

The determinants of faculty patenting behavior: Demographics or opportunities?[☆]

Pierre Azoulay^{a,*}, Waverly Ding^{b,1}, Toby Stuart^{c,2}

^a MIT and NBER, Sloan School of Management, 50 Memorial Drive, E52-555, Cambridge, MA 02142, United States

^b University of California, Haas School of Business, 545 Student Services, #1900, Berkeley, CA 94720, United States

^c Harvard Business School, Rock 211, Soldiers Field, Boston, MA 02163, United States

Received 16 June 2005; accepted 19 May 2006

Available online 16 January 2007

Abstract

We examine the individual, contextual, and institutional determinants of faculty patenting behavior in a panel dataset of 3862 academic life scientists. Using discrete time hazard rate models and fixed effects logistic models, we find that patenting events are preceded by a flurry of publications, even holding constant time-invariant scientific talent and the latent patentability of a scientist's research. Whereas previous research emphasized that academic patenters are more accomplished on average than their non-patenting counterparts, our findings suggest that patenting behavior is also a function of scientific opportunities. This result has important implications for the public policy debate surrounding academic patenting.

© 2007 Elsevier B.V. All rights reserved.

JEL Classification: O31; O32; O33

Keywords: University patenting; Academic science

1. Introduction

In the past few decades, universities and other public-sector research organizations have become more proactive in their efforts to commercialize scientific discoveries (e.g., Jaffe and Lerner, 2001;

[☆] This paper was written for the NBER Conference “Academic Science and Entrepreneurship: Dual Engines of Growth?” Santa Fe, NM, 1 April 2005.

* Corresponding author. Tel.: +1 617 258 9766; fax: +1 617 253 2660.

E-mail addresses: pazoulay@mit.edu (P. Azoulay), wding@haas.berkeley.edu (W. Ding), tstuart@hbs.edu (T. Stuart).

¹ Tel.: +1 510 643 4252; fax: +1 510 643 1412.

² Tel.: +1 617 496 4626; fax: +1 617 495 3817.

Jensen and Thursby, 2001; Thursby and Thursby, 2002). This change has spawned a growing academic literature on university technology transfer, one stream of which has assessed trends in university patenting and the spillover of university science into the private sector (Jaffe, 1989; Mansfield, 1995; Henderson et al., 1998). Underlying the well documented upswing in university patenting has been a sharp increase in the number of individual academic scientists who are listed as inventors on patents. In this paper, we examine the individual, contextual, and institutional determinants of academic patenting in a panel dataset of 3862 academic life scientists.

Past research on academic patenting has proceeded along three distinct tracks. The first track has explored the influence of institutional factors, such as the prestige of the university or the quality of its technology licensing office, on the intensity of patenting at the university level (Di Gregorio and Shane, 2003). The second track has followed a more qualitative approach; the architects of it have shown that there are important differences in the propensity to patent across scientific and technological fields and in the area-specific motivations underlying patenting activity. This work has highlighted the notable difference in the precursors to patenting between the life and physical sciences/engineering (Owen-Smith and Powell, 2001). Researchers in the third track have focused on evaluating statistically the consequences of patenting for the scientific output of individual academic scientists (Breschi et al., 2005; Fabrizio and Diminin, 2005).

Our paper shares with this last group of studies a focus on the individual scientist as the level of analysis, but we concentrate on the antecedents (rather than the consequences) of faculty patenting activity. We adopt this approach because we view the answer to the question “Who Patents?” as a necessary input into the policy debate surrounding the consequences of the academic patenting phenomenon. For example, it seems imprudent to make policy recommendations (e.g., should patent output be taken into consideration in hiring and promotion decisions?) before gaining a clear understanding of who is involved in this activity, at what career stage, and at what type of institution. Furthermore, understanding the determinants of the patenting decision provides a window into the broader phenomenon of academic entrepreneurship, since among academic life scientists at least, applying for a patent is a very strong and robust predictor of the decision to participate in the founding of a biotechnology firm (Stuart and Ding, 2006).

Our study generates a novel set of results, underscoring the benefits of fine-grained longitudinal data at the researcher level of analysis. Although we present a number of findings, we consider two to be most distinctive relative to the existing literature on academic patenting. First, we document that patenting is often accompanied by a flurry of publication activity in the year preceding the patent application, even after accounting for the lagged stock of publications (in hazard rate models) or controlling for scientist fixed effects. This result highlights the fact that academic patenting, rather than merely reflecting the influence of time-invariant demographic factors, also responds to variation in scientific opportunities. Holding life-time scientific achievement constant, we find that surges of scientific productivity, not steady research performance, is most likely to be associated with patenting. We interpret this as suggesting that, at the individual level, the uncovering of new, productive areas of scientific inquiry is an important precursor to the act of patenting.

Our second novel finding concerns the establishment of a relationship between what we construe to be the *latent patentability* of a faculty's research and his/her propensity to patent. While latent patentability typically has been assumed to be unobservable, we are able to devise a patentability score for each scientist in our sample by using keywords in the publications of scientists that have already applied for patent rights as a benchmark for patentable research, and then comparing the research of each scientist in our dataset to this benchmark. Although there is noise in this proxy, it nevertheless quite strongly predicts a patenting event.

In addition to our results concerning the association between the onset of patenting and both scientific opportunity and latent patentability, we also find a pronounced gender difference in the propensity to patent. In a typical specification, we find that women scientists patent at half the rate of men (see Ding et al., 2006, for additional exploration of this difference). We also observe life-cycle effects on the propensity to patent. We find that mid-career academics are much more likely to patent than are younger and older faculty members. Further, we provide suggestive evidence that current institutional affiliation and scientists' social networks influence the onset of patenting. In particular, our analyses uncover two results: first, we find that the presence of co-authors who have patented in the past increases the likelihood of a patent application, and second, we show that scientists are more likely to patent when they are employed at universities with large patent portfolios. These two findings are consistent both with the existence of genuine peer effects in patenting and with the idea that the academic labor market matches scientists with commercial proclivities in similar departments. In addition, the university-level patent effect may reflect the role of a well-functioning technology licensing office (TLO) in either increasing the individual-level rewards or reducing the costs of faculty patenting.

Independent of any specific finding, the general analysis herein is relevant to the broader question of the impact of patenting on the development of academic science. Surveys of university faculty have found rampant concern that patenting is skewing research agendas toward commercial priorities, causing delay in the public dissemination of research findings, and crowding out effort devoted to the pursuit of fundamental knowledge (Blumenthal et al., 1996; Campbell et al., 2002; Krinsky, 2003). Insofar as our results relate to this issue, the finding that patenting follows a flurry of publications suggests to us that the crowd-out hypothesis is unlikely to hold true.³ Although we cannot adjudicate between opposing claims regarding the *effect* of patenting on individual-level or university-level outcomes in the present study, one can construe our results as providing the “first stage” of an econometric analysis of the effect of academic patenting on the rate and direction of scientific progress, an evaluation we are pursuing in other research (Azoulay et al., 2006).

The rest of the paper proceeds as follows. In the next section, we situate our contribution in the large and growing literature on academic patenting and highlight what we regard as outstanding issues that can only be resolved with researcher-level longitudinal data of the kind we analyze. Section 3 describes data sources and the construction of the sample and discusses our econometric approach. Section 4 presents descriptive statistics and reports our results. Section 5 concludes.

2. Who patents?

In recent times, the region of overlap between academic science and commercial markets has experienced significant growth. The expanding interface between these two domains raises myriad questions, ranging from the amount of near-term economic value created by the spillovers of university research, to the emergence of select universities as engines of entrepreneurial activity, to the influence of opportunities to commercialize scientific research on the traditional incentive systems that have governed academic science. Researchers have engaged a variety of these questions, and advancement in our understanding is occurring along many fronts.

Spurred in part by accessible data, many studies have assessed the role of universities as direct sources of commercial innovations, primarily considering the quality and quantity of their

³ However, if scientific trajectories associated with patents exhaust themselves more quickly than those remaining free of associations with the world of commerce, then intertemporal substitution of “basic”, fundamental knowledge by “applied”, patentable output could still be consistent with the patterns we observe in the data.

innovative outputs. For instance, Henderson et al. examine the relative importance of university patents, finding that over time there has been a decline in the positive quality gap separating university patents from those assigned to for-profit firms. Mowery et al. (2001) have investigated the consequence of the policy changes brought about by the Bayh-Dole Act. They challenge the conventional wisdom that Bayh-Dole has accelerated universities' production of patents, showing that the legislation was not a primary factor in explaining the uptick in patenting at three prominent universities. At the level of the university, Thursby and Thursby find that university administrators have become more proactive in pursuing patents and licensing opportunities. Di Gregorio and Shane explore cross-university differences in the formation of start-up companies, discovering that intellectual eminence is a central factor distinguishing the universities that spawn start-up companies.

