The (Changing) Knowledge Production Function: Evidence from the MIT Department of Biology for 1970-2000

Annamaria Conti (Georgia Institute of Technology)

&

Christopher C. Liu (University of Toronto)

July 2013

Executive Summary

Considerable attention has been focused, in recent years, on the role that graduate and postdoc students play in the production of academic knowledge. Using data from the MIT Department of Biology for the period 1970 - 2000, we analyze the evolution, over time, of four fundamental aspects of their productivity: i) training duration; ii) time to a first publication; iii) productivity over the training period; and iv) collaboration with other scientists. We find four main results. First, training periods have increased for later cohorts of graduate and postdoc students. Second, recent cohorts tend to publish their first article later than the earlier cohorts. Third, they are less productive, especially when it comes to first-author publications. Finally, collaborations with other scientists, as measured by the number of coauthors on a paper, have increased. This increase is driven by collaborations with scientists outside of a student's laboratory. We interpret these results in light of the following two paradigms: the increased burden of knowledge that later generations of scientists face and the limited availability of permanent academic positions.

I. Introduction

Knowledge production is considered one of the main determinants of economic growth, as well as one of the main causes of income disparities across countries in the world (Romer, 1990). In the knowledge production function, one of the most critical inputs is knowledge produced by university researchers, who are responsible for more than 70% of all scientific articles (National Science Board, 2008).

Academic knowledge has increasingly become a collective phenomenon. Seminal studies have documented the increase in the size of scientific collaborations, with special focus on the evolution of the geographic dispersion of team members (e.g. Adams *et al.*, 2005; Wuchty *et al.*, 2007). Even though university scientists increasingly collaborate with colleagues outside of their research institutions, an important percentage of scientific research occurs within laboratories (Stephan, 2012b). These laboratories are largely populated by graduate students and postdocs, whose contributions to their laboratory's knowledge stock has been acknowledged in a number of studies (see, for instance, Stephan, 2012b; Conti *et al.*, 2013). These research trainees have coauthored an important fraction of their laboratory's papers and, moreover, have produced a considerable share of the articles published in high-end journals (Black and Stephan, 2010).

In this study we use a unique database that allows us to examine the productivity, training duration, and the collaborative behavior of graduate students and postdocs, as well as and the extent to which these aspects have evolved over time. We interpret the patterns we find in light of two paradigms: the increased burden of knowledge that successive generations of scientists face (Jones, 2009 and 2010) and the limited availability of permanent academic positions (Stephan, 1996; Freeman *et al.*, 2011).

Our data represent an extensive laboratory sample from the MIT Department of Biology, observed from 1970 to 2000. This department has been a major locus of basic and applied discoveries in the life sciences for the latter half of the 20th century. Through the timeframe of our dataset, the scientists working at the MIT Department of Biology made discoveries as varied as the molecular mechanisms underpinning recombinant DNA (e.g., the discovery of splicing and introns), cell death, aging, and the progression of cancer. This work has resulted in six Nobel Laureates and 43 members of the National Academy of Sciences between 1966 and 2000. MIT's Department of Biology has roughly doubled in size, from 27 laboratories in 1966 to 49 laboratories in the year 2000. Given this department's elite status, the findings in this paper may be difficult to extend beyond other elite North American laboratories. With this caveat in mind, we follow in the footsteps of other scholars in trading analytical depth with a focus on an elite setting, rather than speaking to the median laboratory with far more marginal contributions (Azoulay, 2010; Jones, 2010).

We collected relevant information on the graduate students and postdocs who populated these laboratories, including their publication output. For the purposes of this study, we use this information to analyze the evolution over time of four fundamental aspects of their productivity: i) training duration, ii) time to a first publication, iii) productivity over the training period, and iv) collaboration with other scientists.

We find four main results. First, training periods have increased for later cohorts of research trainees. Second, recent cohorts tend to publish their first article later than the earlier cohorts. Third, they are less productive, especially when it comes to first-author publications. Finally, collaborations with other scientists, as measured by the number of coauthors on a paper, have increased. This increase is driven by collaborations with scientists outside of a trainee's laboratory.

The remainder of this study is organized as follows. Section II describes the empirical setting. Section III presents the scientific productivity trends for graduate students and postdocs. Section IV concludes and discusses policy implications.

II. Empirical setting

For the period we study, the MIT Department of Biology generated an Annual Report, which serves as our core data source. The primary purpose of the Annual Report was to (internally) distribute information about the department's scientific activities. As a result, the report includes technical summaries of ongoing projects, as well as a list of publications over the course of the year. From 1966-1989, technical summaries were at the project level, and individuals could contribute to multiple projects. The size of the Annual Report grew in accordance with the size of the department. After the Annual Report reached 629 pages in 1987, summaries were condensed to two pages per laboratory, regardless of size. Unfortunately, starting from 2001, even the summaries ceased to be published and subsequent data have been lost to posterity.

The Annual Report documents a roster of each laboratory's members: we know the names of every individual in each laboratory, as well as the individual's personnel type (e.g., postdoc, graduate student, technician). As a result, we know the characteristics of the department, its laboratories, as well as its individual members over the course of 35 years. Figure 1 provides an example of the data available for each laboratory-year. We know of no other data source that provides as detailed a window into the organization of scientific work as this one.

< Insert Figure 1 about here>

We supplemented this departmental personnel roster with a number of other data sources. To examine scientific outputs, we (hand) collected each laboratory's paper output from Medline. We then matched each publication's author list with our personnel roster to examine the extent to which individual laboratory members contributed to scientific output. As it is exceedingly rare for laboratory members to publish scientific papers without the laboratory head (i.e., PI) as an author, we do not believe we are missing any publications.

Overall, our dataset comprises 1,494 laboratory-years, and 20,324 laboratory member-years that span 1966-2000. Within this dataset, there are 120 professors and 6,938 laboratory members who collectively produced 7,553 journal publications.

We restrict our analysis to the 1970-2000 period, as we found important differences in the way personnel categories were defined before and after 1970. We begin with a description of the laboratories, and their changes over time. We then turn our attention to examine the laboratory members, with a particular emphasis on two major types, postdocs and graduate students, who comprise more than half of our personnel roster.

Within our dataset, the average laboratory has 10 members, of which five are postdocs, three are graduate students, and two are technicians. Staff scientists are rare, but their prevalence has increased over time. As shown in Figure 2, laboratories have grown in size through the latter part of the 20th century, and this increase has been fostered by the number of postdoctoral scientists. There is no change in the number of graduate students or technicians over time, although the number of salaried staff (i.e., technicians and staff scientists) appears to have increased in the late 1990s[1]. Figure 3 presents trends in scientific output for our laboratories. As shown, the average number of articles has steadily increased over time, from an average of four articles per laboratory-year in the 1970s to six articles per laboratory-year in the 1990s. We observe a very similar trend in the number of impact factor-weighted publications.

< Insert Figure 2 about here>

< Insert Figure 3 about here>

We restrict our analysis of laboratory members to trainees – graduate students and postdocs – for the following reasons. First, these individuals make large contributions to

[1] A likely reason why the number of graduate students remained steady over the years is that university departments in the US tend to set a limit to the number of students that can enroll in a PhD program.

a PI's publication output: their purpose is to directly produce scientific publications, rather than to play a supporting role (e.g., technicians). Second, these two types of members are the most predominant personnel types within the laboratory. Together they make up more than half of the laboratory. Third, these two personnel types have been the focus of recent interest in the literature because of their contributions to knowledge and technology production (e.g. Dasgupta and David, 1994; Waldinger, 2010). Lastly, we note that these personnel types are easily and unambiguously identified, rather than more murky categories (e.g., visiting scientists).

Our sample is thus composed of 991 graduate students and 2,427 postdocs. Figures 4a and 4b provide descriptive results of their scientific output. Interestingly, a significant proportion of them (about 35%) did not publish any articles during their training period. Conditional upon having published, the mean number of papers is about three articles for both graduate students and postdocs.

< Insert Figure 4a about here>

< Insert Figure 4b about here>

III. Trends in scientific productivity of graduate students and postdocs

This section explores trends in four major dimensions of the scientific productivity of graduate students and postdocs. First, we look at training duration. Second, we investigate the timing to a first publication. Third, we examine scientific output. Finally, we explore collaboration patterns.

A. Training duration

Postdocs and graduate students are a fundamental input into a laboratory's production function. A quick look at faculty websites convinces one of the importance of their contributions, be it measured by publications, citations, or grants. While there is no doubt that both types of trainees play a large role in expanding a PI's knowledge capital, they fundamentally differ in their distance from the knowledge frontier. In fact, postdocs are closer to the knowledge frontier than graduate students (especially those in their earliest years) and require less supervision from their PIs. Typically, a PI hires postdocs expecting them to come already supplied with skills and knowledge and to work independently (Stephan, 2012b). This is reflected in the fact that postdocs tend to have shorter training periods than graduate students. Indeed, graduate students spend a fraction of their time taking classes and learning how to conduct experiments, while the learning burden for postdocs is definitely lower. Consistent with these facts, Figures 5a and 5b show that the majority of graduate students in our sample completed their training between five and seven years, while postdocs tended to spend between two and four years in a PI's laboratory[2].

