Measuring the Impact of Academic Science on Industrial Innovation:
The Case of California's Research Universities*

Lee Branstetter
Columbia Business School
815 Uris Hall
3022 Broadway
New York, NY 10027
and NBER

This Version: August 2003

* **Acknowledgements**: I wish to thank Jim Adams, Pierre Azoulay, Iain Cockburn, Robert Feenstra, Rebecca Henderson, Marvin Lieberman, David Mowery, Ariel Pakes, and seminar participants at the NBER, Columbia Business School, and the UCLA Anderson Graduate School of Management for useful comments and suggestions. I am especially grateful to Adam Jaffe and Joshua Lerner for detailed feedback on an earlier draft. I also wish to thank a number of academic scientists and industrial R&D managers for providing me with their insights into the process by which knowledge flows from academia to industry. I am indebted to Masami Imai, Hiau-Looi Kee, Changxiu Li, Kaoru Nabeshima, and Yoshiaki Ogura for excellent research assistance. I would like to thank Tony Breitzman and Francis Narin of CHI-Research, Adam Jaffe, and Marie and Jerry Thursby for their help in obtaining the data used in this study. This project was funded by grants from the University of California Industry-University Cooperative Research Program, the NBER Project on Industrial Technology and Productivity, the Japan Foundation Center for Global Partnership, and the National Science Foundation.

I. Introduction

The large investment by the U.S. federal government in academic research in the postwar era has been predicated on the belief that this expenditure would stimulate economic growth in the long run by laying down a scientific foundation upon which inventors could develop useful new technology.[1] Recent studies suggest that the nature of the relationship between academic science and industrial innovation is changing. At least in some fields of science and technology, the positive impact of publicly funded science on private innovation appears to have been increasing in strength in recent years. If this is true, it could have important implications for U.S. science policy and for the prospects for continued technology-driven economic growth.

This paper seeks to contribute to our understanding of this changing relationship in three ways. First, I argue that examining patent citations to academic papers offers a useful window through which the process of knowledge spillovers from science to invention can be viewed. Second, I bring nonlinear regression analysis techniques to bear on a large random sample of 30,000 U.S. utility patent grants. Using these data, I show what kinds of patents cite academic science and how these patterns have changed over time. I also examine the linkage between citations to academic science and the quality of invention. Third, I undertake an econometric "case study" of the changing impact of academic science on innovation by combining comprehensive data on the publications generated by a set of California research universities, the universe of patent citations to these publications, and the universe of potentially citing U.S. patents over the 1987-99 period. Using these data, I document changes in the propensity of patents to cite

[1] For an early and influential statement of this belief, see Bush, 1945.

science while *controlling for* changes in the level and distribution of scientific articles and changes in the level and distribution of potentially citing patents.

To summarize some of my most significant findings, I find that the knowledge spillovers from academic science to invention are highly concentrated in a small subset of technological fields and geographic regions. I show evidence of a positive link in the cross-section between citations to science and invention quality. Finally, I present some preliminary evidence on the relative importance of three changes in explaining the rise in the incidence of patent citations to academic papers: 1) a change in the quantity and distribution across fields of potentially cited scientific publications, 2) a change in the distribution of potentially citing inventors, and 3) a change in the propensity of inventors to cite academic science. The implications of these results and possible extensions are discussed at length in the conclusion.

II. The Link Between Academic Science and Industrial Innovation

Lessons from the Prior Literature

This paper draws on and contributes to a burgeoning literature on the impact of academic science on industrial innovation. While much of my focus will be on the recent research by economists on this topic, I should note that important work by both noneconomists and economists on the relationship between science and technology stretches back several decades.[2] The consensus of this early research was that the relationship between science and technology was generally neither close nor direct. Based on science and technology literature citations studies, Derek De Solla Price (1965)

[2] I thank Marvin Lieberman for pointing to me to some of these early studies. I should note that Schmookler (1966) made important contributions to the early economics literature on this and related subjects. Nathan Rosenberg generated a number of pioneering studies of the economic history of interaction between American universities and industry. See, for example, Rosenberg (1982).

concluded that there was only a weak interaction between the two. This view was generally supported by the Defense Department's ambitious "Project Hindsight" study of the impact of basic scientific research on weapons development, which concluded that the primary impact of science on weapons development came not from recent science at the research frontier, but instead from "packed-down, thoroughly understood, and carefully taught old science," such as that typically presented in textbooks and university courses.[3]

Early researchers noted that there were cases when relatively "new" scientific discoveries quickly found early application in new inventions, leading to a closer coupling between the advance of the scientific frontier, as traced out in the recent scientific literature, and the rapid incorporation of these advances into new *products* or *commercial processes*.[4] However, these deviations from the norm tended to be temporary phenomenon. For instance, Lieberman's (1978) study on the introduction of MOS transistor technology suggests that the linkage between science and technology weakened as the technology matured and the crucial advances in science became embodied in succeeding generations of products.[5]

The recent economics literature has argued that the linkage between "new" science and technology is potentially stronger and more direct than this earlier literature suggested. Case studies and surveys have been used to assess both the magnitude of this impact and the channels through which it flows.[6] These studies suggest that firms

[3] The quoted phrase comes from Sherwin and Isenson (1967).
[4] See, for example, Marquis and Allen (1966).
[5] Darby and Zucker (2003) argue that a similar pattern can be seen in the more recent impacts of biotechnology and nanotechnology on industrial invention.
[6] Important recent studies relying primarily on case study techniques and surveys include Mansfield (1995), Cohen et. al. (1994), Faulkner and Senker (1995) and Gambardella (1995). Rosenberg and Nelson (1994) have also contributed to this literature with a more historical approach.

perceive academic research to be an important input into their own research process, though this importance differs widely across firms and industries.[7] A second stream of recent research has undertaken quantitative studies of knowledge spillovers from academic research. Jaffe (1989) and Adams (1990) were early contributors to this literature. More recently, Jaffe et. al. (1993, 1996, 1998) have used data on university patents and citations to these patents to quantify knowledge spillovers from academic science.[8] This paper will also seek to quantify knowledge spillovers from academic research, and it borrows from the empirical methods introduced in these papers.

A related stream of research has undertaken quantitative analysis of university-industry research collaboration. Contributors include Zucker et. al. (1998) and Cockburn and Henderson (1998, 2000). A number of papers in this literature have studied "start-up" activity related to academic science or academic scientists, such as Zucker et. al. (1998) and Audretsch and Stephan (1996). Finally, several recent studies have examined university licensing of university generated inventions, such as Barnes et al. (1998), Mowery et. al. (1998), Thursby and Thursby (2002), Shane (2000), and Shane (2001).

Collectively, the recent literature has highlighted several key changes that have potentially affected the relationship between academic science and private sector innovation. First, the *quantity* of academic science and its *distribution across fields* has changed over the last three decades, with a substantial shift in federal funding away from the physics-based disciplines that were connected to weapons development and the space

---

[7] While the channels by which firms absorb the results of academic research vary across industries, the Cohen et. al. (1994) study suggests that the formal scientific literature is, on average, an important channel.
[8] Barnes, Mowery, and Ziedonis (1998) and Mowery, Nelson, Sampat, and Ziedonis (1998) have undertaken a similar study for a smaller number of universities.

program and toward the life sciences.[9] Second, the *nature of inventive activity* seems to have changed. Firms in some industries, especially those related to drugs and medical technology, have changed their approach to research in a way that brings them "closer" to academic science. While this is well documented in the context of the pharmaceutical industry,[10] it is less clear to what extent there have been similar changes in other technology-intensive industries. Third, the *institutional environment* in which scientists and inventors interact has changed. Partly, this is the result of public policies designed to encourage the commercialization of university-developed science, such as the Bayh-Dole Act.[11] However, the rise of venture-capital investments in small high-technology firms has arguably made it easier for entrepreneurial academics to commercialize their discoveries. In terms of the data-generating process, this institutional change has brought in a "new" group of patenting entities with a higher propensity to cite academic research than others.[12]

If we are to understand the policy implications of the observed increase in the incidence of patent citations to academic science, it would be obviously helpful to understand the relative importance of these and other factors in explaining the overall increase. A finding that there are more citations simply because there are more publications in fields that have always been highly cited would have quite different implications from a finding that showed a large increase in the propensity to cite science across all classes of inventors and all fields of science. A key aim of the current paper is

---

[9] See Cockburn and Henderson (2000).
[10] See Zucker et. al. (1998) or Cockburn and Henderson (2000), among many other studies.
[11] See Henderson, Jaffe, and Trajtenberg (1998) and Mowery et. al. (1998).
[12] See Kortum and Lerner (1997). In the last few years, of course, the sharp downturn in venture capital funding across a range of technologies has changed the institutional environment yet again.

to try to present evidence on which factors are most important in explaining the overall increase.

Using Patent Citations to Academic Science as Measures of Knowledge Spillovers

This paper will use patent citations to academic papers to measure "knowledge spillovers" between academic science and industrial R&D. In doing so, I am building on the work of Francis Narin and his collaborators, who have pioneered the use of these data in large-sample "bibliometric" analysis.[13] As indicators of knowledge spillovers from academia to the private sector, these data have a number of advantages. The academic promotion system creates strong incentives for academic scientists, regardless of discipline, to publish all research results of scientific merit. As a consequence, the top-ranked research universities generate thousands of academic papers each year. Similarly, inventors have an incentive to patent their useful inventions, and a legal obligation under U.S. patent law to make appropriate citations to the prior art – including academic science. As Figure 1 illustrates, the number of citations to these papers in patents has been growing rapidly for much of the 1990s.

