CHANGES IN UNIVERSITY PATENT QUALITY AFTER THE BAYH-DOLE ACT: A RE-EXAMINATION*

Bhaven N. Sampat
School of Public Policy
Georgia Institute of Technology
Atlanta, GA 30332
bhaven.sampat@pubpolicy.gatech.edu

David C. Mowery
Haas School of Business
U.C. Berkeley
Berkeley, CA 94720
mowery@haas.berkeley.edu

Arvids A. Ziedonis
Wharton School of Business
University of Pennsylvania
Philadelphia, PA 19104
ziedonisa@wharton.upenn.edu

ABSTRACT: The Bayh-Dole Act permitted universities to retain rights to patents resulting from government funded academic research, and encouraged university entry into patenting and licensing. Though it is widely recognized to be a major change in federal policy towards academic research, surprisingly little empirical analysis has been directed at assessing the impacts of Bayh-Dole on the academy and on university-industry research relationships. An important exception is the work of Henderson, Jaffe and Trajtenberg (1998) which examined the impact of Bayh-Dole on the quality of university patents, as measured by the number of times they are cited in subsequent patents. The authors found that the quality of academic patents declined dramatically after Bayh-Dole, a finding that has potentially important policy implications. In this paper, we show that their results are not robust to using a longer stream of citation data, i.e. they are driven by truncation bias. This has important implications not only for thinking about Bayh-Dole, but also for future work using patent citations as economic indicators.

* Stephen Cameron, Rebecca Henderson, Ken Leonard, Frank Lichtenberg, and Richard Nelson provided useful comments on earlier drafts of this paper. The usual disclaimer applies.

1. Introduction

The Bayh-Dole act, passed in 1980, is widely recognized as a major change in federal policy towards academic research. Bayh-Dole allowed universities to retain the rights to patents resulting from federally funded research and encouraged universities to become actively involved in patenting and licensing, activities they had historically avoided. There has been much discussion of the positive and negative impacts of Bayh-Dole on the academy and on technology transfer, but relatively little empirical analysis. An important exception is the work of Henderson, Jaffe and Trajtenberg (1995, 1998a, 1998b), which examined the impact of Bayh-Dole on the quality of university patents, as measured by the number of times they are cited in subsequent patents.[1] The authors found that the quality of academic patents declined dramatically after Bayh-Dole, a finding that has potentially important policy implications.

In this paper, we show that the post-Bayh-Dole "quality decline" disappears when patent-citation data covering a longer period than was available to Henderson and colleagues are used in similar empirical tests. This evidence suggests that the "quality decline" that Henderson et al. found in their analysis may reflect changes in the intertemporal distribution of citations to university patents, rather than a significant change in the number of citations these patents receive.

In Section 2, we briefly discuss the rationale for the passage of Bayh-Dole, and the post-Bayh-Dole growth of university patenting and licensing. In Section 3, we discuss the use of patent citations as economic indicators, and summarize the analysis conducted by Henderson et al. In Section 4, we replicate their results, and show that they are not

[1] See also Jensen and Thursby (2001) and Mowery et al. (2001).

robust to use of more complete data on citations. In particular, we argue that the quality decline is an artifact of the "truncation bias" inherent in the use of patent citation data. In Section 5, we consider why the various statistical techniques used by Henderson et al. to control for truncation were not effective. We conclude in Section 6 with a discussion of implications for the analysis of the effects of the Bayh-Dole Act and for future work using patent citation data.

2. The Bayh-Dole Act and the Growth of University Patenting

American research universities have been an important source of knowledge and technologies useful to U.S. industry since the beginning of the twentieth century (Rosenberg and Nelson 1994). A range of industries critical to American technological leadership, including chemicals, agriculture, pharmaceuticals, and electronics, relied heavily on basic and applied academic research during the twentieth century. The flows of knowledge and technology transfer between U.S. universities and industry occurred via informal channels, including presentations at conferences and publication in journals, as well as more direct channels that include hiring of graduate students, consulting relationships with faculty, and licensing of university patents.[2]

Bayh-Dole created a uniform federal patent policy that allowed universities to retain rights to any patents resulting from government funded research and to license these patents on an exclusive or non-exclusive basis. Before the passage of the Act, universities wishing to retain title to patents resulting from federally funded research had

[2] See Cohen et al. (2002) and Agrawal and Henderson (2002) for discussion of the relative importance of different channels of knowledge and technology transfer.

to negotiate an Institutional Patent Arrangements (IPA) with individual funding agencies or petition for title on an invention by invention basis.

The Bayh-Dole Act was passed in the throes of the "competitiveness crisis" of the 1970s and 1980s, in response to the belief that IPAs and other arrangements for university patenting of publicly funded research results impeded technology transfer and commercialization of these results, thereby weakening U.S. competitiveness. In particular, the framers of Bayh-Dole argued that if universities could not be granted clear title to patents and allowed to license them exclusively, firms would lack the incentive to develop and commercialize university inventions. This argument was based on "evidence" that government owned patents had lower utilization rates than those held by contractors, evidence that Eisenberg (1996) has shown to be faulty.[3]

The Bayh-Dole Act, which passed in 1980 and became effective in 1981, had three primary effects. First, it created a uniform policy on disposal of patent rights resulting from federally funded university research, giving title to university performers of research rather than the funding agencies. Second, it gave universities permission to license government-funded inventions on an exclusive basis, without requiring approval from funding agencies. Third, and perhaps most importantly, it offered strong Congressional endorsement of active university involvement patenting and licensing, activities that they historically had avoided (Mowery and Sampat 2001b, Mowery et al. 2001).In the twenty years since its passage, many universities with little previous

[3] The framers of Bayh-Dole interpreted these data as indicators of the economic impact of university research, neglecting the range of other formal and informal channels through which firms historically benefited from university research. In opening the Senate Bayh-Dole hearings, one of the architects of the legislation, Senator Bayh, cited the low rate of commercialization of government owned patents as evidence of "very little return on the billions of dollars we spend every year on research and development" (Senate Committee on the Judiciary 1979, 2).

patenting experience entered the patenting business, and "incumbent" universities expanded their patenting (Mowery et al. 2001).

Some observers have expressed concern that the incentives created by Bayh-Dole may have led to a shift the content of academic research towards more "applied" work and away from fundamental research, a development that could have detrimental long-term effects (Dasgupta and David 1994). Other observers, however, have praised the Act as an important contributor to the U.S. competitive “revival” of the 1990s and the growth of the “New Economy” that received great attention in the latter half of that decade (AUTM 1996, GAO 1998; Congressional Joint Economic Committee, 2000; OECD, 2001).[4] Yet virtually no quantitative work on the effects of Bayh-Dole was published prior to the the analyses by Henderson, Jaffe, and Trajtenberg.

