
Putting Patents in Context: Exploring Knowledge Transfer from MIT

Author(s): Ajay Agrawal and Rebecca Henderson

Source: *Management Science*, Vol. 48, No. 1, Special Issue on University Entrepreneurship and Technology Transfer (Jan., 2002), pp. 44-60

Published by: INFORMS

Stable URL: <http://www.jstor.org/stable/822683>

Accessed: 04/04/2010 19:11

Your use of the JSTOR archive indicates your acceptance of JSTOR's Terms and Conditions of Use, available at <http://www.jstor.org/page/info/about/policies/terms.jsp>. JSTOR's Terms and Conditions of Use provides, in part, that unless you have obtained prior permission, you may not download an entire issue of a journal or multiple copies of articles, and you may use content in the JSTOR archive only for your personal, non-commercial use.

Please contact the publisher regarding any further use of this work. Publisher contact information may be obtained at <http://www.jstor.org/action/showPublisher?publisherCode=informs>.

Each copy of any part of a JSTOR transmission must contain the same copyright notice that appears on the screen or printed page of such transmission.

JSTOR is a not-for-profit service that helps scholars, researchers, and students discover, use, and build upon a wide range of content in a trusted digital archive. We use information technology and tools to increase productivity and facilitate new forms of scholarship. For more information about JSTOR, please contact support@jstor.org.



INFORMS is collaborating with JSTOR to digitize, preserve and extend access to *Management Science*.

Putting Patents in Context: Exploring Knowledge Transfer from MIT

Ajay Agrawal • Rebecca Henderson

Queen's School of Business, Queen's University, Kingston, Ontario, Canada K7L 3N6
Sloan School of Management, Massachusetts Institute of Technology, Cambridge, Massachusetts 02142
aagrawal@business.queensu.ca • rhenders@mit.edu

In this paper we explore the degree to which patents are representative of the magnitude, direction, and impact of the knowledge spilling out of the university by focusing on the Massachusetts Institute of Technology (MIT), and in particular, on the Departments of Mechanical and Electrical Engineering. Drawing on both qualitative and quantitative data, we show that patenting is a minority activity: a majority of the faculty in our sample never patent, and publication rates far outstrip patenting rates. Most faculty members estimate that patents account for less than 10% of the knowledge that transfers from their labs. Our results also suggest that in two important ways patenting is not representative of the patterns of knowledge generation and transfer from MIT: patent volume does not predict publication volume, and those firms that cite MIT papers are in general not the same firms as those that cite MIT patents. However, patent volume is positively correlated with paper citations, suggesting that patent counts may be reasonable measures of research impact. We close by speculating on the implications of our results for the difficult but important question of whether, in this setting, patenting acts as a substitute or a complement to the process of fundamental research.

(Patents; University Science; Knowledge Transfer; Technology Transfer)

1. Introduction

While there is a widespread belief that publicly funded research conducted at universities has a significant impact on the rate of economic growth, estimating the magnitude and describing the nature of this impact remains extremely difficult. Recent quantitative work in the area has focused particularly on patents as a measure of university "output" (Jaffe 1989, Henderson et al. 1998), on licenses and on the new firms created by licenses (Gregorio and Shane 2000, Jensen and Thursby 1998, Thursby and Thursby 2000), or on patents and licensing considered simultaneously (Mowery et al. 1998). As a logical extension, patent citation data has been widely used in a variety of studies concerning university innovation (Jaffe et al. 1993, 1998, Jaffe and Trajtenberg 1996).

Patent and license data has become particularly important in this context for three reasons. First, the patenting process requires that inventor names, dates, assignee institutions, locations, and detailed descriptions of invention claims be recorded. Such systematically recorded innovation-related details are very rare outside of patent records. Second, innovations that are patented are expected, by definition, to be commercially useful.¹ Third, patenting data has recently become widely available in machine-readable form, and aggressive research programs, such as the one centered at Columbia University (Mowery et al. 2000), coupled with the generous efforts of

¹ This is, of course, not to say that they are commercially successful. In fact, only a very small percentage of patented inventions result in financial success.

AUTM (The Association of University Technology Managers), has recently made much university licensing data available. A focus on patents and licensing as an important mechanism of knowledge transfer from universities to the private sector is thus understandable. However, it is almost certainly incomplete.

Public support of university research is commonly justified on the grounds that the private sector is likely to systematically underfund "basic" or fundamental research because the results are, in general, difficult to appropriate. Thus university research is largely funded on precisely the premise that mechanisms such as patents will be particularly ill-suited to capturing the returns. Professors transfer knowledge through mentoring their students' research, through giving conference presentations, and, most notably, through the free publication of ideas in refereed scientific publications.² If patents characterize only a small proportion of all the work being conducted within the academy, and, even more importantly, if the research that is patented is not representative either of the work being done within the university or of the mode with which it is generally transferred to the private sector, then too great a focus on patenting may seriously misrepresent the nature of the impact of the university on the private sector.

In this paper, we begin to explore this issue by focusing in depth on two departments at MIT, one of the nation's preeminent research institutions. Drawing on in-depth qualitative interviews with the faculty in the Departments of Mechanical Engineering (hereafter ME) and Electrical Engineering and Computer Science (hereafter EECS), coupled with comprehensive quantitative information about each faculty member's patenting and publication behavior, we explore the degree to which patenting is representative of the work being conducted at MIT, of the ways in which it is transferred to the private sector, and of its ultimate impact.

² It is important to note that we are referring to the creation and transfer of *new* knowledge. This refers to knowledge that is generated from laboratory experiments or theory development and is of the type that could be patented or published in science- or engineering-oriented journals. In other words, this does not include common knowledge contained in textbooks and taught to students by professors in regular classes.

