

Working Paper No. E01-304

Market Value and Patent Citations: A First Look

Bronwyn H. Hall

Adam Jaffe and Manuel Trajtenberg

August 2001

Keywords: market value, patent citations, bibliometrics, innovation

JEL Classification: O31, O38

This is a revision of a paper prepared for the Conference in Commemoration of Zvi Griliches' 20 Years as Director of the NBER Program on Productivity and Technological Progress, Cambridge, Massachusetts, March 5 and 6, 1999. As should be clear from the discussion in Section 2, this paper represents but one further step in the research agenda sketched in Zvi's 1979 Bell Journal paper, reported on in the 1984 NBER volume Zvi edited, and continued by Zvi, his students and associates in the ensuing decade. Earlier versions of this paper were presented at the Conference on Intangibles and Capital Markets, New York University, May 15-16, 1998 and the Conference on the Economics of Science and Technology, University of Urbino, Italy, June 5-6, 1998. We have benefited from comments at these conferences and by seminar participants at Keele University, the University of Paris I, the University of Reading, the University of Manchester, University College London, and Oxford University. The data construction effort described in this paper was partially supported by the National Science Foundation, via grants SBR-9413099 and SBR-9320973. We are extremely grateful to Meg Fernando of REI, Case Western Reserve University for excellent assistance in matching the patenting data to Compustat. Correspondence: bhall@econ.berkeley.edu

This paper is available on-line at <http://www.haas.berkeley.edu/groups/iber/wps/econwp.html> or contact University of California Berkeley, IBER, Berkeley CA 94720-3880.

Abstract E01-304

As patent data become more available in machine-readable form, an increasing number of researchers have begun to use measures based on patents and their citations as indicators of technological output and information flow. This paper explores the economic meaning of these citation-based patent measures using the financial market valuation of the firms that own the patents. Using a new and comprehensive dataset containing over 4800 U. S. Manufacturing firms and their patenting activity for the past 30 years, we explore the contributions of R&D spending, patents, and citation-weighted patents to measures of Tobin's Q for the firms. We find that citation-weighted patent stocks are more highly correlated with market value than patent stocks themselves and that this fact is due mainly to the high valuation placed on firms that hold very highly cited patents. We also find that self-citations are worth about twice as much as ordinary citations, especially to smaller firms.

Bronwyn H. Hall

University of California, Berkeley
Nuffield College Oxford
NBER and IFS

Adam Jaffe

Brandeis University and NBER

Manuel Trajtenberg

Tel Aviv University and NBER

Market Value and Patent Citations: A First Look

Bronwyn H. Hall, Adam B. Jaffe, and Manuel Trajtenberg

May 2001

1. Introduction

Micro-level data on patents that include detailed technology field, citations to other patents, number of claims, geographical location, and a variety of other information are increasingly available in machine-readable form. For economists in the field of technical change and innovation, these data have enormous potential: in addition to providing rich technological, geographic and institutional detail, patent data are publicly available for all kinds of research institutions (firms, universities, other non-profits, and government labs) in virtually every country. At a general level, economists have used patents and/or patents weighted by subsequent citations to measure the inventive output of organizations or geographic units; they have used citation intensity or measures related to the nature of citations that an entities patents receive to measure the importance or impact of that entity's inventions; and they have used aggregate flows of citations to proxy for flows of knowledge to investigate knowledge spillovers across organizational, technological and geographic boundaries.

With a few exceptions discussed below, this work has relied on maintained hypotheses that patents are a proxy for inventive output, and patent citations are a proxy for knowledge flows or knowledge impacts. In this kind of work, these maintained hypotheses cannot really be tested, though they may be supported by results that are consistent with strong priors about the nature of the innovation process, and which are internally consistent. In this paper, we seek to strengthen the foundation for the use of patent and patent citation data, by exploring the extent to which firms' stock market value is correlated with their stocks of patents and patent citations. *Our* maintained assumption is that stock market investors hold rational expectations about the extent to which the present value of a firm's future profits varies with its stock of knowledge. Hence evidence that patent-related measures are correlated with market values represents evidence that they are proxies for the private economic value of the firm's knowledge stock. Previous work looking at the relationship between patents and market value suggested

that the extremely skewed nature of the value distribution of individual patents makes firm patent totals very noisy as an indicator of the value of firms' knowledge. In this paper, we explore the extent to which this problem can be mitigated by using citation-weighted patent counts, in the context of a larger and more comprehensive dataset than has been used before.

We begin the paper with a discussion of the meaning of patent citations and a brief survey of prior uses of these data for economic analysis. Then we review what is known about the relationship between patent counts and a firm's value in the financial markets. The next sections of the paper present our data, which is the product of a large-scale matching effort at the NBER and Case Western Reserve University, and the relatively simple "hedonic" model for market valuation that we use. The primary contribution of this paper, estimates of the market value equation that include patent citations, is contained in Section 6; the conclusions contain an extensive discussion of further work and refinements to be implemented in a revision of this paper. Appendices describe the construction of the data, and discuss the important issue of adjusting patent citation data for the truncation inherent in the fact that we cannot observe the entire citation life of patents, with the extent of this truncation increasing for more recent patents.

2. Prior Research using Micro-level Patent Data¹

A patent, as a matter of definition, is a temporary legal monopoly granted to inventors for the commercial use of an invention. In principle, in order to receive this right, the invention must be nontrivial, in the sense that it must not be obvious to a skilled practitioner of the relevant technology, and it must be useful, meaning that it has potential commercial value. If the patent is granted, an extensive public document is created which contains detailed information about the invention, the inventor(s), the organization to which the inventor assigns the patent property right (usually an employer), and the technological antecedents of the invention.² These antecedents, identified as references or citations, include previous patents and other published material that identify or describe aspects of the relevant technology that were previously publicly known. The citations identify "prior art," the practice of which is necessarily excluded from the

¹For more comprehensive literature reviews, see Griliches (1990) and Lanjouw and Schankerman (1999).

²See Appendix B for an example of the front page of such a document, in the form in which it appears on the publicly accessible website of the United States Patent Office (<http://www.uspto.gov>). Note that no public information is currently available for patent applications that are still pending, or for patent applications that are denied by the patent office.

property right granted by the patent. Thus, together with the language of the patent claims—which describe exactly what the patented invention does that has never been done before—the citations help to delimit the property right that the patent represents. As will be discussed further below, the cited patents can be identified by the inventor herself, by a search conducted by the inventor’s patent attorney, or by the patent examiner who reviews the application for the patent office.³

The use of patent data in the economic analysis of technological change has a fairly long, if somewhat unsatisfactory history, which stretches back to the pathbreaking analyses of Schmookler (1966) and Scherer (1965). The availability of information from the U.S. patent office in machine-readable form in the late 1970s spurred greater interest in econometric analyses using these data; much of the resulting early work is reported in Griliches (1984).⁴ In the late 1980s, patent citation information began to be available in computerized form, which led to a second wave of econometric research, utilizing the citation information to increase the information content of the patent data themselves, as well as to investigate an additional set of questions related to the flow of knowledge across time, space and organizational boundaries.⁵

2.1. Patent Citations

Viewed optimistically, patent citations can be seen as providing direct observations of technological impact and knowledge spillovers, in that one technological innovation explicitly identifies several others as constituting the technological state-of-the-art on which it builds. Unfortunately, this optimistic view is somewhat clouded by the reality that there is substantial “noise” in the patent citations data. The nature and extent of such noise depends, to some extent, on the purpose to which the patent data are put. Some authors have used these data to explore

³As can be seen in the patent in Appendix B., patents can make citations both to earlier patents and to non-patent publications. The non-patent citations appear in plain text form, and hence are difficult to manipulate electronically. Research that utilizes the non-patent references includes Trajtenberg, Henderson and Jaffe (1997) and Narin et al (1997).

⁴See also Pakes (1986) and Griliches, Hall and Pakes (1987). Also in the late 1970s Mark Schankerman and Ariel Pakes pioneered the use of renewal data from the European patent office to estimate the value distribution of patents. (Pakes and Schankerman 1984, Schankerman and Pakes 1986). (Renewal of U.S. patents was not required before the mid-1980s.)

⁵It is perhaps interesting to chart the effect of computerization on research via the authors’ experience with the acquisition of patent data. In the early 1980s, Trajtenberg collected citations information for hundreds of CT-scanner patents by hand from paper patent documents. In 1989, we paid \$10,000 to a private data firm for citation information on about 10,000 patents. In the mid-1990s, we began construction of a database with citations to about 2.5 million patents, using an NSF grant of about \$100,000. Today, citations to over 3 million U.S. patents are available free from numerous websites, and a CD-ROM is available from the authors with comprehensive information on all patents granted between 1964 and 1996, and all citations made between 1976 and 1996.

questions involving spatial spillovers (e.g., Jaffe, Trajtenberg, and Henderson 1993), knowledge flows among firms in a research consortium (e.g., Ham 1998), and spillovers from public research (e.g., Jaffe and Trajtenberg 1996; Jaffe and Lerner 1999). In using citations as evidence of spillovers, or at least knowledge flows, from cited inventors to citing inventors, it is clearly a problem that many of the citations are added by the inventor's patent attorney or the patent examiner, and may represent inventions that were wholly unknown to the citing inventor. On the other hand, in using citations *received* by a patent as an indication of that patent's importance, impact or even economic value, the citations that are identified by parties other than the citing inventor may well convey valuable information about the size of the technological "footprint" of the cited patent. That is, if a patent stakes out a territory in technology space that is later frequently deemed to abut areas that are patented in the future, this suggests that the cited patent is important, whoever it is that decides that the citation is necessary.

A recent survey of inventors sheds some qualitative light on these issues (Jaffe, Trajtenberg and Fogarty, 2000). Approximately 160 patentees answered questions about their inventions, the relationship of their inventions to patents that were cited by their patents, and the relationship to "placebo" patents that were technologically similar to the cited patents but which were not cited. The cited and placebo patents were not distinguished in the survey questionnaire, although it is possible that the surveyed inventors knew or looked up which patents they actually cited. The results confirm that citations are a noisy measure of knowledge flow, but also suggest that they do have substantial information content. Overall, as many as half of all citations did not seem to correspond to any kind of knowledge flow; indeed, in a substantial fraction of cases the inventors judged that the two patents were not even very closely related to each other.⁶ At the same time, the answers revealed statistically and quantitatively significant differences between the cited patents and the placebos with respect to whether the citing inventor felt that she had learned from the cited patent, when she learned about it, how she learned about, and what she learned from it. Qualitatively, it appears that something like one-half of citations correspond to some kind of impact of the cited invention on the citing inventor, and something like one-quarter correspond to fairly rich knowledge flow, fairly significant impact, or both.

There are also a small number of studies that "validate" the use of citations data to measure

⁶The results also confirmed that the addition of citations by parties other than the inventor is a major explanation for citations that do not correspond to knowledge flow. About 40% of respondees indicated that they first learned of the cited invention during the patent application process.

economic impact, by showing that citations are correlated with non-patent-based measures of value.⁷ Trajtenberg (1990) collected patents related to a class of medical instruments (computerized tomography, or "CAT" scanners), and related the flow of patents over time to the estimated social surplus attributed to scanner inventions. When simple patent counts are used, there is essentially no correlation with estimated surplus, but when citation-weighted patent counts are used, the correlations between welfare improvements and patenting are extremely high, on the order of 0.5 and above. This suggests that citations are a measure of patent "quality" as indicated by the generation of social welfare. Interesting recent work by Lanjouw and Schankerman (1997, 1999) also uses citations, together with other attributes of the patent (number of claims and number of different countries in which an invention is patented) as a proxy for patent quality. They find that a patent "quality" measure based on these multiple indicators has significant power in predicting which patents will be renewed, and which will be litigated. They infer from this that these quality measures are significantly associated with the private value of patents. With respect to university patents, Shane (1999a, 1999b) finds that more highly cited M.I.T. patents are more likely to be successfully licensed, and also more likely to form the basis of starting a new firm. Sampat (1998) and Ziedonis (1998) explore the relationship between citations and licensing revenues from university patents. Harhoff et al (1999) surveyed the German patentholders of 962 U. S. invention patents that were also filed in Germany, asking them to estimate at what price they would have been willing to sell the patent right in 1980, about three years after the date at which the German patent was filed. They find both that more valuable patents are more likely to be renewed to full term and that the estimated value is correlated with subsequent citations to that patent. The most highly cited patents are very valuable, "with a single U.S. citation implying on average more than \$1 million of economic value" (Harhoff, et al 1999).