The majority of the archival work that has looked at the commercial outputs of universities has treated the organization as the level of analysis (notable exceptions include Agrawal and Henderson, 2002; Stephan et al., 2004). Two recent papers (Breschi et al., 2005; Fabrizio and Diminin, 2005) evaluate statistically the effect of academic patenting on publication output in a panel dataset of individual scientists. These authors start from a sample of patenting academics, which they then supplement with a set of non-patenter controls. This empirical strategy is legitimate if one focuses on the *consequences* of academic patenting for other outcomes of interest, but less so if one is concerned with the *antecedents* of faculty patenting activity.

In this article, we analyze the probability of patenting in a large, longitudinal, random sample of university faculty in the life sciences. Our analysis is guided by an interest in four issues. First, how does the proclivity to patent vary with scientists' gender and experience in the profession? Second, what is the relationship between scientific productivity (measured as papers published) and patenting? Third, are there significant differences across research areas within scientific disciplines in terms of the apparent "patentability" of the work, and is there any evidence to suggest that scientists may be altering their research to move toward patentable research? Fourth, to what extent is the propensity to patent sensitive to the work context of the individual scientist, particularly the level of commercial orientation of a scientist's university and his or her co-authors?

Treating each of these in turn, we first ask, how does the propensity to patent change over the scientific career? Economists and sociologists alike have a long-standing interest in career dynamics in academe, in part because incentives in science vary over the professional life cycle. Two elements of the institutionalized reward system in science are generally thought to be tenure-invariant: the tying of peer recognition to priority in research discovery, and the intrinsic satisfaction garnered from solving vexing problems. However, monetary incentives to produce research output in science do depend on the career stage, and it is well known that the wage-tenure profile in academic science is not steep (Stephan, 1996). Given the shallow slope of post-tenure salary increases, Levin and Stephan (1991) suggest that levels of investment in research should vary over the career life-cycle. In particular, senior scientists with tenured appointments may reallocate some of their effort to consulting and other extra-university income generating opportunities. Therefore, if the widely held assumptions about changing incentives over the career do in fact hold, we should observe that the rate of patenting accelerates in the post-tenure interval.

A countervailing possibility is implied by a growing body of ethnographic research that portrays the increasing acceptance of patenting as a legitimate activity in academic science (Etzkowitz, 1998). If the pendulum has swung to the point that patenting is perceived to contribute to scientists' scholarly reputations and influence, we would expect to observe that, viewing successive cohorts of scientists, patenting occurs with increasing frequency in the early career stage. Consistent with this perspective, Owen-Smith and Powell describe interviews with scientists who have come to

view patents as reaffirmations of the originality of their work and as contributing to their scientific visibility. Recent interview-based accounts thus raise the possibility of a significant shift in the norms and reward system in science, with implications for life-cycle effects in patenting.

Moving on, there is a large literature that connects the gender of a scientist to his or her performance along traditional dimensions of scientific achievement, such as productivity and rates of rank advancement (e.g., Long and Fox, 1995). Numerous studies have documented that women earn comparably lower salaries, garner fewer rank advancements, produce fewer papers, and exit the profession at a great rate than do men. However, little is currently known about gender differences in engagement in the commercial aspects of science. In this article, we estimate the effect of scientists' gender on the rate of patenting. In a related paper (Ding et al., 2006), we incorporate qualitative data to probe the possible causes of the difference we find.

Next, we seek to identify the relationship between scientists' productivity and the likelihood that they patent. Existing evidence suggests that the scientists with the most stellar academic credentials are also the most likely to be involved in commercial endeavors. In particular, Zucker et al. (1998) describe the importance of the geographic location of star scientists in the emergence of the biotechnology industry. They argue that the direct participation of leading academic scientists in early stage biotechnology companies was so important that the locations of star scientists served as geographical constraints on the development of the industry.

The existing literature thus provides reason to expect that patenting is concentrated among the group of eminent scientists. Yet beyond the general association between research output and the likelihood of engaging in market-related activities, identifying more precisely the relationship between the production of papers and patents may adjudicate among the competing mechanisms that could generate the relationship. In particular, if the magnitude of the *stock* of scientists' research output predicts the onset of patenting, it is likely that faculty members' scientific reputations are important considerations in the decision to patent. If this proves to be the case, a plausible explanation is that the prominence of the inventor on a patented technology may influence the university's ability to capitalize on the intellectual property by affecting the probability that potential licensees become aware of and interested in the technology.

Consider instead the implication of a positive relationship between the *flow*, but not the stock, of scientists' research output and the probability that a patent is issued. If the flow of output is the determining factor, we would suggest that scientific opportunities loom large in the transition to patenting. A flurry of scientific output (a high flow of publications) occurs when a scientist unearths a productive domain of research. If patenting is a byproduct of a surge in productivity, we think it reasonable to conclude that a patent is often an opportunistic response to the discovery of a promising research area.

The third issue we consider is how the specific areas of expertise of academic scientists affect the likelihood of patenting. Obviously, there exists heterogeneity across scientists in the potential commercial value of the research they produce. If one needs to account for such differences, it is tempting to argue that the analyst can accommodate them by incorporating scientist fixed effects in the analysis. We believe, however, that this represents just a partial solution given the volume and the diversity of research projects that scientists participate in throughout their careers. We therefore attempt to develop a direct measure of the "patentability" of scientific research. The intuition behind the measure is that knowledge of the research foci of academic scientists who have already patented can be used to identify the domains of science in which research is patentable. With such a measure in hand, we ask two questions. First, does the latent patentability of scientists' research in fact affect the probability of patenting? Second, is it the patentability of the stock or the flow of research outputs that most consequentially influences the propensity to patent?

Fourth, we explore two elements of scientists' work contexts. While it is well established that propensities to patent vary substantially across universities, we do not have a clear sense for the influence of organizational characteristics on the patenting rates of otherwise similar scientists within different universities. Numerous studies suggest that the decision to engage in commercial activity of all sorts is strongly influenced by factors ranging from the norms and culture of an institution vis-à-vis commercial activity to the quality of the university's technology transfer office (Owen-Smith and Powell, 2001; Thursby and Thursby, 2002). Two prevalent considerations are thus the (potentially endogenous) role of a smoothly functioning technology transfer office in encouraging faculty to disclose possibly patentable research findings and, more generally, a pro-commercialization "entrepreneurial culture" at a university. In our analysis, we ask whether university-level variables influence the patent rate net of controls for many observable individual-level characteristics.

A related question concerns the influence of proximate colleagues on the patent proclivities of individual scientists. There are a set of reasons to expect that scientists who work closely with commercially inclined peers will be more likely to pursue commercial applications of their scientific research. Stuart and Ding argue that there are two mechanisms through which colleagues affect the probability that a particular scientist engages in commercial activities. First, peers exert attitudinal influences, in particular shaping the degree to which a given scientist is likely to embrace patenting as both a legitimate undertaking for an academic scientist and as a potential contributor to his or her professional standing. Second, peers convey information that may lower the cost of patenting, such as contacts in the technology transfer office and advice about how to minimize the amount of time consumed in patenting. We thus look for what might be labeled as "peer effects" on the transition to patenting. Specifically, we examine whether scientists who have co-authorship links with patent holders or with researchers employed in the private sector are themselves more likely to patent.

A necessary caveat pertains to the thorny issue of causality. Many of our independent variables, such as publications or latent patentability, could be considered outcomes of interest. Moreover, it would be incorrect to interpret our findings as providing evidence, *inter alia*, that publications and patents are complements or that latent patentability "causes" patent applications. Rather, we have identified correlates of patenting. The conditional correlations we estimate can still be useful insofar as they help to narrow the range of plausible theories regarding the *effect* of academic patenting on scientific productivity. In addition, since our most interesting results pertain to what are in fact lagged dependent variables, the study highlights the need to use correct econometric methodologies to recover causal effects. This is pursued in a companion paper (Azoulay et al.).

3. Data, sample characteristics, and econometric approach

We examine the determinants of faculty patenting behavior in a panel dataset of academic life scientists employed at universities and non-profit research institutes. This area was chosen because the biomedical fields have accounted for the preponderance of university patenting and licensing activity (Mowery et al.). While we have not selected scientists because they have patented, we have sampled from scientific disciplines that we know to have significantly contributed to a vibrant area of technological development. We began by drawing 12,000 doctoral degree recipients from UMI Proquest Digital Dissertations, which lists Ph.D. recipients from more than one thousand universities. In forming the sample, we randomly selected individuals, but only those with Ph.D.s

Table 1
Top 15 scientific disciplines in the sample

UMI subject code	UMI subject description	Frequency
487; 303	Biochemistry	855 (22.2 percent)
306	Biology, general	563 (14.6 percent)
410	Biology, microbiology	466 (12.1 percent)
419	Health sciences, pharmacology	239 (6.2 percent)
490	Chemistry, organic	212 (5.5 percent)
786	Biophysics, general	210 (5.4 percent)
369	Biology, genetics	191 (4.9 percent)
433	Biology, animal physiology	170 (4.4 percent)
982	Health sciences, immunology	167 (4.3 percent)
307	Biology, molecular	102 (2.6 percent)
301	Bacteriology	61 (1.6 percent)
287	Biology, anatomy	54 (1.4 percent)
571	Health sciences, pathology	52 (1.3 percent)
349	Psychology, psychobiology	37 (1.0 percent)
572	Health sciences, pharmacy	34 (0.9 percent)

This table reports the Top 15 disciplines from which our sample was drawn. These disciplines have spawned the greatest number of biotechnology firm founders, scientific advisors and executives. The table also reports the frequency and the proportion of scientists in our sample for each of these 15 scientific disciplines.

in scientific disciplines that have informed commercial biotechnology.⁴ This assures a random sample of Ph.D.s in areas in which academic research may have significant, short-term commercial value.