Our study builds on work by Zucker and Darby and their collaborators (Zucker et al. 1998a, 1998b), who have demonstrated the importance of geographic proximity, research collaborations, and personal relationships in the transfer of knowledge, on the work of Cockburn and Henderson (1998), who focus on coauthorships, and on the work of Branstetter (2000), who focuses on citations to academic papers, as opposed to patents, as indicators of knowledge transfer. However, in contrast to these studies, which in general have focused on a single transfer mechanism in depth, here we attempt to place patenting "in context," exploring its importance relative to other mechanisms of knowledge transfer, particularly journal publications, and the degree to which patenting is representative of knowledge transferred through other channels. This paper is thus most similar in spirit to the work of Cohen et al. (1998). Cohen and his coauthors used extensive interview data to estimate the relative importance of patenting as a mechanism for knowledge transfer from the university. However, whereas Cohen et al. asked their questions of the U.S. manufacturing industry, or the "demand" side of the equation, we complement their work by focusing our inquiries on the university, the "supply" side of the equation, and by supplementing our qualitative work with comprehensive quantitative data on patents, papers, and their citations.

Our results suggest that a focus on patenting as a measure of the impact of university research must be carefully qualified by the recognition that patenting may play a relatively small role in the transfer of knowledge out of the university. As one might expect, for the faculty in our target departments, publishing academic papers is a far more important activity than patenting. In fact, only a small fraction of the faculty patent at all. On average, only about 10–20% of the faculty patent in any given year, and nearly half of the faculty in our sample never filed a patent during the 15-year period under investigation. In contrast, an average of 60% of the faculty publish in any given year and less than 3% never publish over the same period. Indeed, even amongst those faculty that do patent, our informants estimated patents were responsible for as little as 7% of the knowledge that was transferred from their labs to industry, a number

very consistent with the Cohen et al. finding that only about 11% of the information obtained from university research was transferred through patents.

Our analysis also suggests that the channel-of-knowledge flow associated with patents may be quite different from those associated with papers. Branstetter has shown that for the University of California, citations to academic papers far exceed those to academic patents (Branstetter, 2000). We show that the set of firms that most frequently collaborates with MIT faculty on patented research is very different from the set of firms that most frequently collaborates on published research. Moreover, the set of firms that most frequently cites MIT patents is very different from the set that cites MIT papers. Thus, while the patent-related channel-of-information flow out of the university is important, it is by no means the only channel, and it may not be representative of the others.

We then tackle the difficult question of whether patenting activity is a good predictor of publishing behavior. We show that patenting activity is not a good predictor of publishing volumes, but that there is some evidence that those professors who patent more write papers that are more highly cited, and thus that patenting volume may be correlated with research impact.

We close this paper with a brief description of the degree to which our results speak to the related question of whether patenting is a substitute or complement for more "basic" research. Some observers have voiced the fear that as researchers focus more on patenting as a primary means of knowledge transfer, the core goals and values of the university will be compromised. (See, for example, Cohen et al. 1998 and the references therein.) It is difficult to test this idea empirically, but our preliminary results are consistent with the hypothesis that, at least at MIT, patenting is not substituting for more fundamental research activity for the vast majority of the faculty.

We believe that these results are important. As universities defend their public role and governments look to maximize their return on investments in public science, it is important to build as clear a picture of the manner in which universities impact the economy as possible. Our results suggest that a focus on patenting or licensing statistics, in isolation, may

significantly misrepresent the nature of the universities' impact on the economy and that any comprehensive study of the issue must include a focus on the other channels through which university knowledge is transferred to private firms.

2. Data and Methods

This paper draws upon both qualitative and quantitative data. Since this is an exploratory foray, we focus on a single university and on two departments, rather than attempt a broad survey. The heart of the study is an in-depth, quantitative, and qualitative study of professors who are currently on the faculty at MIT in the Departments of Mechanical Engineering (ME) and Electrical Engineering and Computer Science (EECS).

MIT was chosen as the focal university both for reasons of convenience (it is the home institution of one of the authors) and because it is one of the premier research institutions in the United States. In 1998, MIT claimed almost 4% of all the patents given to American universities and received over 1.5% of all federal funding for science and engineering at universities and colleges in fiscal year 1999.³ Moreover, it has historically been firmly orientated towards a goal of having an immediate impact on the world around it. The MIT motto is "hands and mind": MIT was founded as a land grant college, and its leaders have always been concerned about generating value for the economy in which it is embedded.

We chose to focus on the departments of ME and EECS because, after biology, they are the departments that have generated the largest number of patents, and because biology departments have already been quite extensively studied. (See, for example, work by Zucker et al. 1998a, 1998b and more recent work by the same authors, and work by Blumenthal and his collaborators, including Blumenthal et al. 1996.) They are also two of the largest and most vibrant departments at the university, with almost 18% of the Institute's faculty. The data for this study is based on the population of professors who were on the faculty in September 2000 and who generated at least one paper

³ NSF report: Federal Science and Engineering Support to Universities, Colleges, and Nonprofit Institutions: Fiscal Year 1999.

or patent during the period 1983–1997. This includes 154 EECS professors and 82 ME professors for a total of 236. Professors enter the population when they publish their first paper or file their first patent during the period under investigation; not all the professors in the sample were active at MIT throughout the entire period. We chose to focus on the period 1983–1997 because publication data is available in electronic format from 1983 onward and patent data necessitates accommodating two- to three-year lags since we use application dates rather than issue dates (hence data stops at 1997).