In response to the Bayh-Dole Act and other public policy measures, universities have increased the extent to which they patent the research of university-affiliated scientists. They have also increased the extent to which they license these patented technologies to private firms. Nevertheless, it is clear to observers that only a *tiny fraction* of the typical research university's research output is ever patented, and only a fraction of this set of patents is ever licensed. To illustrate this, Figure 2 shows the trends over the 1988-1997 period in several alternative indices of university research output and knowledge spillovers for one university system: the University of California's 9

[13] See Narin et. al. (1997) and Hicks et. al. (2001) for recent examples of this work.

campuses and affiliated laboratories. The figure graphs university patents by issue year (patents), invention disclosures by year of disclosure filing (invention disclosures), new licenses of university technology by date of contract (licenses), the number of citations to previous university patents by issue year of the citing patent (citations to UC patents), and the number of citations to UC-generated academic papers by issue year of the citing patent (citations to UC papers). Throughout the sample period, there are far more citations to UC papers than any other kind of indicator.[14]

This figure suggests that patent citations to academic papers may provide a much broader window through which to observe knowledge spillovers from academic science to inventive activity than any available alternative. But while citations may be easy to count, they are more difficult to interpret. This paper goes beyond simple tabulations of citations to explore their determinants and effects.[15]

Having made the case for the use of patent citations to science as a measure of these spillovers, it is also appropriate to acknowledge the limitations of this measure. Universities also contribute to the advance of industrial technology through the education and training of engineers employed in private firms. Patent citations to academic science are unlikely to be a particularly effective measure of this "general human capital" channel. On the other hand, American universities have played this role of human capital generation for decades, and it is unlikely that the operation of this channel has changed so much as to be the principal driver of a closer connection, if any, between academic science and industrial invention. It is also true that university faculty members engage in

---

[14] Data for Stanford reveal a similar picture.

[15] Other recent studies using data on patent citations to scientific papers include work by Fleming and Sorenson (2000, 2001) and Lim (2001). Agrawal and Cockburn (2002) examine the impact of academic science in industrial innovation in three technological fields, although they do not use data on patent citations to academic science.

formal and informal consulting with industry – and have done so for decades. To the extent that the nature of this consulting involves advising industrial scientists and engineers on the import of recent scientific discoveries, my measure of patent citations to science is likely to be positively correlated – perhaps highly so – with this consulting activity. To the extent that this consulting involves advising industrial scientists and engineers on well-established principles, findings, algorithms, or techniques (i.e., "old science"), patent citations to science are unlikely to be highly correlated with it.[16] Thus, patent citations to science can be viewed as a reasonable measure of the incorporation of recent science into inventive activity. To the extent that change along this particular dimension of university-industry interaction is of interest, patent citations to science are likely to be a useful indicator.

III. Evidence from the Random Sample

Citations Patterns in the Random Sample

I start with a random sample of 30,000 utility patents granted over the 1981-1996 period, approximately 4,500 of which make at least one citation to "science."[17] The sample is large enough that changes in the distribution of patents and patent citations to science in the sample should be reflective of changes in the underlying sample. I focus here on obtaining econometric estimates of the *conditional* impact of various attributes of citing patents on the propensity to cite, holding others constant. The nature of the data suggests the use of a negative binomial specification, since most patents make no references to science but small numbers of patents make numerous references to

---

[16] Agrawal and Henderson report results based on a survey of university faculty at MIT indicating that consulting arrangements are an important channel of knowledge flow to industry.
[17] These data were purchased from CHI Research, Inc. At the moment, budgetary limitations preclude expansion of the random sample past 1996.

science.[18] Independent variables of interest include dummy variables for the (application) year of the patent cohort, the technology category of the patent, the category of organization to which the patent is assigned, and a crude measure of geographic proximity between the region in which the (first) inventor of the patent is located and the region(s) in which academic science is produced.

There are several ways in which I can define "academic science" and thus measure patent citations to it. The most comprehensive such measure is to consider all nonpatent citations which appear to be to "scientific documents" (including conference proceedings and technical manuals) as citations to "science." A narrower measure would be to count all references to articles in SCI-indexed academic journals. This database tracks articles appearing in many of the most influential peer-reviewed journals across all major scientific disciplines, and it may correspond more closely to the output of "academic science." A still narrower measure would count only references to *university*-authored papers in tracked journals. This distinction is useful because, in some scientific disciplines, large corporate R&D labs and public science agencies generate a substantial contribution to "academic science," publishing in the same journals as their university-affiliated peers.[19]

Table 1 presents empirical results based on a negative binomial specification. The three columns correspond to the different definitions of "academic science" described above. The first five rows present the coefficients on dummy variables equal

---

[18] An alternative logit specification, in which the dependent variable was a dummy indicating whether or not the patent in question made any citations to academic science, yielded results qualitatively similar to those shown in Table 1. Specification tests suggested that the more flexible zero-inflated negative binomial model did not fit the data significantly better than the standard negative binomial specification.

[19] An important data limitation to the data used here is that, in most specifications focusing on geographic proximity, I only examine citations made to authors based in the United States. Results presented in the next section utilize data on cited authors worldwide.

to one if the patent assignee falls into one of the five listed categories: university, non-profit R&D organization (many of these are research hospitals), U.S. government agency (i.e., NASA), foreign (foreign firms, individuals, and government agencies are all placed in this category), and "other" (the largest fraction of which are U.S. individuals). The reference category here is private firms. It is immediately clear that universities, nonprofit R&D organizations, and U.S. government agencies are all more likely to cite academic research than are firms. This differential gets generally more pronounced as one restricts the definition of what constitutes academic science. That being said, the vast majority of citing patents (76%) are generated by firms.[20]

The next set of dummy variables corresponds to the technology class of the citing patent. Using a taxonomy developed by Adam Jaffe and Manuel Trajtenberg, I have aggregated the primary patent classes of the U.S. Patent and Trademark Office patent classification system into six groups – chemicals, communications/computers, drugs/medical, electronics/electrical machinery (not directly computer related), mechanical devices, and a catch-all "other" category which constitutes the reference group in these regressions.[21] Patents in the drugs/medical category stand out as being disproportionately likely to cite. This differential effect gets stronger as I narrow the definition of academic science across columns. The chemicals category ranks second in terms of likelihood of citing.

"Science center" is a dummy variable equal to 1 if the patent inventor is located in one of the top 100 U.S. counties in terms of generation of scientific publications. This

[20] It is also possible to control for "self-citation" by excluding from the data counts of made by an entities' own patents to papers generated by authors affiliated with that entity. Regressions run on data purged of such self-citation generate results qualitatively similar to those presented here.
[21] I thank Adam Jaffe for providing this taxonomy in electronic form. Note that there are several hundred primary patent classes.

variable is positive and statistically significant at conventional levels. This suggests that patents are more likely to cite when an inventor is located in a region with a high level of scientific research. However, this does not necessarily constitute evidence of the geographic localization of knowledge spillovers.

Following Jaffe et. al. (1993), I use a different approach to this question that explicitly controls for the skewed distribution of research activity across U.S. counties. I match each of the citing patents in my random sample with a nonciting "control" patent issued on the same date in the same patent class as the citing patent. Let p_c be the probability that a citing patent is generated in the same county as that in which the cited "science source" is located. Let p_0 be the corresponding probability for a randomly drawn control patent. I test for "geographic localization of knowledge spillovers" using the following test statistic:

$$t = \frac{\hat{p}_c - \hat{p}_0}{\sqrt{[\hat{p}_c(1-\hat{p}_c) + \hat{p}_0(1-\hat{p}_o)]/n}} \qquad (1)$$

where the two terms in the numerator are the sample proportion estimates of p_c and p_0. The null hypothesis that $p_c = p_0$ is easily rejected at conventional levels.[22]

All regression specifications are run with patent application year cohort effects. While the coefficients are not shown in Table 1, the results from Table 1, column 1 are graphed out in Figure 3, along with the 95% confidence bounds. What is evident from this graph is a pronounced rise in the tendency of patents to cite over time, controlling for the increase in university patenting and changes in the distribution of patents over classes with different tendencies to cite science. It is interesting to compare the shape of this

[22] This test was conducted using both the "state" and the "county" as the regional unit of analysis. The t-statistic of the difference in ratios was 11.27 for state-level comparisons, 11.03 for county-level comparisons.

graph to Figures 1 and 2. When one looks further back into the past, it seems that the most pronounced increase in the *conditional* likelihood of citation came in the early 1980s rather than in the 1990s, as would be suggested by the unconditional distribution of citations over time. While there has been an increase in the conditional probability tendency to cite science across later cohorts of patents in the 1990s, it may be that much of the "spike" in citations so visible in Figures 1 and 2 has been driven by the widely documented increase in patenting in the health care related technologies.[23]

Does Citation of Academic Science Make Inventions Better?

The discussion of trends in the citations data above is of limited interest unless the knowledge spillovers indicated by these citations are actually enhancing the research productivity of the firms and other organizations that receive them. Are innovators learning from academic science in such a way that they are able to produce more inventions than they otherwise could or better inventions than they otherwise could? Alternatively, does the information generated by academic science allow them to invent in areas in which they could not work without the pre-existing foundation of academic science on which to build?

It is very difficult to establish the technological *dependence* of a particular invention on a cited scientific article without engaging in an in-depth study of the invention and extensive interviews with its inventors.[24] However, I can seek to measure whether or not patented inventions that cite UC or Stanford academic science are

---

[23] See Hicks et. al. (2001) for evidence on the increase in "biomedical" patenting.