Given our focus on the life sciences, one might question whether our results generalize to other academic fields, such as mechanical or electrical engineering. One should note, however, that our definition of life sciences is expansive. For example, our data include scientists holding Ph.D.'s in chemistry, chemical engineering, materials engineering, plant biology, veterinary sciences, and food science. The life sciences, broadly construed, represent such a large slice of the academic patenting phenomenon that the issue of generalizability does not loom particularly large.⁵

Next, we obtained scientists' publication records from the ISI's *Web of Science* database. Because the Web of Science includes authors' affiliations, we were able to identify Ph.D. graduates who pursued careers outside of academe. After removing individuals who (i) had no publications in any post-graduate year, (ii) published exclusively under corporate affiliations, or (iii) exited academe early in their careers,⁶ we were left with 3862 scientists, all of whom we know to have

⁴ To identify the scientific disciplines that have been most important to biotechnology, we coded the educational backgrounds of the Ph.D. holding, university-employed scientific advisory board members of all publicly traded biotechnology firms. The source of information on scientific advisors' degrees was the IPO prospectuses of the 533 U.S.-based biotechnology firms that have been filed with the U.S. Securities and Exchange Committee. We then stratified the random draw from UMI to correspond to the disciplines and Ph.D. grant years of firms' scientific advisors. For example, 22 percent of biotechnology company scientific advisors hold biochemistry Ph.D.s; we drew a corresponding proportion of biochemists into our sample. Table 1 lists the Top 15 disciplines from which scientists in our sample are selected.

⁵ In a related paper, one of the authors assembled a dataset of "superstar" academic patenters, who were defined to be U.S.-based academics with more than 17 patents between 1976 and 2004 (this corresponds to scientists above the 99th percentile of the patent count distribution). Among the 544 such scientists, he found only 138 (25.37 percent) who did not fit our definition of "life scientists".

⁶ Ph.D.s with academic affiliations lasting less than 5 years were dropped from the dataset to exclude post-doctoral fellows who later moved to jobs in industry.

been employed at research institutions. Each scientist is observed from the year after he or she earned a Ph.D. until 1999, unless the individual exited academia.⁷ The final panel contains 54,737 person–year observations between 1968 and 1999.

3.1. Variables

A brief description of the patenting process in academia is useful to interpret the results we will present. The process begins when a faculty member discloses an invention to the university's technology transfer office (TTO).⁸ The commercial potential of this invention is then evaluated by the TTO, which may decide to seek patent rights on the invention. Usually concurrently, the TTO will market the innovation to potential licensing partners in industry. A typical licensing agreement specifies a 40 percent royalty rate for the individual faculty inventor, to be assessed on the gross licensing revenues the invention accrues.

The patents of the academic scientists in our data were assembled from the NBER patent database (Hall et al., 2001). To identify academic patenters, we matched the scientists in our dataset to the list of inventors in the NBER patent database. Matches were performed on the basis of last and first names, and we used information on assignee (university) and geographic region to eliminate false matches.⁹

Within the sample of 3862 scientists, we found 473 (12.2 percent) patenters who were listed on 1372 patents. Out of these patents, 342 were assigned to corporate entities (of which 31 were co-assigned to a university and a corporation), even though the inventors of interest were academically affiliated based on information revealed in other patent applications filed by the inventor or in publication records. Most of these corporate patents have multiple inventors, and a university scientist could be listed as one of the inventors for his advice during the process of invention. A typical example is Richard J. Lagow, who obtained a Ph.D. in inorganic chemistry from Rice University in 1970 and subsequently held professorships at MIT and the University of Texas Austin. Lagow began patenting in 1973, and his patents have been assigned to MIT, University of Texas, and Exflur Research Corporation. Among the 31 patents for which Exflur is the assignee and Lagow is an inventor, 28 involved multiple inventors and 3 listed Lagow as the sole inventor. Based on the data sources available to us, it is not possible to determine Lagow's exact role in developing these inventions and what type of arrangement Lagow has with University of Texas, but from the titles and abstracts of the Exflur patents it is clear that the patented inventions are based on knowledge closely related to Lagow's

⁷ We assume a researcher has exited academia when he or she fails to publish for five consecutive years or when the scientist begins to publish almost exclusively under a corporate affiliation. In either case, we censor observation in the year in which a scientist last publishes under a university affiliation.

⁸ Faculty members are contractually obligated to disclose potentially commercializable discoveries developed on university premises to the TTO; in theory (but see below) they do not have the option to patent university-originated discoveries without going through the official channels. The average TTO received 78 invention disclosures in 2003 but filed only 40 new patent applications (AUTM, 2003). Of course, these numbers vary widely across institutions depending on whether involvement with the world of commerce corresponds to a well-established culture within the institution.

⁹ Because we know the affiliations of the scientists in our data, we do not face the daunting name-matching challenges described in Trajtenberg (2004). We are able to rule out false positives by insisting that both scientists' names and affiliations match the inventor and assignee fields in the patent data. When scientist and inventor names matched but institutions did not, we compared the subject matter of patents to the content of scientists' papers to form a judgment about whether a match was valid. In making this assessment, we incorporated additional information, such as the location of the inventor and whether or not a given patent had multiple inventors.

research. Therefore, our data suggests that a non-trivial portion of faculty patenting activity may occur without the official involvement of their employing university's technology transfer office.

For each scientist in our data, we explore two dependent variables: the transition to first patent application and a dummy variable indicating whether the researcher applied for at least one patent in a given year.

3.1.1. Research output and latent patentability

We create three measures of scientists' research output. From the *Web of Science* we computed annual paper publication counts for each scientist. We count all papers on which a scientist is listed as an author.¹⁰ While this seems a good proxy for the rate of a given scientist's output, we would also like to measure the content of this output. We do this in two different ways. First, we use the affiliation data available from *Web of Science* to identify all instances in which a scientist wrote a paper that was co-authored with one or more individuals in a corporate research and development lab. We assume that papers co-authored with researchers in industry are more likely to be of an applied nature, and thus we consider publishing with co-authors in industry as an indicator of the degree to which scientists are engaging in commercially oriented research.

Second, it would be desirable to account directly for differences among scientists in the inherent "patentability" of their research. To construct such a measure, we have used the title words in scientists' publications to identify the areas in which they have conducted research, and then applied weights to these areas based on an (endogenous-to-the-sample) measure of the extent to which other scientists working in these areas have patented their discoveries. Intuitively, we use the publications of scientists who have already applied for patent rights as the benchmark for patentable research, and then compare the research of each scientist in our dataset to this benchmark to generate a research patentability score for each scientist-year. Specifically, the research patentability score for scientist i in year t is defined as:

$$\text{PATENTABILITY}_{it} = \sum_{j=1}^J w_{j,t-1}^i \frac{n_{ijt}}{\sum_k n_{ikt}}$$

where $j = 1, \dots, J$ indexes each of the scientific keywords appearing in the titles of the journal articles published by scientist i in year t ,¹¹ n_{ijt} is the number of times each of the keywords j has appeared in scientist i 's articles published in year t , and w_{jt}^i is a weight for each keyword that measures the frequency with which word j is used in the titles of articles published by scientists who have entered the patenting regime in year t or earlier, relative to those who have not entered the patenting regime as of year t (the calculation of w_{jt}^i is detailed in the data Appendix A available on the JEBO website). Intuitively, the patentability of a scientist's research can change because of a change in the direction of the research of that scientist or because other patenters' research increasingly comes to resemble that of the scientist. The former effect is captured by the ratio

¹⁰ In other words, we treat sole-authored and co-authored papers as equivalents, but restricting the set of papers to those where the focal scientist appears first or last in the authorship list generates results substantively similar to those we present below.

¹¹ We relied on title words in journal articles instead of journal- or author-assigned keywords because the *Web of Science* database did not begin to include keyword descriptors until 1992. However, the titles of biomedical research papers typically indicate the research area and the methodology used in the paper. We find high overlap between title words and keywords in the papers for which both are available.