< Insert Figure 5a about here>

< Insert Figure 5b about here>

After establishing that postdocs have shorter training periods than graduate students, an interesting question becomes whether the length of training has changed over time. There are at least three reasons that might lead one to think that training periods have increased in recent years. The first is that as knowledge cumulates, earlier cohorts of trainees face an increased educational burden than older ones (Jones, 2009 and 2010). Second, it is also possible that the recent cohorts of postdocs and graduate students tend to stay longer in their positions because of the increased mismatch between the supply of trainees and the availability of permanent academic positions (Stephan, 1996; Freeman *et al.*, 2011). Finally, one cannot exclude the possibility that the increased pressure on PIs to publish and apply for grants has led them to impose longer training periods on their students (Freeman *et al.*, 2011). Figure 6 shows the evolution of training periods for graduate students (in red) and postdocs (in blue) over the period 1970-1995. We exclude the last years, since students who enrolled in these years might not have completed their training by the end of our sample period. In line with previous studies[3], we find that

[2] It is possible that postdocs have worked in more than a PI's laboratory before they are offered a faculty position. However, from discussions with MIT PIs, as well as from an examination of a sample of CVs, it is evident that, at least for the period we examine, this is less the case for MIT postdocs.

[3] See, for instance, the findings by the National Research Council (1990), Tilghman (1998), Jones (2009), Jones and Weinberg (2011), and Freeman *et al.* (2011).

training periods for recent cohorts of students tend to be about one year longer than those for the earliest cohorts. The training period increases from three to about four years for postdocs, and from five to six years for graduate students.

< Insert Figure 6 about here>

To better assess the evolution of training periods over time, we estimate a Poisson regression model[4] in which we relate the training duration of graduate students and postdocs to whether these trainees had enrolled during the following periods: *i*) 1970-1979; *ii*) 1980-1989; and *iii*) 1990-1995. The distribution of students across periods is reported in Table 1.

Table 2 presents the regression results. For each category of trainee, we first include PI fixed effects (column I) to control for PI characteristics that might affect training duration. Second, we add enrollment year fixed effects (column II) to control for year-specific factors, such as economic conditions, which might affect the opportunity costs for potential trainees to invest in education, and hence their average quality (Boehm and Watzinger, 2011; Shu, 2012). We first describe the results for graduate students, and then for postdocs.

As Table 2 shows, when we include PI fixed effects, the coefficients of the dummies for whether a graduate student had enrolled during the 1980-1989 and the 1990-1995 periods have a positive coefficient, although the coefficient is statistically significant only for the first dummy. These results provide some evidence that later cohorts of students take longer to complete their PhD than earlier cohorts (cohorts who enrolled during the 1970-1979 period). In the second column, we add enrollment year fixed effects, and the significance of the coefficients declines, although the coefficient of the 1980-1989 dummy remains statistically significant. We find similar results for postdocs. The coefficients of the 1980-1989 and 1990-1995 period dummies are positive

[4] We use robust standard errors.

and statistically significant in the baseline regression with PI fixed effects, but the significance is reduced once we introduce enrollment year fixed effects.

To summarize, the results in this section suggest that training periods have increased in recent years for both graduate students and postdocs. While we cannot precisely disentangle the mechanisms behind these trends, we nevertheless believe that increasing challenges imposed on trainees, in terms of increased educational burden or reduced availability of permanent academic positions, may play an important role.

< Insert Table 1 about here>

< Insert Table 2 about here>

B. Time to first publication

In this section, we focus on the time it takes trainees to publish their first article. If one thinks of the time interval between trainee enrollment in a PI's laboratory and their first publication as the time it takes them to acquire the necessary knowledge for deriving publishable findings, then this interval can be used as a measure for their distance to the existing knowledge frontier.

We argued in the previous section that postdocs and graduate students fundamentally differ with regard to their position relative to the knowledge frontier. We also showed that training periods are shorter for postdocs than for graduate students. Consistent with this finding, we also observe that postdocs tend to publish the results of their research at a faster rate than graduate students[5]. Figure 7 shows Kaplan-Meier estimates of the time to a first publication for postdocs and graduate students. The probability of publishing a paper in each training year is higher for postdocs than for graduate students. This holds true even when we focus exclusively on *first-author*

[5] We count trainees' publications up until two years after they leave their PI's laboratory to take into account the fact that there are lags between the time trainees end a research project and the time at which the results of the project are published.

publications, which we take as a proxy for those projects to which trainees have given their greatest contribution[6].

< Insert Figure 7 about here>

Once more, we are interested in the evolution of time to a first publication over our sample period, for both graduate students and postdocs. If the knowledge burden for the more recent cohorts is larger than that for the oldest ones, then we should expect that the time it takes to publish a first article is the longest for graduate students than for postdocs. Another reason to expect such a trend is that, over time, the population of scientists has increased disproportionately relative to the population of scientific journals, and there are grounds to suspect that it has become increasingly difficult to publish. By way of example, the percentage of papers published in the journal *Nature* fell from 11% in 1997 to 8% in 2011, despite the increase in the number of available volumes[7]. As for the journal *Science*, submissions from countries such as China, South Korea, and Turkey have increased by 46% in the past ten years, with no corresponding increment in the number of published articles (Franzoni *et al.*, 2011)[8].

Figures 8 and 9 display Kaplan-Meier estimates of the time it takes to publish a first article, distinguishing between the following periods: *i*) 1970-1979; *ii*) 1980-1989; and *iii*) 1990-2000. They provide evidence that the probability of publishing a paper at any given period is higher for the oldest cohorts than for the most recent ones.

< Insert Figure 8 about here>

< Insert Figure 9 about here>

---

6 For the sake of brevity, we do not show the results for first-author publications, but they are available upon request.

7 Data are available from http://www.nature.com/nature/authors/get_published/index.html.

8 These are countries that have implemented incentives for rewarding scientists who submit papers in high-end journals (Franzoni *et al.*, 2011).

What we need to understand is whether these trends persist once we take into account PI characteristics or other factors, such as economic conditions at the time of enrollment, which are also deemed to affect the time to a first publication. For this purpose, we estimate a series of Cox proportional hazard models in which we model the hazard to publish a first article as a function of our period dummies, as well as PI and enrollment year fixed effects. The results for graduate students are presented in Table 3, while those for postdocs are in Table 4. Standard errors are clustered around the PI.

We first focus on the results for graduate students, distinguishing between the time to a first publication and the time to a first *first-author* publication. Estimates are presented in terms of their effect on the odds of publishing a first paper: a coefficient smaller (larger) than one, reflects a negative (positive) effect. When we only include PI fixed effects, both the dummy for the 1980-1989 period and that for the 1990-2000 period have statistically significant coefficients. The magnitudes suggest that the hazard of publishing a first paper, for graduate students who enrolled in the 1980-1989 period, is 0.8 times the hazard of those who enrolled in the 1980-1989 period, and it is 0.7 times the hazard of graduate students who enrolled during 1990-2000. We obtain similar results by focusing solely on first-author publications. Once we include enrollment year fixed effects, the magnitude of the coefficients remains below one, but the coefficients are no longer significant.

< Insert Table 3 about here>

< Insert Table 4 about here>

The results for postdocs are quite similar to those for graduate students. Relative to postdocs who enrolled during 1970-1979, the hazard of publishing a first paper appears to be lower for postdocs who started in the 1980-1989 period, and it is lowest for those who started during 1990-2000. Once we introduce enrollment year fixed effects, results remain significant only for the last period dummy.

Overall, we provide some evidence that the time to a first publication has increased, and this result is strongest for trainees in the most recent decade. In general, these results seem to be consistent with the findings that training periods have increased over time and might suggest that, at least in part, recent cohorts of trainees use their extra training time to achieve first, publishable results.

C. Publication trends

In this section, we turn our attention to trends in the publication production of graduate students and postdocs. The question we want to explore is whether recent cohorts of graduate students and postdocs have become less productive than older ones. Indeed, if one posits that with the accumulation of knowledge recent cohorts of scientists face a larger learning burden (Jones, 2010), or that the mismatch between the supply of scientists (and their papers) and the availability journal space has increased over time, then we should observe a declining trend in the publication output of graduate students and postdocs, once we control for their training duration.

To investigate this hypothesis, we estimate count regression models[9] in which we relate publication outputs that graduate students and postdocs had produced during their training, as a function of whether their enrollment year falls within the 1970-1979, 1980-1989, or 1990-1995 periods. As for the analysis of training durations, we exclude the latest years because graduate students and postdocs who enrolled in these years might not have completed their training by the end of our sample period. In an initial specification, we control for training duration and PI fixed effects; we then add enrollment year fixed effects. We distinguish between the total count of trainee publications, the impact factor-weighted count, and the count of *first-author* publications. The results for graduate students are displayed in Table 5, while those for postdocs are presented in Table 6.