2.2. Market Value and Patents

Until recently, research that uses patents in the market value equation (in addition to or in place of R&D) has been somewhat limited, primarily because of the difficulty of constructing firm datasets that contain patent data. Most of the work shown in Table 1 and described here has been done by Griliches and his coworkers using the database constructed at the NBER that

⁷We are not aware of any studies that validate (by reference to non-patent data) the use of citations to trace knowledge flows. Since it is hard to find other measures or proxies for knowledge flows, this kind of validation is inherently difficult.

contained data on patents only through 1981. This dataset did not include information on the citations related to the patents. The other papers in the table use a cross section constructed by Connolly et al. for 1997 of Fortune 500 companies, and datasets involving UK data, one of which uses innovation counts rather than patents.

[Table 1 about here]

When patents are included in a market value equation, they typically do not have as much explanatory power as R&D measures, but they do appear to add information above and beyond that obtained from R&D, as one would expect if they measure the "success" of an R&D program. Griliches, Hall, and Pakes (1987) show that one reason patents may not exhibit very much correlation with dollar-denominated measures like R&D or market value is that they are an extremely noisy measure of the underlying economic value of the innovations with which they are associated. This is because the distribution of the value of patented innovations is known to be extremely skew, i.e., a few patents are very valuable, and many are worth almost nothing. Scherer (1965) was one of the first to make this point, and it has recently been explored further by Scherer and his co-authors (Scherer 1998; Harhoff, et al 1999). Therefore the number of patents held by a firm is a poor proxy for the sum of the value of those patents and we should not expect the correlation to be high. If the number of citations received by a patent is indicative of its value, then weighting patent counts by citation intensity should mitigate the skewness problem and increase the information content of the patents. As will be shown below, the distribution of citations is also quite skewed, suggesting perhaps that it can mirror the value distribution.

Shane (1993) regresses Tobin's Q for 11 semiconductor firms between 1977 and 1990 on measures of R&D stock, patent stock, and patent stock weighted by citations and finds that the weighted measure has more predictive power than the unweighted measure, entering significantly even when R&D stock is included in the regression; that is, there is independent information about the success of R&D in the weighted patent count measure. Weighted patent counts are also more highly correlated with the R&D input measure than simple patent counts; this implies that ex ante more effort was put into the patents that ultimately yielded more citations. An implication of this finding is that the citations may be measuring something that is not just the luck of the draw; the firms may have known what they were shooting at. Austin (1993a, 1993b) finds that citation-weighted counts enter positively but not significantly (due to small

sample size) in an event study of patent grants in the biotechnology industry. This means they add a small amount of information beyond the simple fact of a patent having granted. Many of the important Austin patents were applied for after 1987, which makes this study subject to serious truncation bias (discussed below). In addition, it uses the 3-day CAR (cumulative abnormal return) at the time of the patent grant as the indicator of economic value; this is an underestimate of the actual value of the patent because there is substantial informational leakage before . In fact, the work surveyed in Griliches, Pakes, and Hall (1987) typically finds that patent counts by application date are more tightly linked to market value than counts by granting date.

The present research project was inspired partly by the tantalizing results in these earlier small-scale studies that citation-weighted patent counts might prove a better measure of the economic value of innovative output than simple patent counts. Rather than working with a single product or market segment, we have assembled data on the entire manufacturing sector and the patenting and citing behavior of the firms within it, taking advantage of the computerized databases now available from the United States Patent Office. The remainder of the paper discusses our dataset construction and some preliminary results.

3. Data

Because of the way in which data on patents are collected at the United States Patent Office, matching the patents owned by a firm to firm data is not a trivial task. Firms patent under a variety of names (their own and those of their subsidiaries) and the Patent Office does not keep a unique identifier for the same patenting entity from year to year. To construct our dataset, a large name-matching effort was undertaken that matched the names of patenting organizations to the names of the approximately 6000 manufacturing firms on the Compustat files available to us and to about 30,000 of their subsidiaries (obtained from the Who Owns Whom directory). As described in Appendix A, the majority of the large patenting organizations have been matched to our data, or we have established why they will not match (because they are foreign or non-manufacturing corporations). However, budget constraints mean that our matching is primarily a snapshot exercise conducted using 1989 ownership patterns; we have limited evidence using patents in the semiconductor industry that this leads to some undercounting of patents for some firms.⁸ Thus the precise results here should be viewed with some caution: they are unlikely to

⁸See Hall and Ham (1999).

change drastically, but they may be affected by a slight undercount of the firms' patents.

The firm data to which we have matched patenting information is drawn from the Compustat file. The full sample consists of over 6000 publicly traded manufacturing firms with data between 1957 and 1995. After restricting the sample to 1965-1995, dropping duplicate observations and partially-owned subsidiaries, and cleaning on our key variables, we have about 4800 firms in an unbalanced panel (approximately 1700 per year). The variables are described in somewhat more detail in the appendix and the construction of R&D capital and Tobin's q is described fully in Hall (1990). For the purposes of this paper, we constructed patent and citation-weighted patent stocks using the same declining balance formula that we used for R&D, with a depreciation rate of fifteen percent. Our patent data goes back to 1964, and the first year for which we used a patent stock variable in the pooled regressions was 1975, so the effect of the missing initial condition should be small for the patent variable. However, we only know about citations made by patents beginning in 1976.

Figure 1 shows the fraction of firms in our sample in a given year who reported R&D expenditures to the SEC, the fraction that year who applied for a patent that was ultimately granted, and the fraction who have a nonzero stock of patents in that year.⁹ The increase in R&D reporting between 1969 and 1972 is due to the imposition of FASB rule no. 2, which mandated the public reporting of "material" R&D expenditures. The fall in the later years in the number of firms with patent applications is due to the fact that we only know about patent applications when they are granted, and our grant data end in 1996. The fact the share of firms with patent applications and with R&D spending is approximately the same is only a coincidence: although there is substantial overlap, these samples are not nested. 19 percent of the firms have R&D stocks and no patents while 13 percent have patent stocks but no R&D.

[Figure 1 about here]

We want to construct a citation-weighted stock of patents held by the firm, as a proxy for its stock of knowledge. This requires that we have a measure of the citation intensity for each patent that is comparable across patents with different grant years. The difficulty is that, for

⁹The stock of patents is defined using a declining balance formula and a depreciation rate of 15 percent, by analogy to the stock of R&D spending:

$$PS_t = 0.85PS_{t-1} + P_t \tag{3.1}$$

each patent, we only observe a portion of the period of time over which it can be cited, and the length of this observed interval varies depending on where the patent’s grant date falls within our data. For patents granted towards the end our data period, we only observe the first few years of citations. Hence, a 1993 patent that has gotten 10 citations by 1996 (the end of our data) is likely to be a higher citation-intensity patent than one that was granted in 1985 and received 11 citations within our data period. To make matters worse, although our basic patent information begins in 1964, with respect to citations we only have data on the citations made by patents beginning in 1976. Hence for patents granted before 1976, we have truncation at the other end of the patent’s life—a patent granted in 1964 that has 10 citations between 1976 and 1996 is probably more citation-intensive than one granted in 1976 that has 11 citations over that same period.

Our solution to this problem is to estimate from the data the shape of the citation-lag distribution, i.e. the fraction of lifetime (defined as the 30 years after the grant date) citations that are received in each year after patent grant. We assume that this distribution is stationary and independent of overall citation intensity. Given this distribution, we can estimate the total (20-year lifetime) citations for any patent for which we observe a portion of the citation life, simply by dividing the observed citations by the fraction of the *population* distribution that lies in the interval for which citations were observed. For patents where we observe the “meat” of the distribution (roughly years 3-10 after grant), this should give an accurate estimate of lifetime citations. For other patents, particularly where we observe only the first few years of patent life, this will give a very noisy estimate of lifetime citations. Many patents receive no citations in their first few years, leading to a prediction of zero lifetime citations despite the fact that some patents with no citations in the first few years are eventually cited.¹⁰

The details of the estimation of the citation lag distribution and the derived adjustment to citation intensity are described in Appendix D. Figure 2 shows the ratio of total citations to total patents for the firms with patents in our data, both uncorrected and corrected for truncation. The raw numbers decline beginning in about 1983, because citations are frequently made more

¹⁰Another issue is that the number of citations *made* by each patent has been rising over time, suggesting a kind of “citation inflation” that renders each citation less significant in later years. It is hard to know, however, to what extent this increased intensity is an artificial artifact of patent office practices, and the extent to which it might represent true secular changes in patent impact. In this paper we choose not to make any correction or deflation for the secular changes in citation rates, with the cost that our extrapolation attempts become somewhat inaccurate later in the sample. For further discussion of this point, see Appendix D. For an attempt to econometrically separate such effects, see Caballero and Jaffe, 1993.

than 10 years after the original patent is issued, and these later citations are unobserved for patents at the end of the data period. The truncation-corrected citation intensity is flat after about 1988 and then begins to rise again. Recall that we date the patents by year of application so that a patent applied for in 1988 was most likely granted between 1989 and 1991 and hence effectively had only 4-6 years to be cited. In addition, the citing patents were also less and less likely to have been observed as we reach 1995 and 1996. Because of the increasing imprecision in measuring cites per patent as we approach the end of our sample period, our pooled regressions focus first on the 1976-1992 period, and then on the subset of years between 1979 and 1988.

[Figure 2 about here]

Figure 3 shows the total citation and patenting rates per real R&D spending for our sample. The patent counts are adjusted for the application-grant lag and the citation counts are shown both corrected and uncorrected. Although the earlier years (1975-1985) show a steady decline in patenting and citation weighted patenting per R&D dollar, one can clearly see the recent increase in patenting rate beginning in 1986-87 that has been remarked upon by other authors (Kortum and Lerner 1998, Hall and Ziedonis 2001). However, the yield begins to decline in about 1993, two years before the end of our sample, mostly because real R&D increases during that period. The corrected patent citation yield also begins to increase in 1986-87 but does not decline quite as much as the patent yield, reflecting an increase in citations per patent in the early to mid-nineties.

[Figure 3 about here]

Figure 4 provides some evidence on the skewness of the distribution of citations per patent. In this figure we plot a distribution of the number of citations received by each of the approximately one million patents that we have assigned to manufacturing corporations. Fully one quarter of the patents have no citations, 150,000 have only one, 125,000 have two, and 4 patents have more than 200 citations. Fitting a Pareto distribution to this curve yields a parameter of 1.8, which implies that the distribution has a mean but no variance. However, a Kolmogorov-Smirnov or other distributional test would easily reject that the data are actually Pareto.

The most cited patents since 1976 are a patent for Crystalline silicoaluminophosphates held by Union Carbide Corporation (227 citations through 1995) and a patent for a Transfer

imaging system held by Mead Corporation (195 citations through 1995).¹¹ In Appendix B, we show detailed information for the first of these patents obtained from the USPTO website. It is apparent from the list of citations that the patent is important because the compound it describes is used as a catalyst in many processes. This single example suggests already that a high citation rate may be correlated with the value of a patent right, because such a product is useful both directly (via sales to other users) and in licensing and cross-licensing.¹²

[Figure 4 about here]

4. Equation Specification and Estimation Results

4.1. The Market Value Equation

We use a specification of the firm-level market-value function that is predominant in the literature: an additively separable linear specification, as was used by Griliches (1981) and his various co-workers. The advantage of this specification is that it assumes that the marginal shadow value of the assets is equalized across firms. The model is given by

$$V_{it} = q_t(A_{it} + \gamma_t K_{it})^{\sigma_t} \quad (4.1)$$

where A_{it} denotes the ordinary physical assets of firm i at time t and K_{it} denotes the firm's knowledge assets. Both variables are in nominal terms.

Taking logarithms of both sides of equation 4.1, we obtain

$$\log V_{it} = \log q_t + \sigma_t \log A_{it} + \sigma_t \log(1 + \gamma_t K_{it}/A_{it}) \quad (4.2)$$

In most of the previous work using this equation, the last term is approximated by $\gamma_t K_{it}/A_{it}$, in spite of the fact that the approximation can be relatively inaccurate for K/A ratios of the magnitude that are now common (above 15 percent). In this formulation, γ_t measures the shadow value of knowledge assets relative to the tangible assets of the firm and $\sigma_t \gamma_t$ measures their absolute value.

¹¹These two patents are the third and fifth most cited overall. The other 3 in the top 5 were taken out before 1976, so they are not contained in the online database provided by the U.S. Patent Office (<http://patents.uspto.gov>).