$n_{ijt}/\sum_k n_{ikt}$, the latter by the weights $w_{j,t-1}^i$. Because the benchmark in year $t-1$ is used to weight title words in year t , year-to-year changes in the research patentability score will reflect only actions of the scientist (through their choices of title keywords), rather than contemporaneous changes in the benchmark.¹²

Finally, to capture the idea that the inherent patentability of past research might still influence the current propensity to patent, we compute a depreciated stock of the research patentability score using a perpetual inventory model. Through the impact of the depreciation rate δ ,¹³ this formulation captures the fact that the recent substantive research orientation of a scientist's research should influence current behavior more strongly than scientific trajectories that unfolded in the more distant past:

$$\text{STOCK_RP}_{it} = (1 - \delta)\text{STOCK_RP}_{i,t-1} + \text{FLOW_RP}_{it} = \sum_{\tau=0}^t (1 - \delta)^{t-\tau} \text{FLOW_RP}_{i\tau}$$

3.1.2. Demographic, university, and network attributes

Following a number of studies of the determinants of scientists' productivity, we were also able to construct many covariates to account for individual and institutional factors that may influence rates of publication and patenting. We coded gender based on first names. In the few instances of gender-neutral names for which we could not determine sex based on web searches, we assumed that the focal scientist was male. To account for life-cycle effects (Stephan), we include the number of years since a scientist earned his or her Ph.D. Because the time involved in publishing scientific research varies across fields, the regressions include a full set of dummies for researchers' dissertation subject areas. Some of the regressions control for time-invariant quality differences among researchers through the inclusion of scientist fixed effects. In specifications without fixed effects, we enter a dichotomous measure of the quality of a scientists' Ph.D. degree granting institution, which is a dummy variable indicating whether or not a scientists' doctoral program was ranked in the Top 20. Specifically, we collected Gourman Report rankings for all institutions in our dataset. Gourman ranking are available at the field level and were issued for the first time in 1980. We assigned universities their original rating for all years prior to 1980 (and updated them every other year for the subsequent period). In addition, we compute a count of the number of patents held by a scientist's Ph.D. granting university during the 5 years prior to his or her degree granting year to control for any possible imprinting effects.

We also include a number of employer-level variables that may influence scientists' patenting. These covariates are updated each year and when scientists change employers. First, given the existing evidence that prominent universities are more likely to be involved in commercial activities, we include a quality rank dummy variable analogous to the one constructed for Ph.D. granting institutions. Second, we use the AUTM surveys to create a technology transfer office (TTO) dummy variable, which is set to one in all years in which a scientist's employing university

¹² Previous researchers have developed other measures of proximity in technological space. For instance, Jaffe (1986) used a cosine-based measure to assess the proximity between the R&D portfolio of any given pair of firms. While this approach works well for measuring technological distance between dyads, it is not well suited to our setting since we need to measure the distance between the scientific trajectory of any given scientist relative to that of a benchmark *group* of (patenting) scientists.

¹³ We set δ equal to 0.15, the Griliches constant used by many innovation researchers on whose work this paper builds. We verified that our core results are not sensitive to this arbitrary choice.

has an active TTO. Third, a university's stock of patents (excluding those of the focal scientist) is entered in the model, among other things to control further for institutional differences in support for patenting.

Finally, to capture the patenting proclivities of our scientists' co-authors, we measure both the number of co-authors and whether the co-authors have applied for patents. We are able to identify patenting behavior only for co-authors who are also members of our sample. Since the set of scientists analyzed here are drawn randomly from the population, this limitation should not introduce bias, although the resulting count is clearly a noisy proxy for the underlying concept. Furthermore, to distinguish the co-author peer effect from the influence of peers at the same institution, we exclude co-authors who are also co-workers when creating these two variables.

3.2. Econometric considerations

Estimating the determinants of faculty patenting behavior requires a procedure that accommodates the discrete nature of the event. Since our interest lies in analyzing the dynamics associated with the onset of patenting in scientific careers, we employ discrete-time hazard rate models (Cox, 1972; Myers et al., 1973; Allison, 1982). The use of discrete-time models (as opposed to continuous-time models such as the Cox) is motivated by the fact that our failure time variable displays multiple events within each time period. For a researcher i during experience interval t , let the discrete-time hazard rate be $p_{it} = \Pr[T_i = t | T_i \geq t, X_{it}]$, where T_i is the time at which researcher i experiences an event and X_{it} a vector of covariates. We use a logistic regression function to link the hazard rate with time and the explanatory covariates:

$$\ln \left[\frac{p_{it}}{1 - p_{it}} \right] = \alpha_t + \beta' X_{it}$$

where α_t is a set of experience interval dummies. In practice, we estimate a simple logit of the decision to apply for a patent, where the observations corresponding to years subsequent to the first event have been dropped from the estimation sample. In this model, the factor change in the probability that a scientist will experience an event relative to probability of not experiencing an event is computed as $\exp(\beta)$. Though this differs from the Cox model (in which the hazard rate multipliers are to be interpreted relative to a non-parametric baseline hazard), the logit model estimates converge to the Cox estimates as the time interval approaches 0 (Thompson, 1977).

These models largely rely on between-scientist variation in the covariates to identify the determinants of the first transition to patenting. A complementary approach is to consider how within-scientist changes in covariates influence the propensity to patent. We do so by estimating so-called "fixed-effects" logit models by conditional maximum likelihood (Chamberlain, 1984). In contrast to our implementation of the standard logits, this approach analyzes the careers of patenting scientists in their entirety rather than just until the year of first patent application. In other words, we treat patenting as a repeatable event in the fixed-effects logit regressions. There is, however, a countervailing cost in the fixed-effects approach—it drops all observations corresponding to scientists who never patent.¹⁴ We believe that together, the discrete-time hazard models and the fixed effects logit models provide a comprehensive picture of the academic patenting phenomenon.

¹⁴ Conditional maximum likelihood estimation requires some variation in the dependent variable to condition out the individual scientist effects. Because scientists who have never patented have no variation on the outcome variable, they must be dropped from the analysis.




Fig. 1. Distribution of patent count for patenting scientists.

4. Results

We begin with some basic descriptive information about the sample. In Fig. 1, we plot the distribution of patents for the 473 (12.2 percent) patenting researchers in the sample. The histogram illustrates a rapid drop off after one; most patenters are listed on one or two patents throughout their career, and very few scientists in our data receive more than 10 patents. Thus, even when we restrict our attention to patenting scientists, the act of patenting never becomes routine for the majority. Fig. 2 displays the distribution of scientists’ total publication counts, broken out by their patenting status. Consistent with the conventional wisdom that patenting is concentrated among




Fig. 2. Distribution of publication count for patenting and non-patenting scientists.

Table 2
Descriptive statistics

	Mean	S.D.	Min.	Max.
Time-varying (58,562 person–year observations)				
Experience	10.201	7.122	1	32
Patent flow dummy	0.017	0.131	0	1
Publications flow	1.729	2.379	0	35
Publications stock	17.563	26.759	0	386
Research patentability flow	0.022	0.049	0	4.173
Research patentability stock	0.111	0.142	0	4.201
Collaboration tie with company scientists	0.235	0.424	0	1
Average number of identified co-authors per paper	0.122	0.244	0	10
Identified co-authors have patents	0.039	0.193	0	1
Employer graduate school in Top 20	0.231	0.422	0	1
Employer has TTO	0.488	0.500	0	1
Employer patent stock	71.80	145.18	0	2189
Calendar year	1986	7.741	1968	1999
Time-invariant (3862 observations)				
Female	0.211	0.408	0	1
Ph.D. university graduate school in Top 20	0.308	0.462	0	1
Ph.D. university 5-year patent stock	18.983	40.906	0	566
Scientist has one or more patents	0.122	0.328	0	1

the group of academically productive scientists, the distribution for the patenter subsample is much less skewed than that of the non-patenter subsample.

4.1. Descriptive statistics

Table 2 presents the summary descriptive statistics for variables used in our analysis. Table 3 reports, by scientists' patenting status, the mean research and employer characteristics measured at five career stages. The data in this table are quite suggestive in that they show, across all variables and all experience intervals, a difference between patenters and non-patenters. The table shows that researchers who have sought and received patent rights for their discoveries are more productive at each career stage: they publish on average 35 percent more research papers as those who have not yet patented. Scientists who have applied for patent rights are closer to commercial research than their non-patenting counterparts, as indicated by the fact that they have collaborated more often with researchers in the private sector. Likewise, the intrinsic patentability of their research is higher across all years of professional experience. Although this relationship is intuitive, it suggests that our journal article title-word-based measure of research patentability may in fact capture some of the patentability in scientists' research. Finally, patenters are more likely to work at universities where patenting activity is intensive (as measured by current employer patent stock); the presence of a TTO does not appear to make a difference in these univariate comparisons; and they are more likely to have co-authors that have themselves patented.

Fig. 3 displays the distribution of patenting events over time. Although we observe an uptick in the years following Bayh-Dole, it is also clear that patenting activity was taking place even before the adoption of the Act. This is consistent with the findings of Mowery et al. In addition, the figure demonstrates a relatively persistent increase in patenting in the academic life sciences, which has been documented elsewhere.