< Insert Table 5 about here>

< Insert Table 6 about here>

[9] We estimate Poisson models with robust standard errors.

Irrespective of the model we estimate, we find that graduate students who enrolled in the period 1990-1995 tend to be less productive than their colleagues who enrolled during 1970-1979. One might wonder whether this effect is driven by the fact that fewer graduate students are publishing in later years or that, on average, recent cohorts are publishing fewer articles. In an attempt to disentangle the two explanations, we estimate a linear probability model in which the dependent variable is an indicator that takes a value of one if graduate students had published at least one article during their training. The results are displayed in the last column of Table 5. As shown, once we control for enrollment year fixed effects, none of the period dummy coefficients appear to be statistically significant, suggesting that the most plausible explanation for our results is that recent cohorts of graduate students are publishing fewer articles.

We obtain similar results once we focus on postdocs. The most recent cohorts of postdocs appear to be less productive than the ones who enrolled during 1970-1979. However, once we control for enrollment year fixed effects, the 1990-1995 period dummy appears to have a statistically significant impact only on the number of first-author publications. When we analyze the probability of publishing at least one paper, we find this time that postdoc cohorts who enrolled during 1990-1995 have a lower probability of publishing than cohorts who enrolled during 1970-1979.

Can we conclude from these results that recent cohorts of trainees have become less productive? By looking at publication trends reported in Figures 10 and 11, we do not observe well-defined, declining trends for the yearly counts of publications and impact-factor weighted publications. This seems to suggest that the negative coefficients we found for the last period dummies capture a *temporary* decline in productivity, confined to the period 1990-1995, and not a general declining trend in these publication counts. On the contrary, for the yearly count of first-author publications, there seems to be a declining trend for both graduate students and postdocs.

< Insert Figure 10 about here>

< Insert Figure 11 about here>

In analyses we do not present here for the sake of brevity, we attempted to analyze whether the decline in the number of first-author publications was correlated with larger time intervals between publications, for publications subsequent to the first. For this purpose, we estimated hazard models for publishing a second first-author paper, conditional on having published a first, and for publishing a third first-author paper, conditional on having published a second. Because we have yearly data, we cannot analyze the time interval between two papers published in a same year. With this caveat in mind, we find that the time intervals between publications subsequent to the first are not larger for the most recent cohorts of trainees. This seems to suggest that the decline in in the number of first-author papers for these trainees could be explained by the fact that: *i*) trainees take longer to publish a first article; and *ii*) they publish fewer publications per year.

To summarize, the results from this section lead us to infer that when we measure trainee productivity by their count of first-author publications, recent cohorts of trainees appear to be less productive than older ones.

D. Collaboration trends

We have analyzed the training period and productivity trends of postdoc and graduate students in light of the challenges that recent cohorts of scientists face relative to older ones. The question remaining to be answered is whether trainees, together with the other categories of scientists, have reacted to these challenges by working in larger teams.

The benefits of teamwork have been extensively discussed by the economics literature and include, among others, output gains derived from labor specialization (Becker and Murphy, 1992), and from the circulation of new ideas among team members (Adams *et al.*, 2005). In the economics of science, scholars have found that scientists increasingly work in teams (Zuckerman and Merton, 1973; Wuchty *et al.*, 2007)[10], and

[10] See also Agrawal and Goldfarb (2008) and Forman and Van Zeebroeck (2012).

that team size has expanded over time (Adams *et al.*, 2005), largely due to an intensification of multi-university collaborations (Jones *et al.*, 2008).

Figure 13 reports trends in the average number of coauthors per paper, distinguishing between postdocs and graduate students. In line with previous studies, we observe that, for both categories of trainees, the average number of coauthors per paper has increased over time, from approximately 1.5 at the beginning of the 1970s, to approximately 3.5 by the second half of the 1990s. Interestingly enough, we also observe that the increased collaboration size was mainly driven by an increase in the number outside laboratory coauthors.

< Insert Figure 12 about here>

Overall, this suggests that trainees, similar to other categories of scientists across a broad range of disciplines, are increasingly working in teams, and that these teams increasingly encompass authors from outside the trainees' laboratories.

IV. Conclusions and policy implications

A. Summary

While knowledge production is considered one of the main determinants of economic growth, there is no doubt that academic knowledge is one of the most decisive inputs in the knowledge production function, representing by far the largest source of codified knowledge.

This study focuses on the contributions to academic knowledge by postdocs and graduate students. Using data from the MIT Department of Biology from 1970 to 2000, we look at the evolution of four fundamental aspects of their productivity: *i*) training duration, *ii*) time to a first publication, *iii*) productivity over the training period, and *iv*) collaboration with other scientists.

We identified four main results. First, training periods have increased for later cohorts of research trainees. Second, recent cohorts tend to publish their first article later than the earlier cohorts. Third, they are less productive, especially when it comes to first-author publications. Finally, collaborations with other scientists, as measured by the number of coauthors on a paper, have increased. This increase is driven by collaborations with scientists outside of a trainee's laboratory.

B. Interpreting the results

What are the mechanisms that drive our results? Our findings are consistent with Jones' educational burden story (Jones, 2009, 2010), which states that, as knowledge accumulates, future generations of scientists require a greater effort to stand on a giant's shoulders. Hence, they can either make a greater effort or specialize in a narrower field and collaborate with other scientists. Our first three results –longer training periods, longer times to publish, lower productivity for later trainee cohorts– could be interpreted as an indication that the knowledge burden has increased. The final result of trainee collaboration provides an indication that these cohorts have become increasingly specialized.

While the educational burden story is indeed a compelling explanation, we nevertheless think that other mechanisms might also be playing a role. One of these is the mismatch between the supply of trainees and the availability of post-training academic positions that scholars have observed in recent decades (Stephan, 2012a; Freeman *et al.*, 2011). Data from the NSF-NIH Survey of Graduate Students & Postdoctorates in Science and Engineering, shows that enrollment into PhD life science programs has increased by 80% between 1972 and 2005[11]. While we do not have information on the availability of post-training positions, it is nonetheless unlikely that the supply of these positions has increased at the same pace. In support of this view, Tilghman (1998) reports that life science employment in "permanent" positions in academia or, more generally, in research laboratories, declined from 87% in 1975 to 73% in 1995. These market imbalances can provide an additional explanation for the longer training periods we observe in the most

[11] Data is available from https://webcaspar.nsf.gov/.

recent cohorts who are forced to prolong their trainee status until permanent positions become available. We also should note that longer training periods certainly benefit and are encouraged by PIs. In fact, their compensation is increasingly assigned according to the rules of a tournament model in which trainee contributions have become key to making discoveries, first (Freeman *et al.*, 2011).

If market frictions were to be responsible for longer training periods, should we also expect them to explain the lower productivity of recent trainee cohort and their increased propensity to work in collaboration with other scientists? Is it plausible to think that market disequilibria last for decades? Why is the market not redirecting the excess supply of trainees to other fields?

To answer the first question, one might consider that the excess supply of scientists has led to an increase in academic journal submissions, without a corresponding increase in the number of publications. If there is an excess supply of submissions, then the direct consequence is that publishing becomes more competitive, which might explain the lower productivity of recent trainee cohorts. Moreover, specialization and collaboration become ways of dealing with market disequilibria, and one wonders whether the reduction in recent cohort productivity could have been even more accentuated had recent trainees not worked with other scientists. This mechanism is not necessarily in contrast with the educational burden explanation, rather, it offers a complementary perspective. In fact, market imbalances might act as a stimulus for scientists to expand the knowledge frontier so as to be able to publish, thus increasing the burden on future generations.

While the mechanism we have highlighted seems to be plausible, one cannot exclude the possibility that the mismatch between the supply of trainees and the availability of academic positions might have led the most brilliant students to shy away from careers in life science; thus, the increase in training periods and the lower productivity of the most recent cohorts is a reflection of their lower quality. The introduction of enrollment year dummies was meant to control for the impact of temporal economic conditions on the average quality of enrolled students, but clearly this is only an imperfect control.

To answer the second and third questions on the duration of market imbalances, we should refer to studies by Freeman *et al.* (2011) and Stephan (2012b) and mention that, increasingly, PhD programs in life science (among others) tend to be populated by foreign students. Indeed, while domestic students might be discouraged from continuing their studies in life science PhD programs, these remain attractive to foreign students not only because of their prestige (especially in the case of the MIT Department of Biology), but also because salary differentials between their country and the US are typically large. Clearly, if the average salary of a PhD holder in Italy is about 2,000 USD per month, then Italian students will be attracted by a graduate degree in the US because by the end of their studies, they will earn more than what they would have earned in their home country. To verify that the proportion of foreign graduate students at the MIT Department of Biology has increased over time, we examined the first and last names of the trainees in our sample, and codified those who had an Asian last name, as well as those with an Italian or French first *and* last name. As a result, we find that the proportion of Asian, Italian, or French students has increased from 17% in 1970 to 27% in 1995 (see Figure 12). While these figures are suggestive, they represent only a conservative estimate of the population of foreign students at the MIT Department of Biology.

