixteen percent of the patents in CT make reference to an earlier 'related US application', i.e. they are designated as 'continuations' of previous applications, which may have failed, or given rise to other (related) patents. Unfortunately, I do not have the actual dates of those earlier applications, only the fact of a reference to them.

¹⁶ As quoted above, the application in the foreign country has to be made within a year of the home country filing date; and indeed, some 80% of the

TABLE 3
Lags Between Foreign and US Applications, by Years

Year	% of patents with foreign priority	Mean Lags (in months)		
		All Patents	Patents to Firms in CT	Patents with foreign priority
1972	1.00	41.0	41.00	41.0
1973	1.00	41.3	41.3	41.3
1974	0.76	17.1	21.0	22.4
1975	0.48	6.0	5.8	12.6
1976	0.56	7.7	8.1	13.7
1977	0.64	6.8	7.2	12.6
1978	0.56	8.4	8.9	14.9
1979	0.54	7.1	4.9	13.1
1980	0.65	8.0	9.8	12.2
1981	0.60	8.1	6.6	13.6
1982	0.42	3.2	3.0	7.6
Average	0.56	7.8	8.2	14.1
Weighted* Average	0.63	11.5	12.7	14.1

* Weighted by the number of citations of each patent.

lag simply implies that the innovation process underlying those patents, stretches back at least 14 months longer than what could be inferred just from the US application date. The lag is also consistent with the observation that patenting requirements are more exacting in the US than in foreign countries, since often times the innovations had not been developed enough at the time of the foreign application, for them to meet US standards. Now, because the standards differ, and moreover, because the composition of the foreign priority countries changed significantly over the period studied (the UK ceding its initial dominant place to Germany and Japan), it does not make sense to date all patents according to the **earliest** date appearing on them. Thus, they will still be dated according to the US application date, so as to be able to interpret consistently whatever lag between patent counts and ΔW is found, in light of the implied (uniform) standard. One can then superimpose on it the foreign-US application lag (which, as table 3 reveals, varies by year), thus gaining some notion of the **minimal overall span** of the innovation process in CT.

Finally, it should be clear that the most serious limitation in this respect stems from the fact that the period studied is very short, not allowing for the estimation of any sort of lag **structure**, let alone of a changing lag. Thus, the findings referring to the lag between patents and ΔW , obtained by maximizing pairwise correlations over the time dimension, should

patents with foreign priority were filed in the US just before a year elapsed. The rest have a bridging 'related US application', i.e. within a year of the filing abroad a patent was indeed applied for in the US, but it either failed, or gave rise to another patent; in either case, the present patent, by being designated as a 'continuation' of the related US application, can still claim the foreign priority date.

be regarded as tentative.

Now to the statistical results: table 4 presents the correlations between alternative versions of simple and weighted patent counts, and (ΔW , TW), with the former variables lagged between 0 and 6 months¹⁷. The first and most important finding is that, in effect, weighted patent counts are correlated with the value measures of innovation, whereas simple counts are clearly not, in all the cases considered. Thus, the first hypothesis is confirmed, and decisively so. Second, the correlations increase substantially, as we narrow-down the scope of the (weighted) counts to the patents granted to firms in CT. This supports my prior and implies, as suggested, that the patents awarded to other assignees did not quite become a 'public good' (i.e. appropriability was not nil).

Third, the correlations peak when the patent counts are lagged just one quarter, declining monotonically as the lag increases (this is true for lags far beyond the 6 months shown in table 4)¹⁸. Superimposing on it the mean foreign-US application lag of 12.7 months (for patents to firms in CT, weighted by citations - see table 3), one obtains a minimum innovation span (or minimum 'lead time') of 16 months. This may strike as rather short

¹⁷Notice that, even though the variables refer to **yearly** figures, the lag can be varied by **monthly** increments, since the patent data is virtually continuous in time. It is also important to note that, since the ΔW series begins in 1973, I just added the 1972 (first) patent to the patent count of 1973, i.e. the ΔW figure for 1973 refers to the **first** CT scanner marketed, and hence it obviously corresponds to the initial patents in the field, including the very first.

¹⁸However, recall from the previous footnote that the 1972 patent was simply added to the 1973 patents in computing the correlations. Thus, the first lag was actually longer (about one year long), and the overall lag would increase from 3 to 4 months if one averages that first lag with the rest.

TABLE 4
Correlations of Simple and Weighted Patent Counts
with ΔW , TW

(a) with Weighted Counts

Lags	all patents		patents to firms in CT	
	ΔW	TW	ΔW	TW
Contemporary	0.509 (0.13)	0.587 (0.07)	0.616 (0.06)	0.626 (0.05)
3 months	0.513 (0.13)	0.635 (0.05)	0.685 (0.03)	0.755 (0.01)
4 months	0.480 (0.16)	0.600 (0.07)	0.677 (0.03)	0.744 (0.01)
6 months	0.317 (0.37)	0.466 (0.17)	0.495 (0.15)	0.605 (0.006)

(b) with Simple Counts

Lags	all patents		patents to firms in CT	
	ΔW	TW	ΔW	TW
Contemporary	-0.162 (0.65)	0.032 (0.93)	-0.087 (0.81)	0.093 (0.80)
3 months	-0.198 (0.58)	0.006 (0.99)	-0.076 (0.83)	0.131 (0.72)
6 months	-0.283 (0.43)	-0.090 (0.81)	-0.175 (0.63)	0.027 (0.94)

Significance levels in parentheses

(although it is not clear against what standard one ought to judge it), but if so it would certainly be consistent with the intense technological rivalry that characterized the evolution of CT scanners. Note also from table 3 that the overall lag drops from close to five years in 1973, to one and a half year in 1974, hovering around one year from then on. Figure 2, tracing the 'geneological tree' of the first patents taken by G. Hounsfield, the inventor of CT, throws light on those initial long lags, and confirms the stringency of the US application requirements vis a vis those in the UK¹⁹. Closely related, figure 2 implies also that the innovations contained in those first patents were more 'general' (and hence more 'important' indeed), in the sense that it took longer to embed them in demonstrably working - and useful - systems that could comply with the US requirements for patentability.