3. Patent Citations and the Post-Bayh-Dole "Quality Decline"

Henderson and colleagues (1995, 1998a, 1998b), (hereafter, HJT) used information on patents assigned to universities from 1965-1988 and citations to those patents observed until 1992 to examine changes in the characteristics of patents issued to U.S. universities before and after Bayh-Dole. In this section, we discuss the use of patent citations as economic indicators, and review HJT's main results.

[4] The belief that Bayh-Dole "worked" in the United States has led other industrialized countries to try to emulate it. Aa recent report by the Organization of Economic Cooperation and Development (OECD) notes "[I]n nearly all OECD countries there has been a marked trend towards transferring ownership of publicly funded research from the state (government) to the (public or private) agent performing the research. The underlying rationale for such change is that it increases the social rate of return on public investment in research" (OECD 2002, 48).

3.1 Patents and Patent Citations

In the United States, in order to be patentable, an invention ("a new and useful process, machine, manufacture, or composition of matter" (35 USC § 101), must be shown to be both "novel" (35 USC § 102) and "non-obvious" (USC 35 § 103). Under the novelty criterion, an invention cannot be patented if it was known or previously used. The non-obviousness criterion means that an invention cannot be patented if "the differences between the subject matter sought to be patented and the prior art are such that the subject matter as a whole would have been obvious at the time the invention was made to a person having ordinary skill in the art to which the subject matter pertains" (35 USC § 103a).

In order to assess whether an invention meets these standards, patent examiners at the United States Patent Office (USPTO) must identify the prior art, i.e. information related to the invention that was known or already in use. For each issued patent, prior patents against which the novelty and non-obviousness of the patent were evaluated are listed in a section entitled "References Cited, U.S. Patent Documents" on the front page of the patent.

Economists began using patent citations and patent citation-based data as measures that were superior to simple patent counts as measures of innovative output (see Griliches 1990 for a review). The large variance in the economic and technological significance of individual patents rendered patent counts extremely noisy indicators of the innovative output of a firm or a government programs. But weighting patents by the number of times they are cited in subsequent patents provided a better measure of the economic and technological importance of these patents (See Trajtenberg, 1990, for one of the first applications of this measure).

Economists have used citation counts as proxies for two different measures of the value of a patented invention. One is the patent's "technological importance," based on the assumption that when patent A cites patent B, A is drawing on the knowledge embodied in B, or that B is a technological antecedent of A (Trajtenberg et al 1997; Caballero and Jaffe 1993; Jaffe and Trajtenberg 1999). Citation of patent B by many subsequent patents suggests that numerous inventions build upon the knowledge embodied in patent B, i.e. patent B has generated significant technological spillovers. Based on this logic, a significant stream of research has used citation counts to patents to assess the "importance" of patented inventions, based on the rationale that inventions that generate a higher level of spillovers are more economically or technologically important (see Jaffe 1998 for a review).

Economists have also used citation counts as proxies for measures of the private value of the invention to the patentholder in studies that have analyzed the relation between the market value of a firm and its citation-weighted patent stock (Hall et al. 2000; Shane and Klock 1997; Austin 1994). In empirical work on a different measure of the private value of patents, citation weights have been used in analyses of the probability that a given patent is litigated (Lanjouw and Schankerman 2000).

3.2 The HJT analysis of "quality decline" in U.S. university patents

The previously cited work by HJT concludes that expanded patenting by U.S. universities after Bayh-Dole was accompanied by a decline in the quality of these patents, as measured by citations. During the 1970s (before the passage of the Bayh-Dole Act), according to HJT,university patents were significantly more likely to be cited than a 1% random sample of patents. But after 1980, this difference diminished, and its statistical signficance disappears after the mid-1980s.

The meaning of this change in citations to U.S. university patents is open to multiple interpretations. Some authors have interpreted it as evidence that research priorities within universities may have shifted towards "applied research" after Bayh-Dole (e.g. Foray and Kazancigil 1999).[5] Implicit in this interpretation is the assumption that patents based on "basic" research would generate more spillovers, and hence citations, than "applied" research (cf. Trajtenberg et al. 1997).

However, Henderson and colleagues (1998) interpret these results as suggesting that there was *not* a shift towards applied research after Bayh-Dole, arguing that

> "… the Bayh-Dole Act and other related changes in federal law and institutional capability have not had a significant impact on the underlying rate of generation of commercially important inventions at universities. Universities either did not significantly shift their research efforts towards areas likely to produce commercial inventions, or, if they did, they did so unsuccessfully" (Henderson et al. 1998).

The logic underlying this interpretation is apparently that a "successful" shift towards applied research would yield patents more heavily cited by industry.

The fact that different scholars have drawn the opposite conclusions from the same results suggests that interpreting the economic meaning of patent citations can be difficult, a topic to which we return below.

---

[5] Henderson and colleagues also suggested this possible interpretation of the "quality decline" in a draft version of their paper, Henderson et al. (1995).

Although the interpretation of changes in the rate of citation to U.S. university patents after Bayh-Dole is itself ambiguous, the considerable attention that this result has attracted suggests that its robustness merits further analysis for three related reasons. First, as Henderson et al. noted, their analysis was done shortly after the passage of Bayh-Dole, and the full effects of this legislation may not yet have been revealed in the data on university patenting and citations to these patents. Second, the authors measure quality via citation counts, which are subject to truncation: a smaller portion of total citations is observable for more recent patents. This fact, combined with the fact that the quality decline is observed only for the most recent patents in their sample, suggests that the issue is worth revisiting with the benefit of a longer time span. Third, several recent papers (Mowery and Ziedonis 2001, Mowery, Sampat, and Ziedonis 2002) find little evidence of a decline in the citation intensity of patents issued to academic patenters, although these analyses employ different control samples and methodologies.

4. The Quality Decline Revisited

In light of the vulnerability of patent citation data to the truncation problems mentioned above, this section examines whether the quality decline reported in Henderson et al. (1998) is robust to the use of a longer span of citations than was available to those authors. In order to isolate the possible effects of truncation, we use a longer stream of citations, but otherwise structure our analysis to be as similar to that of HJT as possible. Specifically, we use a dataset consisting of all university patents applied for between 1975 and 1988 and granted before 1992, identical to the university patent dataset used by HJT for this period.[6] We also follow HJT in constructing a 1% random sample of all U.S. patents granted during the same period. Although we unfortunately do not have access to the exact control sample used by HJT, this control sample is sufficiently similar to theirs for the purposes of this analysis.[7] Third, we collect counts of all citations issued by the end of 1999 to the university patents and control sample patents. This information is extracted from the NBER Patent Citation database, described in Hall et al. (2001).