Table 3
Scientists' mean research and employer characteristics at five career stages by patent application status

	Experience = 5		Experience = 10		Experience = 15		Experience = 20		Experience = 25	
	Yes ^a	No ^a	Yes ^a	No ^a	Yes ^a	No ^a	Yes ^a	No ^a	Yes ^a	No ^a
Research publication flow	1.563 (1.648)	1.290 (1.801)	2.524 (2.945)	1.821 (2.228)	3.208 (3.276)	2.036 (2.629)	3.513 (4.029)	2.215 (2.888)	3.395 (4.023)	2.179 (2.955)
Research publications stock	6.760 (5.971)	5.832 (6.668)	19.066 (16.753)	14.996 (14.819)	35.389 (28.251)	24.429 (23.490)	50.974 (40.143)	37.227 (34.069)	74.386 (60.078)	48.098 (45.535)
Research patentability flow	0.024 (0.028)	0.016 (0.050)	0.043 (0.133)	0.023 (0.047)	0.037 (0.032)	0.027 (0.032)	0.047 (0.048)	0.032 (0.041)	0.037 (0.029)	0.036 (0.038)
Research patentability stock	0.078 (0.074)	0.052 (0.091)	0.178 (0.183)	0.113 (0.142)	0.230 (0.181)	0.157 (0.130)	0.289 (0.166)	0.209 (0.144)	0.293 (0.122)	0.245 (0.175)
Has collaboration ties with firm scientists	0.156 (0.365)	0.088 (0.283)	0.355 (0.480)	0.237 (0.425)	0.528 (0.500)	0.357 (0.479)	0.693 (0.462)	0.463 (0.499)	0.772 (0.421)	0.520 (0.500)
Identified co-authors have patents	0.042 (0.201)	0.013 (0.114)	0.072 (0.260)	0.030 (0.170)	0.120 (0.326)	0.046 (0.210)	0.136 (0.344)	0.084 (0.277)	0.149 (0.358)	0.108 (0.311)
Employer graduate school ranks in Top 20	0.323 (0.470)	0.264 (0.441)	0.313 (0.465)	0.219 (0.413)	0.250 (0.434)	0.200 (0.400)	0.197 (0.399)	0.181 (0.385)	0.175 (0.382)	0.170 (0.376)
Employer has TTO	0.531 (0.502)	0.384 (0.486)	0.620 (0.487)	0.486 (0.500)	0.694 (0.462)	0.595 (0.491)	0.719 (0.450)	0.688 (0.463)	0.825 (0.382)	0.738 (0.440)
Employer patent stock	107.4 (206.8)	53.6 (136.7)	159.4 (307.3)	64.6 (133.7)	143.0 (224.1)	75.9 (116.4)	134.4 (185.1)	110.2 (155.1)	172.3 (238.6)	120.8 (163.7)
Observations	96	3610	166	2429	216	1621	228	1072	114	519

This table reports the mean research and employer characteristics measured at five different stages in scientists' career: the 5th, 10th, 15th, 20th and 25th year after the scientist was granted a Ph.D. Within each career stage, the table is further broken out by whether a scientist has ever applied for a patent (e.g., if a scientist applied for a patent between the 16th and 20th year after he was granted a Ph.D., he contributed to the mean values of the "no" category of experience: 5, 10 and 15, and to the mean values of the "yes" category of experience: 20 and 25).

^a Scientist has at least one patent application.




Fig. 3. Distribution of patenting activity over time. This figure plots the number of patent applications filed in each year and the proportion of scientists who have filed for one or more patent applications in each year.




Fig. 4. Unconditional hazard of first patent application by Ph.D. cohort. This figure plots the kernel-smoothed estimate of the unconditional hazard of first patent application for three cohorts of scientists.

Fig. 4 displays, for the first decade of scientists’ careers, the unconditional hazard of first patent application against experience (as measured by years since graduation) for three non-overlapping cohorts of scientists: those who received their Ph.D. between 1967 and 1975, those who earned their degree between 1976 and 1985, and those who matriculated between 1986 and 1990. The figure is suggestive. It shows that over successive cohorts, the probability of patenting in an early career stage has increased, and in the latest cohort of life scientists in our data, it is increasing at a greater rate.¹⁵

¹⁵ The decline in the unconditional hazard for the third cohort after the 5th year of experience is caused by the gradual censoring of the patent data. Specifically, the NBER patent database contains data on patents issued until 1999. Because




Fig. 5. Patent co-inventorship patterns over professional experience. A “solo” inventor patent is one that lists a single inventor. A “three or more inventor” patent is one that lists at least three inventors. Proportions represent the percent of all first-time patents applied for by academic scientists in a given post-Ph.D. experience year that belong to either of these two categories.

The observed increase in early career patenting is likely caused by multiple factors. One possible explanation for the greater incidence of patenting among early career scientists is that in recent years, post-doctoral fellows are more likely to be listed as co-inventors on patents emanating from the research conducted in their advisors’ labs. Examining data from the 1970s and 1980s, [Stephan and Ma \(2005\)](#) report that there has been an increase during this period both in the proportion of scientists who begin their careers as post-doctoral fellows and in the duration of these fellowships. We cannot explore this possibility directly because our data neither allow us to distinguish post-docs from regular faculty, nor do they identify advisor/post-doc pairings. However, we are able to document general trends in the incidence of patent co-inventorship.

[Fig. 5](#) presents the proportion of all first-time patents in our data that list (i) a sole inventor (dashed line) or (ii) three or more inventors (solid line), plotted against the number of years since the patenting academic scientist received his or her Ph.D. degree. The figure demonstrates a clear negative trend in scientists’ proclivities to receive sole invented patents over their careers and a slightly positive trend in the incidence of multiple-inventor patents over the career. These data alone do not permit us to rule out firmly the possibility that early career patenting is somehow associated with changes in the profession-wide duration and prevalence of post-doctoral fellowships, but it is evidently the case that the life-cycle trend is from sole to multiple-inventor patents, not vice versa.

Although other explanations are feasible, we think it likely that part of the cause of the increase in slope in the early career hazard of patenting observed in [Fig. 4](#) reflects the fact that, over the three decades spanned by our data, patenting has come to be recognized as a legitimate form of scientific output in the academic life sciences. If this is the case, in future years we may observe

our measure of patenting is dated to the time of the filing of an application for a patent that eventually issues, the final years of our data contain fewer patenting events because we do not observe patents that were applied for prior to 1999, but did not issue until after this year. This same censoring process explains the apparent decline in patenting in the final few years displayed in [Fig. 3](#).

that further increases in academic patenting will occur in the early career years as patents come to be viewed as an important means of disseminating research findings and therefore of building a professional reputation.

4.2. *Discrete-time hazard rate models*

We now present results from the discrete-time hazard rate regressions. The results can be found in Table 4. Model (1) includes the variables often thought to be associated with academic patenting, but without the paper count and the patentability variables. All models control for (unreported) Ph.D. subject areas and calendar year dummies. The results confirm the patterns that were already apparent in the descriptive statistics. We find evidence that controlling for the number of co-authors, scientists with at least one patenting co-author are more likely to patent. We caution readers against interpreting this correlation as evidence of patenting peer effects, as it could merely reflect assortative matching among scientists along some other dimension correlated with patenting. In the baseline specification, we also find an effect of co-authorship with corporate researchers on the likelihood of first patent application. At the mean of the other covariates, having co-authored with researchers in industry increases the predicted probability of patenting by 25 percent.¹⁶

The female scientist dummy variable has a sharply negative and highly statistically significant effect on the likelihood of patenting. We estimate that the odds that a women scientist will become a first-time patenter are 49 percent the odds of a male scientist (see Ding et al.).

Unlike the individual-level covariates, the impact of employer-related variables is mixed. We do find a robust effect for the patent stock of the current employer. Like the positive effect of having co-authors with patents, this result raises the possibility of peer effects, but it also could be that employment at a high-patent-intensity university reflects the (otherwise unmeasured) influence of well-functioning technology licensing office, or that it derives from a matching process by which patent-oriented scientists are more likely to receive and accept offers from patent-focused employers. In contrast to the patent activity of the current employer, we do not find an effect for the intensity of patenting at the university where the scientist earned his/her doctorate. Likewise, the effect of the presence of a technology licensing office is also negligible. This could be due to the fact that this organizational innovation diffused quite rapidly among Tier-1 universities following passage of the Bayh-Dole Act.