Returning to the basic finding of this section (i.e., that weighted patent counts are highly correlated with ΔW), we can now (re)interpret the distribution of citation counts across patents, as an implied distribution of the value of innovations. As shown in table 5, the observed distribution fits well the received wisdom on this matter (see for example Shenkerman and Pakes [1984], and Pakes [1986]): it is very skewed, with almost half the patents never cited (and hence of little ex post value), and a 'lucky few' being worth a great deal²⁰. Thus, and contrary to Scherer's pessimistic note to the effect

¹⁹ Notice that the first US application actually failed, even though it was filed one year after the first application in the UK. It then took Hounsfield more than two years of further development (in the course of which a first working prototype of a CT scanner was installed in a London hospital), in order to win a US patent.

²⁰ Campbell and Nieves [1979] present the distribution of citations for all US patents issued from 1791 to 1978 (over four million patents), whereas Narin [1983] shows the distribution corresponding to 13,264 chemical and

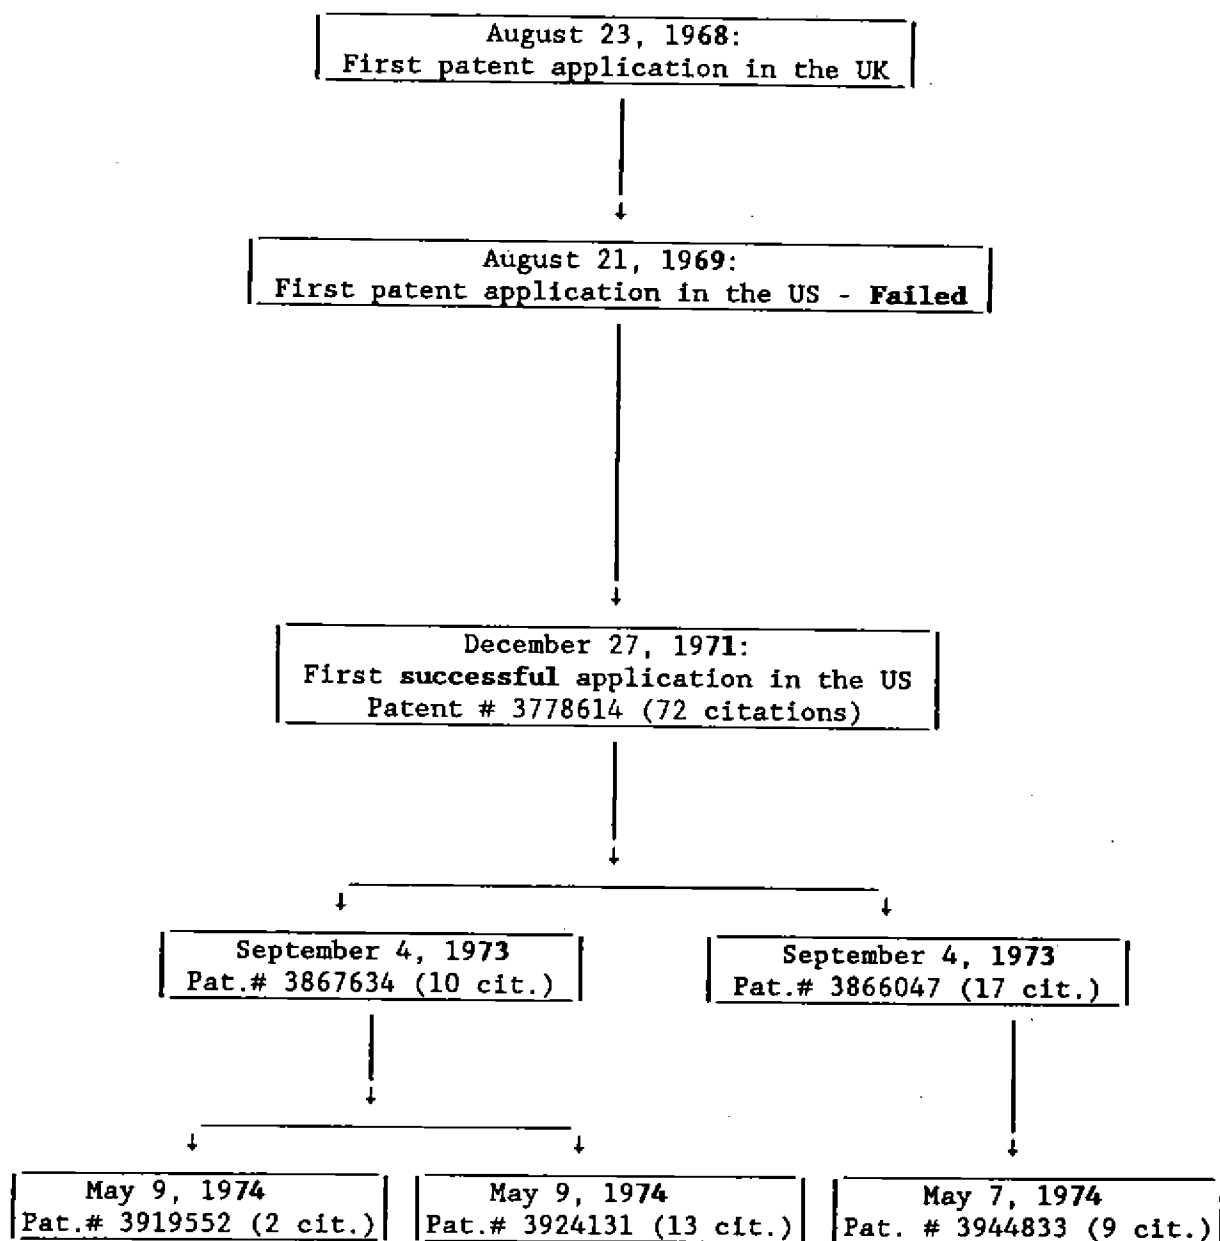
innovations per se, than of the size of the market for the products embedding those innovations. This is a reassuring result, since we expect technological-related factors to be dominant in the patenting and citing processes.