For purposes of comparison, for each potentially cited patent we create two citation counts: the number of citations issued by the end of 1992, which covers the period included in the HJT analysis, and the number of citations issued by 1999. Like HJT, we include self-citations in these counts, but our main conclusions are not sensitive to inclusion or exclusion of self-citations.

Our basic analysis is similar to that in HJT. We regress the number of citations to those patents on application year dummies for years 1975-1988,[8] patent class dummies for each of the patent classes spanned by the university and control patents, and application year dummies interacted with a dummy variable indicating whether the patent is a university or control sample patent. Specifically, we estimate equations of the form:

$$Citations = \sum_t \left[\alpha_t APP_t + \beta_t (APP_t * UNIV)\right] + \sum_c \lambda_c CLASS_c + \varepsilon$$

---

[6] Rebecca Henderson kindly provided us with access to the list of university patents used in their analyses.

[7] We checked the sensitivity of our basic results to the particular control sample employed by conducting the analysis discussed below using 5 different 1% control samples (constructed using STATA's "sample" command). The results from these analyses are basically similar to those reported below, and are available from the authors upon request.

[8] HJT use patents applied for from 1965 onwards. Unfortunately, it is difficult to obtain citation information for pre-1975 patents, so we begin this analysis with patents applied for in 1975. Since the "quality decline" is observed after the mid-1980s, this data limitation does not affect the substance or the conclusions of the present exercise.

where CITES gives the number of citations to patent i, APP_t is a dummy variable taking on the value of 1 if patent was applied for in year t (t=1975 …1988), UNIV is a dummy variable taking on the value of 1 if patent i is assigned to a university, and $CLASS_c$ is dummy variable taking on the value of 1 if the patent class is c[9]. The coefficients on the interaction terms, β, are estimates of the mean differences in the number of citations to university patents and patents from our random sample for a given application year, controlling for technological field effects. Like HJT, we estimate this equation using ordinary least squares.

The main results are shown in Table 1. The first column of the Table reproduces results from Table 4 of Henderson et al. (1998). The second column shows coefficients from a similar regression using the control sample described above. Importantly, for the estimation reported in the second column, the dependent variable is citation counts as observed in 1992, as in HJT. The point estimates are similar, and qualitative results (sign and significance of coefficients) are similar to those in column 1. The quality of university patents relative to the controls declines after Bayh-Dole, and the difference in citations to these two groups of patents becomes statistically insignificant after the mid-1980s. Indeed, university patents are actually cited less frequently than those in the control sample for patents applied for in 1987 and 1988, though this difference is not statistically significant.

TABLE 1 HERE

But when when "citations to 1999" (rather than "citations to 1992") is used as the dependent variable, the results change dramatically (column 3 of Table 1). University

[9] Our university and control patents span 394 patent classes.

patents are more highly cited than the controls for patents applied for throughout the 1975-88 period, and the difference is highly statistically significant. There is little evidence in column 3 of Table 1 of a post-Bayh-Dole "quality decline" in U.S. university patents.[10] This is seen more clearly in Figure 1, which plots the coefficients on the interaction term from both regressions, as well as the coefficients from HJT's analysis.

FIGURE 1 HERE

The difference between the results is attributable to our inclusion of an additional seven years of citation data for the same group of patents as those examined by HJT. The HJT analysis included only 4-7 years of citation data for the patents applied for by universities and entities in our control sample from 1985-1988, but we observe 11-13 years of citations for these patents. The shorter time span of patent citations apparently does not provide an accurate signal of the (relative) quality of later patents, or at least a signal comparable to that obtained for the earlier patents in the 1975-88 sample. In other words, the patent "quality decline" reported by HJT was an artifact of truncation bias.[11] In the next section, we discuss the effects of truncation in more detail.

[10] To test this more formally, we regressed citations on a pre-Bayh-Dole and post-Bayh-Dole dummy, each interacted with a university dummy, as well application year and patent class dummies. When using "citations to 1992" as the dependent variable, the estimated coefficient on the pre-Bayh-Dole*University dummy is 1.8 and the estimated coefficient on the post-Bayh-Dole*University dummy is .55. Based on an F-test of equality of coefficients, we can reject the hypothesis that the coefficients are equal. However, when using "citations to 1999" as the dependent variable, the pre- and post-Bayh-Dole coefficients are 3.2 and 2.8 respectively, and we cannot reject the hypothesis that they are equal. In other words, based on citations observed until 1999, there is no statistically significant post-Bayh-Dole change in the quality of university patents relative to the controls.

[11] In addition to the analysis reported above, we estimated quantile regressions to explore changes in "quality" over time. Least-squares regressions estimate conditional means, but quantile regressions estimate conditional quantiles, and allow for examination of the differences between citations to university and control patents within each quantile. This is useful for two reasons. First, it is well known that patents in the upper tail of the patent value distribution account for the bulk of the value of patent portfolios (Harhoff and Scherer, 1999). Consequently, it is useful to the examine effects of Bayh-Dole on the upper tail of the distribution of citations. Second, one argument proposed by HJT for their "quality decline" results is an increase in patenting of "marginal" inventions by universities--in response to the reduction in the costs of patenting effected by Bayh-Dole, U.S. universities reduced the threshold level of invention quality above

5. Truncation Bias

Truncation of citation information could affect measures of patent citation intensity in at least three ways. First, consider the case where university patents are ultimately more heavily cited than those in a random sample of corporate patents, and the distribution of proportion of citations received over time is identical for both sets of patents. A "quality decline" based on patent citations could appear under these circumstances simply because the cumulative difference between the number of citations to university and control patents takes time to appear—the real divergence in citation-intensity for more recently issued patents falls outside of the time span covered by the analysis. A second possibility arises if university patents are more heavily cited than the control patents, but on average, citations to university patents occur later after issue than is true for the patents in the control sample. Here too, one could observe a decline in the number of patents citing university patents relative to the number citing the controls towards the end of the 1975-1988 period, simply because of the differences in the time profile of these citations between the two groups of patents. A third possibility (a variant of the second) is that the fraction of citations to university patents observed early decreases relative to the fraction of control sample patents observed early, i.e. the

which they would file for patent protection. Despite the fact that the "quality decline" disappears with the addition of additional years of citation data, the threshold hypothesis may still be valid. Accordingly, we estimated quantile regressions for q=.1, .25, .5,.75, and .9. The results (not reported) show no evidence of a post-Bayh-Dole decrease in the quality of university patents (relative to the controls) in any of the quartiles. Indeed, there is some evidence in these results of an *increase* in quality in the upper quantiles of our university patent sample. These results should be interpreted with caution, however, since in order to get the quantile regression algorithm to converge, however, we needed to aggregate the 394 patent class dummies into five broad technological categories. The quantile regressions thus do not include complete controls for differences in citation intensity across patent classes.

intertemporal distribution of citations for one or both of these groups of patents changes during the period covered by the data.