¹²See Somaya and Teece (1999) for an interesting discussion of the IP value creation choice between production or licensing.

The coefficient of $\log A$ is unity under constant returns to scale or linear homogeneity of the value function. If constant returns to scale holds (as it does approximately in the cross section), the log of ordinary assets can be moved to the left hand side of the equation and the model estimated with the conventional Tobin's q as the dependent variable, as we do here. The intercept of the model can be interpreted as an estimate of the logarithmic average of Tobin's q for manufacturing corporations during the relevant period. Thus our estimating equation becomes the following:

$$\log V_{it}/A_{it} = \log Q_{it} = \log q_t + \log(1 + \gamma_t K_{it}/A_{it}) + \delta_t D(K_{it} = 0) \quad (4.3)$$

where the last term is included to control for the overall level of Q when either R&D or patents are missing. Theory does not give much guidance for the specification of intangible capital stocks and it is not clear how we should specify an equation for K that incorporates patents and citation-weighted patents as measures of intangible assets in addition to R&D stocks. There are at least two possible approaches. In our first approach, we simply assume that R&D stocks, patent stocks or citation stocks are all just different measures of the same thing, and compare their performance in a $\log Q$ equation like equation 4.3, one at a time. This is the simplest way to validate our measures and compare their performance, but it implies that each stock (R&D, patent, or citation) is just another indicator of the same underlying concept, the knowledge assets of the firm. In our second approach to the problem, we ask what incremental value patents add in the presence of R&D stocks, and similarly for citations in the presence of patent stocks.

4.2. Citation-weighted patent stocks

The central problem we face in estimation is how to model the stock of intangible assets that is associated with the patents owned by a particular firm. We know that firms apply for patents for a variety of reasons: to secure exclusive production marketing rights to an invention/innovation, to obtain a currency that can be used in trading for the technology of other firms, to serve as a benchmark for the productivity of their research staff, and so forth. We also know that firms in different technology areas have substantially different propensities to patent. For the valuation function, we want a measure of the "book" value of the knowledge capital owned by the firm. That is, ideally we would like to know the cost in current prices of reproducing the knowledge this firm has of how to make new products today and how to undertake future innovation.

When we use past R&D expenditure to proxy for the book value of knowledge capital, we are implicitly assuming that a dollar is a dollar, i.e., that each dollar spent on research generates the same amount of knowledge capital. The reason one might want to use patents as a proxy for knowledge capital is because a patent could represent the "success" of an R&D program. That is, some of the R&D undertaken by the firm produces "dry holes" and although the knowledge gained by doing that research may have some value, such R&D should not be weighted equally with successful innovation-producing R&D in our measure of knowledge capital. Our problem is that to the extent that patents are used as engineer productivity measures and as a currency for technology licensing exchanges, some of the patents held by a firm may represent the same kind of "dry hole," in the sense that they document technological avenues that turn out not to be productive. More generally, it is clear from the work cited earlier that the private value distribution of the patent right is extremely skewed, making simple counts a noisy measure of value.

The idea of using subsequent citations to a patent as a measure of the patent's value rests on the argument that valuable technological knowledge within the firm tends to generate patents that future researchers build on (and therefore cite) when doing their own innovation. The example we gave earlier, the highly cited patent for Crystalline silicoaluminophosphates applied for by Union Carbide in 1984 (and subsequently granted), suggested that this could be the case. From the abstract and the citing patents it is clear that this class of chemicals has widespread use as a catalyst in chemical reactions, which doubtless creates value for the holder of the patent.

Appendix C presents the details of the construction of our citation-weighted patent stocks. Because citations can happen at any time after a patent is applied for,¹³ a natural question is whether we should use citation weights based on all the citations to patents applied for this year and earlier, or whether we should use only citations that have already occurred. That is, do the citations proxy for an innovation value that is known at the time the patent is applied for, or do they proxy for the future value of the innovation, for which the current market value of the firm is only an unbiased forecast? We attempt to explore this question by dividing our stocks into past and future. First, we construct the "total" citation stock for a given firm as of a given date, based on the number of citations made through 1996 to patents held by the firm

¹³There is at least one citation in our sample that is over 50 years old, to a patent that was applied for in 1921 and granted in 1992! Such very long grant lags usually are the result of the "continuation" process allowed by the patent rules, under which an inventor can file a modified patent application that retains the application date of the original.

as of the given date (depreciated). Then, we construct the "future" stock, as the difference between the total citation stock as of the date, and the stock based only on citations that were actually observed by the given date. The latter variable represents the future citations that will eventually be made to patents already held by the firm.

4.3. Basic Results

In Table 2, we show the results from running a "horse race" between R&D stocks, simple patent stocks, and citation-weighted patent stocks on data pooled across two subperiods (1976-1984 and 1985-1992).¹⁴ As others (including some the present authors) have found before, R&D stock is more highly correlated with market value than either patents or citations, but it is also clear that citation-weighted patent stocks are more highly correlated than patents themselves (compare the R-squares).

[Table 2 about here]

Comparing the coefficients in these equations is somewhat difficult, because the units are not the same. The coefficient of the R&D stock/assets ratio is in units of dollar for dollar, i.e., market value per R&D dollar, whereas that for the patent stock/assets ratio is in units of millions of dollars per patent. One possibility is to normalize the patent coefficients by the average or median patent per million R&D dollars or citation per million R&D dollars in the sample. Because of the presence of many zeros and the skewness of both the patent and citation distributions, neither measure is very robust, so we have used the ratio of the total patent stock or total citation stock to the total stock of R&D for all firms, rather than the average of these ratio across firms. For the firms in the first period these numbers are 0.62 (that is, approximately 1.6 million 1980 dollars per patent) and 4.7 (that is, about 210,000 1980 dollars per citation). Using this method, the marginal shadow value of a patent (measured in R&D dollars) for this period is approximately 0.37 million 1980 dollars and the marginal shadow value of a citation (again, measured in R&D dollars) is about 0.50 million 1980 dollars. These numbers can be directly compared to the R&D coefficient of 1.75. The magnitudes suggest substantial downward bias from measurement error in the patents or citations variable and from the use of an *average* patent per R&D yield for normalization. It is noteworthy that the

¹⁴These estimates are computed holding γ_t and δ_t constant across the subperiods for simplicity. The R-squared graph shown later is based on estimates that allow the coefficients to vary over time.

citation coefficient is somewhat higher, and that the difference in explanatory power is more marked for the firms that patent.

Although the results in Table 2 are somewhat encouraging, the extremely oversimplified equation we are using here is likely to obscure much that is of interest. In the next few sections of the paper we explore various ways of looking at this relationship in more detail. But first we examine how it has changed over time.

Figures 5a and 5b show the R-squared from the same simple Tobin's q regression on R&D stock, patent stocks, and citation-weighted patent stocks, estimated year-by-year between 1973 and 1993. Figure 5a shows the result for the R&D-performing firms and Figure 5b for the patenting firms. While neither patents nor citation-weighted patents have as great an explanatory power for the market to book ratio as R&D during the earlier years, by 1984-1986 the citation-weighted patents are doing about as well as R&D, especially when we focus on patenting firms, though this is partly because the explanatory power of R&D has declined.¹⁵

It is noteworthy that the date at which the explanatory power of citation-weighted patents converges to that of R&D in Figure 5b coincides roughly with a number of events that led to an increase in patenting activity during the mid-eighties, such as the Kodak-Polaroid decision. One interpretation is that patenting and citation behavior changed around this time because of changing litigation conditions. This might be explored further by looking more closely at which firms are making and receiving the citations. That is, does fear of litigation lead firms to cite others' patents more carefully in order to fence off their own technology?

[Figures 5a and 5b about here]

Because of the inaccuracy of our citation measures post-1990 and because the shadow value of our measures seems to change over time, in the remainder of the paper we focus on one ten-year period in the middle of our sample where the data are the most complete, and where the valuation coefficients do not change dramatically in Figures 5a and 5b, the 1979-1988 period. We also confine the sample to observations on the firms that have obtained at least one patent between 1975 and 1988 (the firms in the final columns of Table 2).

¹⁵See Hall (1993a,b) for year-by-year measures of the market value of R&D investments.

4.4. Explorations(1): Do citations add information?

A second, more informative way to look at the valuation problem is to hypothesize that although patents are clearly correlated with R&D activity at the firm level, they measure something that is distinct from R&D, either "success" in innovative activity, or perhaps success in appropriating the returns to such activity. This suggests that we should include the yield of patents or citations per R&D dollar as a separate variable in the equation. When interpreting the coefficient, it is important to note that R&D is a nominal quantity while patents are "real," so part of what we see is the changing real price of R&D over the sample. It is also likely that the expected yield of patents per R&D varies by industry, although we do not allow for this in the present paper.

In Table 3, we explore two variations of our base equation. In column 2, we add a patent yield (the ratio of patent stock to R&D stock) to the equation that already has R&D stock and find that it is significant and has a small amount of additional explanatory power.¹⁶ In column 3, we add the firm's average cites per patent to the equation to see if the citation rate has any impact on market value above and beyond that due to R&D and patenting behavior. This variable is quite significant and its coefficient is fairly large. To interpret the results, we use the following expression for the semi-elasticity of market value or Q with respect to the citation-patent ratio:

$$\frac{\partial \log Q}{\partial C/P} = \frac{\gamma_2}{1 + \gamma_0 K/A + \gamma_1 P/K + \gamma_2 C/P}$$

where K is the R&D stock, P is the patent stock, and C is the citation stock. In the table below we show this quantity together with the corresponding quantity for the patents-R&D stock ratio, evaluated at a range of values of the independent variables:

	Mean	Median	Ratio of Totals	Standard Deviation
$K/A(\$M/\$M)$	0.31	0.10	0.15	0.57
$P/K(1/\$M)$	0.64	0.0	0.47	3.87
$C/P(1/\$M)$	4.90	3.88	7.48	6.87
Denominator	1.669	1.336	1.604	
$\frac{\partial \log Q}{\partial C/P}$	0.031	0.038	0.032	
$\frac{\partial \log Q}{\partial P/K}$	0.018	0.023	0.019	

¹⁶Note also that the dummy for not doing R&D is now significantly positive. This occurs because the patent yield variable (P/K) has been set to zero when the firm has no R&D stock. The interpretation is that the average market value effect of being a firm with patents that does no R&D is approximately 7 percent. There are 2,934 observations with no R&D in the current year; 1,960 also have no stock of R&D either.

Thus an increase of one citation per patent is associated with a three percent increase in market value at the firm level. This is a very large number and may be consistent with the "million dollar" citations reported by Harhoff et al. (1999). The value of additional patents per R&D is somewhat lower: an increased yield of one patent per million dollars of R&D is associated with a two percent increase in the market value of the firm. Note that the statistics in the table above make it clear that the ratios are far from normally or even symmetrically distributed, which suggests that some exploration of the functional form of our equation might be useful. We present a simple version of such an exploration later in the paper.

[Table 3 about here]

4.5. Explorations(2): When do citations add information?

Table 4 shows the results of an investigation into whether there is a difference between the market valuation of past and future citations. The answer is a resounding yes (see columns 2 and 5 of Table 4). Whether we include the citation stock alone (column 2) or use the full model with R&D and patents (column 5), the coefficient of a stock based only on future citations is equal to or greater than the coefficient of the stock based on all citations, and the coefficient of the past citation stock is negative and marginally significant or insignificant. The apparent implication is that future citations are more correlated than past citations with the expected profitability of the patent right.

[Table 4 about here]

Because the two stocks, past and future, are highly correlated measures of the same underlying quantity, this finding does not necessarily imply that citations are worthless for forecasting the value of the knowledge assets associated with patents or the expected profit stream from those assets. The past citation stock of a firm could be an excellent forecast of the future citations that are expected for its patent portfolio, even though it is not quite as good as knowing the future citations when predicting the firm's market value. To explore this idea, we decomposed the total citation stock into the part predicted by the past citation stock and the part that is not predicted:¹⁷

$$\begin{aligned} K^C(t) &= E[K^C(t)|K^{PPC}(t)] + K^C(t) - E[K^C(t)|K^{PPC}(t)] \\ &= \widehat{K}^C(t) + e_{K^c}(t) \end{aligned}$$

¹⁷We also included a full set of time dummies in the conditioning set.

where $K^C(t)$ denotes the total citation weighted patent stock at time t and $K^{PPC}(t)$ denotes the patent stock weighted by the citations received as of time t (see Appendix C for details on construction of these variables). The results of including the citation-assets and citation-patent ratios partitioned in this way are shown in columns 3 and 6 of Table 4. In both cases, the coefficient of the unexpected portion of the total citation stock is less than the coefficient of the future citation stock in the preceding column and the coefficient of the predictable portion of the total citation stock is significantly positive. In the citations/assets version (where we have no other information besides citations), the predictable portion of citations has more predictive power than the unpredictable part. The converse is true in our preferred specification (column 6). Thus, although future citations are a more powerful indicator of the market value of the patent portfolio held by these firms, past citations clearly also help in forecasting future returns.