There are important policy implications arising from the interpretations of our results. We will discuss them below.

C. Policy implications

Regardless of the reasons for the observed trends, it is important to note that the *costs of science* have increased (Jones, 2011). These are costs for the *individuals*, who have to endure longer training periods and greater uncertainty regarding their future prospects, as well as costs for the *society*, which cannot recuperate the returns from its investments. The natural question, therefore, becomes how to reduce these costs. As previous scholars have highlighted (Jones, 2011; Stephan 2012a), costs can be reduced by ensuring that graduate students and postdocs receive effective training. This, in turn, improves the trainee amount and quality of learning during their training periods. To increase effectiveness one could begin by alleviating teaching charges for trainees, thereby ensuring that the majority of their time is dedicated to research.

It is also very important that trainees receive adequate supervision by their PIs or other laboratory senior members. Supervision should not only encompass knowledge transfer *sensu stricto*, but also direction toward research domains that are palatable for the rest of the scientific community or toward colleagues who might bring value to a collaboration, as well as advice on future career prospects. To achieve this goal, incentives are fundamental, and PIs should start being seriously evaluated based on the placement of their students. While we increasingly see PIs including information about the careers of their students on their websites, this should become a common practice, as well as a criterion for evaluating their impact (Stephan 2012a).

Another fundamental reason for reducing the costs of science is to avoid discouraging the most brilliant students from undertaking graduate studies in life science. Building upon this comment, improving the effectiveness of training should also be accompanied by measures that guarantee trainees adequate career prospects, as well as adequate salaries during and after the training period. After all, if studying finance entails shorter training periods, better career prospects and higher salaries, why should the most brilliant minds study life science?

Expanding on trainee career prospects, it is also important for PIs and their departments to maintain solid links with the industry sector. In fact, given the increasing mismatch between the supply of graduate students and the availability of positions in academia, students graduating in life science (and the society at large) need to be offered concrete opportunities to enjoy the returns of their investment by finding positions in industry commensurate to the investment they have made. While some schools and departments are heading in this direction (Cyranoski *et al.*, 2011), much remains to be accomplished.

We conclude with a few remarks on the increased participation of foreign students in US graduate and postdoctoral programs. In our view, hiring foreign students imposes three fundamental challenges. The first is the selection of brilliant minds. While this is admittedly less of a problem in the case of the MIT Department of Biology, which typically receives applications from the best schools around the world, it becomes a serious concern for other departments that need to spend considerable time and resources in evaluating applications from less renowned schools. These costs can be alleviated by systematically gathering information on the backgrounds of foreign students and circulating it among faculty members. Second, training students is costly, but training foreign students can be even costlier. This is because foreign students often need time to learn the language of their host countries and acclimate to their cultures and customs. Hence, in order for a society to fully enjoy the returns of their investment, it is essential for foreign students, once they complete their studies, to be given the opportunity to remain in their host country and to find positions in which they can apply the knowledge they have acquired during their studies. Finally, allowing students from other countries raises the question of whether these countries should be compensated for their initial investment on these students. While previous studies have analyzed the benefits that these countries enjoy in terms of increased knowledge flows (Kerr, 2008), it would be interesting to see whether these gains outweigh the initial investment costs.

References

Adams, J. D., Black, G. C., Clemmons, J. R., & Stephan, P. E. (2005). Scientific Teams and Institutional Collaborations: Evidence from US universities, 1981–1999. *Research Policy*, 34(3), 259-285.

Agrawal, A., & Goldfarb, A. (2008). Restructuring Research: Communication Costs and the Democratization of University Innovation. *American Economic Review*, 98(4), 1578-1590.

Azoulay, P., Zivin, J. G., & Wang, J. (2010). Superstar Extinction. *Quarterly Journal of Economics*, 125(2), 549-589.

Becker, G. S., & Murphy, K. M. (1992). The Division of Labor, Coordination Costs, and Knowledge. *The Quarterly Journal of Economics*, 107(4), 1137-1160.

Black, G. C., & Stephan, P. E. (2010). The Economics of University Science and the Role of Foreign Graduate Students and Postdoctoral Scholars. In *American universities in a global market* (pp. 129-161). University of Chicago Press.

Boehm, M., & Watzinger, M. (2011). The Sectorial Allocation of Talent: Evidence from the Market for Economists. Working Paper.

Cyranoski, D., Gilbert, N., Ledford, H., Nayar, A., & Yahia, M. (2011). The PhD Factory. *Nature*, 472(7343), 276-279.

Conti, A., Denas, O., & Visentin , F. (2013). Knoweldge Specialization in PhD Students Groups. Working Paper.

Dasgupta, P., & David, P. A. (1994). Toward a New Economics of Science. *Research Policy*, 23, 487–521.

Forman, C., & Van Zeebroeck, N. (2012). From Wires to Partners: How the Internet has Fostered R&D Collaborations within Firms. *Management Science*, 58(8), 1549-1568.

Franzoni, C., Scellato, G., & Stephan, P. (2011). Changing Incentives to Publish. *Science*, 333: 702-703.

Freeman, R., Weinstein, E., Marincola, E., Rosenbaum, J., Solomon, F. (2011). Competition and Careers in Biosciences. *Science*, 294, 2293-2294.

Jones, B. F., Wuchty, S., & Uzzi, B. (2008). Multi-University Research Teams: Shifting Impact, Geography, and Stratification in Science. *Science*, 322(5905), 1259-1262.

Jones, B. F. (2009). The Burden of Knowledge and the Death of the Renaissance Man: Is Innovation Getting Harder. *Review of Economics and Statistics*, 76(1), 283-317.

Jones, B. F. (2010). Age and Great Invention. *Review of Economics and Statistics*, 92(1), 1-14.

Jones, B. F., & Weinberg, B. A. (2011). Age Dynamics in Scientific Creativity. *Proceedings of the National Academy of Sciences*, 108(47), 1-5.

Jones, B. F. (2011). As Science Evolves, How Can Science Policy?. In *Innovation Policy and the Economy, Volume 11* (pp. 103-131). University of Chicago Press.

Kerr, W. R. (2008). Ethnic Scientific Communities and International Technology Diffusion. *The Review of Economics and Statistics*, 90(3), 518-537.

National Science Board (2008). Science and Engineering Indicators. Arlington, VA: NSB.

National Science Foundation (2005). Industrial Research and Development Information System, Table H-19, www.nsf.gov/sbe/srs/iris/start.cfm.

Romer, P. M. (1990). Endogenous Technological Change. *Journal of Political Economy*, 71-102.

Shu, P. (2012). The Long-Term Impact of Business Cycles on Innovation: Evidence from the Massachusetts Institute of Technology. Working Paper.

Stephan, P. (1996). The Economics of Science. *Journal of Economic Literature*, 34(3), 1199-1235.

Stephan, P. (2012a). Perverse Incentives. *Nature*, 484, 29-31.

Stephan, P. (2012b). How Economics Shapes Science. Cambridge MA: Harvard University Press.

Tilghman, S. (1998). Trends in the Early Careers of Life Sciences. Report by the Committee on Dimensions, Causes, and Implications of Recent Trends in the Careers of Life Scientists, National Research Council. Washington DC: National Academy Press.

Waldinger, F. (2010). Quality Matters: The expulsion of Professors and the Consequences for PhD Student Outcomes in Nazi Germany. *Journal of Political Economy*, 118(4), 787-831.

Wuchty, S., Jones, B. F., & Uzzi, B. (2007). The Increasing Dominance of Teams in Production of Knowledge. *Science*, 316(5827), 1036-1039.

Figure 1: Personnel composition of Professor Baltimore's laboratory.

Professor:	David Baltimore
Visiting Scientists:	Samuel Latt and Richard Van Etten
Postdoctoral Associates:	Brygida Berse, Mark Feinberg, Michael Lenardo, Jing-Po Li, Shiv Pillai, Louis Staudt and Xiao-Hong Sun
Postdoctoral Fellows:	Raul Andino, Patrick Baeuerle, Andre Bernards, Lynn Corcoran, Sunyoung Kim, Towia Libermann, Ricardo Martinez, Mark Muesing, Cornelis Murre, Jacqueline Pierce, Stephen Smale, Didier Trono, Anna Voronova and Astar Winoto
Technical Assistants:	Ann Gifford, Carolyn Gorka, Patrick McCaw, Michael Paskind and Gabrielle Rieckhof
Graduate Students:	George Daley, Peter Jackson, Marjorie Oettinger, David Schatz and Dan Silver
Undergraduate Student:	Anna Kuang

Figure 2: Number of laboratory's personnel by type.