HJT examine whether their results are robust to the first two of these possible truncation dynamics. First, they re-estimate regression (1) in log-linear form, regressing the log of citations on the same explanatory variables listed above in the following specification:[12]

$$\textit{Log Citations} = \sum_t \left[\alpha_t APP_t + \beta_t (APP_t * UNIV)\right] + \sum_c \lambda_c CLASS_c + \varepsilon$$

In this case, the coefficient on the university/year interaction term estimates the proportionate difference between citations to university patents and citations to the controls*, ceteris paribus*. If an equal proportion of citations is truncated for the university and controls at any point in time, in each year the proportionate difference estimated here would equal the proportionate difference were citations observed at the end of time. The authors report that this proportionate difference becomes insignificant after the mid-1980s, suggesting that if the citation frequency distributions have the same shape, then truncation is not driving the observed quality decline.

It is more difficult to check for the effects of truncation if the university and control-patent samples display different intertemporal distributions of citations, e.g. if citations to university patents come later on average than those to other patents. To examine this possibility, HJT created a dataset with one observation for each year from the application date of each of the patents in their datasets. The authors regressed the

[12] This requires dropping all patents with zero citations.

number of citations received by a patent in a given year on patent class dummies, application year dummies, the university/application year interaction terms, and dummies for the lag between the application years of the citing and cited patent, with an interaction term allowing for different lag coefficients for the university and control sample patents. Specifically, they estimated an equation of the form:

$$Citations, Year\ j. = \sum_t [\alpha_t APP_t + \beta_t (APP_t * UNIV)] + \sum_c \lambda_c CLASS_c + \sum_j [\gamma_j LAG_j + \phi_j (LAG_j * UNIV)] + \varepsilon$$

where $LAG_j=1$ if the lag between the application year of patent i and the potential citing year τ (that is, τ-t) is j years, j=0, 1, 2 … 15.

In this formulation, the coefficients on the university/application year interaction terms (β) give the difference in the number of citations between the university and control sample patents in an application year, controlling for the predicted number citations, based on the average citation lag structure for each sample. HJT report that this exercise reveals quality dynamics that are quite similar to those estimated from their other regressions, with university quality advantage diminishing after Bayh-Dole, and disappearing after the mid-1980s.

The approach above controls for differences in average citation lags for university and control patents, based on those observed in the data. But it does not "control" for truncation if the lag structure itself is changing over time, i.e. if citations to university patents come increasingly later over time, or citations to control patents increasingly earlier. This possibility provides an explanation for the differences between their results and those reported above.

Evidence that citations to university patents are arriving increasingly later (relative to controls) in more recent application years is provided in Figures 2-4, which display plots of the cumulative distribution of citations received in eleven year “windows” following the date of application, for the university and control patents applied for respectively during 1977-1980, 1981-1984, and 1985-1988 . The value of the cumulative distribution function at t=x years after the patent application ($0 \leq x \leq 11$) is the proportion of total citations during the 11-year period that are observed after x years. The eleven-year window is employed to enable us to examine the same "span" of citations for all patents and identify shifts in the intertemporal distribution of citations within that span.

The data employed by HJT allowed only four years of citation for the “newest” patents in their analysis (those applied for in 1988 by universities or corporations). Figure 2 shows that the 1977-1980 cohort of patents in our control sample accumulate 16.3% of their “total” citations (i.e., total citations accumulated within 11 years) during the first four years following their application date. The identical cohort of university patents, however, accumulate12.6% of their “total” citations during the first four years following their application data. Figure 3 shows that in the 1981-1984 period, these respective proportions are 15.1% and 9.4% respectively, and in Figure 4, covering 1985-88, these proportions are 18.7% and 12.0% respectively. Visual inspection of these figures suggests that the vertical distance between them is increasing not just at t=4 years, but also more generally for more recent cohorts of patents. That is, citations to university patents are arriving increasingly later (relative to those in the control sample) in more recent cohorts.

FIGURES 2-4 HERE

The data in Figures 2-4 do not control for technological field effects, and we accordingly estimated two models to analyze changes in the distribution of citation lags for university and control-sample patents. First, we calculated the average lag to forward citations for each cited patent, using all citations within 11 years of the application date, and estimated the following equation[13]:

$$Average\ Citation\ Lag = \sum_t \left[\alpha_t APP_t + \beta_t (APP_t * UNIV)\right] + \sum_c \lambda_c CLASS_c + \varepsilon$$

where the dependent variable is the average lag to all citations to patent i, within the 11 year window. Second, we performed a similar analysis using median citation lags:

$$Median\ Citation\ Lag = \sum_t \left[\alpha_t APP_t + \beta_t (APP_t * UNIV)\right] + \sum_c \lambda_c CLASS_c + \varepsilon$$

The least squares estimates of the differences in average and median citation lags (the β coefficients) are shown in Table 2. In 8 of the 9 years before 1984, there are no significant differences between the university and controls in the (average or median) lengths of lags to "forward" citations. In each after 1984, both the mean and median forward lags are significantly longer for the university patents.

TABLE 2 HERE

These results confirm that citations to university patents are indeed arriving increasingly later relative to citations to the control-sample patents. Thus, there *is* a change in the characteristics of post-1980 university patents. However, this change does not reflect a decline in the average number of citations to university patents (relative to

[13] Importantly, while the lags in Figures 2-4 were calculated across each citing-cited pair, the "average" and "median" lags used in these regressions are calculated at the level of the cited patent.

citations to those in the control sample), but rather a change in the intertemporal distribution of citations.

Interpretation of this shift is difficult, since the meaning of patent citations generally is open to different interpretations. Lanjouw and Schankerman (1999) find that early citations (those observed within 5 years of application date) are highly correlated with their measures of the importance and economic value of a patent that are taken directly from the application (i.e. number of claims, number of backward citations, number of countries in which patent protection is sought). Later citations are less strongly correlated with these "time zero" quality measures. Lanjouw and Schankerman suggest that this finding may be interpreted in two ways. The first is that later citations may represent "citing the classics" and have little to do with economic value of inventions. The second is that later citations represent the ultimate market success of an invention, which takes time to manifest itself and is not revealed in the early quality measures. The latter interpretation is supported by the results of Hall, Jaffe, and Trajtenberg (2000), who find that unanticipated future citations (those not predictable by early citations) are the most important citations in predicting the market value of the patentholding firm.