[24] In a series of interviews with cited academics and citing firms in which I presented both parties with a list of patent citations to the work of a particular academic, it was often quite easy, based on the titles/abstracts of the patents, to identify a technological linkage between the cited paper and the citing patent. Obviously, it is difficult to draw sweeping generalizations from a small number of interviews. A brief summary of this fieldwork component of the project is available from the author upon request.

systematically "better" than patents that do not. The micro literature on patents has suggested several measures of patent "quality" – quantitative features of the patent document – that have been demonstrated to be positively correlated with the *ex-post* commercial and technological importance of the patent. Three such measures include counts of *ex-post* (or "forward") citations, counts of claims contained in the patent document, and a measure of "generality" proposed by Henderson, Jaffe, and Trajtenberg (1998). This latter measure is a quantitative index of the diversity of technological fields across which *ex-post* citations occur. An invention whose citations come from multiple technological fields can be thought of as having a more "general" impact than an invention whose citations come from a single technological field. The formal definition of the index is

$$Generality_i = 1 - \sum_{k=1}^{N_i} \left(\frac{Nciting_{ik}}{Nciting_i} \right)^2 \qquad (2)$$

where the numerator in the expression measures the number of citations to patent *i* coming from patent class *k*, while the denominator measures the total number of citations to patent *i* across all classes.

Table 3 presents the results of regressions in which these three measures of quality are the dependent variable, a dummy variable indicating patents which cite academic research is the chief independent variable of interest, and I use as controls measures of the patent cohort (application year) and technological field. The results in

Table 3 suggest that patents citing academic research are significantly better according to all three indices of patent quality.[25]

However, in this context, it is very difficult to interpret this result in a *causal* way. Are patents that cite academic research "better" because they cite, or do they tend to cite academic research more frequently because they are "better"? At this level of aggregation, it is difficult to determine which interpretation is correct.[26]

IV. Evidence from California Research Universities

Evidence from a Citations Function Approach

As I noted in the introduction, much can potentially be learned by examining changes in citations while controlling for changes in the population of potentially cited papers and in the population of potentially citing patents. While it would be impractical to do this for the universe of academic publications and U.S. patents, it has been possible for me to link data on the universe of SCI-indexed academic publications generated by the campuses and affiliated research units of the University of California, Stanford University, Caltech, and the University of Southern California, the universe of patent citations made to these publications over the 1983-1999 (grant year) period, and the universe of potentially citing U.S. utility patents granted over that same period.

Restricting the sources of science to a relatively small number of universities based in a single state brings with it obvious disadvantages. Nevertheless, a study of this kind in the context of California is of particular interest because of the substantial growth

---

[25] After completing the first draft of this paper, it was brought to my attention that Sorenson and Fleming (2001) have also documented a positive relationship between patent quality and academic citations, using a smaller sample drawn from two years.

[26] Fleming and Sorenson (2001) question the interpretation that the higher level of citations received by patents citing academic science is indicative of a higher level of patent quality. They find that citations are higher for patents citing any kind of publication, including classes of publication with limited scientific content.

of the state's relative importance in national innovative inputs and outcomes. Over the past twenty years, the geography of innovation within the United States has changed substantially. As Hicks et. al. (2001) document, California has dramatically increased its share of domestically generated U.S. patents, and it has been the leading center of venture capital-backed entrepreneurial activity. One of the reasons given for California's innovative ascendancy is the high quality of the state's academic science base, to which locally based firms are believed to have preferential access. The approach taken below will actually allow for a test of the hypothesis that location within the state provides preferential access to spillovers from California academic science.

An examination of some features of the raw data illustrate why the approach taken in this section may be useful. As Figure 4 illustrates, the majority of publications generated by the UC system in 1999 was concentrated the life sciences. This is reflective of national trends, and the preponderance of publication in these fields has been a feature of the data for much of my sample period. If I were to find that citations to academic science are dominated by the life sciences and medicine, this could simply reflect the greater volume of publication in those fields. To put it simply, there are many more relevant articles to cite.

Figure 5 presents a brief look at changes in patenting across aggregated fields of technology over time, where time is measured by the year in which the patent is granted. It is clear that patenting has been growing overall – but that growth has been particularly rapid in the categories of computers and communications and drugs, with level of patenting in these fields rising roughly 6-fold and almost four-fold, respectively, over the course of my sample period. Any investigation of the impact of academic science on invention needs to control for changes, such as these, in the distribution of invention across fields. The finding of increase in patent citations to life science articles could simply reflect the explosion of patenting in drug-related technology categories.

The empirical framework I use for this analysis borrows from the work of Caballero and Jaffe (1993) and Jaffe and Trajtenberg (1996, 2002). In this framework, I model the probability that a particular patent, *p*, applied for in year *T*, will cite a particular article, *a*, published in year *t*. This probability is determined by the combination of an exponential process by which knowledge diffuses and a second exponential process by which knowledge becomes obsolete.

This probability is referred to in the work of Jaffe and Trajtenberg (1996) as the *citation frequency*. It is a function of the attributes of the citing patent (P), the attributes of the cited article (a), and the time lag between them (T-t). It can be rendered in notation as

$$p(a,P) = \alpha(a,P)\exp[-\beta_1(T-t)][1-\exp(-\beta_2(T-t)] \qquad (3)$$

Attributes of the citing patent that I incorporate into my analysis include the application year, the technical field (based on the primary technology class assigned by the patent examiner), the type of entity owning the patent (based on the identity of the assignee), and the geographic location of the patent, based on the address of the inventor. Attributes of the cited article that I consider include the publication year, the scientific field of the article, and the institution with which the authors were affiliated at the time of publication.

Given these data, one could sort all potentially citing patents and all potentially cited articles into cells corresponding to the attributes of articles and patents. The expected value of the number of citations from a particular group of patents to a particular group of articles could be represented as

$$E[c_{tcelTSL}] = (n_{TSL})(n_{tcel})\alpha_{tcelTSL} \exp[-(\beta_1)(T-t)][1-\exp(-\beta_2(T-t))] \quad (4)$$

which can easily be rewritten as

$$\frac{E[c_{tcelTSL}]}{(n_{TSL})*(n_{tcel})} = \alpha_{tcelTSL} \exp[-(\beta_1)(T-t)][1-\exp(-\beta_2(T-t))] \quad (5)$$

This is what Jaffe and Trajtenberg (1996) refer to as a *citations function*. If one adds an error term, then this equation can be estimated using nonlinear least squares. The estimating equation is thus

$$p_{tcelTSL} = \alpha_t \alpha_c \alpha_e \alpha_l \alpha_T \alpha_S \alpha_L \exp[-(\beta_1)(T-t)][1-\exp(-\beta_2(T-t))] + \varepsilon_{tcelTSL} \quad (6)$$

where the dependent variable measures the likelihood that a particular patent in the appropriate categories of application year (t), technology class (c), institutional type (e), and location of the citing patent's inventor (l) will cite an article in the appropriate categories of scientific field (based on the scientific content of the article) (S), a particular campus (L), and publication year (T). The α's are multiplicative effects estimated relative to a benchmark or "base" group of patents and articles. In this model, unlike the linear case, the null hypothesis of no effect corresponds to parameter values of unity rather than zero.

I estimate various versions of (6) using the nonlinear least squares estimation routine of the STATA software package. When doing so, I weight the observations by the square root of the product of potentially cited articles and potentially citing patents corresponding to the cell, that is

$$w = \sqrt{(n_{tcel})*(n_{TSL})} \quad (7)$$

This weighting scheme should take care of possible heteroskedasticity, since the observations correspond to "grouped data," that is, each observation is an average (in the corresponding cell), computed by dividing the number of citations by $(n_{tcel})*(n_{TSL})$.

This approach allows us to examine changes in citation patterns over time controlling for differences in the intensity of citation of science across different industrial technology classes, changes in the distribution of patents across technological fields, and changes in the distribution of scientific articles across scientific fields. Regression results from a version of (6) run on the full sample are given in Table 3. Using the parameter values from this regression, it is also possible to graph out the double exponential function implied by our parameter estimates, giving us a sense of how the "citedness" of a particular group of articles by a particular group of patents changes over time. This is graphed out for our "base case" in Figure 6. The base case in this regression corresponds to patents assigned to firms, where the first inventor resides in the U.S. outside the state of California. The base patent application period is 1981-1987, and the base publication period is 1981-1985. The base science category is biology, the base patent category is chemistry, and the base institution is Stanford University.

The shape of the curve graphically demonstrates the first key result of this section – namely that citations to academic science are somewhat localized in time. Citations to science appear almost immediately after article publication, and the citation function peaks at a lag of about four years after article publication. These lags are measured with respect to the application date of the patent, implying rapid spillovers of knowledge from science into industrial invention. While the estimated lag structure demonstrates that papers continue to receive some citations even at relatively long lags, the citation frequency declines steadily after the peak lag. For the "base category," the estimated citation frequency drops below the level for a single year at a lag length of about twelve years.

The similarities between my methodology and that of Jaffe and Trajtenberg (1996, 2002) invite an informal comparison of my results to theirs. My findings are not directly comparable to theirs, because I date patents according to date of application rather than date of grant. Nevertheless, the general patterns of growth and decay of citations to academic papers over time seem to be broadly similar to those of citations to other patents.

The second key result of this section is the finding of striking differences in the incidence of citation across fields of academic science over time. Note that the citation function specification controls for the number of "citable papers" within these science categories over time, as well as the number of potentially citing patents across fields of technology, so the coefficients on science categories are akin to a "per-paper" measure of technological fertility. The coefficients in Table 3 suggest that a paper in the "biomedical research" field is *nearly 38 times* more likely to be cited in a patent than a paper in the base category of biology. Papers in "chemistry" and "clinical medicine" are nearly five times as likely to be cited as a biology paper, while papers in the other science categories are substantially less likely to be cited. This differential is illustrated in Figure 7, where the double exponential function for "biomedical research" is graphed out relative to the base category for general "biological sciences."[27]

[27] In results available upon request, I specified an "academic production function" for the university systems studied in this section of the paper, in which the output measure was the count of publications generated in a scientific field by a particular campus in a particular year. This was regressed on measures of "inputs" to the research process, including various measures of R&D funding, post-doctoral students,

Again, an informal comparison with the results of Jaffe and Trajtenberg is useful. These authors allow the technological fertility of different patent classes to vary, but constrain the propensity to cite to be the same across patent classes, so that my measures of technological fertility are not directly comparable to theirs. Nevertheless, I find a much more skewed pattern of citations across science classes than they find across technology classes. Jaffe and Trajtenberg find that the "drug and medicines" category is about 1.4 times as "fertile" as the base category of "other patents." My estimated gap between the most and least fertile categories of science is much wider. To put this another way, the distribution of citations to science is much more narrowly concentrated within particular categories of science than the distribution of citations to patented technologies.