For this 15-year period, we collected comprehensive data about each paper and patent generated by every faculty member in the sample as well as every paper or patent that cited these patents or papers. Our final data set includes information about 640 patents and 5,132 papers assigned to the sample faculty, plus data about the 6,074 patents that cite these patents, data about the 727 patents that cite these papers, and data about the 49,975 papers that cite these papers. Paper data was collected from the Institute of Scientific Information’s Science Citation Index,⁴ and patent data was collected from the US Patent and Trademark Office database.⁵

We supplemented this quantitative data with qualitative interview data. We requested a face-to-face interview with every faculty member in either department who had ever been an inventor on a patented technology that was licensed from MIT’s Technology Licensing Office (TLO). This group was selected because it was assumed that they would be the most familiar with the patenting and licensing process due to their direct experience; in 1999, this was 39% of the faculty in both departments. Of those faculty members, 74% agreed to meet with us, resulting in an interview sample size of 68.

3. Results

3.1. Sample Characteristics

Table 1 presents basic descriptive statistics about the faculty members who agreed to be interviewed, as

Table 1 Descriptive Statistics for Professors Interviewed for Qualitative Research Compared to Those for Total Sample Population and for Professors from ME Compared to EECS

	Total Population	Interviewed Faculty	Mechanical Engineering	Electrical Engineering and Computer Science
<i>N</i>	236	68	82	154
Publications				
Mean	21.7	24.1	22.2	21.5
Standard Deviation	31.7	31.0	31.4	32.0
Median	11	13.5	12.5	10
Maximum	223	169	223	200
Minimum	0	0	0	0
Patents				
Mean	2.7	6.1	2.5	2.8
Standard Deviation	5.0	7.2	4.5	5.3
Median	1	4	1	1
Maximum	36	36	28	36
Minimum	0	1	0	0
Years at MIT				
Mean	11.6	12.3	11.1	11.9
Standard Deviation	4.6	4.0	4.6	4.6
Median	14	14.5	13.5	15
Maximum	15	15	15	15
Minimum	1	1	1	1

contrasted with the entire population of which they are a part. The professors that we interviewed publish slightly more than the population mean, patent a great deal more, and were active for slightly longer during the period under investigation. Table 1 also contrasts the publication and patenting records of the faculty from the two departments. 35% of the 236 professors studied were from ME. ME professors have slightly higher average rates of publishing, slightly lower rates of patenting, and were active for slightly less time than their colleagues in EECS. However, these differences are statistically insignificant and data for the two departments is aggregated for purposes of the analyses that follow.


3.2. Patenting as One Mechanism Amongst Many

The recent increase in university patenting, especially since the passage of the Bayh-Dole Act in 1980, has been well documented. Figure 1 presents total patents

⁴ <www.webofscience.com>.

⁵ <www.uspto.gov>.

Figure 1 University Patenting Over Time



assigned to several research universities from 1983–1997.⁶ At the aggregate level, these numbers are consistent with a substantial increase in patenting as a mechanism of university knowledge transfer, as much of the existing literature suggests. For example, while there were only 26 patents assigned to MIT in 1983, there were 130 in 1997, a 400% increase.

Figure 2 shows patenting and publishing rates over time for our sample in particular. Two things are immediately apparent. First, as one would expect, publishing is a much more important activity than patenting, at least as measured by count data. While the average faculty member publishes between 1.5 and 2.0 papers a year, they only produce about 0.25 of a patent, or roughly one every four years. Second, while there is some evidence of an increase in patenting rates (faculty in our sample move from filing roughly 0.18 to 0.28 patents per year⁷), publishing rates were also increasing significantly over the

period.⁸ The ratio of patents to publications increased from 0.11 to 0.13 over the period, but it rose from a very low base.⁹

Note that these results raise the issue of possible sample selection bias. Recall that our sample consists only of those professors who were on the faculty at MIT in September 2000. If “stronger” faculty stay while “weaker” faculty leave, the apparent increase in both patenting and publication rates over the period might be merely an artifact of sample construction. In order to explore this issue, we compared our sample against the entire population of EECS and ME professors in terms of both paper and patent output for 9 of the 15 years in our sample.¹⁰ Table 2 compares the average paper-to-patent ratio of the population to our sample for each of the nine years for which we have data. There is no systematic difference between

⁶ The data for this figure and all figures related to patents were generated from the USPTO patent database. Also, patent application dates, rather than issue dates, are used in this graph and throughout the remainder of the paper.


⁷ These are three-year averages and aggregated over ME and EECS (i.e., the average for 1984 is calculated using values from 1983–1985).

⁸ Note that publication and patenting rates in 1983 and 1984 are almost certainly overstated, since faculty members only enter the sample when they publish or patent. Thus, there is by definition a lower fraction of “sleeping” faculty in the early years.

⁹ It is important to note that these results are likely to be inconsistent with those for the life sciences where, in some cases, patenting rates per professor have increased substantially.

¹⁰ Population data was collected for the years 1986, 1987, 1988, 1989, 1991, 1992, 1993, 1994, and 1997. These were the years for which we were able to obtain copies of the MIT catalogue, our source of professors currently on the faculty.

Figure 2 Papers and Patents per Professor



the sample and the population during this period (z -statistic = 0.37). Most importantly, only a small fraction of the faculty patent at all. Figure 3 shows the percentage of faculty members in any given year who publish, patent, or license. Patenting and licensing is essentially a minority activity. On average, only about 10–20% of the faculty patent in any given year and 3–7% license an invention. In contrast, while more than 50% of the sample publish at least one paper in any given year and less than 3% have never published, nearly half of our sample have never patented at all!