Model (2) adds two variables to the specification: a scientist's count of publications in year $t - 1$, and a cumulative stock of publications up to year $t - 2$. We consider these variables to represent the centerpiece of our analysis, as they help us to adjudicate between competing explanations of the association between scientific productivity and patenting. In the first and all subsequent regressions, only the flow variable is significant, suggesting that patenting is accompanied by a flurry of scientific activity. At the mean of the data, each additional research publication increases the researcher's odds of entering the patenting regime during the next year by 10 percent; a one

¹⁶ In Table 4, the hazard of patenting appears to start increasing from 5th year after Ph.D. and peak during the 9th to 15th year after Ph.D. relative to the baseline category (23rd to 29th years after Ph.D. grant). We caution the reader that this trend reflects our decision to limit the analysis to the first patenting event. Because we drop scientists from the data once they have patented, we would expect more negative duration dependence as only those scientists who have not yet patented prior to an experience interval remain in the risk set during that interval. In Table 5, when we included repeated patenting events in the fixed effects specifications, the positive effect of experience [5, 8] over the baseline experience category disappears.

Table 4
Discrete-time hazard models of patenting

	(1)	(2)	(3)	(4)
Experience [1, 4]	0.253 (0.272)	0.290 (0.277)	0.329 (0.280)	0.301 (0.278)
Experience [5, 8]	0.438 (0.245) [†]	0.450 (0.250) [†]	0.499 (0.253)*	0.453 (0.252) [†]
Experience [9, 15]	0.652 (0.233)**	0.635 (0.237)**	0.678 (0.239)**	0.637 (0.237)**
Experience [16, 22]	0.623 (0.229)**	0.604 (0.228)**	0.628 (0.228)**	0.606 (0.227)**
Female	−0.707 (0.153)**	−0.672 (0.153)**	−0.674 (0.153)**	−0.673 (0.153)**
Collaboration tie with company scientists _{t−1}	0.208 (0.106)*	0.065 (0.113)	0.075 (0.113)	0.068 (0.114)
Average number of identified co-authors per paper _{t−1}	0.456 (0.113)**	0.441 (0.114)**	0.434 (0.115)**	0.442 (0.114)**
Identified co-authors have patent _{t−1}	0.532 (0.194)**	0.341 (0.207) [†]	0.350 (0.207) [†]	0.342 (0.207) [†]
Ph.D. university graduate school in Top 20	0.103 (0.103)	0.090 (0.103)	0.091 (0.103)	0.091 (0.103)
Ph.D. university 5-year patent stock	0.001 (0.001)	0.001 (0.001)	0.001 (0.001)	0.001 (0.001)
Employer graduate school in Top 20	−0.005 (0.120)	−0.011 (0.119)	−0.009 (0.119)	−0.010 (0.119)
Employer has a TTO _{t−1}	0.043 (0.117)	0.016 (0.118)	0.016 (0.118)	0.015 (0.118)
Employer patent stock _{t−1} × 100	0.087 (0.033)**	0.086 (0.033)*	0.085 (0.033)*	0.086 (0.033)*
Research publication stock _{t−2}		−0.001 (0.002)		
Research publication stock _{t−3}			0.001 (0.003)	
Research publication stock _{t−3}				−0.001 (0.003)
Research publication flow _{t−1}		0.093 (0.020)**	0.104 (0.022)**	0.098 (0.023)**
Research publication flow _{t−2}			−0.030 (0.027)	−0.040 (0.028)
Research publication flow _{t−3}				0.033 (0.027)
Constant	−6.351 (0.398)**	−6.461 (0.402)**	−6.497 (0.404)**	−6.465 (0.402)**
Log-likelihood	−2563.20	−2548.48	−2547.83	−2547.08
Wald χ^2	319.33	359.47	362.09	364.97
Model d.f.	44	46	47	48
Pseudo- R^2	0.06	0.06	0.06	0.06
	(5)	(6)	(7)	(8)
Experience [1, 4]	0.342 (0.279)	0.334 (0.279)	0.355 (0.282)	0.325 (0.280)
Experience [5, 8]	0.473 (0.251) [†]	0.445 (0.251) [†]	0.437 (0.250) [†]	0.464 (0.251) [†]
Experience [9, 15]	0.642 (0.237)**	0.624 (0.237)**	0.615 (0.236)**	0.633 (0.238)**
Experience [16, 22]	0.603 (0.228)**	0.596 (0.228)**	0.591 (0.227)**	0.594 (0.228)**
Female	−0.664 (0.153)**	−0.663 (0.153)**	−0.661 (0.153)**	−0.664 (0.153)**

Collaboration tie with company scientists _{<i>t</i>-1}	0.052 (0.113)	0.045 (0.113)	0.039 (0.113)	0.049 (0.113)	0.056 (0.114)
Average number of identified co-authors per paper _{<i>t</i>-1}	0.426 (0.117)**	0.422 (0.118)**	0.427 (0.119)**	0.424 (0.118)**	0.437 (0.141)**
Identified co-authors have patent _{<i>t</i>-1}	0.306 (0.206)	0.307 (0.205)	0.307 (0.205)	0.295 (0.206)	0.291 (0.208)
Ph.D. university graduate school in Top 20	0.090 (0.103)	0.088 (0.103)	0.087 (0.103)	0.088 (0.104)	0.088 (0.104)
Ph.D. university 5-year patent stock	0.001 (0.001)	0.001 (0.001)	0.001 (0.001)	0.001 (0.001)	0.001 (0.001)
Employer graduate school in Top 20	-0.015 (0.120)	-0.016 (0.120)	-0.016 (0.119)	-0.020 (0.120)	-0.022 (0.120)
Employer has a TTO _{<i>t</i>-1}	0.009 (0.118)	0.008 (0.118)	0.007 (0.118)	0.008 (0.118)	0.012 (0.118)
Employer patent stock _{<i>t</i>-1} × 100	0.089 (0.034)**	0.090 (0.034)**	0.090 (0.034)**	0.090 (0.034)**	0.090 (0.034)**
Research publication stock _{<i>t</i>-2}	-0.001 (0.002)	-0.001 (0.002)	-0.001 (0.003)	-0.001 (0.002)	-0.001 (0.003)
Research publication flow _{<i>t</i>-1}	0.083 (0.023)**	0.082 (0.023)**	0.081 (0.023)**	0.079 (0.022)**	
Research patentability stock _{<i>t</i>-2}	0.402 (0.287)			0.356 (0.301)	0.356 (0.300)
Research patentability stock _{<i>t</i>-3}		0.116 (0.325)			
Research patentability stock _{<i>t</i>-3}			0.104 (0.341)		
Research patentability flow _{<i>t</i>-1}	1.522 (0.739)*	1.475 (0.746)*	1.496 (0.745)*		
Research patentability flow _{<i>t</i>-2}		1.212 (0.698)†	1.246 (0.698)†		
Research patentability flow _{<i>t</i>-3}			-0.140 (0.960)		
High research patentability flow _{<i>t</i>-1}				0.311 (0.113)**	0.322 (0.113)**
Publication flow _{<i>t</i>-1} × experience [1, 4]					0.005 (0.028)
Publication flow _{<i>t</i>-1} × experience [5, 8]					0.109 (0.034)**
Publication flow _{<i>t</i>-1} × experience [9, 15]					0.061 (0.032)†
Publication flow _{<i>t</i>-1} × experience [16, 22]					0.081 (0.036)*
Publication flow _{<i>t</i>-1} × experience [23, 29]					0.116 (0.055)*
Constant	-6.446 (0.414)**	-6.382 (0.419)**	-6.340 (0.423)**	-6.474 (0.404)**	-6.609 (0.439)**
Log-likelihood	-2545.98	-2545.05	-2544.77	-2543.98	-2543.26
Wald χ^2	363.98	365.27	365.42	373.44	378.62
Model d.f.	49	51	53	48	52
Pseudo- R^2	0.06	0.06	0.06	0.06	0.06

Notes: (1) Number of observations: 54,737; number of researchers: 3862; number of first patenting events: 473. (2) For all researchers in the sample, only observations on or before the year of the first patenting event or censoring have been used (i.e., for all researchers that have patented, the observations after the year of their first patent application were not used in the analysis). (3) All models control for Ph.D. subject areas and calendar year dummies. (4) Experience [23, 29] is the base category. (5) A dummy variable indicating whether the researcher has zero publication in year $t - 1$ is included in Models (5)–(7), though not reported in the table; a dummy variable indicating whether the researcher has zero publication in year $t - 2$ is included in Models (6) and (7), though not reported in the table; a dummy variable indicating whether the researcher has zero publication in year $t - 3$ is included in Model (7), though not reported in the table. (6) Robust standard errors in parentheses, clustered by scientist. (7) † significant at 10 percent; *significant at 5 percent; **significant at 1 percent.

standard deviation increase (2.38) in the flow of research publications is associated with a 25 percent increase in the likelihood of patenting relative to the baseline rate. In Models (3) and (4), we explore further the timing of this flurry by using more flexible specifications for the distributed lag of publications. In Model (3), we include the flow of publications in year $t - 2$ and the stock up to year $t - 3$. In Model (4), we include the flow of publications in year $t - 3$ and the stock up to year $t - 4$. In both cases, only the coefficient for the 1-year lagged variable is significant; in other words, a 1-year lag appears to capture accurately the timing of the publication flurry associated with patenting.