If later citations are essentially noise, then the fact that university patents are cited later is consistent with a decline in the "quality" of these patents after Bayh-Dole. But if these late citations are more revealing of the ultimate importance of inventions, their exclusion from empirical analyses of this issue may have serious consequences. An assessment of this finding thus requires more reflection and empirical work on how and why citations might be related to the private value of patents.

But as we noted above, other scholars have interpreted patent citations as measures of "spillovers" from university research, using citation counts to measure the magnitude of such spillovers. In this interpretation, the lengthening citation lags associated with university patents means that the patent-embodied knowledge flows from universities are incorporated into future patents more slowly in the wake of Bayh-Dole. If "science-based" patents take more time to be cited in subsequent patents, then this longer lag in the citation of university patents might reflect some tendency for U.S. universities to patent "science" rather than technology in the aftermath of Bayh-Dole. Any such tendency may be undesirable from a social welfare perspective, if preservation of a scientific "commons" is important for scientific progress (see Eisenberg and Heller 1998, Eisenberg and Nelson 2001). This warrants further investigation.

6. Conclusions

The results of this analysis suggest that there has been no decline in the "quality" of university patents after Bayh-Dole, if we measure quality by the total number of citations to patents. Rather, the quality decline observed by HJT was a result of truncation of the citation data. The fact that other scholars (Mowery and Ziedonis 2001; Mowery, Sampat and Ziedonis 2002) do not observe a quality decline similar to that reported by HJT appears to reflect their use of patent-citation data covering a longer period of time than that available to HJT.[14]

[14] Mowery et al. (2002) also do not observe a quality decline, though they use only citations observed within 5 years of the issue date of patents. Given a median application-grant lag of 2 years, this yields approximately seven years of citation data, which may be enough. That analysis uses a control sample that differs from that used by HJT and excludes self-citations, which may also be responsible for the differences in results.

The sensitivity of these results to truncation and the difficulties in controlling for truncation in the face of shifts in citation lags also suggest that the use of patent-citation data to evaluate relatively recent policy changes is problematic and any results must be interpreted with great care and caution.

The "meaning" of patent citations is itself subject to multiple interpretations, and these different interpretations make it difficult to draw strong conclusions from the changes that we have observed in the intensity and timing of citations to university patents.[15]

Finally, the most important questions relating to the effects of the Bayh-Dole Act on U.S. university research and technology transfer cannot be answered solely with patent citation data. Has Bayh-Dole affected the incentives of academic researchers? Does the introduction of commercial incentives into university research threaten the norms of academe? Have universities begun to patent and license "science" that they previously disseminated freely? Are patents on academic research outputs necessary to facilitate technology transfer? These are important questions that cannot be addressed with patent and patent citation data alone. In addition to paving the ground for subsequent work using patent citation data, another "spillover" from the paper by Henderson and colleagues is that it has focused attention and stimulated research on the broader issues relating to the effects of Bayh-Dole, an important contribution indeed.

[15] Several scholars (Lanjouw and Schankerman 1999; Hall et al. 2000; Jaffe et al. 2000; Sampat and Ziedonis 2000) have begun to explore the economic significance of different types of patent citations, research that should prove illuminating.

6. References

Agrawal, Ajay and Rebecca Henderson. 2002. "Putting Patents in Context: Exploring Knowledge Transfer from MIT" *Management Science* 48(1): 44-60.

Cohen,Wesley, Richard Nelson, and John Walsh. 2002. "Links and Impacts: The Influence of Public Research on Industrial R&D" *Management Science* 48(1): 1-23.

Eisenberg, Rebecca, and Michael Heller. 1998. "Can Patents Deter Innovation? The Anticommons in Biomedical Research." *Science* 280:698-701.

Foray, Dominique, and Ali Kazancigil. 1999. "Science, Economics and Democracy." in *UNESCO World Conference on Science*.

Hall, Bronwyn H., Adam Jaffe, and Manuel Trajtenberg. 2000. "Market Value and Patent Citations: A First Look." *NBER Working Paper W7741* .

Henderson, Rebecca, Adam Jaffe, and Manuel Trajtenberg. 1995. "Numbers Up, Quality Down? Trends In University Patenting 1965-1992." *Stanford CEPR Draft Manuscript* .

Henderson, Rebecca, Adam Jaffe, and Manuel Trajtenberg. 1998. "University Patenting Amid Changing Incentives for Commercialization." in *Creation and Transfer of Knowledge*, edited by G. Barba Navaretti, P. Dasgupta, K.G. Mäler, and D. Siniscalco. New York: Springer.

Henderson, Rebecca, Adam-B Jaffe, and Manuel Trajtenberg. 1998. "Universities as a Source of Commercial Technology: A Detailed Analysis of University Patenting, 1965-1988." *Review of Economics and Statistics* 80:119-27.

Jensen, R. and M. Thursby, 2001. "Proofs and Prototypes for Sale: The Licensing of University Inventions," *American Economic Review* 91, 240-259.

Lanjouw, Jean, and Mark Schankerman. 1999. "The Quality of Ideas: Measuring Innovation With Multiple Indicators." *NBER Working Paper 7345* .

Mowery, David, Richard Nelson, Bhaven N. Sampat, and Arvids Ziedonis. 2001. "The Growth of Patenting and Licensing By U.S. Universities: An Assessment of the Effects of the Bayh-Dole Act of 1980." *Research Policy* 30:99-119.

Mowery, David, and Bhaven Sampat. 2001a. "Patenting and Licensing University Inventions: Lessons from the History of the Research Corporation." *Industrial and Corporate Change* 10:317-355.

Mowery, David, and Bhaven Sampat. 2001b. "University Patents, Patent Policies, and Patent Policy Debates: 1925-1980." *Industrial and Corporate Change* 10:781-814.

Mowery, David, Bhaven N. Sampat, and Arvids Ziedonis. 2002. “Learning to Patent: Experience and the Quality of University Patents 1980-1994.” *Management Science*.

Mowery, David, and Arvids Ziedonis. 2000. “Numbers, Quality and Entry: How Has the Bayh-Dole Act Affected US University Patenting and Licensing?” *Draft Manuscript*.

National Science Foundation. 1996. *Science and Engineering Indicators*. Washington D.C.: USGPO.

Sampat, Bhaven, and Arvids Ziedonis. 2000. “Patent Citations and the Value of University Patent Rights.” *Draft Manuscript* .