Figures 4a and 4b expand on this point by illustrating the distribution of professors in terms of patenting and publishing frequency, respectively. Notice how different these distributions are. Not surprisingly, given the results of Figure 3, the distribution of patenting faculty is heavily skewed to the left. 44%

of the professors have never been an inventor on a patent, less than 15% have been granted more than 5 patents, and less than 6% have been granted more than 10 patents. While the distribution of publishing faculty is also skewed, it has much less mass to the left, and a significant tail at the far right. 14% of the faculty have published more than 35 papers, while 5% have published more than 100. Given these data, it is perhaps not surprising that even those faculty with considerable patent portfolios and/or licensing experience often dismissed the idea that patenting or licensing activity could be used as an important measure of their activities:


I don't think these [patent counts] tell you very much. I don't care too much for patents. I wouldn't have even bothered to patent most of these things that are on your list. Most of those were patented by scien-

Table 2 Comparison of Sample to Population in Terms of Paper:Patent Output Ratios for Select Years

	1986	1987	1988	1989	1991	1992	1993	1994	1997
Paper/patent ratio (sample)	9.83	8.85	7.73	6.48	5.92	9.63	7.50	6.93	7.74
Paper/patent ratio (population)	9.60	8.66	8.01	7.57	6.68	12.23	9.23	6.51	4.93

Note. These mean values were calculated by dividing the total number of papers by the total number of patents for all professors for a given year, not by averaging the ratio for each individual professor. This method was used because many professors have no patents in a given year, thus resulting in a zero denominator for their individual paper:patent ratio. As a result of using this method to construct the mean, standard errors to describe the distribution cannot be calculated and thus are not reported here.

Figure 3 Percentage of Faculty Publishing, Patenting, and Licensing



tists from Japanese firms who were visiting my lab for 6, or 12, or 18 months. That's why I am listed as a coinventor. They file these patents to show their companies that they are doing work here, but I don't think they really intend to do anything with them. I certainly haven't received a penny from any of these patents. (EECS professor, interview, February 17, 1999)

You can't just look at the patents. Many people don't even care about patents. The patent system is too slow for them. Look at this list [shows list of over 30 companies founded from MIT inventions]. Only a very small handful of these have a patent [from MIT]. And most of these [companies on the list] are, or will be, world-changing companies. There's not a very strict patenting culture here, but we do support and encourage world-changing companies [at this lab]. (EECS professor, interview, March 1, 1999)

Another piece of evidence that patenting and licensing may account for a surprisingly small share of the knowledge that reaches the private sector is shown in Figure 5, which summarizes the results of one of the questions that we explored during our qualitative interviews. Each interview sought to understand how and why the faculty member worked with individuals and firms outside of MIT. In each case, we worked step by step through the respondent's CV, asking about each paper and patent, how it came to be written, and what impact it had had on the private sector.


Table 3 presents our results, and contrasts them with those obtained by Cohen et al. (1998).¹¹ First, notice how relatively unimportant faculty members believe patenting and licensing activity to be: the mean response is that it carries only 6.6% of the information transferred out of the university. Second, notice also how surprisingly consistent the two sets of results are. Recall that the Cohen et al. results were obtained by asking US manufacturing firms how important they considered various knowledge transfer channels from the university to their industry to be, while ours were obtained by asking individual professors. Both view patents and licensing as relatively unimportant (6.6% versus 11.6%), while both sources list publications as around 18%, and informal channels ("consulting" and "conversations") as around 31% of the information that is transferred.¹²

There are, of course, potentially significant limitations associated with our interview data since professors may perceive channels that involve direct interaction with firms that use their knowledge to be

¹¹ Note that we have rescaled the Cohen et al. results to be compatible with our own, by normalizing the total scores to sum to 100%.

¹² Again, it is important to note that there are some exceptional manufacturing sectors in the Cohen et al. study, namely pharmaceuticals, which do consider patents a very important knowledge transfer channel.

Figure 4 (a) Publication Frequency ($n = 236$)
 (b) Patent Frequency ($n = 236$)




relatively more important purely because they are more salient. Faculty might have a tendency to overestimate the relative importance of channels such as consulting and informal conversations and to underestimate the importance of more indirect channels such as patents and publications.¹³ On the other hand,

¹³ However, it is important to note that the patents and licenses channel may not be as “indirect” as papers. Agrawal (2000) reports that approximately two-thirds of his sample of patented inventions licensed from MIT involved direct interaction between the inventor and the licensing firm.

recall that we interviewed only those faculty that had patented and licensed at least one invention. One might expect this group of faculty to overestimate the importance of patents relative to the mean.

Taken together, these results are consistent with much prior research, and with the hypothesis that patenting and licensing constitute a relatively small channel for the transfer of knowledge from academia to the private sector. This in itself is reason enough to think carefully about the degree to which the results obtained from analyzing university patenting behav-

Figure 5 Perception of Relative Importance of Knowledge Channels
 (n = 68)



ior is representative of the nature of knowledge flows out of the university.

However, if it is the case that patenting is *broadly representative*—that is, if the firms that collaborate with professors on patents and that cite faculty patents are the same firms that engage in other channels of access—then a focus on patenting is more likely to provide a useful lens through which to view the impact of the university on the economy, even if patenting represents a relatively small proportion of total knowledge transferred. We turn to this question next.

3.3. Different Firms, Different Channels

In this section we present a preliminary analysis of patterns in collaboration and citation for our sample. Specifically, we examine the degree to which firms that collaborate on or cite MIT patents are the same as those that collaborate on or cite MIT papers. Our results suggest that there is significant variation in terms of the particular firms that employ the various channels. Figures 6a and 6b illustrate the degree to which the firms that use MIT patents are also those that use or reference MIT papers.