5.c The Second Hypothesis

To recall, the second hypothesis is that simple patent counts would be good indicators of the level of innovative activity in a field, as reflected for example in R&D expenditures. The relationship between patents and R&D has been intensively scrutinized in past research²⁵, and the results appear to be quite uniform, centering around the following 'stylized facts': a) there is a strong statistical association between patents and R&D expenditures; b) this relationship appears to be mostly contemporaneous; and c) R&D explains a great deal of the cross-sectional variance in patenting, but not much of the variation along time. Much of this research has been done on short panels of firms' data, leading Griliches to conclude, in summing-up, that "...patents are a good indicator of differences in inventive activity across firms, but that short term fluctuations in their numbers within firms have a large noise component in them" (Griliches[1984], p.3).

The second hypothesis can be seen as extending those results to the time dimension, i.e. it also postulates a close association between patents and R&D, but within a given field and along time, rather than across firms and/or industries. I resort again to simple correlations, and explore in detail the possible existence of lags. The main findings, drawn from Table 7, are: First, there is indeed a high correlation between yearly patent counts and R&D, and a

FIGURE 2
The First Patents in CT, by G. Hounsfield



The arrows indicate that the lower patent(s) has been designated as a 'continuation' of the preceeding patent document.

TABLE 5
Distribution of Patents According to Number of Citations

Number of citations	Number of patents	percent of patents	cumulative percent
0	215	47.1	47.1
1	78	17.1	64.3
2	54	11.8	76.1
3	35	7.7	83.8
4	21	4.6	88.4
5	10	2.2	90.6
6	15	3.3	93.9
7	8	1.8	95.6
8	3	0.7	96.3
9	3	0.7	96.9
10	2	0.4	97.4
12	1	0.2	97.6
13	2	0.4	98.0
14	1	0.2	98.2
16	1	0.2	98.5
17	2	0.4	98.9
19	1	0.2	99.1
20	1	0.2	99.3
21	1	0.2	99.6
25	1	0.2	99.8
72	1	0.2	100.0

$$(4) \quad \alpha_1^* = \underset{\alpha}{\text{Max Corr}} [WPC_t(\alpha), \Delta W_t] \quad , \quad \text{and} \quad \alpha_2^* = \underset{\alpha}{\text{Max Corr}} [WPC_t(\alpha), TW_t]$$

The answers emerge clearly from table 6: $\alpha_1^* = 1.30$, and $\alpha_2^* = 1.10$, i.e. there are in fact 'increasing returns' to citations, and they manifest themselves more strongly in the context of the relationship of patent counts with ΔW , rather than with TW . Note that these results are robust, in that they obtain also in the absence of a lag, and when using counts of all patents rather than counts of patents to firms in CT²³. Moreover, and as figure 3 shows, the highest correlations arrived at can be taken indeed as **global** maxima. In other words, the WPC's based upon patents to firms in CT, lagged 1 quarter and using as exponents $\alpha_1 = 1.30$ and $\alpha_2 = 1.10$, dominate all other cases along the three dimensions considered here.

The finding that $\alpha_1^* > 1$ is in itself an important one²⁴: first, it provides further evidence to the effect that WPC's convey a great deal of information on the value of innovations disclosed in patents; in particular, it means that the **marginal** informational content of the WPC's increases with

each patent, and not in the aggregate (weighted) patent count.

²³ The actual maximized values of the exponents are somewhat higher in the latter case: $\alpha_1^* = 1.40$, and $\alpha_2^* = 1.30$. This may be related to the fact that, as already mentioned, the patents of firms in CT received more than their proportional share of citations.

²⁴ And it came as a surprise: originally I thought that there might be **diminishing** returns to citations, i.e. that citations may be given 'too generously', and hence that their marginal informative value would be low and declining. Had that been the case, the role of WPC's as indicators of innovations would have been weakened.


TABLE 6
Correlations of WPC with ΔW , TW: Searching
for Non-Linearities

Exponent α	Patents to Firms in CT				All patents	
	Contemporary ΔW	TW	Lagged 3 months ΔW	TW	Lagged 3 months ΔW	TW
0.80	0.455 (0.19)	0.543 (0.10)	0.512 (0.13)	0.653 (0.04)	0.329 (0.35)	0.503 (0.14)
0.90	0.538 (0.11)	0.590 (0.07)	0.601 (0.07)	0.711 (0.02)	0.419 (0.23)	0.570 (0.09)
1.00	0.616 (0.06)	0.626 (0.05)	0.685 (0.03)	0.755 (0.01)	0.513 (0.13)	0.635 (0.05)
1.10	0.680 (0.03)	0.642 (0.05)	0.754 (0.01)	0.777 (0.008)	0.605 (0.06)	0.687 (0.03)
1.20	0.721 (0.02)	0.635 (0.05)	0.798 (0.006)	0.770 (0.009)	0.684 (0.03)	0.719 (0.02)
1.30	0.738 (0.01)	0.607 (0.06)	0.813 (0.004)	0.736 (0.02)	0.738 (0.02)	0.720 (0.02)
1.40	0.730 (0.02)	0.560 (0.09)	0.800 (0.006)	0.677 (0.03)	0.760 (0.01)	0.689 (0.03)
1.50	0.703 (0.02)	0.501 (0.14)	0.766 (0.01)	0.606 (0.06)	0.751 (0.01)	0.634 (0.05)
1.60	0.663 (0.04)	0.436 (0.21)	0.718 (0.02)	0.527 (0.11)	0.719 (0.02)	0.562 (0.09)

Significance levels in parentheses.

FIGURE 3

Correlations between WPC and ΔW , for various exponents



□ All patents, no lags.

+ Patents to firms in CT, no lags.

◇ Patents to firms in CT, lagged 3 months.

the number of citations. Second, it implies that the variance in the value of patents is larger, and that the distribution of those values more skewed, than what could be inferred from the simple count of citations (recall table 5).