Continuing in this theme, I can allow different categories of patented technologies to display different propensities to cite science. Relative to the base category (chemicals), drug/medicine patents are 2.6 times more likely to cite science, whereas all other categories are substantially less likely to cite science. The typical patent in the least likely-to-cite category, mechanical patents, is only about 1% as likely to cite science as the typical chemical patent. Again, the estimated gap between technology categories in citation propensity is quite substantial. Note that these estimated propensities control for the number of patents in these categories over time, so that these coefficients are properly interpreted as an estimate of the differential "per-patent" propensity to cite science.

Taken together with the result on differences in fertility across science classes, this suggests that the aggregate trends in patent citations to science are driven largely by "biotech" patents citing "bioscience" papers. While there is certainly growing citation

graduate students, etc. The results suggest that the higher "productivity" of the biological sciences is not driven purely by the increase in R&D funding in that field.

activity outside this nexus, citations to date have been highly concentrated within it. The existing literature has pointed out that university patenting and licensing activity have been concentrated in biotechnology. To the extent that the concentration of citations across fields reflects the real underlying distribution of knowledge spillovers across fields, these results would seem to imply that the much discussed shift in federal R&D funding toward the life sciences is actually a step toward improving the impact of R&D spending on industrial invention. However, I cannot, at this point, rule out that the concentration of citation activity in the bioscience-biotech nexus also reflects field-specific differences in citations practices.

I have seen that the citation function results suggest that knowledge spillovers from academic science to industrial invention are concentrated in time and technology space. These results also provide evidence of concentration in geographic space. Citing patents are assigned to three categories based on their recorded addresses: California inventors, U.S. inventors outside California, and non-U.S. inventors. U.S. inventors outside California are the base category, so the coefficients imply that California-based inventors in a given technology class are nearly three times more likely to cite California academic science. The evolution of this differential over time is graphed in Figure 8, which compares the predicted citation frequencies for California-based inventors to those of the base category at different lag lengths. Non-U.S. inventors are only about half as likely to cite California science as is the base category. The U.S. / non U.S. differential propensity to cite implied by the coefficients of Table 3 is broadly comparable to international differences in knowledge flows documented by Jaffe and Trajtenberg.

The intranational localization of knowledge spillovers implied by the California effect seems large. California's share of national patenting has grown substantially over the course of my sample period (reflecting among other things, the regional concentration of venture capital funding), but the citation function approach controls for that. However, the current specification arguably does not control well for regional clustering of industrial R&D within the particular niches of the broad technology categories I have employed. A finer disaggregation of patent classes would likely attenuate the measured degree of localization. Furthermore, as can be seen in Figure 9, it is still the case that large numbers of citations are made by inventors far from California. In fact, one sees a "bicoastal" concentration of citations, reflecting the clustering of U.S. innovative activity in the Northeast and the West Coast.

I have also looked at patenting by different categories of assignees: firms, public science institutions (universities, research institutes, and research hospitals), and a grab-bag category of "other institutions" in the non-profit sector. Assignment of a patent to one of these categories is based on the typography of assignees developed in the NBER patent citation database. Relative to the base category of firms, public science institutions are nearly four times as likely to cite academic science, and "other institutions" are almost twice as likely to cite academic science. This is unsurprising, given the connection that is likely to exist between academic science and academic patenting. Because these institutional categories accounted for a small fraction of total U.S. patenting, even by the end of my sample period, it is still the case that the vast majority of patent citations to California academic science are made by the patents of industrial firms. This reality notwithstanding, it is important to control in a study like this for the impact of university patenting on patent citations to science, and this breakdown by assignee category helps to accomplish that goal.

I also included a set of cited institution effects, to get a sense of differences across institutions in citedness. The actual coefficients may be of limited interest to readers not based in California, but they provide an interesting lesson in the utility of the citation function approach. When one does simple data tabulations, the institution with the largest number of patent citations to its academic research – by a considerable margin – is UC-San Francisco. When one controls for changes in the distribution of papers across fields, UCSF's average level of citedness over time drops below that of the base institution, Stanford. In other words, UCSF's high number of citations completely reflects its specialization in the bioscience disciplines. The institution that seems to be a standout with science field controls in place is Caltech. Although the scale of its academic output is limited (reflecting its small size), and concentrated to some extent outside the fields where the connections between academic science and industrial invention seem to be the strongest (Caltech has no medical school), it has a proportionately greater impact on industrial invention than Stanford. Within the UC institutions, the campus with the highest degree of citedness, controlling for levels of academic publications across fields and their changes over time, is UC-San Diego.

Having incorporated fixed effects associated with the cited field of science, the cited institution, the citing field of technology, and characteristics of the citing inventor/assignee, I can also make some inference about changes in citation patterns over time across fields. Perhaps the most interesting finding here is that the propensity to cite academic science is evidently growing over time. This can be seen by examining the pattern of coefficients on the citing year cohort terms. They generally increase from the "base category" of 1981-87. Note that I have explicitly controlled for the fact that academic publications in the heavily cited branches of science have become more numerous and that there has been an increase in patenting in fields that heavily cite academic science. These results are consistent with the view that there has been *a change in the nature of invention* such that inventors now draw more heavily on academic science.[28]

That being said, we find some evidence of a decline in citation propensity in the most recent period. Controlling for changes in the volume and distribution of patents and publications, the average per-patent propensity to cite science seems to have declined somewhat in the late 1990s from the peak levels of the mid-1990s. While still more than 50% higher than the base period, the finding of a decline in citation propensity raises an immediate question about the permanence of recent growth in the measured linkage between academic science and industrial technology. Are recent trends beginning to reverse themselves? Could we be seeing a replay of the kind of cycle of interaction between science and technology identified by earlier researchers, in which, once a set of significant scientific discoveries is effectively assimilated by industrial inventors, there is a decoupling of the formerly strong relationship between technology and frontier academic science?

[28] Of course, it is also possible that these coefficients simply reflect a change in citation practices rather than an actual change in knowledge spillovers. Jaffe and Trajtenberg also find an increase in propensity to cite prior patents that rises fairly steadily over time, and they attribute some of this to advances in information technology that make prior art easier to find. I find that the fraction of citations made to academic science is going up – in other words, science citations are increasing even more quickly than citations to prior patents – but I cannot, at this stage, definitely rule out the alternative interpretation that much of this change is driven by changes in citation practices. However, anecdotal evidence from conversations with intensively citing firms and highly cited academic scientists strongly suggests that at least part of the measured increase in citation propensity is attributable to an increasingly close connection between science and innovation, especially in the biotech arena.

Interpretation of this measured decline is clouded by two problems in the data. The first is the issue of the so-called "spike patents," which is discussed at some length in the most recent edition of *Science and Engineering Indicators*.[29] In order to bring the U.S. patent system into compliance with the set of international intellectual property rights standards embodied in the Trade-Related Intellectual Property Rights (TRIPs) agreement that was part of the charter of the World Trade Organization (WTO), the U.S. Patent and Trademark Office changed the effective period of monopoly granted to U.S. patent holders from 17 years after the grant date to 20 years from the filing date. This change took effect for patents filed after June 8, 1995. Previously rejected patents re-filed after this deadline would also be subject to new rules. Applications submitted to the U.S. PTO more than doubled in May and June of 1995, and these applications carried an unusually large number of citations to science. This surge in patenting seems to have been driven in part by a rush to file as much as possible under the "old rules." The increase in citations to science seems to have been driven in part by uncertainty out of what was appropriate description of the prior art and a desire to avoid having to refile under the new rules. Patents applied for in this period were issued gradually over the next few years – dramatically increasing the average citations to science in the overall data. Once the last of these applications was processed, the rate of citation fell to something closer to earlier levels.

The most recent edition of *Science and Engineering Indicators* notes that average citations to science per patent continue to increase through the late 1990s when one removes these "spike patents" from the data, though the growth rate in average citations slows. However, that data tabulation does not control for the continuing shift in the

---

[29] This issue is also discussed in Hicks et. al. (2001).

distribution of patenting towards frequently citing categories of technology. I also removed the so-called "spike patents" from my sample, and re-estimated the citation function. My citation function approach controls for the continued changing distribution of patents across technology classes in the 1996-1999 period – and doing still produces the finding of a *decline* in citation propensity in the most recent period.

However, this measured decline could be an artifact of another feature of the data for which it will be harder to control. Note that we measure patents by application year cohort. Because it takes time for patent applications to be processed, and because we only have data on patent applications that are eventually granted, we observe a truncated sample of the patents applied for in the 1996-1999 period. If there is any connection between the science citations contained in a patent application and the length of time it takes the U.S. PTO to evaluate the application, then this could bias the measured citation propensity downward.

Even if the per-patent propensity to cite science has not declined in the most recent period, one needs to put its ability to explain overall trends in the data into perspective. It is certainly true that the data reject the imposition of the constraint that per-patent citation propensities have not changed over time. Imposition of this constraint causes a significant degradation in the fit of the model to the data. But the degradation in model fit generated by this constraint is small relative to the degradation in fit generated by imposing the constraint that the relative propensity of different patent classes to cite science is the same or the constraint that the relative citedness of different categories of science is the same. In other words, changes in the distribution of patenting across technologies and changes in the distribution of publications across fields explain much more of the total variance in patent citations to science than does changes in per-patent citation behavior over time.