Figure 6a shows that more firms collaborate on papers than patents and that the fraction of firms

that engage both channels is quite small. Specifically, 58% of the firms that collaborated with this set of professors did so by writing papers together, but did not write patents together. Even amongst the 20 firms with the highest number of paper collaborations, which account for 83% of the total number of paper collaborations, 14 of these firms did not collaborate on any patents. Thus, any examination of patent collaborations would miss a substantial fraction of those firms that engage in the type of collaborative research that results in journal publications.

Figure 6b illustrates that while there are more firms that cite MIT patents than papers, the set of firms that cite papers is not a strict subset of the former. In fact, 24% of all citing firms never cite MIT patents. Specifically, 11% only write papers that cite MIT papers and 13% write patents that only cite MIT papers, not patents. These firms would be missed in analyses that only include firms that cite MIT patents. Even amongst the 20 firms with the highest number of citing papers, which account for 67% of the total number of citing papers, six of these firms did not write any citing patents. Thus, once again, an examination of patent citations misses a substantial fraction of those firms that engage in the type of research that results in the writing of papers that cite MIT papers.

Collectively, the data presented in Figures 6a and 6b support the idea that different firms employ quite different channels for gaining access to MIT-produced knowledge. These results are preliminary, but they are consistent with the hypothesis that a focus on those firms that cite or that collaborate in writing MIT patents may not accurately represent the set of firms that gain knowledge from MIT.

3.4. Patenting Activity as a Predictor of Publishing Behavior

We next focus on the degree to which patenting activity is a good predictor of publishing activity or impact. On the quantitative front, Figure 7a shows a scatter plot of total patents versus total publications, where the unit of observation is the professor. There is no clear relationship between the two, and the plot illustrates the great diversity of behavior across the faculty. Figure (7b) shows a similar scatter plot where the data has been age-adjusted such that the total

Table 3 Distribution of Perceived Importance of Various Modes of Knowledge Transfer—Qualitative Interviews (Agrawal 2000) vs. Questionnaire Results (Cohen et al. 1998)

	Estimate the portion of the influence your research has had on industry activities, including research, development, and production that was transmitted through each of the following channels:	How important are the following sources to Industrial R&D
	% Total (Standard Deviation)	% Total that responded at least "moderately important" (3 on 4-point Likert scale)
	Agrawal Interview 2000	Cohen et al. 1998, normalized to equal 100
Patents and licenses	6.6 (5.6)	11.6
Publications	18.5 (17.3)	17.4
Consulting	25.1 (18.4)	13.7
Conversations	6.3 (6.8)	17.5
Cosupervising	9.4 (10.2)	7.7
Recruiting/hiring	16.8 (12.5)	8.5
Conferences	5.2 (5.6)	14.6
Research collaborations	12.1 (10.8)	9.1

paper and patent output has been divided by the number of years each professor was active during the period under investigation. Still, no clear relationship is evident. If anything, the plot might suggest a negative correlation between patenting and publishing behavior, with a few individuals publishing heavily but not patenting and a few patenting heavily but not publishing. However, as we will show, the relationship is not statistically significant and is in fact positive when patents are compared to paper citations.


Table 4a extends this analysis by showing correlation coefficients for a variety of flow measures of patenting and publishing behavior. While it is reasonable to assume that in most cases a patent and a paper written in the same year will be measured in the same year since we use patent application dates

and paper publication dates, we include three one-year lag variables for each measure to capture any systematic variations from this assumption.¹⁴ While there is clearly correlation across publishing behavior over time (0.76, 0.73, and 0.62 are the correlation coefficients of $paper_{(t)}$ with $paper_{(t-1)}$, $paper_{(t-2)}$, and $paper_{(t-3)}$, respectively), as well as patenting behavior over time (0.49, 0.41, 0.30), there is very little evidence that patenting and publishing behavior are correlated with each other (0.01, -0.003, -0.004, and -0.02 are the correlation coefficients of $paper_{(t)}$ with $patent_{(t)}$, $patent_{(t-1)}$, $patent_{(t-2)}$, and $patent_{(t-3)}$, respectively).

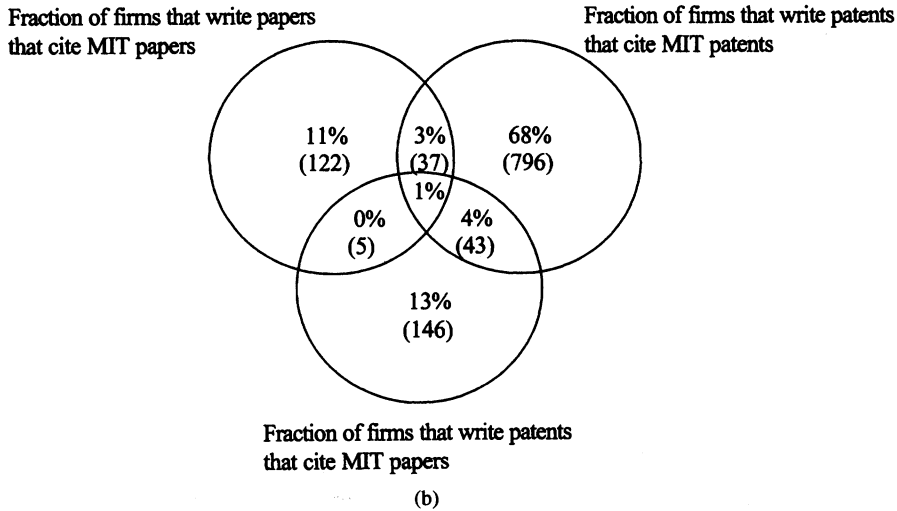
¹⁴ Most science and engineering publications have a publication cycle that is less than one year from the time of receiving the first draft. This is in contrast to many areas in the social sciences where the lag is often two to three years.