Now to the finding that $\alpha_1^* > \alpha_2^*$. Recall that ΔW is a measure of the gains to the representative consumer from improvements in the set of available products, and therefore it amounts to a snapshot valuation of the innovations underlying those improvements. TW , on the other hand, multiplies ΔW by the number of consumers that benefit from the innovation, at present and in the future. Thus, the fact that $\alpha_1^* > \alpha_2^*$, and that $\text{Corr}[WPC(\alpha_1^*), \Delta W] > \text{Corr}[WPC(\alpha_2^*), TW]$, means that citations are more informative of the value of innovations per se, than of the size of the market for the products embedding those innovations. This is a reassuring result, since we expect technological-related factors to be dominant in the patenting and citing processes.

5.c The Second Hypothesis

To recall, the second hypothesis is that simple patent counts would be good indicators of the level of innovative activity in a field, as reflected for example in R&D expenditures. The relationship between patents and R&D has been intensively scrutinized in past research²⁵, and the results appear to be quite uniform, centering around the following 'stylized facts': a) there is a strong statistical association between patents and R&D expenditures; b) this relationship appears to be mostly contemporaneous; and c) R&D explains a great deal of the cross-sectional variance in patenting, but not much of the

²⁵ Many of the papers in Griliches[1984] have to do in one way or the other with this issue; extensive references to previous works can also be found there.

variation along time. Much of this research has been done on short panels of firms' data, leading Griliches to conclude, in summing-up, that "...patents are a good indicator of differences in inventive activity across firms, but that short term fluctuations in their numbers within firms have a large noise component in them" (Griliches[1984], p.3).

The second hypothesis can be seen as extending those results to the time dimension, i.e. it also postulates a close association between patents and R&D, but within a given field and along time, rather than across firms and/or industries. I resort again to simple correlations, and explore in detail the possible existence of lags. The main findings, drawn from Table 7, are: First, there is indeed a high correlation between yearly patent counts and R&D, and a much weaker one between R&D and patents weighted by citations; thus, the second hypothesis is amply confirmed. Second, the degree of association peaks when R&D is lagged just 5 months, supporting previous findings of short 'gestation lags'. Third, the correlations are slightly higher for counts of all patents than for patents to firms in CT, suggesting some degree of spill-overs from the R&D done by manufacturers of CT to other assignees.

In order to strengthen the notion that patent counts are to be seen as indicators of innovative activity in the broad sense, and not just as proxies for R&D, I computed also the following correlations (all are contemporaneous, SPC stands for simple patent counts; the data are from table 2):

Cor(SPC, no. of firms in the CT market) = 0.858
(0.0007)

Cor(SPC, no. of new scanners introduced in the market) = 0.813
(0.002)

Cor(SPC, no. of new adopters) = 0.913
(0.0002)

TABLE 7
Correlations Between Patent Counts and R & D

Lags	all patents		patents to firms in CT	
	SPC	WPC	SPC	WPC
None	0.869 (0.0002)	0.609 (0.05)	0.843 (0.001)	0.525 (0.097)
3 months	0.919 (0.0001)	0.591 (0.04)	0.912 (0.0001)	0.495 (0.102)
4 months	0.924 (0.0001)	0.582 (0.05)	0.914 (0.0001)	0.491 (0.105)
5 months	0.933 (0.0001)	0.577 (0.05)	0.918 (0.0001)	0.483 (0.112)
6 months	0.921 (0.0001)	0.543 (0.07)	0.903 (0.0001)	0.450 (0.142)
1 year	0.831 (0.0008)	0.248 (0.44)	0.794 (0.002)	0.152 (0.638)

Significance levels in parentheses.

SPC: Simple Patent Counts; WPC: Weighed Patent Counts.

The first two reflect the fact that competition in this technologically progressive market was driven by rivalry in innovation; the third correlation has to do with the impact of innovation on diffusion, i.e., it indicates that the market expanded as the technology improved (see Trajtenberg [1983] for a detailed account of this relationship).

6. THE PATENTS → R&D → PATENTS LINK, AND FURTHER EXTENSIONS

So far the discussion has been confined to simple statistical associations, in order to shed light on the proper use of patent statistics as indicators of innovation. Here I venture into some causal links that may help grasp the dynamics of the innovative process, and suggest two possible extensions.

Eventhough hard evidence is difficult to come about, it has been widely observed that important innovations often generate a flurry of further innovative activity, that brings in turn a host of minor improvements. In fact, new products embedding truly novel technologies are usually crude and lacking at first, but gradually improve over time as a result of further research efforts²⁶. Thus, and returning to patent indicators, one would expect to find a causal link going from WPC to R&D, probably with a substantial lag. Moreover, since as shown above simple patent counts (SPC) follow R&D after a 5 month lag, we would also expect to observe a link between WPC and SPC, the lag between them being the sum of the waiting period between WPC and R&D, and the

²⁶ Indirect support to this contention can be found in the often-noted fact that consumers are well aware of this process, and their ensuing 'technological expectations' appear to play a key role in their decision of when to purchase a new or rapidly advancing product.

gestation lag. In other words, I have in mind a sort of self-propelling innovative cycle, by which important innovations (as reflected in high WPC) bring about further R&D aimed at implementing and perfecting them, and this in turn results in further patents. Some of the later may prove to be important innovations (and thus collect many citations), opening-up a new phase in the cycle.

Table 8 shows the correlations between WPC, R&D and SPC, allowing for various lags: first, there is indeed a strong correlation between lagged WPC and R&D, peaking with a lag of 9 months; second, the correlation between WPC and SPC is highest when the two are 14 months apart, this corresponding exactly to the sum of the WPC → R&D 9-month lag, and the R&D → SPC 5-month gestation lag. Notice also that the correlations between lagged WPC and R&D are systematically higher than those between SPC and R&D, and hence the chain of events is clearly of the form WPC → R&D → SPC, and not SPC → R&D → SPC.