While the propensity to cite relative to the base period rises, albeit not monotonically, with later cohorts of patents, there is little evidence that the relative per-paper citedness of different cohorts of papers is rising over time.[30] However, this statement requires some qualification. The nature of the lag structure suggests that the lags between the appearance of a scientific article and its initial impact on patenting are generally short. Controlling for these lags, one sees some decline in the estimated citedness of more recent cohorts of publications. Recall, however, that this coefficient has an interpretation of "citedness per paper." The volume of publications has been trending upward over time – on average, by the end of my sample period, publication levels had expanded by over 50% relative to the beginning of the sample. Thus, even though the per-paper measure of citedness has modestly declined, the expanding volume of publication has more than compensated, pushing up the total number of citations to more recent cohorts of papers.

Given the extent to which aggregate numbers of citations are driven by biotech, I break the data into a biotech-only subsample and a subsample from which papers in bioscience and patents in biotechnology are excluded. This partition of the data allows me, at least in principle, to examine changes in the bioscience-biotechnology nexus in some detail. Then, I can separately estimate the key parameters of the citations function for the "non-bio" subsample, such that the parameters of obsolescence and diffusion –

---

[30] This result parallels that of Jaffe and Trajtenberg, who also find no evidence of an increase in citedness over time once they control for lags between cited and citing patents.

constrained here to be the same across fields of science -- are not driven by observations in the bioscience-biotech nexus.[31]

When partitioning the data in this manner, the field categories must obviously be adjusted. In the biotech-only subsample, the "drugs and medical" patent category is subdivided into drugs, surgery/medical instruments, and "other drug and medical." The bioscience categories are now clinical medicine, "other biotech," and biochemistry/biophysics/molecular biology. The other categories remain as before. Examination of a regression on the biotech subsample demonstrates that patent citations are concentrated in drug patents and the most heavily cited science category is the "other biotech" field. Despite its identification with information technology, Stanford is the most highly cited institution.[32] Compared to the full sample, the geographic bias toward California inventors and the institutional bias toward public science assignees are evidently less pronounced. One sees a significant rise in propensity to cite, relative to the base period, in the mid-1990s. While the propensity to cite in the most recent period seems to have actually *fallen* relative to the base period, the differences are statistically indistinguishable from zero. Throughout the 1990s, patenting in these most intensely citing categories was rising rapidly. In these regression results, we also see an increase (albeit a possibly temporary one) within those categories in measured per-patent propensity to cite science. While the coefficients on the paper publication year cohort

[31] Jaffe and Trajtenberg (1996, 2002) allowed the obsolescence parameter to vary across fields of technology, but constrained the diffusion parameter to be the same. In results not shown in this draft, I estimated a variant of the baseline specification that allowed obsolescence to vary in this manner. This produced estimates of varying obsolescence across fields that are similar to those found by Jaffe and Trajtenberg, and this did not change the qualitative pattern of the other multiplicative category fixed effects. However, given the preponderance of citations in the bioscience-biotech nexus, it seemed to make sense estimate citations functions separately for a bio-tech only subsample and a non-biotech sample, allowing both obsolescence and diffusion to vary over these samples.

[32] Note that I have drawn the institutional boundary of Stanford University to include its medical school.

dummies suggest a modest decline in per-paper citedness, the increase in publication volumes in these fields more than compensates for this. The measured obsolescence parameter is slightly smaller, while the measured diffusion parameter is larger, relative to the overall sample.

Interestingly, the non-biotech subsample generates a significantly different pattern of results. The aggregate patent classes used are computers and communications (IT), general electronics, mechanical inventions, chemicals (the base category), and a catch-all "other" category. Science aggregates are engineering and technology, physics, chemistry, and a catch-all "other science" category. The other categories remain as before.

Within "non-biotech," the IT patent classes cite science most frequently, displaying a propensity to cite that is nearly 13 times as high as the base category. General electronics patents are more than 6 times as likely to cite science, while mechanical patents are three times as likely. Articles in the physics fields are nearly 22 times more likely to be cited than base category articles. The physics aggregate includes some fields that relate to semiconductors and advanced materials. The engineering/technology aggregate (which includes computer science) is the next most highly cited, with a citedness per paper that is about 8 times greater than the base category. The rest of the sciences are significantly less likely to be cited.[33]

[33] Note that chemistry seems substantially less significant in this subsample than it did in the overall sample, where it was an important source of cited papers and citing patents. This difference seems to stem from the significant interaction between chemistry and the life sciences. Many citations made by chemical patents are evidently made to articles in the life sciences. Likewise, many citations received by chemical papers come from patents in the bioscience-biotech nexus. In other words, the chemical field can be viewed as being on the border of the bioscience-biotech nexus, and excluding papers and patents in this nexus reduces the measured importance of chemistry in patent citations to science.

In a striking contrast with earlier results, geographic localization seems to be much higher in this subsample. California-based inventors display a much higher likelihood of citing California science than the base (non-California U.S.) category of inventors. While intra-national localization is apparently higher, international localization is lower – the tendency of non-American inventors to cite California science is nearly 75% as high as that of non-California American inventors. This pattern of results could very well reflect the increasing geographic concentration of the U.S. information technology industry in California, as well as strong growth by inventors based outside the United States (particularly East Asia) in patenting in IT-related classes.

Another contrast with earlier results is a much higher propensity (relative to firms) for patents generated by public science institutions to cite science. Public science institutions are more than 25 times as likely to cite science as are firm patents, controlling for patent class. The category of "other institutions," is less likely to cite science in these fields, corresponding to the less significant role played by this category of assignee in non-biotech patenting.

The patterns suggested by the coefficients on patent application year cohorts and paper publication year cohorts also suggest patterns that differ from those in previous regressions. Controlling for changes in the volume and distribution of publications and patents, later cohorts of patents display a markedly lower per-patent propensity to cite science. The volume of patents in these categories has grown substantially, but not enough to fully offset the estimated decline in per-patent propensity to cite. On the other hand, more recent cohorts of papers are more likely to be cited, although the estimated differential relative to the base category is not always statistically significant.

Publications have grown in non-biotech sciences, particularly in physics and engineering technology, reinforcing the impact of higher citedness of the later paper cohorts. If the evidence from the bioscience-biotech nexus provides support for the notion that inventors have moved closer to academic science, the evidence here is perhaps more consistent with the view that academic science has moved closer to invention.

Finally, the estimated obsolescence coefficient is substantially higher than the overall sample, while the diffusion parameter is lower. This implies that citations to science in these categories arrive more quickly, decay more rapidly, and peak at a lower level. Viewed in this context, the results on the paper publication year cohort effects are even more striking. Even controlling for these shorter lags, the incidence of citation has risen for more recent cohorts of papers.

V. Conclusions and Extensions

Relative to other indicators of knowledge flow from academia to the private sector, citations to academic papers are relatively numerous, rich, and widely available across campuses and scientific disciplines. Quite simply, there is a great deal of information to be obtained from this source, and the existing literature has only begun this process. The beginning sections of the paper described some of the basic lessons of these data, based on a representative random sample of U.S. patents. While the logit and negative binomial regressions described therein can control for changes in citations driven by changes in the distribution of patenting across technologies, I was unable, in that context, to control for changes in the underlying distribution of potentially cited science.

At the risk of sacrificing some generality, the citations function estimates focused on California research universities allow me to control for changes in both patenting and publication. While still preliminary, the citation function results offer some suggestive insights into the changing impact of academic science on industrial innovation. First, this impact is quite concentrated in time, geographic location, and, especially, in technology space. There is a widely held perception that the links between academic science and industrial innovation have increased across a broad range of disciplines. This perception, of course, reflects a notion of connection between science and innovation that is much broader than the one used in this paper. However, if one focuses on the incorporation of frontier scientific research into industrial invention as evidenced by patent citations, then the evidence presented here suggests that link is, to a great extent, a phenomenon within the bioscience-biotech nexus. Patents within this nexus have been much more likely to cite science – and increasingly so over time, although the most recent period suggests a possible decline in citation propensity. In other words, citation patterns suggest that, at least for a while, inventors in this nexus moved closer to academic science. This, coupled with the growth in patenting in bioscience-citing classes relative to other categories of patents, explains a significant portion of aggregate increase in patent citations to science. Papers within this nexus have not shown an increased citedness over time, but the increase in publication in these areas is so large as to be an additional driving factor in increased citations to science.

Discussions with industry experts suggest that the increase in citations has been driven by a shift in the focus of research in the biomedical industries, the entry of academic scientists into the marketplace through start-up enterprises, and an expansion of patents to encompass classes of innovations such as genetically modified organisms and "research tools." In principle, one could attempt to quantify the degree to which changes in patent citations to science in biotechnology are statistically associated with each of these changes. Such an investigation is the subject of ongoing research. As noted in the draft, in the most recent period, we find some evidence of a modest decline in per-patent citation propensity, raising an intriguing question about the permanence of recent growth in the linkage between science and technology in this nexus. Further exploration of this apparent decline is also the subject of ongoing research.

Outside the bioscience-biotech nexus, one sees evidence – albeit less pronounced – of a secondary nexus in the IT-related disciplines. One finds a concentration of citation activity by IT-related patents to IT-relevant scientific disciplines. However, the citation patterns in this nexus differ from that in "bio" nexus in a number of dimensions. There is no evidence of rising propensity to cite science – in fact, the data suggest a rather striking monotonic decline in per-patent propensity to cite for more recent patent cohorts. On the other hand, there has been a rapid increase in publication volumes in the most highly cited fields and there is evidence that these more recent paper cohorts are more frequently cited, even controlling for the shorter lags between publication and patent citation. On average, inventors outside of biotech/pharmaceuticals have not moved substantially closer to academic science – the measured increase in citations largely reflects changes in the distribution of patenting across fields and, potentially, the increased relevance to industrial technology of more recent cohorts of scientific papers.