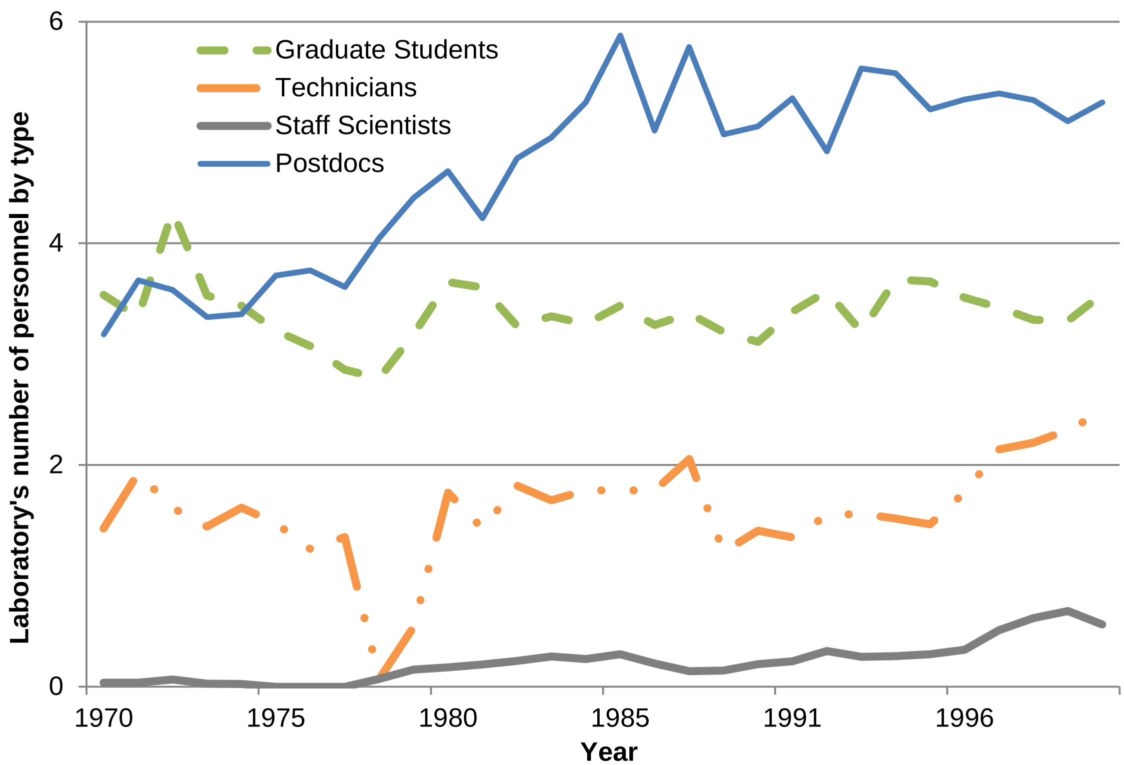

Figure 3: Number of laboratory's publications and impact factor-weighted publications