The process just described has a strong 'supply-push' flavour, and evokes Schumpeterian notions of innovation-induced cycles. It should be clear, however, that within the narrow context of this paper I cannot undertake to weigh this view versus its main contender, namely, that of Schmookler[1966] upholding the centrality of demand-inducement mechanisms. At the same time, Schmookler's results cannot be seen as excluding, or negating, the findings here: to begin with, Schmookler had at his disposal only simple patent counts (that reflect just the level of innovative activity, and not the magnitude of innovation) and hence could not really address the 'supply-push' story²⁷.

²⁷One cannot take too seriously the tests that Schmookler performed in chapter IV of his book, using various series of 'important innovations'; moreover, it seems that Schmookler himself doubted the appropriateness of

TABLE 8
The Patents → R & D → Patents Chain

Lags	Correlation of R & D with		Correlation of SPC with	
	WPC	SPC	WPC	SPC
+ 3 months	0.591 (0.04)	0.919 (0.0001)		
none	0.609 (0.05)	0.855 (0.0008)	0.701 (0.02)	1.00
- 4 months	0.711 (0.01)	0.834 (0.001)	0.784 (0.004)	0.969 (0.0001)
- 6 months	0.796 (0.003)	0.820 (0.002)	0.835 (0.001)	0.922 (0.0001)
- 7 months	0.819 (0.002)	0.810 (0.0025)	0.870 (0.0005)	0.917 (0.0001)
- 8 months	0.861 (0.0007)	0.806 (0.003)	0.854 (0.0008)	0.872 (0.0005)
- 9 months	0.866 (0.0006)	0.772 (0.005)	0.853 (0.0009)	0.843 (0.001)
-12 months	0.854 (0.0008)	0.674 (0.02)	0.875 (0.0004)	0.727 (0.01)
-14 months	0.774 (0.005)	0.611 (0.05)	0.935 (0.0001)	0.706 (0.02)
-16 months	0.697 (0.02)	0.543 (0.08)	0.903 (0.0001)	0.649 (0.03)

Significance levels in parentheses.

Furthermore, he studied the innovation process in conventional, well-established sectors, whereas the case here refers to a radically new product, for which an element of demand creation is undeniable.

In sum, the above results should be regarded for now as suggestive, and be seen as part of recent efforts to shed new light on the classic supply-push demand-pull debate (e.g. Beggs [1984], Griliches et al [1986], and Gort and Wall [1986]).

Now to the intended extensions: the first has to do with the question of whether patent data are also informative at a more disaggregated level (e.g. firms within a product class), and the extent to which they may be indicative of the private value of innovations (recall that ΔW refers to their social value). The idea is to have a panel of sales by firms over time, with own patents, and patents by everybody else, as independent variables (I do have all the data needed in the context of CT scanners). The basic model would be that of Spence [1984], and hence one of the interesting issues that could be investigated is the degree of appropriability (i.e. the coefficient of others' patents). The second extension refers to spill-overs across sectors: a possible source to study this elusive issue may be in cross-sectoral citations, that is, references to patents in the 'source' sector, appearing in patents belonging to the field that benefits from the spill-overs. In the case of CT scanners, for example, one would look into citations to CT patents, made in patents belonging to more recent imaging technologies, such as NMR (also called MRI), ultrasound, and positron emission tomography.

those series and ensuing tests.

7. CONCLUDING REMARKS

In light of the results presented above, there is indeed room to believe that patent data hold a significant potential for research in economics. A key element lies with the search techniques used here: with their help one can unlock the wealth of information contained in the patent file, and overcome the classification problem, at least for case studies or small panels.

A distinguishing feature of the approach put forward, is that the units of analysis are narrowly-defined product classes (very close to the economist's notion of markets), rather than firms or SIC categories. Aside from the obvious advantages for analysing issues in industrial organization, this shift in units carries a major additional benefit: patent counts by product classes appear to be good indicators in the ~~time~~ dimension, and not just cross-sectionally as is the case with counts of firms' patents²⁸. This may be of great significance, in view of the fact that innovation is in essence a dynamic phenomenon. Closely related, patent data can be easily obtained all the way to the very beginning of a product class, whereas the gathering of conventional industry data usually starts only when a sector is well established. Thus, and as shown for CT scanners, patent counts and citations may play an important role in studying the very emergence of new markets, which seems to be the period when most of the innovative activity takes place (quite clearly, studies focused on mature industries are very likely to miss the bulk of the innovative segment). Moreover, the why's and

²⁸ As Griliches et al [1986] point out in summing-up previous research, both the fact that the R&D budgets of firms are typically stable over time, and that most firms take a small but highly variable number of patents per year, make it very difficult to trace innovative activity over time on the basis of firms' data.

how's of cross-sectional results regarding the structural characteristics of mature sectors (e.g. concentration, entry barriers, etc), cannot really be understood but in light of how those sectors evolved into their observed equilibrium; again, the type of patent data used here seems to be particularly well-suited to trace that process.

The results having to do with the lag structure (e.g. from patents to ΔW), underscore the importance of correctly dating patents, and call for extra caution in interpreting such findings. The relatively long lags between foreign and US applications suggest that the filing of patents in the US occurs fairly late along the innovation process; this fact may help explain the pervasive result found in the literature, of a contemporaneous relationship between patents and other indicators, such as R&D, the value of firms, etc.

The close association of patents with R&D raises once again some questions as to the proper use (and interpretation) of either variable in empirical research. In accord with the prevailing view, the results here show that simple patent counts are certainly not to be regarded as a measure of innovative output: that requires the weighting by citations. Yet, patents are not quite akin to an input either (as R&D is), since they also reflect to some extent 'effort', and a modicum of technological success.

In order to understand their role, it may be helpful perhaps to think of patents as working papers in economics (and hence of economic departments as firms, and fields in economics as industries): papers are 'produced' roughly in proportion to the number of faculty²⁹, as patents are with respect to R&D.