Figures 6 and 7 show the coefficients that result when both past and future citation-weighted patent stocks or predictable and unpredictable citation-weighted stocks are included in the same Tobin's q equation year-by-year for the 1973-1993 period. In Figure 6, we see that the future citation-weighted patent stock is clearly preferred over the past and that the latter has a coefficient that is zero or negative. When we separate the citation stock via the orthogonal decomposition of predictable based on the past versus unpredictable, we find that both enter, but that the unpredictable portion has a higher shadow value in the equation and that the predictable portion behaves more or less like the total stock (Figure 7).¹⁸

[Figure 6 about here]

[Figure 7 about here]

4.6. Explorations(3): How do citations add information?

Our working hypothesis is that citations are an indicator of the (private) value of the associated patent right, and are therefore correlated with the market value of the firm because investors value the firm's stock of knowledge. For this reason, it is of interest to explore the question of the precise shape of the citation valuation distribution: does the fact that a firm's patents yield fewer citations than average mean that its R&D has been unproductive? How does the valuation change for firms with patents that have very high citation yields of the sort we saw in the inset portion of Figure 4? Table 5 explores these questions. We broke the average citation

¹⁸Interpretation here is a bit dicey. These coefficients ought to be normalized in some way to put them on a common ground.

stock per patent stock variable up into 5 groups: less than 4, 4-6 (the median for firms with patents), 6-10, 10-20, and more than 20 (see Table 5 for details). The groups are unequal partly because we were interested in the tail behavior. We then included dummy variables for four of the five groups in the valuation regression (the left-out category was 0-4 citations per patent).

[Table 4 about here]

The results are quite striking. For firms with less than the median number of citations per patent (6), it makes no difference how far below the median they fall; firms with 4-6 citations per patent have no higher value than firms with less than 4 (the left-out category). However, firms that average more than the median number of citations per patent have a very significant increase in market value, and one that appears to be approximately linear. The most dramatic effect is for the 573 observations (143 firms) with a stock of more than 20 cites per patent: the market value of these firms is 54 percent higher than would be expected given their R&D capital and their patent stock.

Further investigation of these 143 firms revealed the following: they are concentrated in computing, office equipment, semiconductor, and electronics (270 observations on 66 firms), pharmaceuticals and medical instruments (155 observations on 41 firms), and to a lesser extent, in textiles and apparel (36 observations on 7 firms), and machinery (33 observations on 8 firms). They include both quite small (so they have a very few highly cited patents) and medium to large firms (such as Intel, Compaq Computer, Tandem Computer, Alza Corp, and Signal Companies). It appears that the larger firms are primarily in the electronics sector, broadly defined, while those in the pharmaceutical sector that average a high citation rate are more likely to be smaller biotechnology firms. It should be kept in mind that we are focusing here on a period that spans the period during which profound changes in patenting strategy took place in some industries after the creation of the Circuit Court of Appeals and the well-known Kodak-Polaroid decision of 1986 - see Hall and Ziedonis 2001 *inter alia*, for discussion of this point).

4.7. Explorations(4): Does it matter who does the citing?

Some of the citations that appear in patents of corporations are to patents that were assigned to the same firm as the citing patent. We refer to these citations as "self citations." Because citations are generated by a complicated process involving the inventor, the patent attorney and the patent examiner, it is not clear a priori what interpretation to give to these self-citations.

One possibility is that they appear simply because patents of the same firm are well-known to the parties, or because of an inventor's desire to acknowledge colleagues. If so, then self-citations ought to be economically less significant than other citations. On the other hand, firms citing their own patents could be a consequence of the cumulative nature of innovation and the "increasing returns" property of knowledge accumulation, particularly within a narrow field or technology trajectory. Self-citations could suggest that the firm has a strong competitive position in that particular technology and is in a position to internalize some of the knowledge spillovers created by its development. This will imply both that the firm has lower costs because there is less need to acquire technology from others, and that they will earn higher profits or rent from their activities without encouraging entry. The presence of self-citations may thus be indicative of successful appropriation by the firm of cumulative impacts, while citations by others might indicate cumulative impacts that are spilling over to other firms. If so, then the *private* value of self-citations should be greater than that of other citations. Thus examining the relative value of self-citations to other citations pushes further our examination of the economic "information content" of citations.

In order to investigate this question, we add self-citation measures to our valuation equation. We measure the self-citation propensity in two ways: 1) the share of citations that are self-cites; and 2) the ratio of the stock of self-cites to the patent stock. Because the self-cites are already included in the citation stock, when we add the latter variable to the regression, its coefficient will represent the premium or discount associated with self-cites. We also include a dummy for having no self-cites.

Figure 8 shows the distribution of the share of the stock of citations that can be attributed to self-citations for firms in our sample. Clearly it is quite skew towards zero, but there are a considerable number of firms whose stock of citations has a significant share of their own patents. The median share for patent-holding firms is 6 percent, with an interquartile range of 2 to 14 percent. About 7 percent of the observations have less than 0.1% self-citations and for 1.5% of the observations more than half their citations are to themselves.

[Figure 8 about here]

Columns 2 and 3 of Table 6 show the results of estimating our preferred specification with these variables included (column 1 shows the baseline specification). Both forms of the self-cite variable have highly significant and positive coefficients in the market value equation. In

column 2, an increase of 0.2 in the share of self-citations is associated with a 5 percent increase in market value. In column 3, if each the patents held by the firm acquires an additional cite from another entity, market value increases 5 percent; if that citation is made by the firm itself, market value rises 10 percent. In both cases, having no self citations increases market value, a fact which implies some nonlinearity in the relationship. It is possible that having cites, but a low number of self-cites is associated with being in an active but very crowded technology area, whereas firms with no self-cites are also firms with few patents that compete in a different arena.

[Table 6 about here]

In column (4) we investigate the question of firm size and self-citation. Large established firms will have more self-cites simply because they have been in business a long time and have large R&D and large legal departments. For these firms, it is possible that self-citation is a less important indicator of value in the presence of the other patent-related measures. Column (4) confirms this hypothesis: the value-relevance of both self-cites and not having any self-cites declines with size. To help interpretation, in Figure 9 we show the self-citation effects as a function of firm size together with the size distribution for our firms, on the same horizontal logarithmic scale.

[Figure 9 about here]

Both effects are negative for large firms and positive for small firms. For a firm with \$10 million in sales, one more self-citation per patent adds about 16 percent to market value. For firms with sales greater than \$100 million, there is essentially no change in market value from having more self-citations per patent. On the other hand, for a firm with \$10 million in sales, having no self-citations means your market value is higher by 32%. For a firm with \$10 billion in sales, having no self-citations *subtracts* about 22 percent from market value. The point at which the impact of having no self-citations changes from positive to negative is at about \$500 million in sales. Thus the meaning of self-citation is very different for small and medium-sized firms, many of whom are younger vis-a-vis older target firms. For smaller firms, both having a self-citation rate or having no self-citations is positive for market value. For larger firms, the self-citation rate has little or no effect on market value and having no self-citations is negative.

5. Conclusion and Suggestions for Further Research

This paper is a "first look" at these data. We find that augmenting firms' patent counts with citation intensity information produces a proxy for the firms' knowledge stocks that is considerably more value-relevant than the simple patent count itself. It remains true that, for most of the time period, patent-related measures cannot win a "horserace" with R&D as an explainer of market value. But this should not surprise us. As emphasized by Sam Kortum in his comments on this paper, even if citations are a reasonably informative signal of success, this does not mean they will be more correlated with value than R&D, because optimizing firms will increase their R&D in response to success. The citation stock is also associated with significant incremental market value *after* controlling for firms' R&D. Indeed, firms with very highly cited patents (more than 20 cites per patent), the estimates imply almost implausibly large market value differences, on the order of a 50% increase in value, relative to firm with the same R&D and patent stocks but with the median citation intensity

In addition to this confirmation that citations do contain useful incremental value-relevant information, we also have two intriguing findings regarding when citations are most valuable: 1) market value is correlated, to a significant extent, with the portion of eventual citations that cannot be predicted based on past citations; 2) market value is positively correlated with the share of citations to a firm's patents that the firm itself has made. The first finding suggests that the market "already knows" much about the quality of inventions, which will ultimately be confirmed by the arrival of future citations that are "unexpected" in the sense of unpredictable based only on past *citation* information. This result clearly requires further exploration. First, it would be useful to explore the use of a functional form or normalization that would allow the relative value of past and future citations to be compared more directly, rather than just asking which adds more to the R-squared. In addition, one could ask how many years' worth of citations does one have to see to know most of what citations will eventually reveal. Is 10 years enough? What fraction of what you "know" (in the sense of correlation with market value) by knowing the lifetime citations do you "know" after 5 years? Also, one could explore whether this result is driven by the tail of the distribution, which we know is associated with much of the value. In other words, to what extent is it possible to predict that a patent will ultimately get >20 citations based only on the first few years' patents? Is it this difficulty of predicting the really big winners that makes the unpredictable portion of the citations total so

important?

The second finding, that self-citation is largely positive for value, opens up a very interesting avenue of research. The self-citation variable gives us a window into technological competition, in the sense that it may inform us about the extent to which firms have internalized knowledge spillovers, or the strength of their competitive position vis-a-vis other firms in their industry. Future work should explore the ways in which this finding varies by industry and technology field, and the meaning of the size relationship that we found.

Other variations on the results include more exploration of the shape of the citations-value relationship. Has the importance of highly cited patents changed over time with changes in the patent regime? As noted, the firm-years with a citation intensity above 20 include both small and large firms. It would be useful to sort out whether these are different from each other, and also the extent to which the results relating to an *average* citation intensity of more than 20 are themselves driven by a few patents in the extreme tail. Here again, it would be useful to explore other functional forms.¹⁹

In addition to these variations on the themes already struck, there are other aspects of citation behavior that ought to be value-relevant. One possibility is *generality*. Trajtenberg, Jaffe, and Henderson have proposed a measure of "generality," defined as (1 minus) the Herfindahl Index of concentration of citations over patent classes. The idea is that citations that are spread over a larger number of technological fields are more "general", and vice versa. In terms of impacting the market value of firms, though, one could hypothesize the following: for firms that concentrate in narrow fields of activity, more generality is bad, since the firm will not be able to appropriate the spillovers to other fields. For conglomerates, the opposite may be true. Thus, we could compute the average generality of patents for firm j in year t , and interact this variable with a dummy for whether or not the firm is a conglomerate. This may be tricky; see the strategy literature on diversification that occurs in order to exploit an innovation resource base (Silverman 1997). We may need to normalize generality as well, since the measure depends on the number of citations. This suggests both conceptual difficulty in separating the effects of citation intensity and generality, and also a complex truncation problem in the generality

¹⁹One issue that we have explored little in this paper but which deserves attention is the variation in measurement error across measures based on widely varying numbers of patents and citations. Because of the count nature of the underlying data, measures based on few citations or patents are inherently "noisier" than those based on a large number. See Hall (2000) for a discussion of this issue in the context of a concentration index based on patent counts.

measure itself.

6. Bibliography

Albert, M. B., D. Avery, F. Narin, and P. McAllister. 1991. "Direct Validation of Citation Counts as Indicators of Industrially Important Patents," *Research Policy* 20: 251-259.

Austin, David. 1994. "Patent Citations and Appropriability," Washington, DC: Resources for the Future Discussion paper No. 94-12 (February).

Austin, David. 1993a. "An Event Study Approach to Measuring Innovative Output: The Case of Biotechnology," *American Economic Review* 83 (May): 253-258.

_____. 1993b. Ph. D. thesis, Department of Economics, University of California at Berkeley.

Blundell, Richard, Rachel Griffith, and John van Reenen. 1995. "Market Share, Market Value, and Innovation in a Panel of British Manufacturing Firms." London: University College London Discussion Paper 95/19 (October).