Trajtenberg, Manuel. 1990. “A Penny for Your Quotes: Patent Citations and the Value of Innovations.” *Rand Journal of Economics* 21:172-87.

Rosenberg, Nathan, and Richard Nelson. 1994. “American Universities and Technical Advance in Industry.” *Research Policy* 23:323-348.

Table One: University-Control Sample Mean Difference in Citations, Over Time

Application Year	Dependent Variable: Citations to 1992 (from HJT 1998, Table 4)		Dependent Variable: Citations to 1992		Dependent Variable: Citations to 1999	
1975	2.54	***	2.48	***	3.21	***
	(0.35)		(0.35)		(0.77)	
1976	1.82	***	1.67	***	2.20	***
	(0.34)		(0.35)		(0.76)	
1977	1.31	***	1.42	***	3.02	***
	(0.34)		(0.35)		(0.77)	
1978	2.04	***	2.33	***	4.50	***
	(0.34)		(0.34)		(0.75)	
1979	1.13	**	1.34	***	2.41	***
	(0.31)		(0.32)		(0.71)	
1980	1.91	***	1.82	***	4.01	***
	(0.31)		(0.31)		(0.67)	
1981	1.68	***	1.93	***	3.98	***
	(0.31)		(0.32)		(0.69)	
1982	0.96	***	1.49	***	3.87	***
	(0.31)		(0.31)		(0.68)	
1983	0.97	***	1.14	***	2.89	***
	(0.3)		(0.31)		(0.67)	
1984	0.47	*	0.36		2.22	***
	(0.28)		(0.29)		(0.64)	
1985	0.4		0.68	**	3.41	***
	(0.28)		(0.28)		(0.62)	
1986	0.06		0.06		2.98	***
	(0.27)		(0.27)		(0.6)	
1987	-0.07		-0.11		1.90	***
	(0.25)		(0.25)		(0.56)	
1988	-0.08		-0.16		1.89	***
	(0.24)		(0.24)		(0.53)	
Year Dummies	Sig.		Sig.		Sig.	
Patent Class Dummies	Sig.		Sig.		Sig.	
N	NA		17996		17996	

*Note: *** denotes significance at the 1% level, ** at the 5% level, and * at the 10% level*

Standard Errors in Parenthesis

Table Two: University-Control Sample Mean Difference in Average and Median Lags to Forward Citations (Within 11 Year Windows)

Application Year	**Dependent Variable: Average Time to Forward Citations (within 11 year window)**		**Dependent Variable: Median Time to Forward Citations (within 11 year window)**	
1975	-0.093		-0.127	
	(0.13)		(0.15)	
1976	-0.038		-0.067	
	(0.13)		(0.15)	
1977	0.336	**	0.295	**
	(0.13)		(0.15)	
1978	0.123		0.187	
	(0.12)		(0.14)	
1979	-0.113		-0.135	
	(0.12)		(0.13)	
1980	0.086		0.073	
	(0.11)		(0.13)	
1981	0.152		0.183	
	(0.11)		(0.13)	
1982	0.177		0.156	
	(0.11)		(0.13)	
1983	0.068		0.073	
	(0.11)		(0.12)	
1984	0.221	**	0.239	**
	(0.1)		(0.12)	
1985	0.344	***	0.337	***
	(0.1)		(0.11)	
1986	0.502	***	0.501	***
	(0.1)		(0.11)	
1987	0.419	***	0.481	***
	(0.09)		(0.1)	
1988	0.366	***	0.395	***
	(0.09)		(0.1)	
Year Dummies	Sig.		Sig.	
Patent Class Dummies	Sig.		Sig.	
N	15299		15299	

*Note: *** denotes significance at the 1% level, ** at the 5% level, and * at the 10% level*

Standard Errors in Parenthesis

Figure One: Coefficients on University*Application Year Interaction Term, By Application Year