²⁹Of course, that does not say much about the magnitude of the

The fact is that it does not take much to get a patent, once the firm has an established R&D facility going, as it does not take much to write a working paper. Still, a larger number of patents presumably indicate, *ceteris paribus*, that more research efforts have been invested by the R&D staff (as more papers would suggest that the faculty is 'trying harder'). Thus, it seems that patent counts can be regarded as a more 'refined' measure of innovative activity (*vis a vis* R&D), in the sense that they incorporate at least part of the differences in effort, and net-out the influence of 'luck' in the first round of the innovative process. On the pragmatic side, good R&D data are much more difficult to obtain than patent data, and the latter have a wider coverage. Moreover, patent data are richer and 'finer', in that they are practically continuous in time, and can be further classified by a variety of criteria. Thus, there is plenty of room to expand the use of patent counts, lessening at the same time the dependancy upon hard-to-get R&D data.

The results that hold the greatest promise are those related to the use of citations, since they offer a quantitative indicator for a key variable that had virtually none, namely, the value of innovations. The marked skewness in the distribution of those values appears to be no longer an impediment to the use of patent data, but rather the main source of their usefulness. This as well as all previous conclusions are expressed in a qualified manner, for the obvious reason that they are based upon the findings from a case study; hopefully, future research will bring-in more supportive evidence, and further demonstrate the attractiveness of the proposed indicators.

contribution to economics: for that purpose one would need information on whether and where the working papers get published, the number of citations that they receive over time, etc.

REFERENCES

- Beggs, J.J., "Long-Run Trends in Patenting". in Griliches, Z. (ed.), **R&D, Patents, and Productivity**. Chicago: Chicago University Press, 1984.
- Bresnahan, T., "Measuring the Spillovers from Technical Advance: Mainframe Computers in Financial Services". **American Economic Review**, September 1986, 76, 742-55.
- Campbell, R.S. and A.L.Nieves, "Technology Indicators Based on Patent Data: The Case of Catalytic Converters - Phase I Report: Design and Demonstration". **Battelle, Pacific Northwest Laboratories**, September 1979
- Carpenter, M.P., Narin F. and P. Woolf, "Citation Rates to Technologically Important Patents". **World Patent Information**, 1981, 3, 160-163.
- Carpenter, M.P. and F. Narin, "Validation Study: Patent Citations as Indicators of Science and Foreign Dependence". **World Patent Information**, 1983, 5.
- Cole, J.R. and S. Cole, **Social Stratification in Science**. Chicago, Illinois: University of Chicago Press, 1973.
- Ellis, P., Hepburn G. and C. Oppenheim, "Studies on Patent Citation Networks". **Journal of Documentation**, March 1978, 34, 12-20.
- Gort, M. and R.A. Wall, "The Evolution of Technologies and Investment in Innovation". **Economic Journal**, September 1986, 96, 741-757.
- Griliches, Z., "Hybrid Corn: An Exploration in the Economics of Technological Change". **Econometrica**, 1957, 25, 501-22.
- Griliches, Z. (ed.), **R&D, Patents, and Productivity**. Chicago: University of Chicago Press, 1984.
- Griliches, Z., B.H. Hall and A. Pakes, "Is There a Second (Technological Opportunity) Factor?. Unpublished Manuscript, 1986.
- Lieberman, M.B., "Patents, Learning by Doing, and Market Structure in the Chemical Processing Industries". **International Journal of Industrial Organization**, 1987, 5.
- Line, M.B., "The Half-Life of Periodical Literature: Apparent and Real Obsolescence". **Journal of Documentation**, 1970, 26, 46-52.
- Mansfield, E., Wagner, S. and G. Beardsley, "Social and Private Rates of Return from Industrial Innovation". **Quarterly Journal of Economics**, 1977, 91, 221-40.

- Narin, F., "Evaluative Bibliometrics: The Use of Publication and Citation Analysis in the Evaluation of Scientific Activity". **Computer Horizons, Inc.**, March 1976.
- Narin, F., "Patent Citation Analysis as an Analytical Tool in Marketing Research". **CHI Research/Computer Horizons, Inc.**, September 1983.
- Narin, F. and P. Woolf, "Technological Performance Assessments Based on Patents and Patent Citations". **CHI Research/ Computer Horizons, Inc.**, August 1983.
- Pakes, A. and M. Schankerman, "The Rate of Obsolescence of Patents, Research Gestation Lags, and the Private Rate of Return to Research Resources". in Griliches Z. (ed.), **R&D, Patents, and Productivity**. Chicago: Chicago University Press, 1984.
- Pakes, A., "Patents as Options: Some Estimates of the Value of Holding European Patent Stocks". **Econometrica**, July 1986, 55.
- Office of Technology Assessment and Forecast (OTAF), U.S. Department of Commerce, Patent and Trademark Office, **Fifth Report**. Washington D.C.: U.S. Government Printing Office, August 1975.
- Office of Technology Assessment and Forecast (OTAF), U.S. Department of Commerce, Patent and Trademark Office, **Sixth Report**. Washington D.C.: U.S. Government Printing Office, June 1976.
- Price, D.J. de Solla, **Little Science, Big Science**. New Haven, CT: Yale University Press, 1963.
- Price, D.J. de Solla, **Science Since Babylon**. New Haven, CT: Yale University Press, 1975.
- Scherer, F., "Firm Size, Market Structure, Opportunity, and the Output of Patented Innovations". **American Economic Review**, December 1965, 55, 1097-1123.
- Scherer, F., "Using Linked Patent and R&D Data to Measure Interindustry Technology Flows". in Griliches, Z. (ed.), **R&D, Patents, and Productivity**. Chicago: Chicago University Press, 1984.
- Schmookler, J., **Invention and Economic Growth**. Cambridge: Harvard University Press, 1966.
- Trajtenberg, M., **Dynamics and Welfare Analysis of Product Innovations**. Ph.D dissertation, Harvard University, 1983 (forthcoming, Harvard University Press, 1987).