Caballero, Ricardo, and Adam Jaffe. 1993. "Standing on the Shoulders of Giants: An Empirical Assessment of Knowledge Spillovers and Creative Destruction in a Model of Economic Growth," in O. Blanchard and S. Fischer, eds, *NBER Macroeconomics Annual*, 1993.

Chakrabarti, Alok K., Israel Dror, and Nopphadol Eukabase. 1993. "Interorganization Transfer of Management - An Analysis of Patent Citations for a Defense Firm," *IEEE Transactions on Engineering Management* (February 1993): 91-94.

Chaney, Paul K., Timothy M. Devinney, and Russell S. Winer. 1991. "The Impact of New Product Introductions on the Market Value of Firms," *Journal of Business* 64(4): 573-610.

Cockburn, Iain, and Zvi Griliches. 1988. "Industry Effects and Appropriability Measures in the Stock Market's Valuation of R&D and Patents," *American Economic Review* 78 (May): 419-23.

Cohen, Wesley M., Richard R. Nelson, and John Walsh. 1996. "A First Look at the Results of the Carnegie-Mellon Survey of Industrial R&D in the United States. Pittsburgh, PA: Carnegie-Mellon University. Photocopied.

Connolly, Robert A., and Mark Hirschey. 1988. "Market Value and Patents: A Bayesian Approach," *Economics Letters* 27(1): 83-87.

Connolly, Robert A., Barry T. Hirsch, and Mark Hirschey. 1986. "Union Rent Seeking, Intangible Capital, and Market Value of the Firm," *Review of Economics and Statistics* 68(4):

567-77.

Economist, The. 1993. "Patents: Hidden Agenda," *The Economist* (20 November): 100-101.

Fisher, Franklin M., and John J. McGowan. 1983. "On the Misuse of Accounting Rates of Return to Infer Monopoly Profits," *American Economic Review* 73 (2): 82-97.

Griliches, Zvi. 1961. "Hedonic Prices for Automobiles: An Econometric Analysis of Quality Change," reprinted in Griliches, Zvi (ed.), *Price Indexes and Quality Change: Studies in New Methods of Measurement*. Cambridge, Mass.: Harvard University Press, 55-87.

_____. 1981. "Market Value, R&D, and Patents," *Economic Letters* 7: 183-187.

_____ (editor). 1984. *R&D, Patents and Productivity*, University of Chicago Press.

_____. 1990. "Patents Statistics as Economic Indicators," *Journal of Economic Literature* 92: 630-653.

Griliches, Zvi, Bronwyn H. Hall, and Ariel Pakes. 1991. "R&D, Patents, and Market Value Revisited: Is There a Second (Technological Opportunity) Factor?" *Economics of Innovation and New Technology* 1: 183-202.

Griliches, Zvi, Ariel Pakes, and Bronwyn H. Hall. 1987. "The Value of Patents as Indicators of Inventive Activity," in Dasgupta, Partha, and Paul Stoneman (eds.), *Economic Policy and Technological Performance*. Cambridge, England: Cambridge University Press, 97-124.

Hall, Bronwyn H. 1990. "The Manufacturing Sector Master File: 1959-1987," Cambridge, Mass.: NBER Working Paper No. 3366 (May).

_____. 1993a. "The Stock Market Valuation of R&D Investment during the 1980s," *American Economic Review* 83: 259-264.

_____. 1993b. "Industrial Research during the 1980s: Did the Rate of Return Fall?" *Brookings Papers on Economic Activity Micro* (2): 289-344.

_____. 2000. "A Note on the Bias in the Herfindahl Based on Count Data," UC Berkeley: Photocopied.

Hall, Bronwyn H., and Rosemarie Ham Ziedonis. 2001. "The Patent Paradox Revisited: Determinants of Patenting in the U.S. Semiconductor Industry, 1980-94," *Rand Journal of Economics*, Vol. 32, No. 1.

Hall, Bronwyn H., and Daehwan Kim. 1997. "Valuing Intangible Assets: The Stock Market Value of R&D Revisited." UC Berkeley, Nuffield College, Harvard University, and NBER: work

in progress.

Harhoff, Dietmar, Francis Narin, F.M. Scherer, and Katrin Vopel. 1999. "Citation Frequency and the Value of Patented Inventions," *The Review of Economics and Statistics*, 81, 3, 511-515, August.

Hayashi, Fumio, and Tonru Inoue. 1991. "The Relation between Firm Growth and q with Multiple Capital Goods: Theory and Evidence from Panel Data on Japanese Firms," *Econometrica* 59 (3): 731-753.

Hirschey, Mark, Vernon J. Richardson, and Susan Scholz. 1998. "Value Relevance of Non-financial Information: The Case of Patent Data," University of Kansas School of Business: Photocopied.

Jaffe, Adam. 1986. "Technological Opportunity and Spillovers of R&D: Evidence from Firms' Patents, Profits, and Market Value," *American Economic Review* 76: 984-1001.

Jaffe, Adam, Michael Fogarty and Bruce Banks. 1998. "Evidence from Patents and Patent Citations on the Impact of NASA and Other Federal Labs on Commercial Innovation," *Journal of Industrial Economics*, June 1998

Jaffe, Adam and Joshua Lerner. 1999. "Privatizing R&D: Patent Policy and the Commercialization of National Laboratory Technologies," National Bureau of Economic Research Working Paper No. 7064 (April)

Jaffe, Adam and Manuel Trajtenberg. 1996. "Flows of Knowledge from Universities and Federal Labs: Modeling the Flow of Patent Citations over Time and Across Institutional and Geographic Boundaries," *Proceedings of the National Academy of Sciences*, 93: 12671-12677, November 1996

----- . 1999. "International Knowledge Flows: Evidence from Patent Citations," *Economics of Innovation and New Technology*, 8: 105-136.

Jaffe, Adam, Manuel Trajtenberg and Michael Fogarty. 2000. "The Meaning of Patent Citations: Report of the NBER/Case Western Reserve Survey of Patentees," National Bureau of Economic Research Working Paper No. 7631 (April)

Jaffe, Adam, Manuel Trajtenberg, and Rebecca Henderson. 1993. "Geographic Localization of Knowledge Spillovers as Evidenced by Patent Citations," *Quarterly Journal of Economics* 108(3), August.

Kortum, Samuel, and Joshua Lerner. 1998. "Stronger Protection or Technological Revolution: What is Behind the Recent Surge in Patenting?," *Carnegie-Rochester Conference Series*

on *Public Policy* 48: 247-304.

Lanjouw, Jean O., and Mark Schankerman. 1999. "The Quality of Ideas: Measuring Innovation with Multiple Indicators," National Bureau of Economic Research Working Paper No. 7345 (September)

----- . 1997. "Stylized Facts of Patent Litigation: Value, Scope, and Ownership," London: LSE STICERD Working Paper No. EI/20 (January 1998).

Levin, Richard, Alvin Klevorick, Richard R. Nelson, and Sidney G. Winter. 1987. "Appropriating the Returns from Industrial Research and Development," *Brookings Papers on Economic Activity* 1987 (3): 783-820.

Licht, Georg and Konrad Zos (1996), "Patents and R&D: An Econometric Investigation Using Applications for German, European, and U.S. Patents by German Companies," ZEW, Mannheim: Photocopied.

Megna, Pamela, and Mark Klock. 1993. "The Impact of Intangible Capital on Tobin's Q in the Semiconductor Industry," *American Economic Review* 83: 265-269.

Narin, Francis, Kimberly S. Hamilton and Dominic Olivastro. 1997. "The Increasing Linkage between U.S. Technology and Public Science," *Research Policy*, 26, 3, 317-330

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

Pakes, Ariel. 1986. "Patents as Options: Some Estimates of the Value of Holding European Patent Stocks," *Econometrica* 54: 755-784.

Pakes, Ariel, and Mark Schankerman. 1984. "The Rate of Obsolescence of Patents, Research Gestation Lags, and the Private Rate of Return to Research Resources," in Griliches, Zvi (ed.), *R&D, Patents, and Productivity*, Chicago: University of Chicago Press.

Sampat, Bhaven N. 1998. "Cite-Seeing: Patent Citations and Economic Value," New York: Columbia University. Photocopied.

Schankerman, Mark. 1998. "How Valuable is Patent Protection? Estimates by Technology Field," *Rand Journal of Economics* 29: 77-107.

Schankerman, Mark, and Ariel Pakes. 1986. "Estimates of the Value of Patent Rights in European Countries during the Post-1950 Period," *Economic Journal* 97: 1-25.

Scherer, F. M. 1998. "The Size Distribution of Profits from Innovation," *Annales d'Economie et de Statistique* 49/50: 495-516.

Scherer, F. M. 1965. "Firm Size, Market Structure, Opportunity, and the Output of Patented Innovations," *American Economic Review* 55: 1097-1123.

Scherer, F. M., Dietmar Harhoff, and Joerg Kukies. 1998. "Uncertainty and the Size Distribution of Rewards from Technological Innovation," University of Mannheim. Photocopied.

Schmookler, Jacob. 1966. *Invention and Economic Growth*, Cambridge, Mass.: Harvard University Press.

Shane, Hilary. 1993. "Patent Citations as an Indicator of the Value of Intangible Assets in the Semiconductor Industry." Philadelphia, PA: The Wharton School. Photocopied (November).

Shane, Scott. 1999a. "Selling University Technology," mimeo

Shane, Scott. 1999b. "Technological Opportunities and New Firm Creation," mimeo

Trajtenberg, Manuel. 1989. "The Welfare Analysis of Product Innovations, with an Application to Computed Tomography Scanners," *Journal of Political Economy* 97 (2): 445-479.

----- . 1990. "A Penny for Your Quotes: Patent Citations and the Value of Innovations," *Rand Journal of Economics* 21 (1): 172-187.

Trajtenberg, Manuel, Rebecca Henderson, and Adam Jaffe. 1997. "University versus Corporate Patents: A Window on the Basicness of Invention," *Economics of Innovation and New Technology* 5: 19-50.

Waugh, Frederick V. 1928. "Quality Factors Influencing Vegetable Prices," *Journal of Farm Economics* 10(2): 185-196.

Wildasin, David. 1984. "The q Theory of Investment with Many Capital Goods," *American Economic Review* 74 (1): 203-210.

Table 1
Market Value - Innovation Studies with R&D and Patents

Study	Country (industry)	Years	Functional Form	Other Variables	R&D Coeff	R&D Stock Coeff	Patent or Innov Coeff	Comments
Griliches 1981	US	1968-74	Linear (Q)	Time & Firm dummies, [log Q(-1)]		1.0-2.0	.08 to .25 ?	units appear to be 100 pats
Ben-Zion 1984	US	1969-76	Linear (V)	Ind dummies, Investment, Earnings	3.4 (0.5)		.065 (.055)	No time dummies?
Jaffe 1986	US	1973, 79	Linear (Q)	Time & tech dummies, C4, mkt share, Tech pool, interactions		7.9 (3.3)		3SLS even higher
Connolly, Hirsch, Hirschey 1986	US	1977	Linear (EV/S)	Growth,risk,age,Mkt share,C4,Adv, Union share, Ind dummies	7.0 (0.8)		4.4 (0.6)	Unexpected patents
Cockburn, Griliches 1988	US		Linear (Q)	Industry appropriability (Yale survey)		.9-1.2	0.1	patent coef. Is insignificant
Griliches, Pakes, Hall 1987	US							
Connolly, Hirschey 1988	US	1977	Linear (EV/S)	Growth, risk, C4, Adv	5.6 (0.6)		5.7 (0.5)	Bayesian estimation
Hall 1993a	US	1973-91	Linear (V)	Assets, Cash flow, Adv, Gr, time dummies	2.5-3.0 (.8)	0.48 (.02)		By year also
Hall 1993b	US	1972-90	Linear (Q)	time dummies	2.0-10.0	0.5-2.0		By year; LAD; absolute coeff
Johnson, Pazderka 1993	US							
Thompson 1993	US							
Megna, Klock 1993	semi- conductors	1977-90	Linear (Q)	Rivals R&D and patents		0.82 (0.2)	0.38 (0.2)	Patent stock
Blundell, Griffith, van Reenen 1995	UK	1972-82	Linear (V)	Time dummies,Assets,Mkt share			1.93 (.93)	Innovation counts
Stoneman, Toivanen 1997	UK	1989-95	Linear (V)	Assets,Debt,Growth,Mkt share,investment, Cashflow, time dummies, Mills ratio	2.5 (1.5)		insig.	Selection correction; by year