Figure 6 (a) Many Firms That Collaborate on Patents Are Not the Same As Those That Collaborate on Papers
 (b) Many Firms That Cite Patents Are Not the Same As Those That Cite Papers

271 firms collaborate on papers or patents with our sample of MIT professors.



1163 firms cite papers or patents authored by our sample of MIT professors.



Note. Percentage may not sum to exactly 100% due to rounding.

Similarly Table 4b presents correlation coefficients for stock measures of patenting and publishing behavior including totals and averages. While the correlation between total patents and papers is measurable (0.10), this is largely due to the variance in the number of active years across professors. When this factor is controlled for by taking patent and paper output averaged over years, the coefficient is much smaller (0.04).


Tables 5 and 6 present regression analyses designed to explore this issue more systematically. Table 5

focuses on the question of the degree to which the level of patenting activity is predictive of the volume of publication activity, and presents analyses of the general form:


$$\text{publication behavior}_{it} = \alpha + \beta \text{ patenting behavior}_{it} + \delta x_{it} + \epsilon_{it}, \quad (1)$$

where i is the group index for professor, t is the index for year, x_{it} is a vector of control variables, and ϵ_{it} is

Figure 7 (a) Patent Versus Paper Output—1983–1997 ($n = 236$)
 (b) Patent Versus Paper Output per Year—1983–1997 ($n = 236$)



(a)



(b)

Table 4a Correlation Matrix: Patenting and Publishing (Flow Measures)

	Paper _t	Paper _{t-1}	Paper _{t-2}	Paper _{t-3}	Patent _t	Patent _{t-1}	Patent _{t-2}	Patent _{t-3}
Paper _t	1.00							
Paper _{t-1}	0.76	1.00						
Paper _{t-2}	0.73	0.75	1.00					
Paper _{t-3}	0.62	0.70	0.70	1.00				
Patent _t	0.014	0.006	0.031	0.04	1.00			
Patent _{t-1}	-0.004	-0.00	-0.01	0.01	0.49	1.00		
Patent _{t-2}	-0.00	-0.00	-0.01	-0.01	0.41	0.45	1.00	
Patent _{t-3}	-0.02	-0.00	-0.01	-0.02	0.30	0.39	0.43	1.00

AGRAWAL AND HENDERSON
Putting Patents in Context

Table 4b Correlation Matrix: Patenting and Publishing (Stock Measures)

	Total Papers	Total Patents	Total Years	Papers per Year	Patents per Year
Total papers	1.00				
Total patents	0.10	1.00			
Total years	0.34	0.17	1.00		
Papers per year	0.95	0.06	0.13	1.00	
Patents per year	0.02	0.90	-0.04	0.04	1.00

Table 5 Publication Behavior As a Function of Patenting Activity—Fixed Effects Models

Dependent Variable	Papers _(t)	Patents _(t)	Papers _(t)	Patents _(t)	Depreciated Stock of Papers _(t)
	(5.1)	(5.2)	(5.3)	(5.4)	(5.5)
No. Observations	2,237	2,105	2,105	2,105	2,784
No. Groups	213	213	213	213	236
Patents _(t)	0.02 (0.10)	-0.02 (0.08)	-0.03 (0.08)		
Patents _(t-1)	-0.03 (0.10)	0.03 (0.09)	0.02 (0.09)	0.22** (0.02)	
Patents _(t-2)	-0.02 (0.11)	0.02 (0.09)	0.01 (0.09)	0.09** (0.03)	
Patents _(t-3)	-0.12 (0.11)	-0.13 (0.09)	-0.14 (0.09)	-0.04 (0.03)	
Papers _(t)				-0.00 (0.01)	
Papers _(t-1)		0.40** (0.03)	0.40** (0.03)	-0.01 (0.01)	
Papers _(t-2)		0.36** (0.03)	0.36** (0.03)	0.01 (0.01)	
Papers _(t-3)		0.08** (0.03)	0.08** (0.03)	0.02 (0.01)	
Years active			0.03* (0.01)	0.01* (0.00)	0.64** (0.03)
Depreciated stock of patents _(t)					-0.12 (0.11)
Intercept	1.98** (0.07)	0.46** (0.08)	0.15 (0.17)	0.08 (0.05)	0.79** (0.27)
R ² within	0.00	0.33	0.33	0.07	0.17
between	0.00	0.91	0.90	0.76	0.05
overall	0.00	0.64	0.63	0.26	0.06

Note. Standard errors in parentheses.

* $p < 0.05$; ** $p < 0.01$.

Table 6 Paper Citations As a Function of Patenting Activity—Fixed Effects Models

Dependent Variable	Depreciated Stock of Paper Citations _(t)	Depreciated Stock of Paper Citations _(t)	Depreciated Stock of Patents that Cite Papers _(t)
	(6.1)	(6.2)	(6.3)
No. Observations	2,784	2,784	2,784
No. Groups	236	236	236
Depreciated stock of patents _(t)	6.00** (0.93)	2.56** (0.90)	0.33** (0.03)
Depreciated stock of papers _(t)	10.52** (0.17)	9.28** (0.17)	0.05** (0.01)
Years active		4.82** (0.27)	0.12** (0.01)
Intercept	-18.43** (1.63)	-49.60** (2.33)	-1.01** (0.07)
<i>R</i> ² Within	0.62	0.66	0.22
Between	0.71	0.70	0.21
Overall	0.66	0.68	0.20

Note. Standard errors in parentheses.