APPENDIX A

Online search and retrieval of patent data from large databases

The search techniques referred to above make use of the following basic elements: (a) Large computerized databases, to which one can gain easy access just with a PC equipped with a modem; (b) Online search facilities, consisting essentially of basic boolean operators such as 'and', 'or', 'not' etc, and a set of rules governing the use of keywords, fields of search, and the like. (c) Online retrieval capabilities, that allows one to unload the selected data into the PC. There are today thousands of databases containing millions of documents, and their use is spreading extremely fast, both in business and in academia (it seems though that economists have been particularly slow in taking advantage of these services).

I have used for the present study the PATDATA database, through BRS. This database includes all patents issued in the US from July 1975 to the present, it is updated weekly, and its current size is in the order of 800,000 patents. Each patent document consists of fifteen 'fields', such as application and issue dates, classification codes, assignee, a descriptive abstract of about 10-30 lines long, etc. It is worth pointing out that DIALOG offers access to about 10 databases on patents, covering most countries and going back in some areas to the sixties.

The search for patents in a particular product field or industry can be done in a variety of ways: using key words pertaining to the product in question that may appear in the title and/or in the abstract, identifying a small set of relevant patent classification codes, locating assignees

(typically firms) that are known to operate in the field, etc. Needless to say, there isn't a well-defined method that would deliver with certainty **all** the patents in a given field, and **only** those. Rather, the search process consists , at least at first, of trial and error: it involves sampling by a given set of initial criteria, examining the abstracts in order to determine whether the sampled patents do belong to the desired category, searching anew according to an updated set of criteria, and so forth. Since it is always possible to re-examine the patents **after** the search so as to eliminate those that do not belong, the dominant **ex ante** concern is to minimize the probability of overlooking in advance patents that may belong to the desired field, subject of course to the researcher's budget constraint. Lastly, once the desired set of patents has been identified and retrieved, one can extract from them the required data, with the aid of a specially designed computer programme (patents come in the form of full-text documents). At this point the actual analysis can begin³⁰.

³⁰Narin [1983] suggested a similar search procedure, aimed at corporate consulting rather than academic research. Moreover, he relied upon proprietary search techniques and data, whereas one of the appeals of the approach here is easy access and widespread availability.

APPENDIX B

Patent Citations: A Statistical Analysis of Truncation and Age Effects

B.1 Testing 'age vs. importance'

I begin by defining a hypothetical citation process by which all patents are of equal importance, and the only differentiating factor is 'age'. The distribution of citations thus generated will be then compared to the actual one, and the maintained hypothesis that all patents are equally important tested with the aid of a χ^2 test.

As a first step, patents are ordered along time according to their application date, and indexed with $i = 1, \dots, N$ ($N=456$); note that i thus indicates the cumulative number of patents in CT applied for up to patent i . Denoting by p_{ij} the probability that patent i will be cited in patent j (for $i < j$), and by r_j the number of references to previous patents in CT appearing in patent j , one can now define:

All patents $1 \leq i < j$ are said to be iso-important if

$$(B1) \quad p_{ij} = \frac{r_j}{j-1} = p_j \quad j = 1, \dots, N$$

Thus, equal importance is taken to mean that all patents applied for up to a point in time, have the same probability of being cited by a subsequent patent. In other words, (B1) means that the citations appearing in patent j are the result of r_j random drawings (without replacement) from a pool containing the $j-1$ patents that preceded it³¹. Noting that (B1) implies also

³¹Clearly, this is not the only possible way of designing a citation process that would render patents of 'equal importance'. Notice, however, that by defining p_{ij} to be independent of the time interval $(j-i)$, I implicitly favor the earlier patents, thus increasing the power of the test. That is, any plausible departure from (B1) would have the probabilities decrease with

time independence (i.e. for any $i < j < k$, p_{ik} is independent of p_{ij}), the expected number of citations of patent i can be computed simply as

$$(B2) \quad C_i^e = E(C_i) = \sum_{j=i+1}^N p_j$$

Obviously, $C_i^e > C_j^e$ for any $i < j$, i.e. older patents will get on average more citations than recent ones, just by virtue of their age. Notice also that p_j has to decrease eventually with j (unless r_j were to increase indefinitely over time, but in fact r_j is quite stable), thus reinforcing the pure 'age effect'. In other words, not only do later patents miss the earlier p_j 's, but those probabilities tend to be the large ones, a fact that further reduces the expected number of citations of recent patents vis a vis older ones.

In order to actually perform the χ^2 test, the data are aggregated by months, since it would be unreasonable to attach any significance (in the sense of differences in C_i^e) to the precise day of application. Indexing by r and t the number of months since 1/72, the observed (or actual) number of citations is $C_t^o = \sum_{i=1}^{n_t} C_i$, where n_t is the number of patents in month t .

Similarly, and redefining (B1) in monthly terms,

$$(B1)' \quad p_r = \sum_{i=1}^{n_r} r_i / \sum_{j=1}^{r-1} n_j$$

and, accordingly,

$$(B2)' \quad C_t^e = n_t \sum_{r=t+1}^T p_r$$

Turning now to the test itself,

(j-i), making the distribution of expected citation more uniform, and hence making it easier to reject the null hypothesis.

$$(B3) \quad \chi^2 = \frac{\sum_{t=1}^{155} (C_t^e - C_t^o)^2}{C_t^e} = 1025 > > 148 = \chi^2_{(111)} \quad \alpha=0.01$$

Thus, the hypothesis that the observed distribution of citations is due just to age is strongly rejected. As is to be expected, the largest discrepancies between actual and expected values occur at the very beginning of the period. In particular, the values for the first patent are $C_1^o = 72$, $C_1^e = 5.96$, and hence $(C_1^e - C_1^o)^2 / C_1^e = 731$, which amounts to 3/4 of the computed χ^2 . Since this first patent can be regarded in many ways as an exception, the test was redone after deleting it, and again the null hypothesis is rejected by a wide margin.