Table 2
"Horse-race" Regressions comparing R&D, Patents and Citations
 U.S. Manufacturing Firm Sample (Cleaned) - 1976-92
 Nonlinear Model with Dependent Variable = log Tobin's q

	All Firms			R&D-Doing Firms			Patenting Firms		
	Period: 1976-1984								
Number of observations	17,111			10,761			10,509		
R&D Stock/Assets	1.736 (.069)			1.741 (.070)			1.754 (.082)		
D(R&D=0)	.029 (.014)			--			0.024 (.017)		
Patent Stock/Assets	0.607 (.042)			0.493 (.042)			0.599 (.041)		
Cite Stock/Assets			0.103 (.006)			0.083 (.006)			0.102 (.006)
D(Pats=0)	0.120 (.013)	0.143 (.013)		0.252 (.019)	0.281 (.018)		--	--	
R-squared	0.228	0.110	0.138	0.249	0.127	0.161	0.231	0.129	0.180
Std. Err.	0.669	0.719	0.707	0.683	0.737	0.722	0.644	0.685	0.665
Ratio of Total Pats or Cites to Total R&D (\$1980M)	0.566			0.550			0.619		
Coefficient scaled by ratio of totals	1.736	0.344	0.450	1.741	0.271	0.354	1.754	0.371	0.487
	Period: 1985-1992								
Number of observations	15,605			10,432			9,718		
R&D Stock/Assets	0.547 (.027)			0.560 (.027)			0.563 (.033)		
D(R&D=0)	0.004 (.015)			--			0.027 (.019)		
Patent Stock/Assets	0.710 (.049)			0.638 (.050)			0.710 (.049)		
Cite Stock/Assets			0.080 (.004)			0.073 (.005)			0.080 (.004)
D(Pats=0)	0.144 (.014)	0.151 (.014)		0.279 (.020)	0.295 (.019)		--	--	
R-squared	0.116	0.067	0.087	0.123	0.077	0.105	0.121	0.103	0.137
Std. Err.	0.748	0.769	0.760	0.764	0.784	0.772	0.729	0.737	0.723
Ratio of Total Pats or Cites to Total R&D (\$1987M)	0.331			0.324			0.352		
Coefficient scaled by ratio of totals	0.547	0.235	0.239	0.560	0.207	0.214	0.563	0.250	0.254

Heteroskedastic-consistent standard errors.

All equations have a complete set of year dummies.

Stocks are computed using 15 percent annual depreciation rate.

Table 3**Effect of Adding Patents and Citations to R&D Regression**

U.S. Manufacturing Firms (Cleaned Sample) - 1979-88 - 12,119 firm-years - 1,983 Firms
 Nonlinear Model with Dependent Variable = logarithm of Tobin's q

Independent Variable	(1) K	(2) K with P/K	(3) P/K and C/P	(4) K with C/K
R&D Stock(K)/Assets D(R&D=0)	1.198 (.056) 0.015 (.018)	1.276 (.061) 0.058 (.019)	1.366 (.076) 0.066 (.019)	1.296 (.061) 0.095 (.019)
Pat Stock/K Cite Stk/Pat Stk Cite Stock/K		0.0268 (.0062)	0.0306 (.0076) 0.0514 (.0039)	0.0092 (.0013)
R-squared	0.218	0.222	0.254	0.231
Standard error	0.687	0.685	0.671	0.681

Heteroskedastic-consistent standard errors in parentheses.

All equations include year dummies.

Stocks are computed using 15 percent annual depreciation rate.

Citation stocks are patent stocks weighted by all the cites they received before 1995 plus an estimate of post-1994 cites, depreciated as of the patent date.

Table 4
Splitting Total Citation Stock into Past and Future
 U.S. Manufacturing Firms (Cleaned Sample) - 1979-88 - 12,119 firm-years - 1,983 Firms
 Nonlinear Model with Dependent Variable = logarithm of Tobin's q

Independent Variable	Cite/Assets			P/K and C/P		
	(1)	(2)	(3)	(4)	(5)	(6)
R&D Stock(K)/Assets D(K=0)				1.366 (.076) 0.066 (.019)	1.267 (.070) 0.063 (.018)	1.136 (.059) 0.060 (.018)
Cite Stock/A	0.1048 (.0050)					
Past Cite Stk/A		-0.0554 (.0166)				
Future Cite Stk/A		0.1437 (.0073)				
Pred. Cite Stk/A			0.1140 (.0054)			
Unpred. Cite Stk/A			0.0770 (.0076)			
Pat Stock/K				0.0306 (.0076)	0.0284 (.0070)	0.0225 (.0057)
Cite Stk/Pat Stk				0.0514 (.0039)		
Past Cite Stk/P Stk					-0.0156 (.0068)	
Future Cite Stk/P Stk					0.0601 (.0046)	
Pred. Cite Stk/P Stk						0.0118 (.0037)
Unpred. C Stk/P Stk						0.0500 (.0030)
R-squared	0.203	0.212	0.207	0.254	0.260	0.260
Standard error	0.694	0.690	0.693	0.671	0.668	0.669

Heteroskedastic-consistent standard errors in parentheses.

All equations include year dummies.

Stocks are computed using 15 percent annual depreciation rate.

Citation stocks are patent stocks weighted by all the cites they received before 1995, depreciated as of the patent date (see the text).

Past citation stocks are stocks of citations that have already occurred as of the valuation date, depreciated as of the patent date.

Future citation stocks are the difference between citation stocks and past citation stocks.

Pred. and unpred. citation stocks are the orthogonal decomposition of citation stocks into the piece predictable from the past and the residual.

Table 5**The Shape of the Citations-Value Relationship**

U.S. Manufacturing Firms (Cleaned Sample) - 1979-88 - 12,119 firm-years - 1,983 Firms
 Nonlinear Model with Dependent Variable = log Tobin's q

Independent Variable	K and C/P	K, P/K, C/P	P/A only
R&D Stock (K)/Assets	0.865 (.047)	0.919 (.051)	
D(R&D=0)	-.006 (.017)	0.032 (.018)	
Pat Stock/A			0.623 (.039)
Pat Stock/K		0.025 (.006)	
4-6 Cites per Patent (3,145 observations)	0.002 (.018)	0.002 (.018)	-.004 (.019)
6-10 Cites per Patent (3,993 observations)	0.095 (.018)	0.095 (.018)	0.117 (.018)
10-20 Cites per Patent (1,997 observations)	0.353 (.023)	0.352 (.023)	0.456 (.021)
>20 Cites per Patent (573 observations)	0.546 (.043)	0.541 (.042)	0.796 (.032)
R-squared	0.252	0.255	0.230
Standard error	0.672	0.671	0.682

Heteroskedastic-consistent standard errors in parentheses.

All equations include year dummies.

Stock are computed using 15 percent annual depreciation rate.

Cite stocks are patent stocks weighted by all the cites they received before 1995 plus an estimate of post-1994 ci

The left-out category is fewer than 4 cites per patent (2,411 observations).

Table 6
Market Valuation of Self-Citations
 U.S. Manufacturing Firms (Cleaned Sample) - 1979-88 - 12,119 firm-years - 1,983 Firms
 Nonlinear Model with Dependent Variable = log Tobin's q

Variable	(1)	(2)	(3)	(4)
K/Assets	1.198 (.056)	1.400 (.080)	1.372 (.077)	1.075 (.068)
D(K=0)	0.015 (.017)	0.067 (.019)	0.070 (.019)	0.007 (.018)
Pat stock/K		0.0316 (.0078)	0.0314 (.0077)	0.0141 (.0057)
Citations per patent (stocks)		0.0538 (.0042)	0.0478 (.0041)	0.0439 (.0038)
Share of self citations (stocks)		0.248 (.096)		
Self citations per patent (stocks)			0.0543 (.0139)	0.0967 (.0160)
Log (S) = log(sales)-mean				-.0488 (.0058)
Log (S)*Self cites per patent				-.0237 (.0054)
D(no self citations)		0.166 (.051)	0.154 (.050)	0.119 (.047)
Log (S)*D(no self)				-.077 (.064)
R-squared	0.218	0.255	0.256	0.267
Standard error	0.687	0.671	0.670	0.665
Log likelihood	-12,644.4	-12,347.3	-12,340.5	-12,243.8
# parameters	2	6	6	9
Chi-squared versus col. (1)		594.2	607.8	801.2
Degrees of freedom		4	4	7
Chi-squared per d.f.		148.6	152.0	114.5

84 observations with 100% self-citations have been excluded.

All citation and patent variables are depreciated stocks.

The share of self citations is the fraction of citations that have been made to the patents assigned to the firm.

The self-citations per patent is the rate of self-citation per patent held (these are also included in citations per patent).

The size variable has its mean removed, so coefficients are relative to the average-sized firm.

The range of log(S) demeaned is -11 to 6.7, or \$10K to \$143B.

Figure 1
US Manufacturing - Cleaned Sample - 4,846 Firms




Figure 2

Citation Counts before and after Truncation Correction





Figure 3
U.S. Manufacturing Sector - 4,846 Firms



— Cites per real R&D —◆— Corrected cites per real R&D - - ○ - - Patents per real R&D

Figure 4 Citation Distribution




Figure 5a
R&D Performing Firms - R-Squared from Tobin's Q Equation




Figure 5b
Patenting Firms Only - R-Squared from Tobin's Q Equation




Figure 6
R&D Performing Firms - Splitting Citations Stocks
into Past and Future




Figure 7
R&D Performing Firms - Splitting Citation Stocks
into Predictable and Unpredictable Components




Figure 8






Figure 9
Self-citation Effects as a Function of Firm Size



7. Appendix A: Data Description

The data we use are drawn from the Compustat files and from files produced by the United States Patent Office. We have included all the firms in the manufacturing sector (SIC 2000-3999) between 1976 and 1995 in a large unbalanced panel (approximately 4800 firms). The firms are all publicly traded on the New York, American, and regional stock exchanges, or traded Over-the-Counter on NASDAQ. For details on data construction, see the documentation in Hall (1990), although we have drawn a new sample from a larger dataset than the file described in that document.

The main Compustat variables that we use are the market value of the firm at the close of the year, the book value of the physical assets, and the book value of the R&D investment. The market value is defined as the sum of the value of the common stock, the value of the preferred stock (the preferred dividends capitalized at the preferred dividend rate for medium risk companies given by Moody's), the value of the long-term debt adjusted for inflation, and the value of short-term debt net of assets. The book value is the sum of the net plant and equipment (adjusted for inflation), the inventory (adjusted for inflation), and the investments in unconsolidated subsidiaries, intangibles, and others (all adjusted for inflation). Note that these intangibles are normally the good will and excess of market over book from acquisitions, and do not include the R&D investment of the current firm, although they may include some value for the results of R&D by firms that have been acquired by the current firm. The R&D capital stock is constructed using a declining balance formula and the past history of R&D spending with a 15 percent depreciation rate.

The patents data have been cleaned and aggregated to the patent assignee level at the Regional Economics Institute, Case Western Reserve University. We have matched the patent assignee names with the names of the Compustat firms and the names of their subsidiaries in the *Who Owns Whom Directory of Corporate Affiliations* as of 1989 in order to assign patents to each firm. In order to ensure that we picked up all important subsidiaries, we also tried to positively identify the unmatched patenting organizations that had more than 50 patents during the period to ensure that we had not missed any subsidiaries. A spot check of firms in the semiconductor industry, which is an industry with lots of new entry during the period, suggests that our total patent numbers are fairly accurate for the period 1975-1995, but that

they are an undercount in the case of some firms (averaging about 5-15% under).²⁰

²⁰See Hall and Ham Ziedonis (1999).