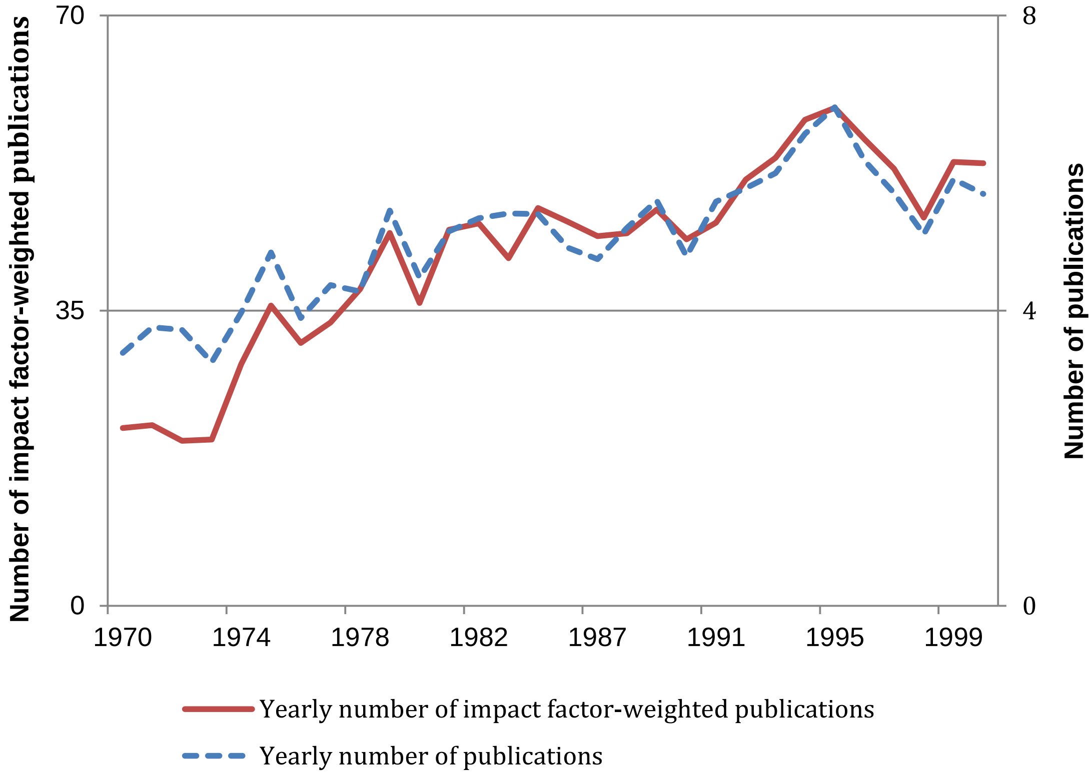

Figure 4a: Distribution of graduate students by their number of papers

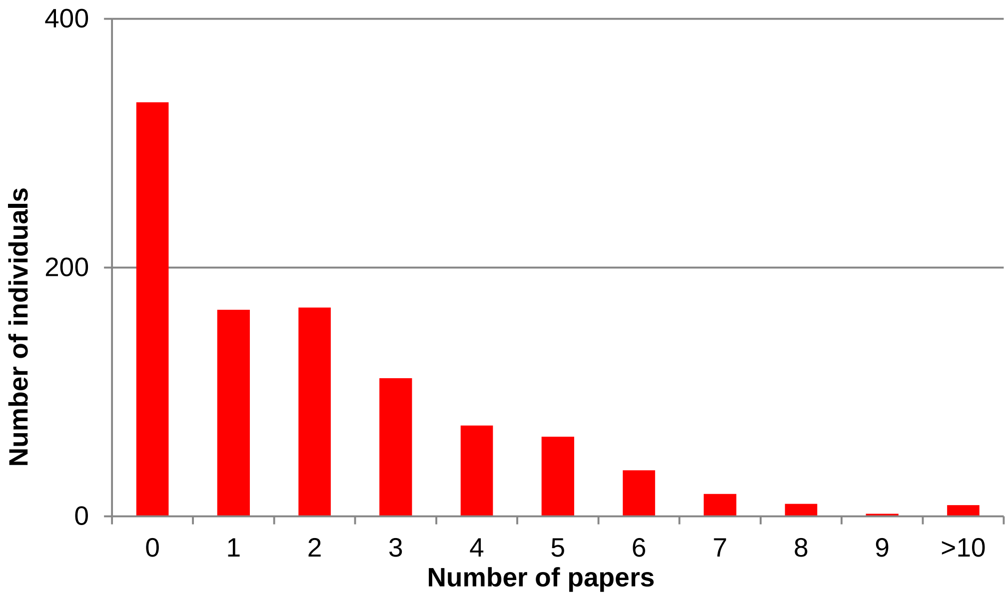

Figure 4b: Distribution of postdocs by their number of papers

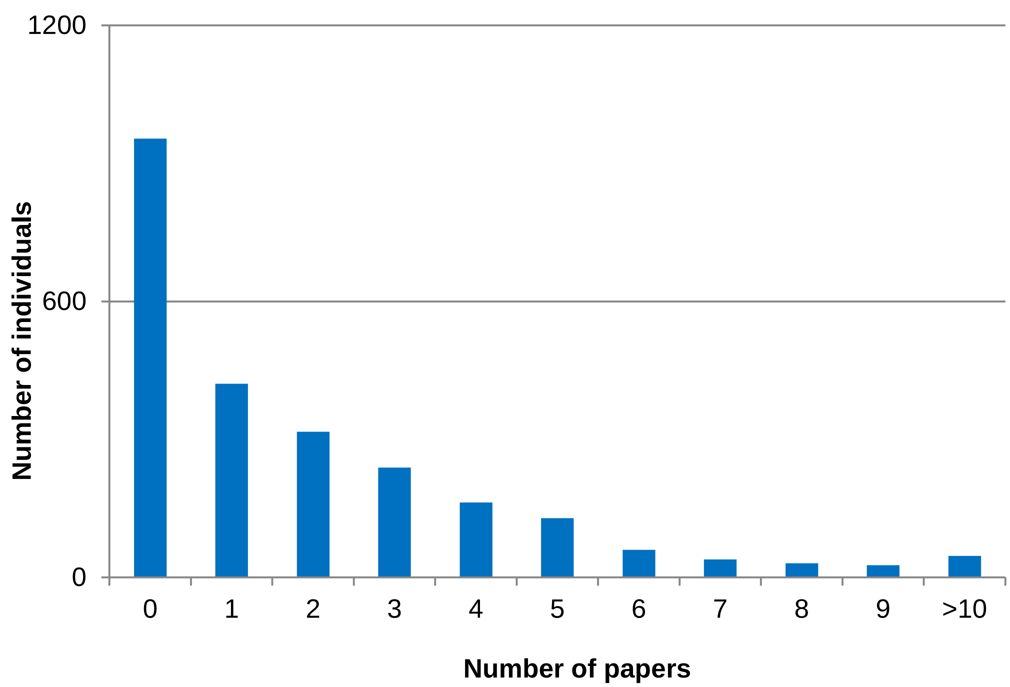

Figure 5a: Distribution of graduate students by their training duration

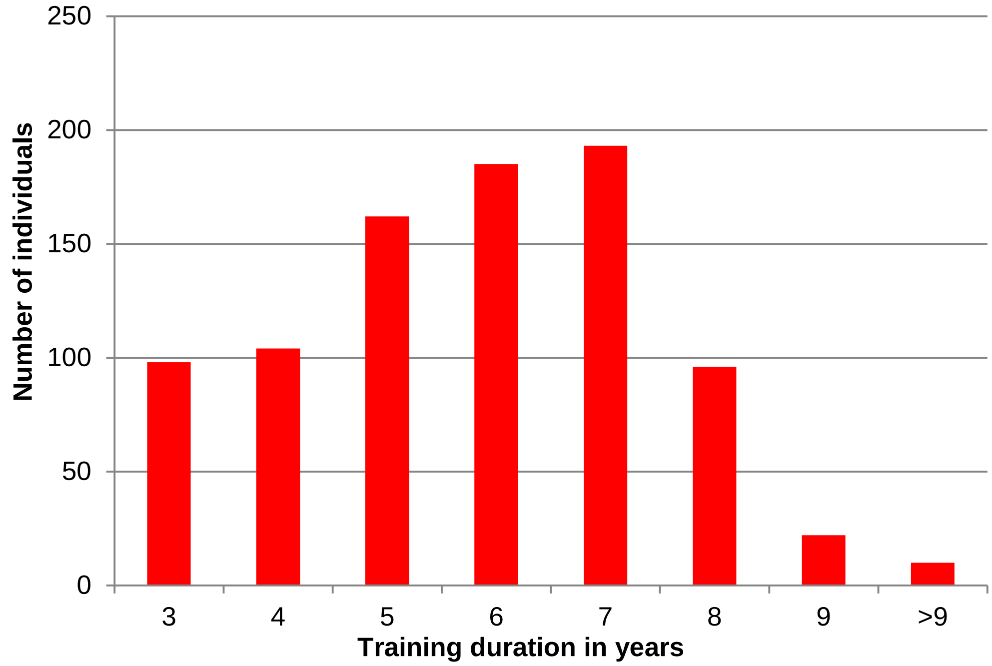

Figure 5b: Distribution of postdocs by their training duration

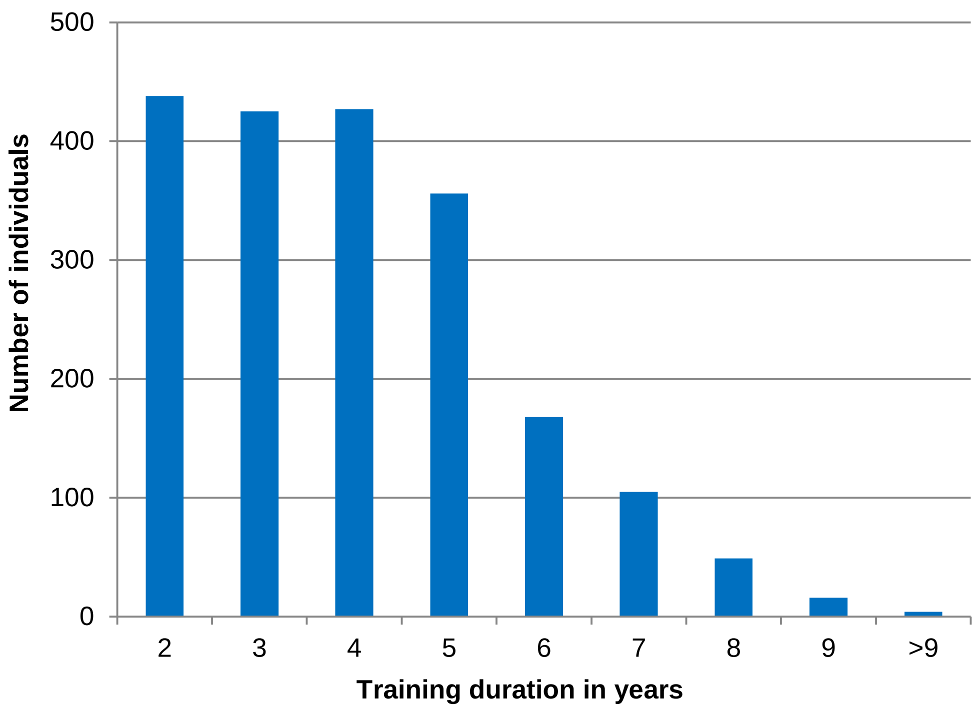

Figure 6: Training duration for graduate students and postdocs over time

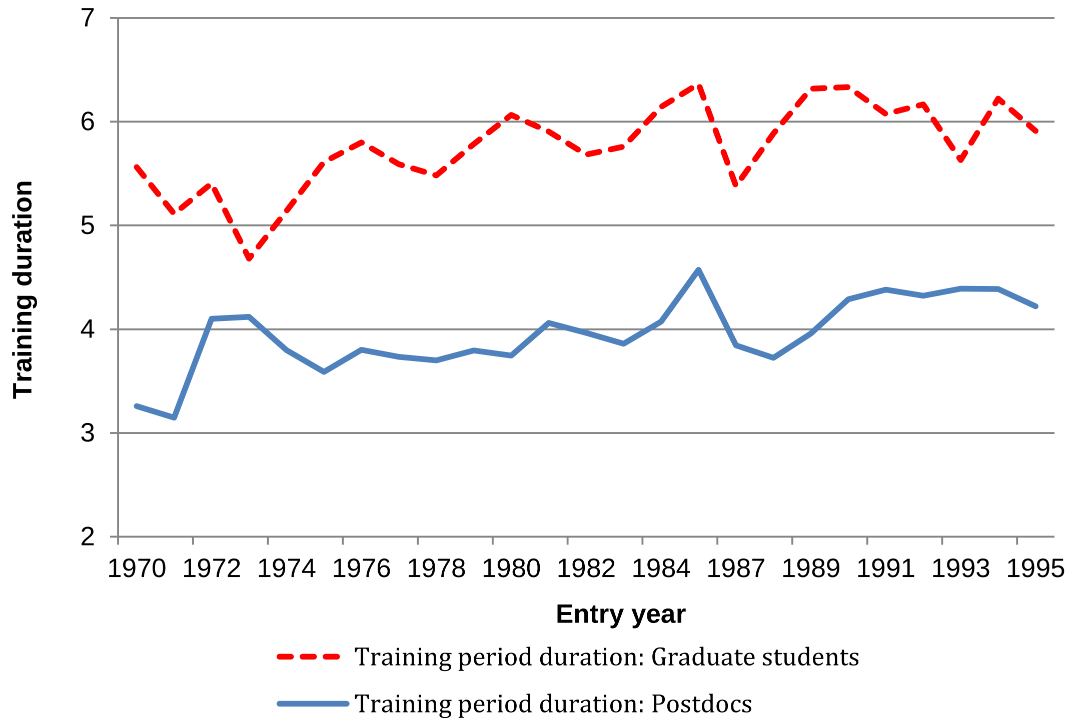

Table 1: Distribution of graduate students and postdocs by enrollment period

	Graduate students	Postdocs
1970-1979	289	560
1980-1989	334	868
1990-1995	247	565
1996-2000	121	434

Table 2: Regression results for graduate student and postdoc training duration

	Graduate students		Postdocs	
	Coeff.	*Coeff.*	*Coeff.*	*Coeff.*
Dummy=1 for 1980-1989 period	0.075**	0.144*	0.065*	0.149
	(0.033)	(0.077)	(0.034)	(0.094)
Dummy=1 for 1990-1995 period	0.055	0.004	0.143***	0.180
	(0.038)	(0.077)	(0.004)	(0.096)
PI FE	YES	YES	YES	YES
Entry Year FE		YES		YES
R^2	0.04	0.04	0.03	0.04
N	870		1993	

Note: We estimated Poisson models. Robust standard errors are in parentheses. *** p < 0.01, ** p < 0.05, * p < 0.10. For these analyses we only consider trainees who had enrolled before 1996.