This conditional correlation strikes us as being an important finding, for it informs the alternative interpretations of the association between scientific productivity and involvement with the world of commerce. In the first interpretation, commercialization activities correspond to attempts by scientists to monetize established reputations and professional status. This view is compatible with a life-cycle model of the decision to patent in which the activity is pursued by highly productive scientists to extract additional value from their reputations for scientific eminence. In the second interpretation, publications and patents are co-occurring outputs that often encode the same set of scientific insights; patents, just like publications, reflect genuine shocks to scientific opportunities. The occurrence of a patenting event is then most likely to be associated with the unearthing of a fruitful new avenue of scientific inquiry. We see the correlation between the onset of patenting and the lagged flow, but not the stock, of publications as much more consistent with the latter interpretation.¹⁷ The plausibility of this interpretation is reinforced by a peculiar aspect of U.S. patent law which grants inventors a 1-year grace period from the date of publication for the filing of a patent application (Merges, 1997, 226). In other words, an academic inventor wishing to maximize the effective life of a patent would apply to the USPTO exactly 364 days after the date of publication, provided that he/she is willing to forego patent protection in foreign jurisdictions.

Using the specification in Model (2) as a benchmark, Models (5) through (8) examine the influence of the latent patentability of the scientist's research on his/her propensity to enter the patenting regime. We perform the analysis to parallel the approach taken in previous models. Model (5) adds the flow of our research patentability score in year $t - 1$ (i.e., based on our endogenous-to-the-sample measure, the extent to which the papers a scientist has published in the previous year are substantively similar to the work previously published by patenting scientists) and the corresponding cumulative stock up to year $t - 2$. At the mean of the data, increasing the patentability score by one standard deviation raises the likelihood of first patent application by 8 percent. Moreover, as can be seen in Models (6) through (8), the conclusion is not substantially altered when using a more flexible functional form to model the distributed lag of the latent patentability score.¹⁸

We draw two conclusions from these results. First, the strong effect of the latent patentability measure suggests that it is indeed possible to use bibliometric data gleaned from journal articles to proxy for inter-person (or intra-person, inter-field, inter-university, etc.) differences in the level of commercial orientation of scientific research. Second, here again, we find that only the flow

¹⁷ This interpretation is also consistent with Murray and Stern's (2005) analysis of paper–patent pairs, but it suggests that this phenomenon is not confined to the single journal whose articles they analyze. Of course, since we do not examine the actual content of patents and papers, we can only provide circumstantial evidence in favor of a substantive linkage between these two forms of output. In practice, it seems likely that patentable claims will be spread over a number of papers revolving around a common theme, some published before, some after the filing of the patent application.

¹⁸ Specifically, in Model (8), we replace the research patentability flow at $t - 1$ with a dummy for observations that lie above the 75th percentile of the research patentability flow variable.

influences the likelihood of patenting. Just as in the case of publications, the onset of patenting appears simultaneous with a change in the content of a scientist's research in a direction that makes it more similar to that of scientists who have already applied for patent rights. However, because it is the flow, and not the stock of this measure that seems to matter, the evidence is consistent with the idea that a patent application does not constitute merely a response to changes in the formal and informal incentives faced by academic scientists over their careers, but also reflects the seizing of opportunities along a novel research trajectory.

Using Model (8) as the benchmark, Model (9) examines whether the pre-application flurry of publications varies in magnitude over the life cycle. We find no evidence of a flurry during the first 4 years of the experience clock. For life scientists, this period typically corresponds to post-doctoral fellowships, before obtaining a first job as an established, independent investigator. In subsequent years, the magnitude of the flurry is in fact quite stable.

In summary, we find that individual rates of patenting respond to scientific opportunities and that patenting coincides with a genuine change in the content of these scientists' research. In contrast, evidence regarding institutional influences on scientists' propensity to patent is mixed. We do not find evidence of "imprinting" of the commercial norms of the university in which a scientist earned his/her Ph.D., but we do find that the current employer's patent stock exerts a positive influence on individual patenting rates. The presence of a TTO does not appear to have a strong effect at the individual-level of analysis, although this variable is of course highly correlated with the university patent stock. Finally, we do find evidence that having patenting co-authors raises the odds of patenting for individual researchers.

4.3. *Fixed-effects logit models*

The results presented above suffer from two limitations. First, they only pertain to the decision to apply for the *first* patent. For a sizable proportion of scientists, patenting is a repeated event, and the determinants of patenting could differ in the group of serial patenters. Moreover, one might object that our result regarding the flurry of publications contemporaneous with patenting assumes that the lagged stock of publications adequately captures differences in talent among scientists. It would be desirable to subject this set of results to a more stringent test. For these reasons, Table 5 replicates Models (1), (4), (7) and (8) in Table 4 using fixed-effects logit models. In these models, patenting is treated as a repeated event, and there are as many observations in the estimation sample as there are person-years for patenting scientists.¹⁹

Model (2) in Table 5 shows that the impact of the 1-year lagged count of publications remains even after accounting for time-invariant talent differences among scientists through fixed individual effects and that the inclusion of additional lags does not modify the result. We interpret this finding as suggesting that within-scientist changes in scientific opportunities influence their likelihood of patenting. Similarly, Model (3) highlights the role of changes in the latent patentability of a scientist's research, which again appear to correlate with patenting events. As in the case of the discrete hazard models, we find an effect only in the year immediately preceding the year of patent application. In Model (4), we summarize the patentability measure with a dummy variable set to one for observations above the 75th percentile of the continuous research patentability variable, which yields similar results. The statistical significance of these effects is weaker than

¹⁹ We also drop the stock variables from the specifications, since they move too slowly to be separately identified from the individual effects.

Table 5
Fixed-effect logit models of probability of patenting

	(1)	(2)	(3)	(4)
Experience [1, 4]	0.060 (0.446)	0.053 (0.446)	0.032 (0.448)	0.061 (0.446)
Experience [5, 8]	0.315 (0.374)	0.295 (0.374)	0.278 (0.376)	0.296 (0.374)
Experience [9, 15]	0.561 (0.280)*	0.524 (0.280)†	0.513 (0.282)†	0.523 (0.280)†
Experience [16, 22]	0.535 (0.188)**	0.500 (0.189)**	0.490 (0.190)**	0.499 (0.189)**
Collaboration tie with company scientists _{t-1}	-0.106 (0.133)	-0.120 (0.135)	-0.121 (0.134)	-0.124 (0.134)
Average number of identified co-authors per paper _{t-1}	0.323 (0.197)	0.303 (0.198)	0.291 (0.199)	0.309 (0.197)
Identified co-authors have patents _{t-1}	0.430 (0.208)*	0.392 (0.209)†	0.379 (0.209)†	0.371 (0.209)†
Employer graduate school in Top 20	-0.195 (0.205)	-0.175 (0.206)	-0.179 (0.206)	-0.173 (0.206)
Employer has a TTO _{t-1}	0.355 (0.138)**	0.358 (0.138)**	0.355 (0.138)**	0.361 (0.138)**
Employer patent stock _{t-1} × 100	0.020 (0.039)	0.019 (0.039)	0.018 (0.039)	0.017 (0.039)
Research publication flow _{t-1}		0.034 (0.017)*	0.030 (0.017)†	0.032 (0.016)*
Research publication flow _{t-2}		-0.003 (0.018)		
Research publication flow _{t-3}		-0.0001 (0.018)		
Research patentability flow _{t-1}			2.138 (1.203)†	
Research patentability flow _{t-2}			-0.164 (1.021)	
Research patentability flow _{t-3}			-0.949 (1.160)	
High research patentability flow _{t-1}				0.199 (0.113)†
Log-likelihood	-2016.65	-2014.37	-2011.87	-2012.87
Wald χ^2	580.56	585.11	590.12	588.11
Model d.f.	19	22	26	21
Pseudo- R^2	0.13	0.13	0.13	0.13

Notes: (1) Number of observations: 9312; number of researchers: 473. (2) All models control for period dummies 1975–1976, 1977–1979, 1980–1982, 1983–1985, 1986–1988, 1989–1991, 1992–1994, 1995–1997, 1998–1999; base category is 1967–1974. (3) Experience [23, 29] is the base category. (4) Dummy variables indicating whether the researcher has zero publication in year $t - 1$, $t - 2$ and $t - 3$ are included in Model (3), though not reported in the table. (5) † significant at 10 percent; *significant at 5 percent; **significant at 1 percent.

in the corresponding “cross-sectional” hazard rate models. This is not surprising since the choice of research topic by individual scientists is likely to exhibit stickiness.