Figure 8.1:

8. Appendix B: Highly Cited Patents

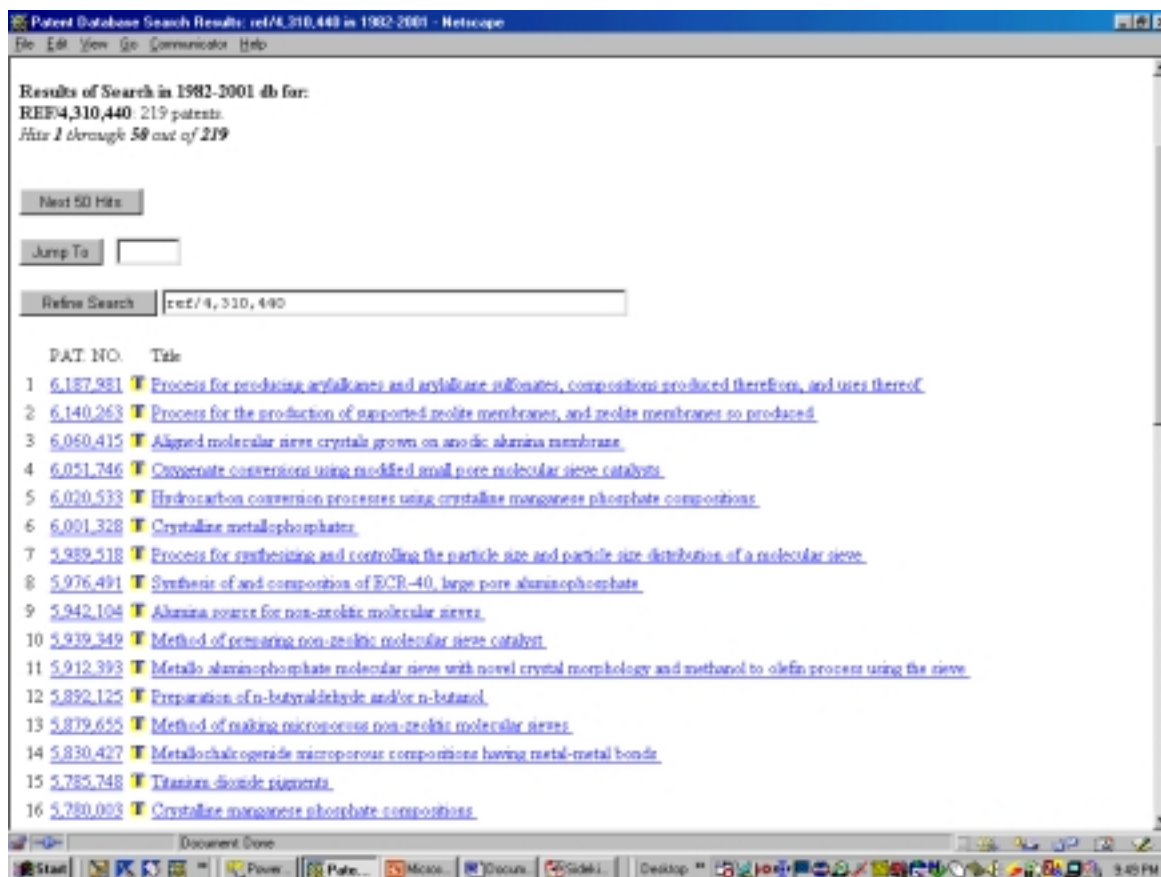


Figure 8.2:

Citations to Patent on Previous Page

9. Appendix C: Citation-weighted patent stocks

Define

$$C(t, s) = \text{number of cites received at time } s \text{ to patents applied for at time } t \quad (9.1)$$

and

$$C(t) = \sum_{s=t}^T C(t, s) = \text{total number of cites to patents applied for at time } t \quad (9.2)$$

Note that in our case $T = 1995$, so the measure of the total number of cites is truncated for all patents, and badly truncated for patents post 1985 or so.

Assume a single depreciation rate for the "private value" of a patent, δ . Then the simplest citation-weighted patent stock treats each patent as if it is worth the number of citations it ever receives and sums these citations over all the patents applied for in a given year to obtain a measure for the increment to the stock of knowledge at time t equal to $C(t)$. Using the standard declining balance formula, the stock of knowledge itself is defined by the following equation:

$$\begin{aligned} K^C(t) &= (1 - \delta)K^C(t - 1) + C(t) & (9.3) \\ &= C(t, t) + (1 - \delta)C(t - 1, t - 1) + (1 - \delta)^2C(t - 2, t - 2) + \dots \\ &\quad C(t, t + 1) + (1 - \delta)C(t - 1, t) + (1 - \delta)^2C(t - 2, t - 1) + \dots \\ &\quad C(t, t + 2) + (1 - \delta)C(t - 1, t + 1) + (1 - \delta)^2C(t - 2, t) + \dots \\ &\quad + \dots \end{aligned}$$

Note that in the above equation, the knowledge associated with a patent is assumed to be created at the time the patent is applied for, and future citations are depreciated as though they occurred at that time. Note also that when we put this stock into a valuation equation, we are implicitly assuming that the market knows the "true" value of the innovation, which is only revealed later to us via the citations.

An alternative possibility is to use citation weights that depend only on citations to the patents that have already occurred, and to depreciate the citations as of the date when they occur, rather than the date of the original patent application:

$$K^{PC}(t) = (1 - \delta)K^{PC}(t - 1) + C(t, t) + C(t - 1, t) + C(t - 2, t) + \dots \quad (9.4)$$

$$\begin{aligned}
&= C(t, t) + (1 - \delta)C(t - 1, t - 1) + (1 - \delta)^2C(t - 2, t - 2) + \dots \\
&\quad C(t - 1, t) + (1 - \delta)C(t - 2, t - 1) + (1 - \delta)^2C(t - 3, t - 2) + \dots \\
&\quad C(t - 2, t) + (1 - \delta)C(t - 3, t - 1) + (1 - \delta)^2C(t - 4, t - 2) + \dots \\
&\quad + \dots
\end{aligned}$$

Comparing equations (9.3) and (9.4) we can see that they differ not only because one is forward-looking (it counts citations at time t that will not be made until later than t), but also because the pattern of depreciation is different. This suggests an alternative to equation (9.4), shown below, that uses only past citations, but depreciates them as of the patent date rather than at the citation date:

$$\begin{aligned}
K^{PPC}(t) &= (1 - \delta)K^{PPC}(t) + \sum_{s=0}^{t-1} (1 - \delta)^s C(t - s, t) & (9.5) \\
&= C(t, t) + (1 - \delta)C(t - 1, t - 1) + (1 - \delta)^2C(t - 2, t - 2) + \dots \\
&\quad + (1 - \delta)C(t - 1, t) + (1 - \delta)^2C(t - 2, t - 1) + \dots \\
&\quad + (1 - \delta)^2C(t - 2, t) + \dots \\
&\quad + \dots
\end{aligned}$$

That is, the innovation is assumed to be valuable when it is made, but we don't learn about it until the citations happen. In the previous equation (9.4), the assumption is that the innovation becomes more valuable each time a citation is made. The advantage of the formulation in (9.5) is that it is nested within (9.3):

$$\begin{aligned}
K^C(t) &= K^{PPC}(t) + C(t, t + 1) + C(t, t + 2) + (1 - \delta)C(t - 1, t + 1) \\
&\quad + C(t, t + 3) + (1 - \delta)C(t - 1, t + 2) + (1 - \delta)^2C(t - 2, t + 1) + \dots \\
&= K^{PPC}(t) + \sum_{s=1}^T \sum_{j=0}^{s-1} (1 - \delta)^j C(t - j, t + s - j) \\
&= K^{PPC}(t) + K^{FPC}(t) & (9.6)
\end{aligned}$$

Thus this formulation will allow us to separate the total citation weighted stock of patents at time t into two components, one that contains only citation information prior to t , and one which contains only citation information at $t+1$ and later. However, it remains unclear whether we have chosen an appropriate depreciation structure.

10. Appendix D: Truncation Bias

We face at least two types of truncation bias in our measures: bias in our patent measures due to the fact that we have patent data only for granting years between 1967 and 1996 and bias in our citation measures because there is a long citation lag and therefore we will not have observed many citations for patents granted during the later part of our sample. The first bias is fairly straightforward and easy to correct for because the application grant lag is fairly stable and has a median of only about two years. The second is both more complex and of substantive interest, since it is unclear whether our forecast of future citations should agree with that of the market. Therefore we may wish to investigate the relationship between current citation stock and future realized citations in somewhat more detail.

In this section of the paper we outline the steps we take to correct for the first type of bias (due to the application-grant lag). Then we describe the measures of citation stocks that we construct for use in our regressions.

10.1. Patent truncation

Figure D.1 shows the average distribution of the lag between application and grant for all U.S. patents issued during the past four decades. The distributions are quite similar across decades, although there does seem to be a net reduction in the lag between the 1960s and the later periods.²¹ Except for the 1960s, 95 percent of patent applications that are eventually granted will be granted by year 3, and 99 percent by year 5. Our measure of patents in a year is the number applied for that are ultimately granted, so our goal is to adjust the granted-application count at both ends of our sample. The fact that the median lag is short means that in making our adjustment we cannot go back before an application data of 1964 (because we only have grants made in 1967 and later) and that we will not be able to adjust the patent counts beyond 1993 (because we only see about half the patents applied for in 1994 due to the grant lags). We describe these adjustments in more detail below.

[Figure D.1 about here]

At the beginning of the sample (1967) we observe some of the patents applied for in 1964-1966, but not the ones that were granted so quickly that their grant date is before 1967. This

²¹The data for the 1990s are based only on the first half of the period and therefore longer issue lags will be truncated, implying that the apparent shortening of the issue lag may be an artifact.

suggests that we might be able to correct our application counts for 1964-1966 (that is, fill in for lags 0 to 3) using weights drawn from the distribution for the 1960s. At the end of the sample, the opposite happens: there are patents applied for between about 1991 and 1996 that are still pending; some of them will be granted eventually, meaning that our counts of successful patent applications for those years are too small. Again, we can scale up the numbers we do have using the grant-lag distribution. Therefore we will compute the following two adjustments to our patent counts:

$$\begin{aligned}\tilde{P}_t &= \frac{P_t}{\sum_{s=67-t}^M w_s} & 64 < t < 67 \\ \tilde{P}_t &= \frac{P_t}{\sum_{s=0}^{94-t} w_s} & 91 < t < 94\end{aligned}\tag{10.1}$$

where P_t is the number of patent applications at time t , M is the maximum issue lag (assumed to be equal to about 10), and the weights w_s are weights constructed from the average lag distributions shown in Figure D.1. In Table D.1 we show the weighting factors we will use (the inverse of the expressions in equations (10.1)). Note the edge effects, which imply that the 1996 data will not be usable, and that the 1964 and 1995 data will have more variance due to estimation error.²²

[Table D.1 about here]

10.2. Citation truncation

As discussed in the text, we seek citation-based measures of patent quality that are comparable across patents. The number of citations that we observe for any given patent may depend on many things. These include:

1. Technological field;
2. Grant date or "vintage" of the patent;
3. The total number of years for which we have data on the patent's citations.

²²For this reason, we make no attempt to use data later than 1993 in the body of the paper (although we do use that data to construct stocks of future citations). The 1964 data is used only to the extent that it enters into the stocks of patents and citations that we construct. The first stock in our regressions is dated around 1973, so the measurement error effect should be quite small (recall that the counts are being depreciated by 15 percent per year).

The first two factors may well contain artificial influences due to variations in patent and citation practice across technologies and across time. For example, the number of citations *made* by each patent rose over much of this period, possibly due to changes in the ease of identifying prior art due to computer technology. Arguably, this "citation inflation" means that a typical citation given in the 1990s is less significant than the typical citation given in the 1970s. Similarly, there may be variations in norms of citation behavior that make a typical citation in a particular technological field more or less significant than the average. In the absence, however, of some external information about these artifactual variations, the only way one could deal with them would be to "take out" time and technology effects, by subtracting from the observed citation counts mean field and time effects. The problem with such an approach is that there is likely to be variation over time and across fields in the *true* importance of the typical patent. Taking out these effects would drastically reduce the variance in the data, probably throwing out a good part of the baby with the bathwater. Therefore in this paper we choose not to make any correction for technology field or secular citation trends.²³²⁴

We do, however, adjust for truncation due to the number of years of citations we actually observe, by fitting a model of the citation lag distribution introduced by Caballero and Jaffe (1993) and developed by Jaffe and Trajtenberg (1996). This model allows the average level of citing activity to vary across fields and times (either because of real or artifactual forces), and thereby estimates a citation lag distribution that is purged of these variations.²⁵

A patent can be cited at any time after it is issued. Although the majority of citations happen in the first 10 years of a patent's life, there is a long tail of citations that can occur into the infinite future.²⁶ Figure D.2 shows the number of citations received at each lag for patents

²³One might believe that secular trends associated with the year of the *citing* patent are associated with changes in citation practice, and hence artifactual, while secular trends associated with the *cited* year represent movements in the true average importance of new inventions. Separating the two, while simultaneously estimating the citation-lag distribution, requires strong functional form assumptions. See Caballero and Jaffe, 1993.