Figure 7: Kaplan-Meier estimates of the time to a first publication: graduate students versus postdocs

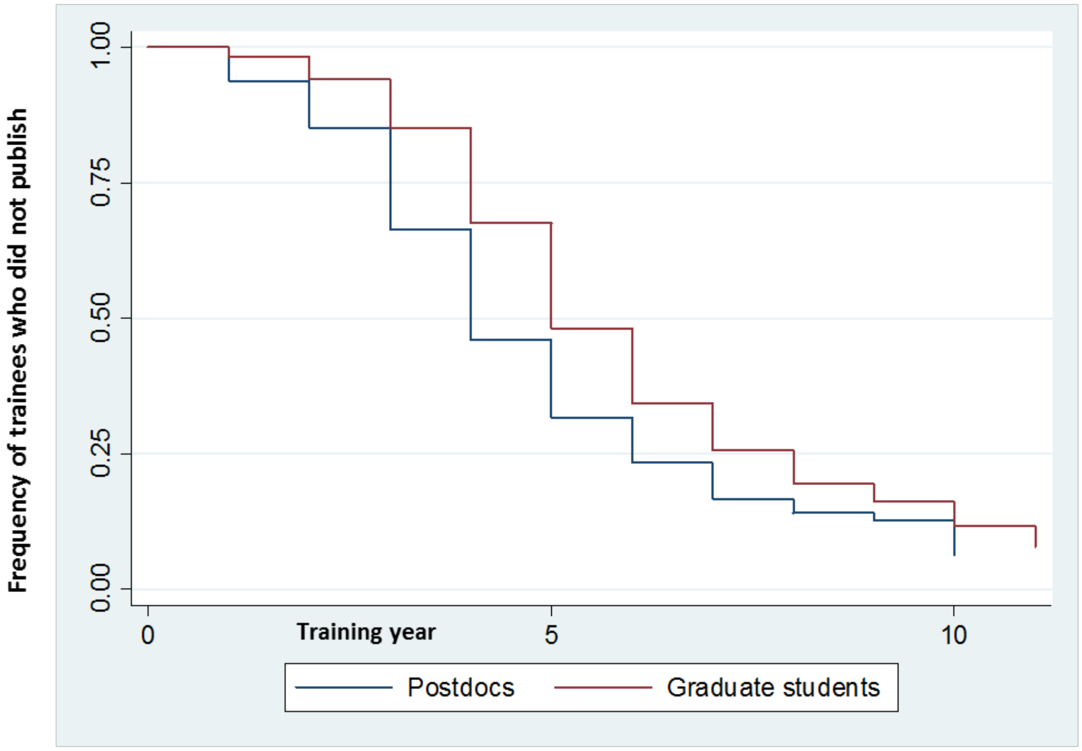

Figure 8: Kaplan-Meier estimates of the time to a first publication: graduate students over time

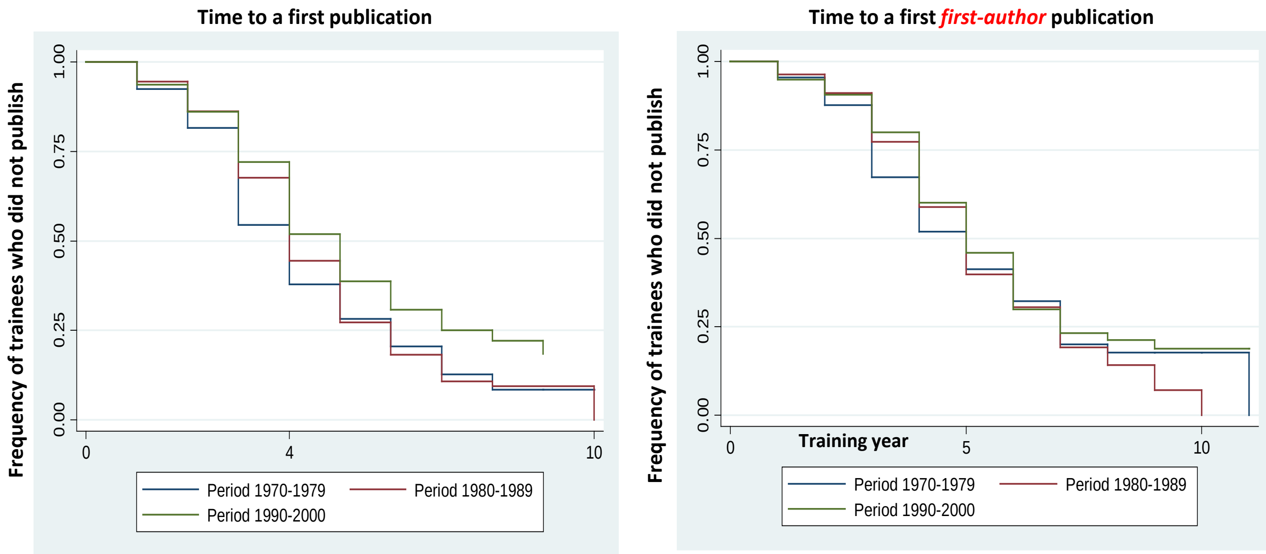

Figure 9: Kaplan-Meier estimates of the time to a first publication: postdocs over time

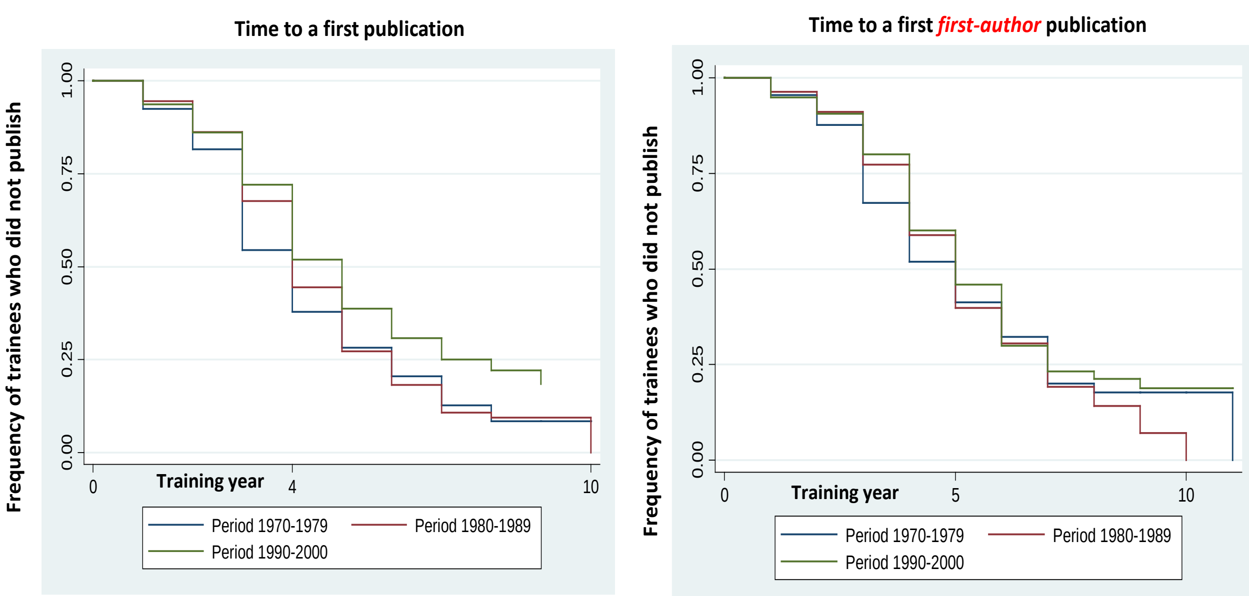

Table 3: Hazard models for the time to a first publication: graduate students over time

	Any publication		First-author publications	
	Hazard Ratios	*Hazard Ratios*	*Hazard Ratios*	*Hazard Ratios*
Dummy=1 for 1980-1989 period	0.768**	0.736	0.718***	0.564
	(0.094)	(0.247)	(0.090)	(0.206)
Dummy=1 for 1990-2000 period	0.650***	0.218**	0.613***	0.224
	(0.103)	(0.145)	(0.081)	(0.204)
PI FE	YES	YES	YES	YES
Entry Year FE		YES		YES
Log likelihood	-4043	-4025	-3380	-3365
N	991			

Note: We estimate Cox proportional hazards models with standard errors clustered around PI. We report hazard ratios. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Table 4: Hazard models for the time to a first publication: postdocs over time

	Any publication		First-author publications	
	Hazard Ratios	*Hazard Ratios*	*Hazard Ratios*	*Hazard Ratios*
Dummy=1 for 1980-1989 period	0.788***	0.876	0.795**	0.838
	(0.055)	(0.259)	(0.075)	(0.290)
Dummy=1 for 1990-2000 period	0.615***	0.197***	0.602***	0.261**
	(0.061)	(0.104)	(0.062)	(0.153)
PI FE	YES	YES	YES	YES
Entry Year FE		YES		YES
Log likelihood	-10479	-10444	-8517	-8488
N	2427			

Note: We estimate Cox proportional hazards models with standard errors clustered around PI. We report hazard ratios. *** p < 0.01, ** p < 0.05, * p < 0.10.