Another finding of the paper is evidence at the patent level that the incidence of citation of academic science is positively associated with measures of invention "quality." While this evidence is consistent with the idea that knowledge spillovers from academia make private inventions better, for reasons discussed at length in the paper, this does not constitute proof that the chain of causality runs from citation to invention quality. Certainly, further analysis is needed at the assignee (firm) level.

Patents do not spring spontaneously into existence – they are created by inventors, most of whom are employed by innovating *firms*. This is perhaps reason enough to conduct analysis at the firm level. There are also econometric reasons for doing so. It is likely that the patents created by a firm possess some common characteristics – it is unlikely that these are realizations of some independent, identically distributed random variable. Secondly, patents, *per se*, possess no panel dimension, because a given patent only appears once in my sample. In contrast, innovative firms generate multiple patents per year. By tracking the patents of a corporate assignee over time, I am able to bring to bear all of the usual fixed effects econometric techniques. This can provide useful leverage in sorting out, for instance, whether citation of academic science actually leads to higher levels of innovative performance. Pursuing such analysis at the assignee level is the focus of current research.

Bibliography

Adams, J., 1990, "Fundamental Stocks of Knowledge and Productivity Growth," *Journal of Political Economy* 98: 673-702.

Adams, J. and Z. Griliches, 1996, "Research Productivity in a System of Universities," NBER working paper no. 5833.

Agrawal, A. and I. Cockburn, 2003, "The Anchor Tenant Hypotheses: Exploring the Role of Large, Local, R&D-Intensive Firms in Regional Innovation Systems," forthcoming in the *International Journal of Industrial Organization*.

Audretsch, D. and P. Stephan, 1996, "Company-Scientist Locational Links: The Case of Biotechnology," *American Economic Review*, Vol. 86, No. 3.

Barnes, M., D. Mowery, A. Ziedonis, 1998, "The Geographic Reach of Market and Nonmarket Channels of Technology Transfer: Comparing Citations and Licenses of University Patents," working paper.

Branstetter, L., 2000, "Is FDI a Channel of R&D Spillovers: Evidence from Japan's FDI in the U.S." NBER working paper.

Bush, V., 1945, *Science – the Endless Frontier: A Report to the President on a Program for Postwar Scientific Research*, Washington, D.C.: U.S. Government Printing Office.

Cameron, A. C. and P. Trivedi, 1998, *The Regression Analysis of Count Data*, Econometric Society Monograph No. 30, Cambridge: Cambridge University Press.

Cockburn, I. and R. Henderson, 2000, "Publicly Funded Science and the Productivity of the Pharmaceutical Industry," paper prepared for the NBER Conference on Science and Public Policy.

Cockburn, I. and R. Henderson, 1998, "The Organization of Research in Drug Discovery,"*Journal of Industrial Economics*, Vol XLVI, No. 2.

Cohen, W., R. Florida, L. Randazzese, and J. Walsh, 1998, "Industry and the Academy: Uneasy Partners in the Cause of Technological Advance," in R. Noll, ed., *Challenges to the Research University*. Washington, D.C.: Brookings Institution

Darby, M. and L. Zucker, 2003, "Grilichesian Breakthroughs: Inventions of Methods of Inventing and Firm Entry in Nanotechnology," NBER working paper. 9825.

Evenson, R. and Y. Kislev, 1976, "A Stochastic Model of Applied Research," *Journal of Political Economy* 84 (2): 265-282.

Faulkner, W. and J. Senker, 1995, *Knowledge Frontiers: Public Sector Research and Industrial Innovation in Biotechnology, Engineering Ceramics, and Parallel Computing*, Oxford: Clarendon Press.

Fleming, L. and O. Sorenson, 2001, “Science as a Map in Technological Search,” working paper.

Gambardella, A, 1995, *Science and Innovation: The U.S. Pharmaceutical Industry during the 1980s*, Cambridge: Cambridge University Press.

Henderson, R., A. B. Jaffe, and M. Trajtenberg, 1998, “Universities as a Source of Commercial Technology: A Detailed Analysis of University Patenting, 1965-1988,” *Review of Economics and Statistics*, 119-127.

Hicks, D., T. Breitzman, D. Olivastro, K. Hamilton, 2001, “The Changing Composition of Innovative Activity in the US – A Portrait Based on Patent Analysis,” *Research Policy* 30, pp. 681-703.

Jaffe, A., 1989, “The Real Effects of Academic Research,” *American Economic Review*, 79 (5), pp. 957-70

Jaffe, A., M. Trajtenberg, and R. Henderson, 1993, “Geographic Localization of Knowledge Spillovers as Evidenced by Patent Citations,” *Quarterly Journal of Economics,* Vol. CVIII, No. 3.

Jaffe, A. and M. Trajtenberg, 1996, “Flows of Knowledge from Universities and Federal Labs: Modeling the Flow of Patent Citations over Time and across Institutional and Geographic Boundaries,” NBER working paper no. 5712.

Jaffe, A., M. Fogarty, and B. Banks, (1998), “Evidence from Patents and Patent Citations on the Impact of NASA and Other Federal Labs on Commercial Innovation,” *Journal of Industrial Economics*, Vol. XLVI, No. 2.

Jensen, R. and M. Thursby, 1999, “Proofs and Prototypes for Sale: The Licensing of University Inventions,” *American Economic Review,* 91, pp. 240-259.

Kortum, S. and J. Lerner, 1997, “Stronger Protection or Technological Revolution: Which is Behind the Recent Surge in Patenting?” working paper.

Lambert, D., 1992, “Zero-inflated Poisson Regression, with an Application to Defects in Manufacturing,” *Technometrics* 34: 1-14.

Lieberman, M., 1978, “A Literature Citation Study of Science-Technology Coupling in Electronics,” *Proceedings of the IEEE*, 36 (1), pp. 5-13.

Lim, K., 2001, "The Relationship between Research and Innovation in the Semiconductor and Pharmaceutical Industries," working paper.

Mansfield, E., 1995, "Academic Research Underlying Industrial Innovations: Sources, Characteristics, and Financing," *The Review of Economics and Statistics* 77: 55-65.

Marquis, D. G. and T. Allen, 1966, "Communication patterns in Applied Technology," *American Psychologist*, 21, pp. 1052-1060.

Mowery, D., R. Nelson, B. Sampat, and A. Ziedonis, 1998, "The Effects of the Bayh -Dole Act on U.S. University Research and Technology Transfer: An Analysis of Data from Columbia University, the University of California, and Stanford University," working paper

Narin, F., K. Hamilton, and D. Olivastro, 1997, "The Increasing Linkage Between U.S. Technology and Public Science," *Research Policy* 197: 101-121.

Narin, F., 1995, "Linking Biomedical Research to Outcomes – The Role of Bibliometrics and Patent Analysis," CHI Working Paper.

Office of Technology Transfer, University of California, 1997, Annual Report: University of California Technology Transfer Program. Oakland, CA: University of California.

Price, D., 1965, "Is Technology Historically Independent of Science? A Study in Statistical Historiography," *Technology and Culture*, 6, pp. 553-568.

Rosenbloom, R. and W. Spencer, 1996, *Engines of Innovation: U.S. Industrial Research at the End of an Era*, Boston: Harvard Business School Press.

Schmookler, J., 1966, *Invention and Economic Growth*, Cambridge, MA: Harvard University Press.

Shane, S., 2000, "Prior Knowledge and the Discovery of Entrepreneurial Opportunities," *Organization Science*, 11, pp. 448-469.

Shane, S., 2001, "Technological Opportunities and New Firm Creation," *Management Science*, 47, pp. 205-220.

Sorenson, O. and L. Fleming, 2001, "Science and the Diffusion of Knowledge," working paper.

Stephan, P., 1996, "The Economics of Science," *Journal of Economic Literature* 34: 1199-1235.

Stern, S., 1999, “Do Scientists Pay to Be Scientists?” NBER Working Paper No. 7410.

Thursby, J. and M. Thursby, 2002, “Who is Selling the Ivory Tower? Sources of Growth in University Licensing,” *Management Science*, 48, pp. 90-104.

Zucker, L., M. Darby, and M. Brewer, 1998, “Intellectual Capital and the Birth of U.S. Biotechnology Enterprises,” *American Economic Review*, 88: 290-306.