²⁴To the extent that the variation across time and fields is artifactual, it represents measurement error in the citation variables that reduces their explanatory power and biases their coefficients. One could construct the citation variables with and without purging citations of field and time effects; comparing the results might yield some insight as to the fraction of field and time variance that is real and the fraction that is artifactual.

²⁵The unbalanced nature of the patent data makes it important to control for field and time effects. For example, the only observations that we have on the longest lags are produced by citations made in years at the end of the data period. If the citation-lag distribution were estimated without controlling for the secular rise in citation rates, the estimated tail of the lag distribution would be too thick, because we only observe long lags from high-citation years.

²⁶Because we date patents by date of application, and grant lags are somewhat variable, we occasionally observe citations with negative citation lags, i.e. the citing patent has an application date before the cited patent. For example, there is one patent in our dataset (#4,291,005, belonging to Calgon Corporation) that was issued in 1992, but had originally been applied for in 1921 and then continued. This patent cited another patent that

in our sample applied for at different dates between 1965 and 1993. Because our citation data only go from 1976 until 1996, our citations are truncated at both ends: for patents applied for in 1965, we have only the citations at lag 10 and later, whereas for patents issued in 1993, our citations are abruptly truncated at lag 3. Thus we need to adjust for both these truncation biases.

[Figure D.2 about here]

In addition, as Jaffe, Trajtenberg, and Henderson (1993) have shown, citations vary considerably by technological fields. To adjust for this fact, we have grouped the 436 3-digit patent classes in our data into 6 major technological fields: drugs and medical, chemical exc. drugs, electrical and electronics, computers and communication, mechanical, and all other.²⁷ Our estimated citation model will include these effects, in order to prevent our estimate of the citation-lag distribution from being distorted by the combination of differences in citation practices across fields and changes in the field mix of patents over time.

We define our dependent variable to be the following:

$$\begin{aligned}
 C_{t,j,s} &= \Pr(\text{cite at lag } s | \text{patent applied for at } t \text{ in field } j) \\
 &= \frac{\# \text{cites at lag } s \text{ to patents applied for at } t \text{ in field } j}{\# \text{patents applied for at } t \text{ in field } j}
 \end{aligned}
 \tag{10.2}$$

where j is the one digit field of the patent, t is the year in which the patent was applied for, and s is equal to $t' - t$, where t' is the application year of the citing patent. A problem one confronts when trying to model the citation lag distribution is the fact that we cannot allow both the citing and cited year effects to change freely over time, because some years (notably those post-1996) are never observed, so their effects will not be identified. We use a slightly more flexible version of the model of Jaffe and Trajtenberg (1996) to solve this problem. Our model of the citation probability is then the following:

$$C_{t,j,s} = \beta_0 \alpha_{\tau(t)} \gamma_{t+s} \beta_j \exp(-\beta_1 \beta_j^S s) (1 - \exp(-\beta_2 \beta_j^D s))
 \tag{10.3}$$

had been issued in 1979, 13 years before Calgon's patent actually issued, but 58 years after the *citing* patent (Calgon's) had been applied for. Such anomalies are rare, and we have chosen to set the cite lag to zero in these cases.

²⁷This classification was performed by Gal Steinberg and Manuel Trajtenberg (Tel Aviv University) and is available from the present authors on request.

where $\tau(t)$ maps the years into five-year cohorts (see the table of results for details). As in Jaffe-Trajtenberg, this model constrains the citation-lag distribution to have the same shape for patents of any year, but allows the shape and level of the distribution to be different for each field and it allows the overall level to vary by 5-year patenting interval and by citation year. The only difference from their model is that in our most general version we have allowed both the diffusion process (the $\beta_2\beta_j^D$ term) and the obsolescence process (the $\beta_1\beta_j^S$ term) to vary by technology field, whereas they held the diffusion constant across field due to difficulties with identification.

Model (10.3) was fit to citation data aggregated by field (6 fields), application year (1963-1993), and citation year (1976-1994), yielding 2616 observations, each of which is the ratio of citations made with lag s to patents in field j applied for in the application year t . The results are shown in Table D.2. There are three sets of columns, the first for a model where diffusion and obsolescence effects do not vary across field ($\beta_j^S = 1 \forall j = 1, \dots, 5$ and $\beta_j^D = 1 \forall j = 1, \dots, 5$), the second for a model where diffusion is the same for all technology fields ($\beta_j^D = 1 \forall j = 1, \dots, 5$), and the third for a model where both diffusion and obsolescence are allowed to vary over technology field. Although allowing diffusion to vary over technology fields produces only a marginal improvement in the R-squared, it is clear from the coefficient estimates that there are substantial differences among fields.

[Table D.2 about here]

The J-T model allows for both cited and citing year effects. Although the cited year effects do not vary much over the sample, the citing year effects show a substantial growth rate, consistent with the "citation inflation" mentioned earlier, plus the increase in the rate of patenting, which generates more citations made each year. Figure D.3 shows the estimated citation lag distribution purged of the fact that the probability of citation is increasing over time for all patents, and that patents in a given cohort may be cited more or less, based on the complete model in the last column of Table D.2. The vertical axis is the relative citation probability; the area under each of the curves is the estimated relative overall citation intensity for a given field; the left-right position of the curve indicates the "speed" of citation in a given field. The most cited patents are those in computers and communications, followed by drug and medical patents. The average citation lag is notably longer for drugs and medical, and the diffusion rate substantially slower for patents in the "other" category. The modal lag ranges from 2.6 years

for mechanical to 5.3 years for other.

[Figure D.3 about here]

We will use the results for the first model (shown on Figure D.3 as column 1 estimates) to estimate the unobserved citation probabilities, since we do not want to purge the data of field effects. We are missing citations made before 1976 to patents issued between 1963 and 1975, and citations made after 1996 for patents issued between 1963 and 1995 (that is, we are missing one year (1996) for patents issued in 1965, two years (1996, 1997) for patents issued in 1966, and 30 years (1997-2026) for patents issued in 1996. Our model allows us to predict the number of these missing citations. Define the fitted lag distribution in the following way:

$$D_s = \exp(-\beta_1 s)(1 - \exp(-\beta_2 s)) \quad s = 0, \dots, 30$$

Then the predicted number of citations at lag s for a patent issued at time t , when we do not observe citations beyond lag $S=T-t$ is the following:

$$C_{t,s} = D_s \frac{\sum_{j=0}^S C_{t,j}}{\sum_{j=0}^S D_j}$$

For example, the predicted number of citations in the year 2000 to a patent issued in 1993 is the following (when we observe citations only through 1996):

$$C_{1993,7} = D_7 \frac{\sum_{j=0}^3 C_{1993,j}}{\sum_{j=0}^3 D_j}$$

There is one problem with this procedure in the later years: because count data has a skew distribution and is bounded below by zero, when zero citations are observed (due to a short period over which to observe them), the expected number of citations that will eventually be observed conditional on this fact is not zero, but positive. We correct for this fact by replacing $\sum_{j=0}^S C_{t,s} / \sum_{j=0}^S D_j$ in the equation above by the following quantity, derived empirically from our data in 1975 (where we see 20 years of citations):

$$E\left[\sum_{j=0}^{20} C_{t,j} \mid \sum_{j=0}^S C_{t,j} = 0\right]$$

That is, we predict the total number of citations that will be observed in the 20 years given that we observed zero citations in the first S years, and then distribute these citations according our citation lag distribution.

To derive the estimated total (20-year) citations for any patent, we simply sum the observed citations from the observed years and the predicted citations based on the above methodology for the unobserved years. Totals based on these "corrected" citation counts are used as indicated in Figures 2 and 3, and were used to construct the citation stocks used in the regressions.

Table D.1
Weighting Factors for Patent Applications

Year	Inverse Weight
1964	2.119
1965	1.229
1966	1.063
1967	1.000
1968-89	1.000
1990	1.000
1991	1.003
1992	1.009
1993	1.034
1994	1.166
1995	2.230
1996	37.461

TABLE D.2
Estimation of Citation Probability based on Jaffe-Trajtenberg Model

Cited Year: 1963-1993 Citing Year: 1976-1994 6 Tech Fields

2616 Observations = $6 \cdot (14 \cdot 19 + (18 \cdot 19) / 2 - 1)$

Dep. Var.=No. of Citations by Citing Year Patents to Patents in Cited Year & Field

	Coefficient	S.E.	Coefficient	S.E.	Coefficient	S.E.
Tech Field Effects (base=other)						
Chemicals excl Drugs	0.8788	0.0130	0.8943	0.0189	0.4443	0.0195
Computers & Communication Equip.	1.6150	0.0184	2.0179	0.0327	1.3016	0.0599
Drugs & Medical	1.3714	0.0164	1.1165	0.0189	0.5396	0.0195
Electrical & Electronics	1.1146	0.0145	1.2660	0.0235	0.6039	0.0238
Mechanical	0.9011	0.0131	0.9068	0.0188	0.4063	0.0160
Citing Year Effects (base=1976)						
1977	0.8920	0.0279	0.8851	0.0220	1.0077	0.0251
1978	0.9961	0.0304	0.9902	0.0240	1.0936	0.0264
1979	1.0821	0.0330	1.0764	0.0261	1.1540	0.0277
1980	1.1507	0.0357	1.1468	0.0281	1.1959	0.0289
1981	1.1751	0.0375	1.1727	0.0296	1.1927	0.0295
1982	1.2540	0.0408	1.2508	0.0322	1.2390	0.0312
1983	1.2530	0.0426	1.2526	0.0337	1.2051	0.0317
1984	1.3981	0.0482	1.3956	0.0381	1.3048	0.0349
1985	1.5637	0.0548	1.5618	0.0434	1.4213	0.0388
1986	1.7444	0.0622	1.7447	0.0494	1.5508	0.0433
1987	1.9992	0.0727	2.0011	0.0578	1.7328	0.0497
1988	2.2563	0.0844	2.2648	0.0674	1.9086	0.0566
1989	2.4626	0.0954	2.4684	0.0761	2.0245	0.0625
1990	2.5883	0.1038	2.5991	0.0831	2.0780	0.0667
1991	2.6691	0.1106	2.6823	0.0886	2.0976	0.0698
1992	2.8730	0.1227	2.8921	0.0986	2.2111	0.0762
1993	2.7362	0.1219	2.7587	0.0981	2.0596	0.0741
1994	1.2983	0.0667	1.3108	0.0538	0.9545	0.0389
Cited Year Effects (base=1963-65)						
1966-70	0.6736	0.0139	0.6723	0.0110	0.8749	0.0174
1971-75	0.6352	0.0125	0.6283	0.0104	0.9295	0.0220
1976-80	0.5687	0.0144	0.5594	0.0120	0.9369	0.0284
1981-85	0.4959	0.0159	0.4860	0.0133	0.9300	0.0352
1986-90	0.4275	0.0173	0.4152	0.0142	0.8935	0.0403
1991-93	0.3332	0.0188	0.3174	0.0148	0.7435	0.0410
Beta1: Obsolescence by Technology Field						
Chemicals excl Drugs			1.0249	0.0204	0.7022	0.0248
Computers & Communication Equip.			1.2731	0.0185	1.1015	0.0268
Drugs & Medical			0.7796	0.0140	0.4559	0.0167
Electrical & Electronics			1.1558	0.0197	0.7960	0.0223
Mechanical			1.0125	0.0199	0.6251	0.0214
Beta1 (Base=Other)	0.1229	0.0018	0.1188	0.0017	0.1250	0.0017
Beta2: Diffusion by Technology Field						
Chemicals excl Drugs					3.5478	0.2699
Computers & Communication Equip.					2.2286	0.1406
Drugs & Medical					3.9224	0.2342
Electrical & Electronics					3.8451	0.2489
Mechanical					4.7339	0.3507
Beta2 (Base=Other)	0.4801	0.0132	0.4608	0.0104	0.1599	0.0063
R-squared	0.904		0.939		0.949	
Standard error of regression	0.0591		0.0472		0.0430	

Figure D.1 Application-Grant Lag Distribution




Figure D.2
Empirical Citation Lag Distribution by Year
HJT Sample of Manufacturing Firms





Figure D.3
Citation Lag Distribution (1976-1994)
JT Methodology with Field-Specific Diffusion and
Obsolescence



—◆— Chemicals excl. drugs - - -◆- Computers and comm. —▲— Drugs and medical - - -■- Electric and electronic
 —□— Mechanical —○— Other - - + - Column 1 est (all fields)

Figure D.4 Citation Lag Distribution by Cited Year