5. Discussion and conclusion

The policy debate regarding interactions between industry and academia in general, and academic patenting in particular, has often taken for granted the idea that patenting represents a fundamental departure from the norms of the “Republic of Science”. According to this view, academic researchers toil in relative obscurity by producing fundamental knowledge up until the time they receive tenure; subsequently, they may monetize their reputation by involving themselves in commercial pursuits. Patents, though not necessarily remunerative in and of themselves, provide academic researchers with visibility and status in the world of commerce, for example by contributing to the likelihood that they are invited to sit on corporate advisory boards (Stuart and Ding).

The findings in this paper challenge the standard account. First and foremost, our results suggest that patents and publications correspond to two types of output that have more in common than previously believed. Indeed, the positive relationship between patent applications and the flow, but not the stock, of publications suggest that patents and papers encode similar pieces of knowledge, a fact exploited by Murray and Stern in their investigation of the anti-commons hypothesis. Second, our results suggest that the academic incentive system may be evolving in ways that accommodate deviations from traditional scientific norms of openness. Many patenting events in our data take place in the early years of scientists’ careers, and the slope of the patent–experience curve has become steeper with more recent cohorts of scientists. This finding dovetails with qualitative accounts that emphasize that patents are becoming *de rigueur* on academic vitas in many institutions and are even considered legitimate forms of research output in promotion decisions.

If the present paper enhances our understanding of the antecedents of academic patenting, much work is yet to be done to comprehend the effects of this now-prevalent practice on the rate of scientific progress. Does applied research (as embodied in patents) crowd-out the fundamental pursuit of knowledge (as measured by publications)? Answering this question is difficult because patenting is a choice variable for scientists and the outcome of a decision that could easily reflect expectations of future scientific productivity. One perspective on the regressions in this paper is that they represent analyses of the selection process whose estimation is necessary to recover causal effects of patenting on scientific output (Azoulay et al.). But our results also alert us to the possibility that the substantive content of post-patent publications might be different from these scientists’ pre-patent output, leading naturally to the study of the effect of patenting on the *direction* of scientific progress. One could envision the use of our measure of latent patentability, whose construction is a contribution of this paper, on the left-hand side of a regression equation to investigate this question. In other words, rather than use inter- or intra-scientist differences in the patentability of research to predict the onset of patenting as we have done, one could analyze how patenting (and other career events, such as promotions) affect the commercial orientation of scientists’ research.

Finally, our findings suggest that social contagion might be an important mechanism through which the practice of academic patenting diffuses among the population of life scientists. The result that scientists whose co-authors patent are more likely to patent themselves is consistent with genuine “peer effects”, but it is also consistent with assortative matching of co-authors along some other dimension correlated with patenting, such as scientific produc-

tivity. Distinguishing between these competing hypotheses remains a valuable goal for future research.

Acknowledgements

We thank Scott Stern, Iain Cockburn, and Paula Stephan for many useful comments, as well as seminar audiences at the University of Toronto, Harvard Business School, and NBER. The usual disclaimer applies.

Appendix A. Supplementary data

Supplementary data associated with this article can be found, in the online version, at doi:10.1016/j.jebo.2006.05.015.

References

- Agrawal, A.K., Henderson, R.M., 2002. Putting patents in context: exploring knowledge transfer from MIT. *Management Science* 48 (1), 44–60.
- Allison, P.D., 1982. Discrete-time methods for the analysis of event histories. In: Leinhardt, S. (Ed.), *Sociological Methodology*. Jossey-Bass, San Francisco, pp. 61–98.
- AUTM Licensing Survey, 2003. In: Stevens, A.J., Toneguzzo, F. (Eds.), *A Survey Summary of Technology Licensing (and Related) Performance for US and Canadian Academic and Non-profit Institutions, and Technology Investment Firms*.
- Azoulay, P., Ding, W.W., Stuart, T.E., 2006. The impact of academic patenting on (public) research output. NBER Working Paper #11917.
- Blumenthal, D., Campbell, E.G., Causino, N., Louis, K.S., 1996. Participation of life-science faculty in research relationships with industry. *New England Journal of Medicine* 335 (23), 1734–1739.
- Breschi, S., Lissoni, F., Montobbio, F., 2005. The scientific productivity of academic inventors. CESPRI Working Paper.
- Campbell, E.G., Clarridge, B.R., Gokhale, N.N., Birenbaum, L., Hilgartner, S., 2002. Data withholding in academic genetics—evidence from a national survey. *Journal of the American Medical Association* 287 (4), 473–480.
- Chamberlain, G., 1984. Panel data. In: Zvi Griliches, Intriligator Michael, D. (Eds.), *Handbook of Econometrics*. North-Holland, Amsterdam.
- Cox, D.R., 1972. Regression models and life tables. *Journal of the Royal Statistical Society Series B* 34 (2), 187–220.
- Di Gregorio, D., Shane, S., 2003. Why do some universities generate more start-ups than others? *Research Policy* 32 (2), 209–227.
- Ding, W.W., Murray, F., Stuart, T.E., 2006. Gender differences in patenting in the academic life sciences. *Science* 313 (5787), 665–667.
- Etzkowitz, H., 1998. The norms of entrepreneurial science: cognitive effects of the new university-industry linkages. *Research Policy* 27 (8), 823–833.
- Fabrizio, K., Diminin, A., 2005. Commercializing the laboratory: the relationship between faculty patenting and publishing. Working Paper. Emory University.
- Hall, B.H., Jaffe, A.B., Trajtenberg, M., 2001. The NBER patent citations data file: lessons, insights and methodological tools. NBER Working Paper.
- Henderson, R.M., Jaffe, A.B., Trajtenberg, M., 1998. Universities as a source of commercial technology: a detailed analysis of university patenting 1965–1988. *Review of Economics and Statistics* 80 (1), 119–127.
- Jaffe, A.B., 1986. Technological opportunities and spillover from R&D: evidence from firms' patents, profits, and market value. *American Economic Review* 76 (5), 984–1001.
- Jaffe, A.B., 1989. Real effects of academic research. *American Economic Review* 79 (5), 957–970.
- Jaffe, A.B., Lerner, J., 2001. Reinventing public R&D: patent policy and the commercialization of national laboratory technologies. *Rand Journal of Economics* 32 (1), 167–198.
- Jensen, R., Thursby, M.C., 2001. Proofs and prototypes for sale: the licensing of university inventions. *American Economic Review* 91 (1), 240–259.

- Krimsky, S., 2003. *Science in the Private Interest: Has the Lure of Profits Corrupted Biomedical Research*. Rowman & Littlefield, Lanham, MD.
- Levin, S.G., Stephan, P.E., 1991. Research productivity over the life-cycle—evidence for academic scientists. *American Economic Review* 81 (1), 114–132.
- Long, J.S., Fox, M.F., 1995. Scientific careers: universalism and particularism. *Annual Review of Sociology* 21, 45–71.
- Mansfield, E., 1995. *Innovation, Technology and the Economy: The Selected Essays of Edwin Mansfield*. Brookfield, Aldershot, UK.
- Merges, R., 1997. *Patent Law and Policy: Cases and Materials*. Michie Co., Charlottesville, VA.
- Mowery, D.C., Nelson, R.R., Sampat, B.N., Ziedonis, A.A., 2001. The growth of patenting and licensing by US universities: an assessment of the effects of the Bayh-Dole Act of 1980. *Research Policy* 30 (1), 99–119.
- Murray, F., Stern, S., 2005. Do formal intellectual property rights hinder the free flow of scientific knowledge? An empirical test of the anti-commons hypothesis. NBER Working Paper #11465.
- Myers, M.H., Hankey, B.F., Mantel, N., 1973. A logistic exponential model for use with response-time data involving regressor variables. *Biometrics* 29, 257–269.
- Owen-Smith, J., Powell, W.W., 2001. Careers and contradictions: faculty responses to the transformation of knowledge and its uses in the life sciences. *Research in the Sociology of Work* 10, 109–140.
- Stephan, P.E., 1996. The economics of science. *Journal of Economic Literature* 34 (3), 1199–1235.
- Stephan, P.E., Gurmu, S., Sumell, A.J., Black, G., 2004. Who's patenting in the university? Evidence from the survey of doctorate recipients. Working Paper.
- Stephan, P.E., Ma, J., 2005. The increased frequency and duration of the postdoctorate career stage. *American Economic Review* 95 (2), 71–75.
- Stuart, T.E., Ding, W.W., 2006. The social structural determinants of academic entrepreneurship: an analysis of university scientists' participation in commercial ventures. *American Journal of Sociology* 112 (1), 97–144.
- Thompson, W.A., 1977. On the treatment of grouped observations in life studies. *Biometrics* 33 (3), 463–470.
- Thursby, J.G., Thursby, M.C., 2002. Who is selling the ivory tower? Sources of growth in university licensing. *Management Science* 48 (1), 90–104.
- Trajtenberg, M., 2004. The name game: using inventors patent data in economic research. Presentation to the NBER Summer Institute.
- Zucker, L.G., Darby, M.R., Brewer, M.B., 1998. Intellectual human capital and the birth of U.S. biotechnology enterprises. *American Economic Review* 88 (1), 290–306.