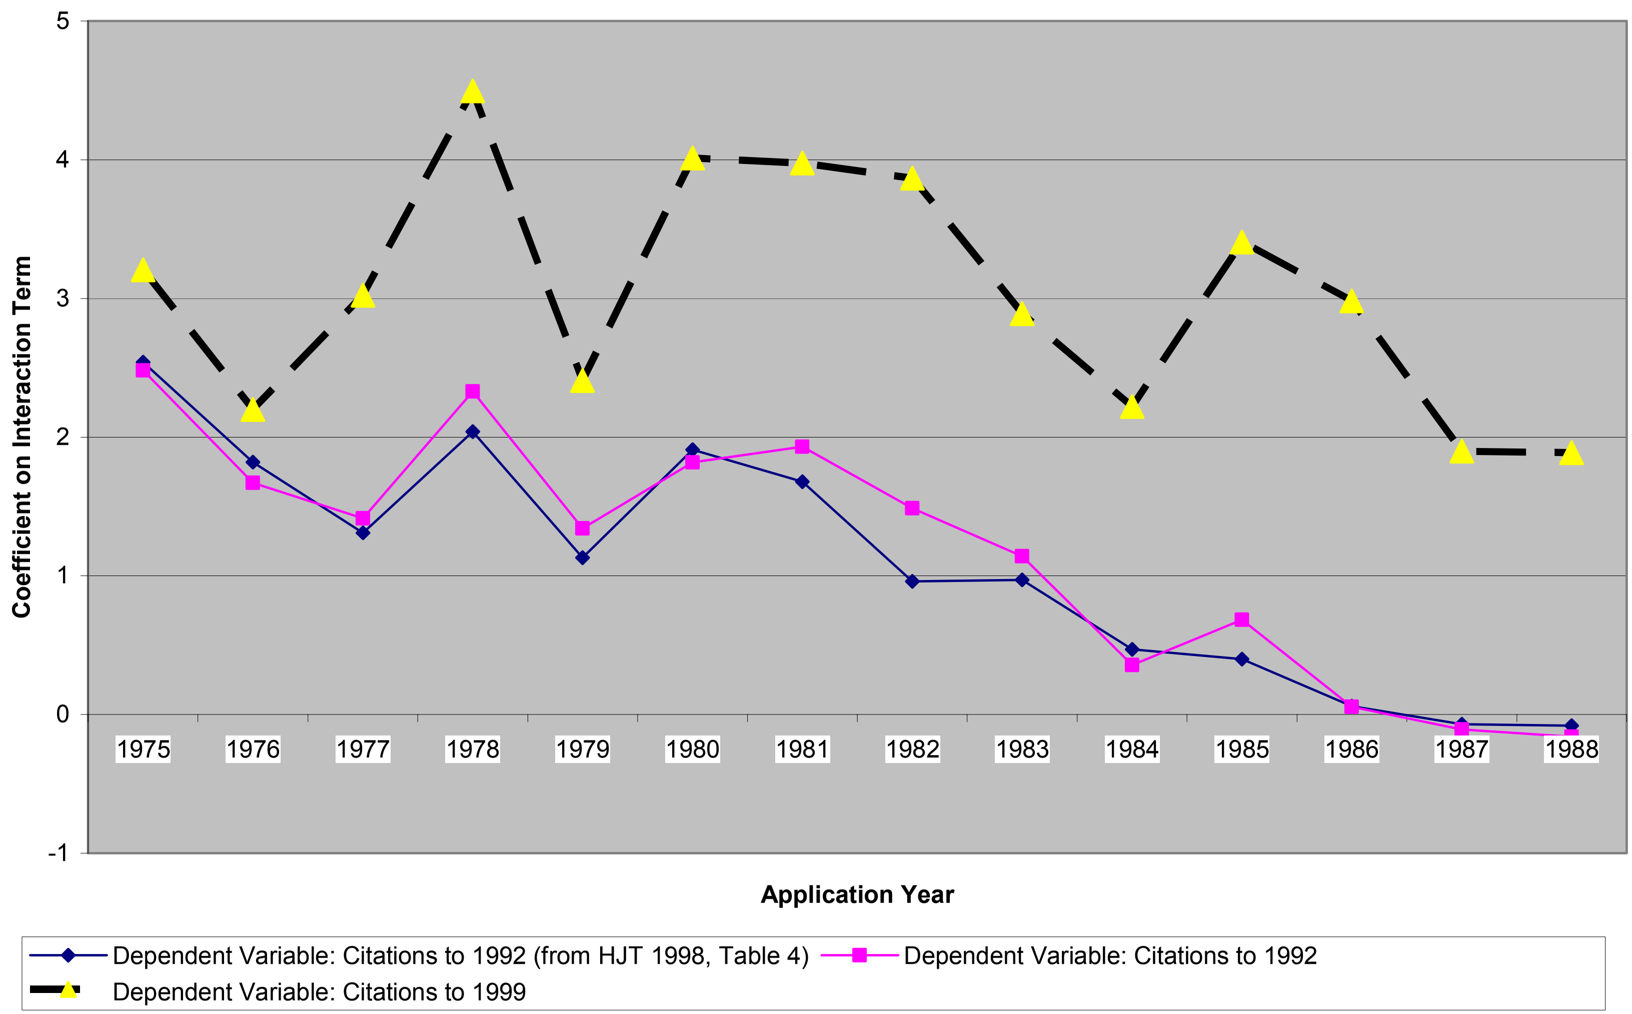

Figure Two: Cumulative Distribution of Citation Lags, University and Control Patents, 1977-1980