B.2 Assessing the truncation bias

Now to the other potential problem in this context, namely, the fact that the - unavoidable - truncation of the data might induce a bias in the citation counts. Of course, the further back in the past the period studied is, the less there would be reason for concern. For a given distance in time, though, the extent of the bias will be determined by the behavior of citation lags, and by the rate of new patent arrivals after the date of search. Citation lags refer to the length of time elapsed between the dates of the citing and of the cited patent: the shorter they are, the less severe the problem will be. Table B.1 presents the distribution of these lags, by year of the cited patents: for example, the 1975-77 patents were subsequently cited 508 times, 12.6% of those citations occurring during the first year following the application date, 37.4% in the course of the second year, etc. the mean lag being almost of 3 years.

TABLE B.1

Frequency Distribution of Citations Lags, by Year of Cited Patent

lag in yrs ^a	1972-74	1975-77	1978-80	1981-82	%	all cumulative %
1	2.0	12.6	7.7	25.5	8.7	8.7
2	15.8	35.6	37.4	75.0	29.4	38.1
3	23.2	25.4	19.8	0	24.0	62.1
4	22.6	15.2	12.1	0	17.2	79.3
5	18.2	4.5	13.2	0	9.9	89.2
6	9.1	3.1	6.6	0	5.4	94.7
7	4.4	2.0	3.3	0	2.9	97.6
8	3.0	1.4	0	0	1.8	99.3
9	1.0	0.2	0	0	0.4	99.8
10	0.7	0	0	0	0.2	100.0
number of citations	297	508	91	4	900	
mean lag (in yrs.)	4.1	2.9	3.2	1.8	3.3	

^aThe lags have been computed on the basis of monthly data, so that a one year lag means the interval 0 - 11 months, a three year lag 24 - 35 months, etc.

Note that the frequency distribution of citation lags for all patents is very skewed to the left, most citations occurring within the first 3-4 years after a patent has been applied for, and the process dwindling down to a trickle after 5-6 years³². In particular, this is true for the distribution of lags of the 1975-77 patents, which is arguably the most 'representative' period in this context. As for the maximum lag, it seems quite certain that it does not exceed 10 years, judging from the evidence of the initial years (1972-74), for which the maximum lag could have been significantly longer (11-14 years).

So far the qualitative evidence seems to indicate that the truncation problem is not too severe; still, we need actual estimates of the biases in order to make a final judgement. Denote by f_r the frequency distribution of citation lags, i.e. if year t patents are to receive (on average) C_t citations per patent, f_r stands for the percentage of those citations to be received after r years (obviously, $\sum_{r=t}^{\infty} f_r = 1$). Likewise, define $c_{tr} = f_r C_t$ and $g_{tr} = c_{tr} / n_r$, where n_r is as before the total number of patents in year r . Now, suppose that because of truncation, one can actually obtain only a fraction h_r of them; then, assuming that g_{tr} is invariant with respect to h_r (i.e., that citations to year t patents are randomly distributed among the n_r patents), the observed average number of citations to year t patents will be: $c_{tr}^0 = g_{tr} h_r n_r = h_r f_r C_t$. Thus, given the sequences $\{h_r, f_r\}$, one can compute for each year the fraction $v_t = \sum_{r=t}^{\infty} h_r f_r$, that is, v_t stands for the percentage of citations that patents in year t can be expected to receive, out of the total

³²Campbell and Nieves [1979] report longer lags for the case of catalytic converters, but theirs refer to all citations (which would have longer lags indeed), rather than to 'within citations' only, as is the case here.

that they would have received had it not been for the truncation of the data.

The figures for h_r are obtained from the granting-application lags shown in section 3, e.g. $h_{83} = 0.76$, $h_{85} = 0.23$, etc. (obviously, for $r \leq 81$, $h_r = 1.00$, and for $r \geq 86$, $h_r = 0$); those for f_r are a slight variation of the citation lags displayed in Table B.1³³. The results are as follows:

Year of Cited patents	v_t	Number of Citations		
		actual	missing (rounded)	fraction missing
up to 75	1.000	491	0	0.00
76	0.998	169	0	0.00
77	0.990	145	1	0.01
78	0.969	55	2	0.04
79	0.930	29	2	0.07
80	0.861	7	1	0.14
81	0.732	3	1	0.33
82	0.527	1	1	1.00

Thus, I do miss a few citations because of the truncation of the data; more importantly, there is as expected a truncation bias, in the sense that there is a smaller fraction of the 'true' number of citations to later patents than to earlier ones. Notice, however, that the **absolute** expected number of missing citations is very small and that, even if the bias was somehow underestimated in those calculations by a factor of 2, the 'true' citations count would still differ only slightly from the count used here. Thus, it is clear that the truncation problem is mostly inconsequential for the

³³The citation lags were computed here as: (year of citing patent - year of cited patent), rather than according to their respective months, as in table B.1. The figures for v_t shown in the table are the averages of two values, one computed on the basis of the distribution f_r corresponding to the 1975-77 patents, and a second on the basis of the f_r 's for all patents.

computations and findings presented above.

Finally, and to press the truncation issue further, I asked the following question: given that the above calculations are done on the basis of averages, could it not be that one or more of the patents applied for in say, 1982, actually turned out to be **very important**, but it went undetected here because there are just 4 years of **granted** patents since? Quite clearly, one cannot rule that out with certainty, but as the following exercise indicates, such possibility is very unlikely: I took the three patents applied for in 1973 (each of which received a large number of citations), and counted the number of citations that they would have been given, if just the patents granted up to the end of 1977 had been available, thus replicating the situation now vis a vis the 1982 patents. The partial count for the first patent was of 7 citations vs. a true count of 11, of 5 vs. 18 for the second, and of 8 vs 21 for the third, i.e., the importance of those patents would have been recognized right away. I repeated the exercise with the latest two patents to receive more than 10 citations each (both were applied for in 1976), obtaining similar results: the restricted count was of 16 citations vs. a total of 20 for one, and of 6 vs. 13 for the other. These findings are important not so much for the statistical analysis, but rather in that they make sure that we get an accurate description of the evolution of the CT scanners over time. In other words, it is very unlikely that some major innovation occurred in the field of CT in the early eighties, and the patent data failed to detect it because of truncation. Instead, the field seems to be lingering on, as Figure 1 shows, giving way to MRI and other rising technologies.