* $p < 0.05$; ** $p < 0.01$.

an error term. Fixed effect models are used to analyze this panel data.¹⁵

In Specification (5.1) (see Table 5), we begin by regressing a simple count of publications on a count of patents, both current and lagged, without including any control variables. None of the coefficients are significant (p -values are 0.82, 0.74, 0.88, and 0.25, respectively). Next, we add lagged publication measures in (5.2) which are highly significant ($p < 0.01$ for all measures). In other words, while the number of papers written three years ago is related to the number of papers written today, the number of *patents* written today or in any of the last three years appears to be unrelated to current paper output. Then, in (5.3), we add a control for the number of years at time t the professor has been active during the period under investigation. One might suspect that this control would eliminate the relationship between current and lagged paper output. However, it does not.

¹⁵ A series of Hausman tests indicates that the hypotheses that individual effects follow a random-normal distribution are rejected for several specifications presented in these tables such that we do not use the random-effects model that would otherwise provide more efficient specifications.

The lagged measures of paper output remain highly significant and the current and lagged measures of patent output remain insignificant.¹⁶ In (5.4), we test the reverse relationship and estimate the effects of paper output on patent output. The results are very similar. While one- and two-year lagged patent measures and the number of years that the professor has been active are positively related to current patent output, none of the paper count variables are significant, and Granger causality tests in both directions suggest that patent and paper outputs are independent. Specification (5.5) relaxes the strict constraint of comparing a single year's publishing output with a single year's patenting output by including stock rather than flow measures.¹⁷ Our core result continues to hold: patenting activity does not appear to be significantly related to publishing activity.

¹⁶ We also test for individual year effects using a specification similar to (5.3) but including dummy variables for each year. The coefficients on year dummies are generally insignificant. This result is not reported in Table 5.

¹⁷ We use a depreciation rate of 20%, which is standard for this kind of analysis (see Henderson and Cockburn, 1996).

The results presented in Table 6 explore the degree to which patenting activity is related to the degree to which a professor's work is *cited*. We estimate specifications of the general form:

stock of paper citations_{it}

$$= \alpha + \beta \text{ stock of patents}_{it} + \delta x_{it} + \epsilon_{it}, \quad (2)$$

where, as above, x_{it} is a vector of control variables and ϵ_{it} is the error term.

These results are more interesting. We include the stock of papers as a control variable in every regression since we expect paper citations to increase as the number of published papers increases. Model (6.1) (see Table 6) suggests that the stock of patents is positively related to the stock of paper citations, even after controlling for the stock of papers. While this effect is reduced by more than half (the coefficient drops from 6.0 to 2.6) after controlling for the number of years the professor has been active (6.2), the relationship is still positive and statistically significant.

Notice that the significance of these results lends additional credibility to the insignificance of our results in Table 5, suggesting that while patent counts are not good predictors of paper counts, or with the *volume* of a faculty member's research, they are correlated with paper citations, or with its *impact*.

Specification (6.3) explores another measure of impact, and regresses the depreciated stock of the patents that cite each professor's papers against their stock of patents and papers. Once again, we find a significant and positive coefficient on patent stock, even when we control for the stock of papers and the number of active years. Collectively, these results suggest that while patent counts are not a good predictor of publication counts, they are a reasonable predictor of the "importance" of a professor's publications, as measured by citations.

3.5. Patents: Substitute or Complement?

Lastly, we turn to the difficult question of the degree to which patenting acts as a substitute or complement to the process of conducting fundamental research. In commercial settings, basic, or "fundamental," research is often considered a substitute for more applied work (Cockburn et al. 2001). Several

observers have worried that a similar dynamic may be at work within universities, and that an increasing focus on the commercial implications of university research may skew university faculty away from the more fundamental work that universities were originally created to produce.

Our qualitative interviews suggest that neither patenting nor publishing is generally the motivation for selecting a particular research agenda. Most faculty members claim that they do not embark on a research program with a particular patent or paper outcome in mind. Rather, they suggest that they are engaged in a research stream that they find interesting and challenging, and that they make patent or publish decisions on a case-by-case basis. This is not to imply that some professors do not have a greater disposition towards patenting than others, but our interviews suggested that the patent-versus-paper question did not seem to drive the direction of research programs, at least in most cases. The following quote is representative:

I don't consciously do patentable research instead of publishable research, or vice versa. First of all, most patentable research is also publishable. Second, when I start working on a research project I have no idea whether it's ever going to result in anything useful, let alone a patented invention or a published journal paper. I work with a number of colleagues and on a variety of research trajectories. When we get on to something that looks like it might be patentable, if we have time, and if we're motivated, we check out whether it's worth patenting. However, it is useful to talk to industry people with real problems because they often reveal interesting research questions—but sometimes they try to steer you towards patenting. Sometimes that research results in something patentable, sometimes not. (EECS professor, interview, April 6, 1999)

Our quantitative results are more ambiguous. On one hand, there are a small group of faculty who appear to patent much more proportionately than their peers (Figures 7a and 7b). On the other hand, if patenting activity was substituting for fundamental research for a majority of the faculty, one might expect publication rates to be negatively correlated with patent counts. Not only have we shown that this is not the case, but we have also demonstrated that

there is some reason to believe that increasing patenting activity is correlated with increased rates of citation to the faculty member's publications. This result might reflect the fact that patent counts are a good measure of the degree to which research results can be immediately applied so that faculty who patent extensively are at increased risk of citation from industry. However, it is important to note that we measure *total* citation counts, academic as well as industrial, so that our result is also consistent with the hypothesis that patenting may actually be a complement to fundamental research.