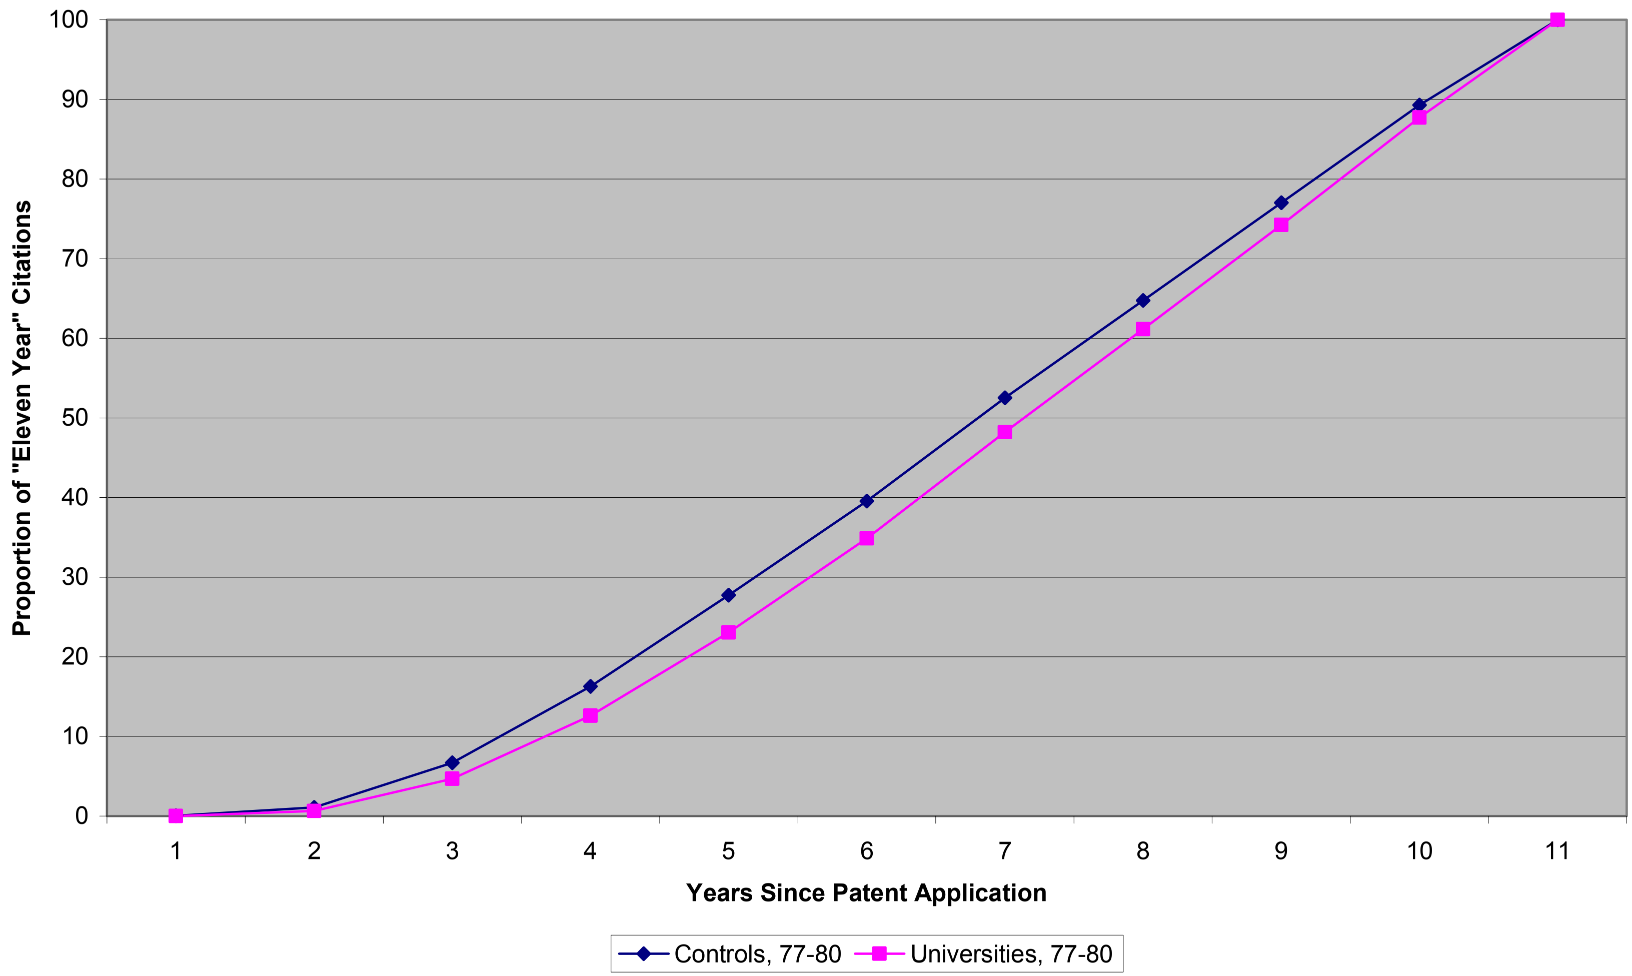

Figure Three: Cumulative Distribution of Citation Lags, University and Control Patents, 1981-1984

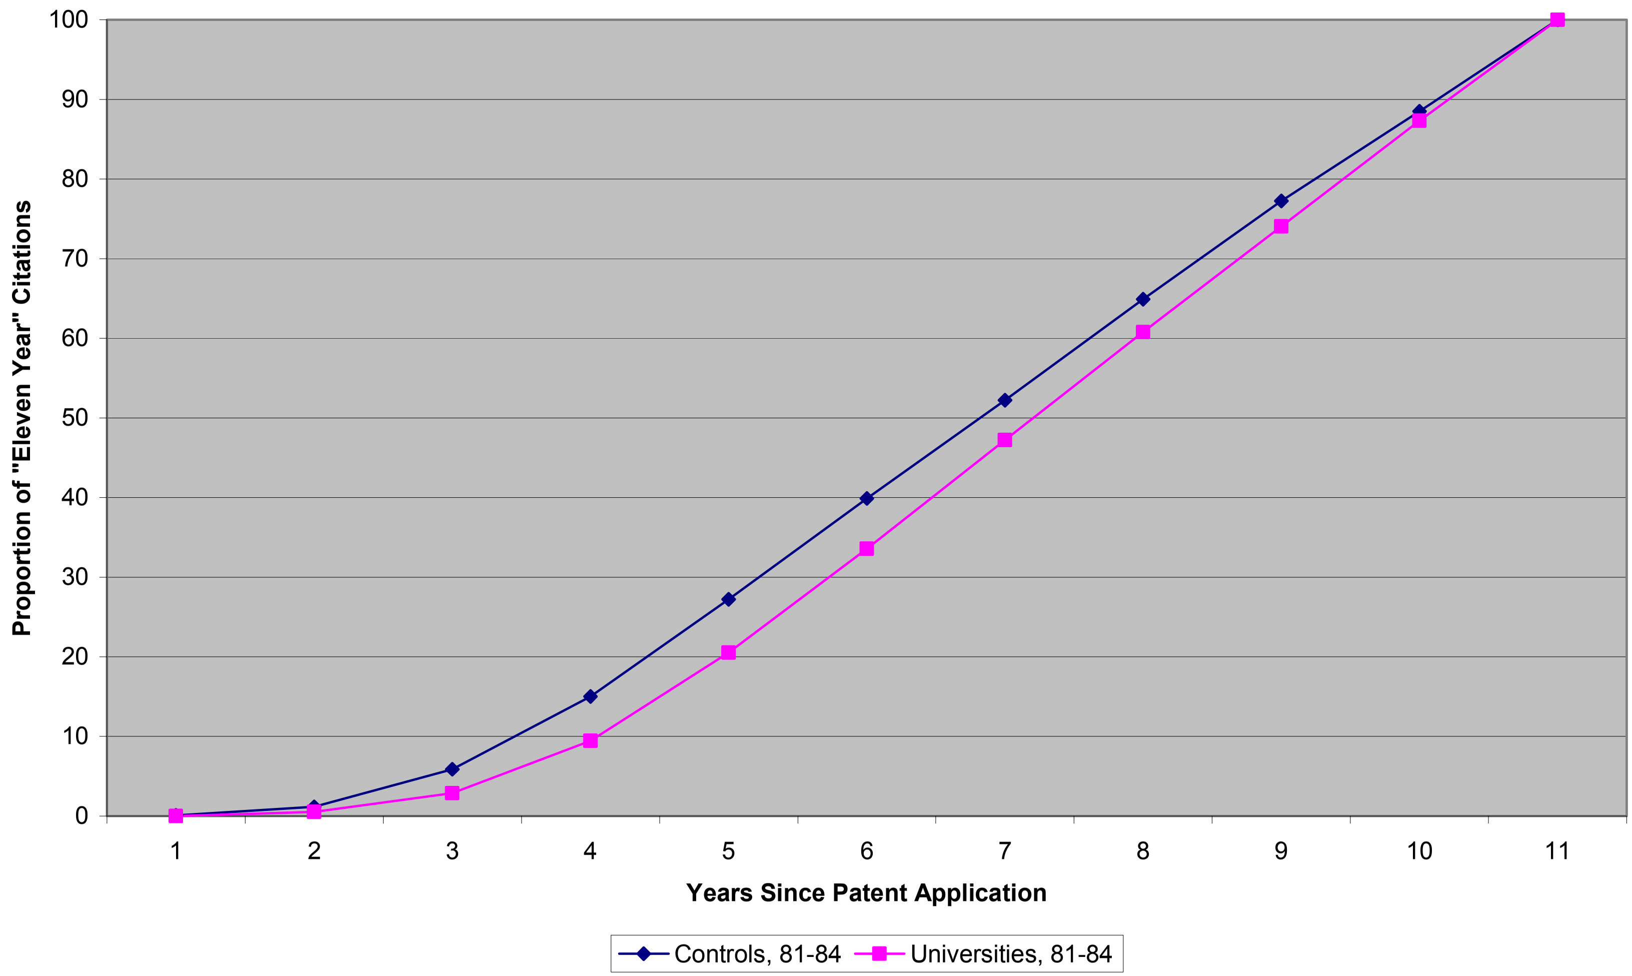

Figure Four: Cumulative Distribution of Citation Lags, University and Control Patents, 1985-1988