Figure 1 Patent Citations to Academic Research

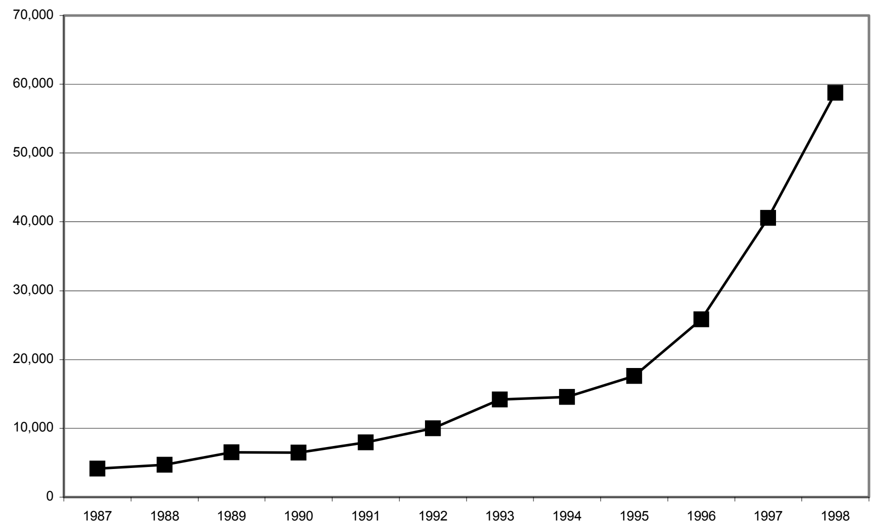

Figure 2 Citations to UC papers vs other indicators

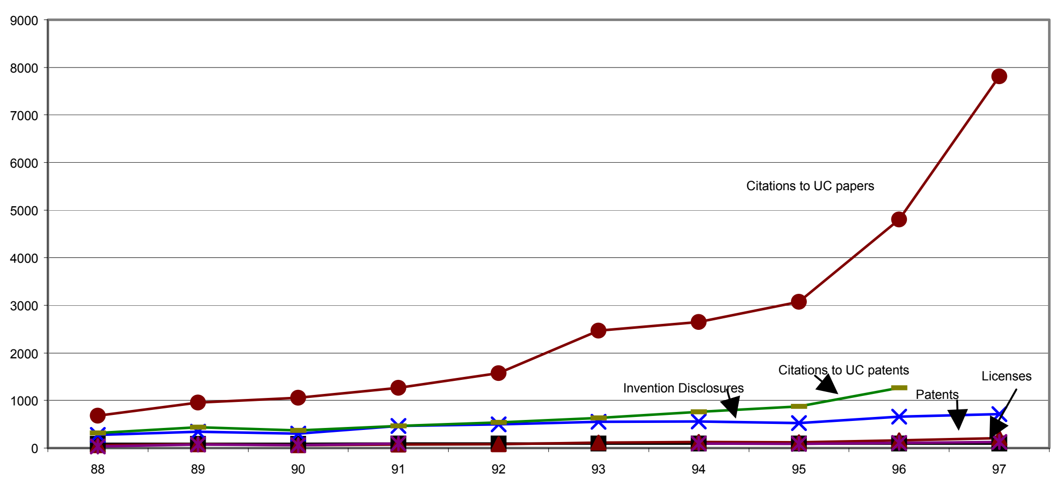

Table 1 Negative Binomial Regressions on the Determinants of Academic Citation

Variable Category	*Variable*	*All citations*	*ISI journals only*	*University-affiliated authors*
Type of Assignee	**Universities**	1.60 (.146)	1.90 (.173)	2.51 (.225)
	Nonprofit R&D organization	1.09 (.319)	1.42 (.377)	1.56 (.495)
	U.S. government agency	.623 (.165)	.758 (.206)	.618 (.290)
	Foreign assignee	-.527 (.048)	-.474 (.063)	-.579 (.095)
	Other	-1.089 (.065)	-1.06 (.088)	-.685 (.125)
Technology Class	**Chemicals**	1.94 (.066)	2.30 (.092)	2.45 (.140)
	Communications/ Computers	1.65 (.074)	1.42 (.105)	1.34 (.161)
	Drugs/Medical	2.90 (.078)	3.46 (.104)	3.97 (.151)
	Electronics	1.42 (.069)	1.61 (.097)	1.41 (.150)
	Mechanical devices	.020 (.075)	-.041 (.113)	-.403 (.194)
Science Center		.436 (.050)	.492 (.064)	.649 (.091)
Application Cohort Effects		Yes	Yes	Yes
Obs		29,876	29,876	29,876
Log-Likelihood		-20,695.575	-12,510.80	-6,641.674

Figure 3 Application Year Cohort Effects, Negative Binomial Regressions

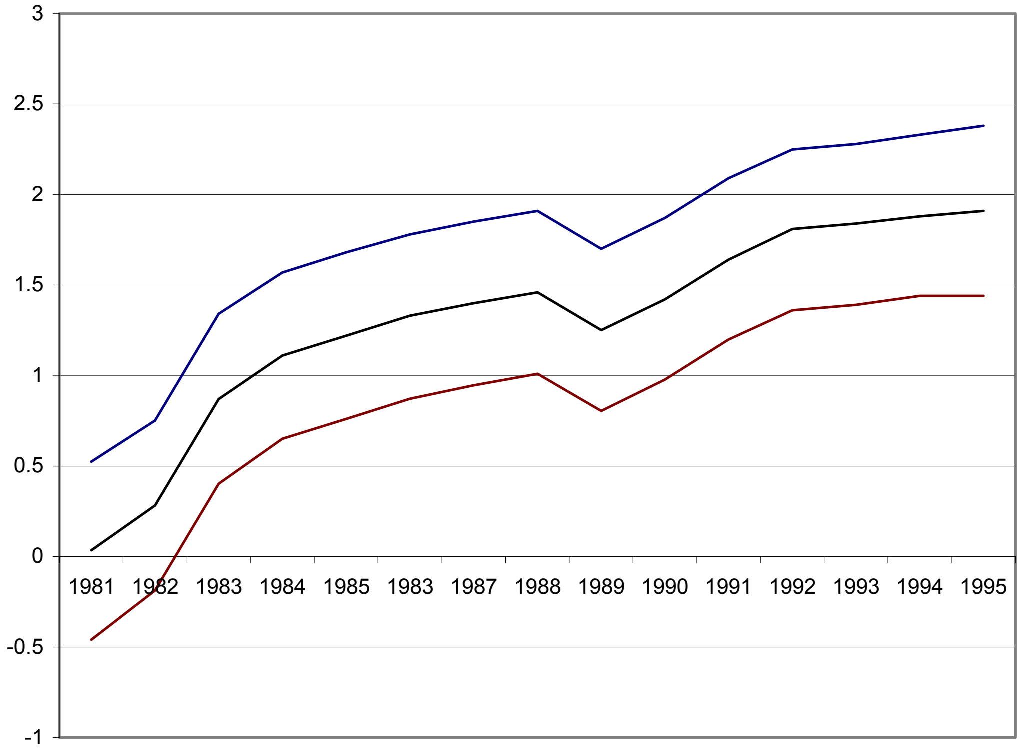

Table 2 Results on Quality Differentials

Variable	*Coefficient*	*Std. Err.*	*Implied Difference*
Claims	2.25	.220	15.1%
Forward Citations	.581	.101	5.5%
Generality	.029	.008	7.5%

Table 1 reports the regression coefficient on a dummy variable identifying patents that cite scientific research. Regressions control for technological field and application year effects. The measures of patent quality in the table are used as the dependent variable in the regression, as in Henderson, Jaffe, and Trajtenberg, 1998.

Figure 4 UC Academic Publishing across Fields, 1999

Source: Institute for Scientific Information

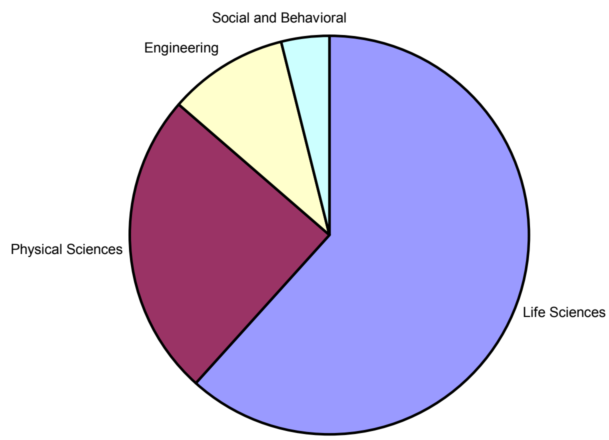

Figure 5 Patent Grants by Technology Category, 1981-1999

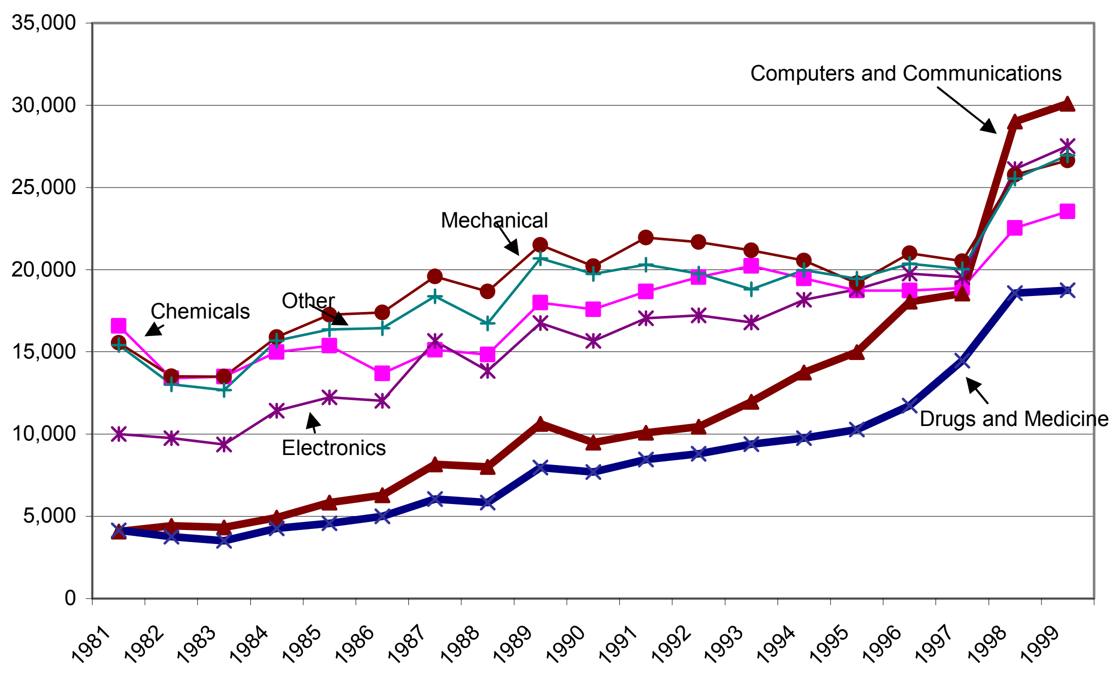

Table 3 Citation Function Results, Full Sample

Variable	Coefficient	T-statistic for H_0: Parameter=1
Computer Communications	0.042641	-101.27
Drugs/medicine	2.580614	65.35
Electronics	0.054423	-109.91
Mechanical	0.013223	-91.96
Other	0.046111	-83.27
Biomedical research	37.58918	7.76
Chemistry	4.612922	6.11
Clinical Medicine	4.880407	6.27
Eng/Technology	0.244953	-6.14
Other Science	0.310768	-6.01
Physics	0.441311	-4.58
Caltech	1.211671	18.34
Berkeley	0.578289	-54.87
Davis	0.408163	-79.11
Irvine	0.441956	-64.1
Los Angeles	0.378945	-89.43
Riverside	0.267284	-74.76
Santa Barbara	0.325639	-57.73
Santa Cruz	0.245496	-60.65
San Diego	1.106468	10.74
Santa Francisco	0.890317	-13.14
USC	0.56083	-48.03
US-CA	2.754931	89.72
Non-US	0.455029	-64.86
Other Institutions	1.720603	30.52
Public Science	3.895451	84.39
App year 88-90	0.839112	-9.43
App year 91-93	1.483676	16.41
App year 94-96	2.141093	22.88
App year 97-99	1.533536	12.35
Paper pub year 86-89	0.937367	-5.41
Paper pub year 90-93	0.867258	-8.19
Paper pub year 94-97	0.771549	-11.01
β_1 **(obsolescence)**	0.258239	0.002268
β_2 **(diffusion)**	3.72E-08	4.73E-09