4. Conclusions

What are the implications of these results for the use of patent-related metrics in studies of university innovation and knowledge transfer? First, they underline the well-established idea that patents are a relatively small channel for the transfer of knowledge out of the university. Echoing Branstetter (2000) and consistent with Cohen et al. (1998), we showed that MIT professors write far more papers than patents, and that many faculty members never patent at all. Moreover, our results suggest that patterns of patent citations may not be representative of wider patterns of collaboration or paper citation: different firms appear to use quite different channels to access knowledge at MIT. They also suggest that patent counts are not useful measures of the overall output of new knowledge, if publication count is taken to be a reasonable measure of such output. These results imply reasonably serious limitations in terms of generalizability across channels and overall knowledge transfer when interpreting results based purely on patent-related data. Second, there is some evidence that patent counts may be correlated with the "impact" of a faculty member's research, at least as measured by paper citations. This result is of significant interest because it suggests that patent data may offer some insight into the impact of university research. Finally, we suggested that our results offer some evidence that, at least at these two departments at MIT, patenting is not substituting for more fundamental research, and that it might even be a complementary activity.

Clearly, much remains to be done. We plan to conduct a much finer-grained analysis of the degree to

which different firms use different channels to access knowledge: Our data set contains information about many thousands of firms, and we hope to use it to explore the degree to which our preliminary analysis is representative of the larger universe. Moreover, we are curious as to why different firms choose to use different channels. Are they significantly different? Do they make quite different use of MIT-generated knowledge?

In addition, we plan to focus in much more depth on heterogeneity in faculty behavior across departments and over time. Do faculty who patent widely and whose patents are widely cited "look different" from their colleagues? Do they work with different types of firms? We are hopeful that these data will allow us to begin to make progress on these and related questions.

Acknowledgments

The authors thank Iain Cockburn, Scott Stern, Scott Shane, Atul Nerkar, and participants of the NBER productivity workshop and of the 'Roundtable for Engineering Research' that was held at the Georgia Institute of Technology (December 2000) for useful comments. They also appreciate the valuable efforts of Alex Oettl, Tin Yau Lee, and Brian Quinlan for their research assistance. This research was partially funded by the MIT Center for Innovation in Product Development under NSF Cooperative Agreement Number EEC-9529140 and the Center for Knowledge-Based Enterprises at Queen's University. Their support is gratefully acknowledged.

References

- Agrawal, A. 2000. Importing scientific inventions: Direct interaction, geography, and economic performance. Doctoral thesis, UBC mimeo, Chapter 2, University of British Columbia, British Columbia, Canada.
- Blumenthal, D. 1996. Relationships between academic institutions and industry in the life sciences—an industry survey. *New England J. Medicine* 334(6) 368–373.
- Branstetter, L. 2000. Measuring the link between academic science and industrial innovation: The case of California's research universities. Unpublished, NBER Summer Institute, Cambridge, MA.
- Cockburn, I., R. Henderson. 1998. Absorptive capacity, coauthoring behavior, and the organization of research in drug discovery. *J. Indust. Econom.* 46(2) 157–182.
- , ———, S. Stern. 2000. Untangling the origins of competitive advantage. *Strategic Management J.* 21 1123–1145.
- , ———, ———. 2001. Balancing incentives: The tension between basic and applied research: MIT mimeo, Cambridge, MA.

- Cohen, W., R. Florida, L. Randazzese, J. Walsh. 1998. Industry and the academy: Uneasy partners in the cause of technological advance. R. Noll, ed. *Challenges to the Research University*. Brookings Institution, Washington, D.C.
- Gregorio, D., S. Shane. 2000. Why do some universities generate more start-ups than others? Unpublished manuscript.
- Henderson, R., I. Cockburn. 1996. Scale, scope and spillovers: The determinants of research productivity in drug discovery. *Rand J. Econom.* 27(1) 32–59.
- , A. Jaffe, M. Trajtenberg. 1998. Universities as a source of commercial technology: A detailed analysis of university patenting, 1965–1988. *Rev. Econom. Statist.* 80(1) 119–127.
- Jaffe, A. 1989. Real effects of academic research. *Amer. Econom. Rev.* 79 957–970.
- , M. Fogarty, B. Banks. 1997. Evidence from patents and patent citations on the impact of NASA and other federal labs on commercial innovation. Working paper 6044, National Bureau of Economic Research, Cambridge, MA.
- , R. Henderson, M. Trajtenberg. 1993. Geographic localization of knowledge spillovers as evidenced by patent citations. *Quart. J. Econom.* 108 577–598.
- , 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. Working paper 5712, National Bureau of Economic Research, Cambridge, MA.
- Jensen, R., M. Thursby. 1998. Proofs and prototypes for sale: The tale of university licensing. Working paper 6698, National Bureau of Economic Research, Cambridge, MA.
- Mansfield, E. 1995. Academic research underlying industrial innovations: Sources, characteristics and financing. *Rev. Econom. Statist.* 77 55–65.
- Mowery, D. C., R. R. Nelson, B. Sampat, A. A. Ziedonis. Academic patent quality and quantity before and after the Bayh-Dole Act in the United States. *Res. Policy* Forthcoming.
- Thursby, J., M. Thursby. 2000. Who is selling the ivory tower? Sources of growth in university licensing. Working paper 7718, National Bureau of Economic Research, Cambridge, MA.
- Zucker, L., M. Darby, J. Armstrong. 1998a. Intellectual capital and the firm: The technology of geographically localized knowledge spillovers. *Econom. Inquiry* 36 65–86.
- , ———, M. Brewer. 1998b. Intellectual capital and the birth of U.S. biotechnology enterprises. *Amer. Econom. Rev.* 88 290–306.

Accepted by David C. Mowery and Scott Shane; received December 2000. This paper was with the authors 9 months for 1 revision.