Table 4, Bioscience/Biotechnology Nexus, R-squared:

Variable	Coefficient	T-statistic for H_0: Parameter=1
Surgery/medical Instruments	0.38233	-71.87
Biotechnology	0.012987	-113.17
Other Drug & Medical	0.014335	-68.08
Clinical Medicine	0.132214	-72.4
Other Biotech	2.338469	29.22
Caltech	0.857317	-7.07
Berkeley	0.436655	-38.01
Davis	0.28297	-49.13
Irvine	0.31106	-39.39
Los Angeles	0.336879	-46.12
Riverside	0.192002	-40.62
Santa Barbara	0.240868	-31.5
Santa Cruz	0.25059	-27.85
San Diego	0.751386	-14.23
Santa Francisco	0.659127	-22.01
USC	0.500215	-26.16
US-CA	2.020528	30.15
Non-US	0.411354	-34.56
Other Institutions	1.381069	10.64
Public Science	1.749713	24.98
App year 88-90	0.760042	-8.28
App year 91-93	1.083733	1.84
App year 94-96	1.340517	4.82
App year 97-99	0.938605	-0.96
Paper pub year 86-89	0.827982	-6.97
Paper pub year 90-93	0.635044	-12.06
Paper pub year 94-97	0.601088	-9.11
		Asymptotic Standard Error
β_1 **(obsolescence)**	0.241896	0.005724
β_2 **(diffusion)**	2.03E-05	8.11E-07

Table 5   Citation Function, Non-biotech Sample

Variable	Coefficient	T-statistic for H_0: Parameter=1
Computers Communications	12.96369	10.38
Electronics	6.535692	9.47
Mechanical	3.27098	7.36
Other	0.297798	-5.65
Eng/Technology	7.674823	5.29
Other Science	0.292769	-4.39
Physics	21.66557	5.84
Caltech	0.259771	-98.2
Berkeley	0.116574	-119.39
Davis	0.415936	-52.22
Irvine	0.014034	-93.23
Los Angeles	0.069184	-114.2
Riverside	0.027343	-73.11
Santa Barbara	0.068202	-111.47
Santa Cruz	0.020609	-71.11
San Diego	0.069723	-103.1
Santa Francisco	0.019566	-118.73
USC	0.150845	-91.43
US-CA	11.51926	16.11
Non-US	0.735649	-3.44
Other Institutions	0.615733	-1.72
Public Science	25.43173	11.54
App year 88-90	0.436647	-55.01
App year 91-93	0.273626	-68.94
App year 94-96	0.124899	-99.31
App year 97-99	0.112237	-69.24
Paper pub year 86-89	1.173268	6.12
Paper pub year 90-93	1.070638	1.26
Paper pub year 94-97	1.255886	1.97
		Asymptotic Standard Error
β_1 (obsolescence)	0.484528	0.005581
β_2 (diffusion)	2.14E-09	4.53E-10

Figure 6 Fitted Citation Frequency (Base Category)

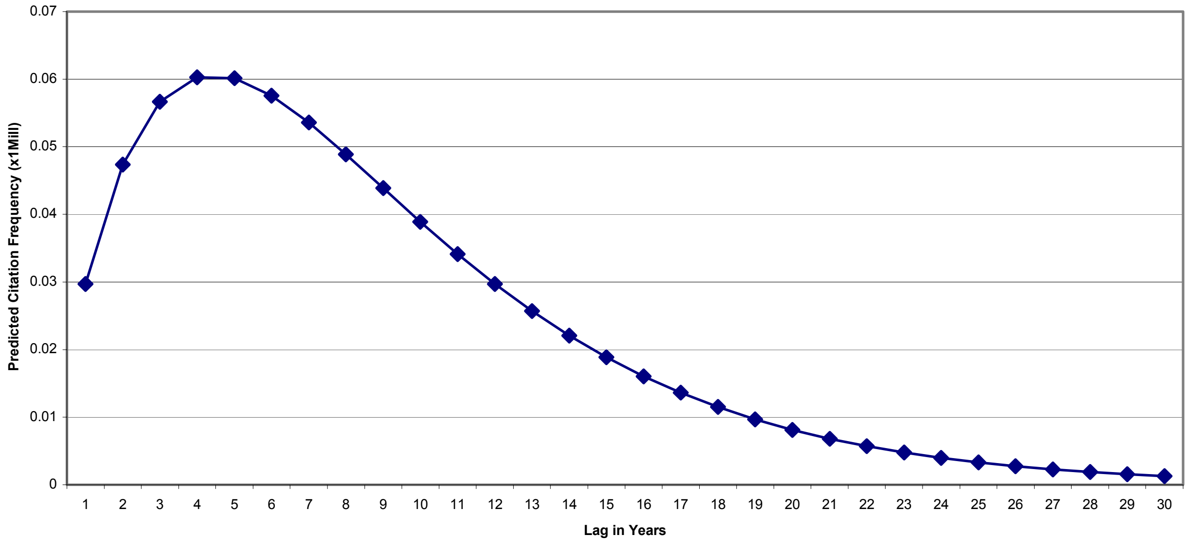

Figure 7   Base Case versus Biomedical Research

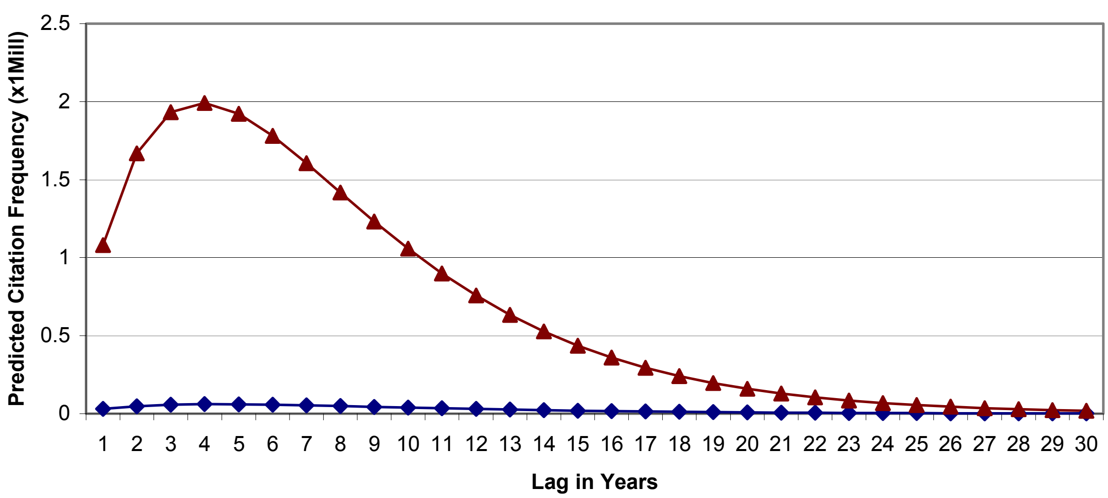

Figure 8 Citations by California-based inventors versus Non-California Inventors

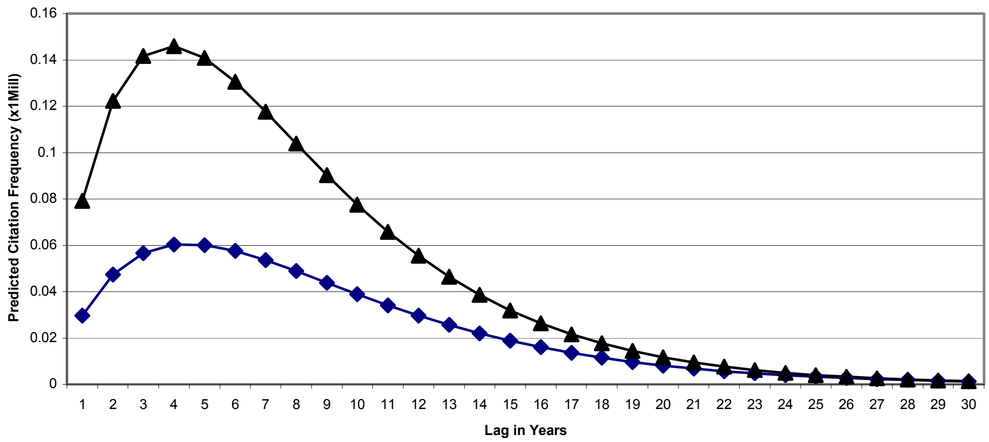

Figure 9   Citations to UC Berkeley Papers, US