Table 5: Regression results for graduate student publications

	# Publications		# Weighted publications		# First-author publications		Probability of publishing	
	Coeff.	*Coeff.*	*Coeff.*	*Coeff.*	*Coeff.*	*Coeff.*	*Coeff.*	*Coeff.*
Dummy=1 for 1980-1989	-0.103	0.137	-0.142	-0.016	-0.241**	0.007	-0.036	0.037
	(0.103)	(0.208)	(0.118)	(0.259)	(0.104)	(0.194)	(0.043)	(0.125)
Dummy=1 for 1990-1995	-0.257**	-0.454*	-0.287**	-0.671**	-0.499***	-0.901***	-0.143***	-0.087
	(0.120)	(0.267)	(0.131)	(0.263)	(0.130)	(0.246)	(0.053)	(0.120)
Duration	YES	YES	YES	YES	YES	YES	YES	YES
PI FE	YES	YES	YES	YES	YES	YES	YES	YES
Entry Year FE		YES		YES		YES		YES
R^2	0.15	0.16	0.29	0.32	0.12	0.13	0.28	0.30
N	870							

Note: Standard errors are in parentheses. For the Poisson models we use robust standard errors, while for the linear probability model we cluster standard errors around PI. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$. For these analyses we only consider trainees who had enrolled before 1996.

Table 6: Regression results for postdoc publications

	# Publications		# Weighted publications		# First-author publications		Probability of publishing	
	Coeff.	*Coeff.*	*Coeff.*	*Coeff.*	*Coeff.*	*Coeff.*	*Coeff.*	*Coeff.*
Dummy=1 for 1980-1989	-0.250***	-0.274	-0.167**	-0.166	-0.255***	-0.382**	-0.036	0.016
	(0.071)	(0.218)	(0.080)	(0.258)	(0.074)	(0.228)	(0.027)	(0.078)
Dummy=1 for 1990-1995	-0.314***	-0.264	-0.251***	-0.167	-0.384***	-0.448**	-0.078**	-0.138*
	(0.086)	(0.240)	(0.096)	(0.279)	(0.089)	(0.228)	(0.033)	(0.081)
Duration	YES	YES	YES	YES	YES	YES	YES	YES
PI FE	YES	YES	YES	YES	YES	YES	YES	YES
Entry Year FE		YES		YES		YES		YES
R^2	0.18	0.19	0.30	0.32	0.14	0.14	0.22	0.23
N	1993							

Note: Standard errors are in parentheses. For the Poisson models we use robust standard errors, while for the linear probability model we cluster standard errors around PI. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$. For these analyses we only consider trainees who had enrolled before 1996.

Figure 10: Publication output of graduate student cohorts

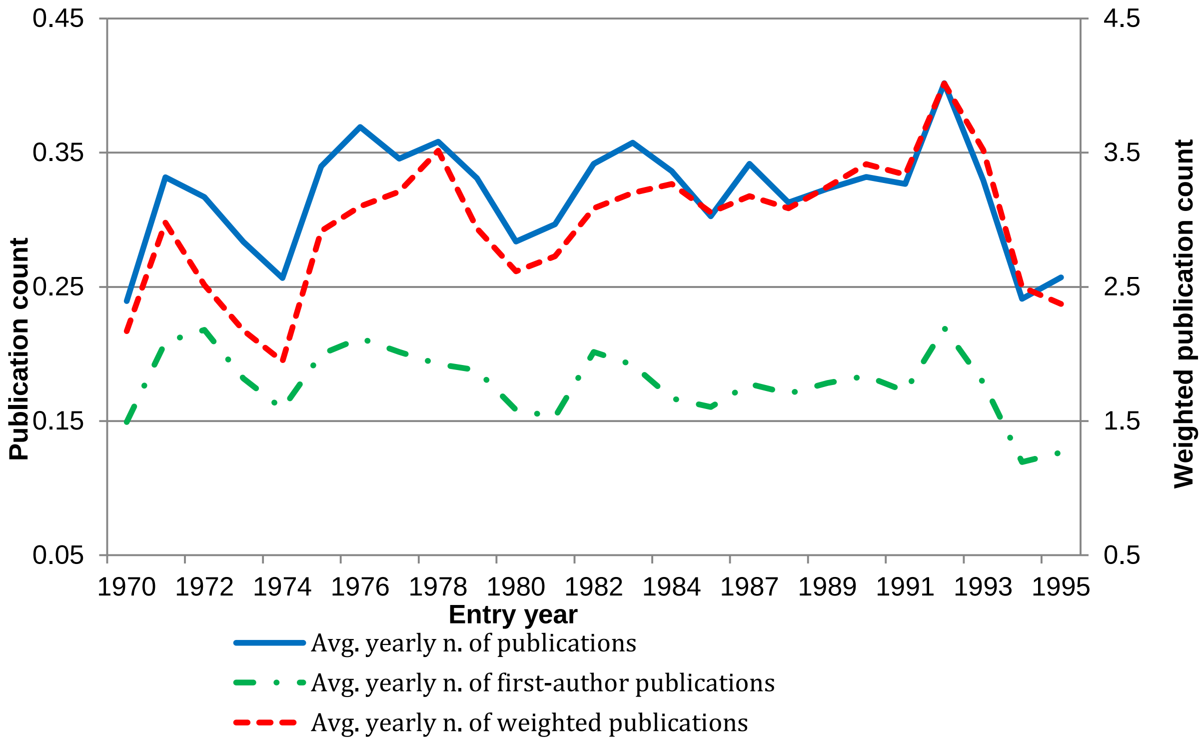

Note: Counts normalized by duration

Figure 11: Publication output of postdoc cohorts

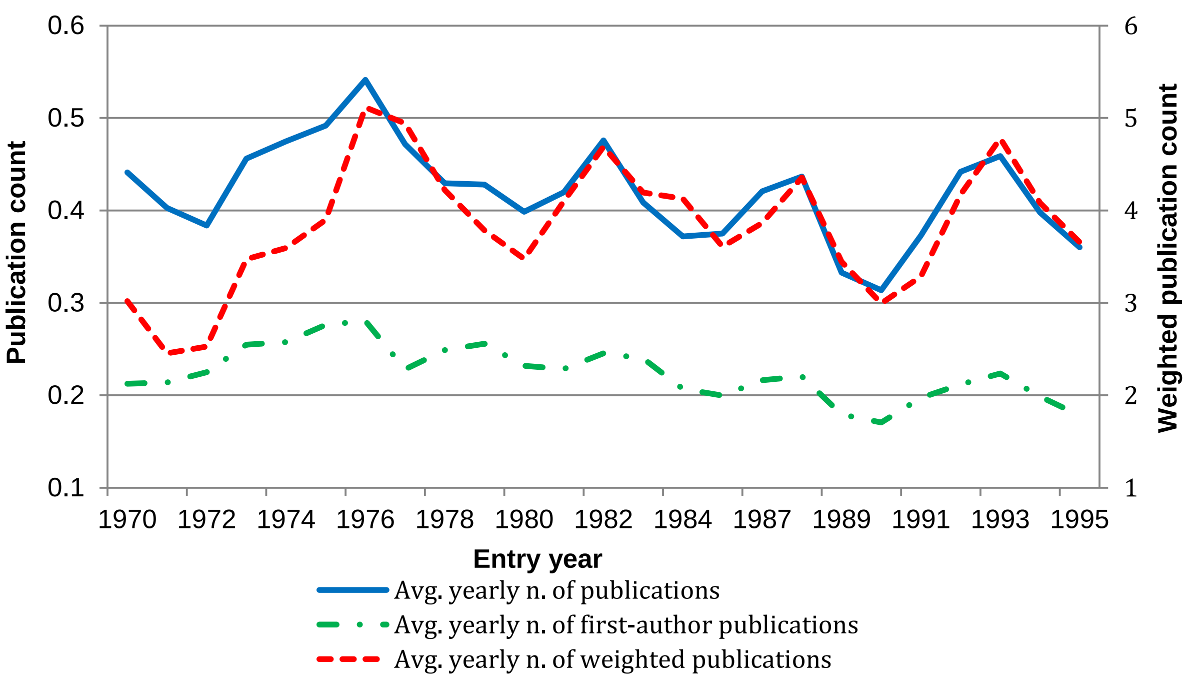

Note: Counts normalized by duration

Figure 12: Average yearly number of coauthors per paper

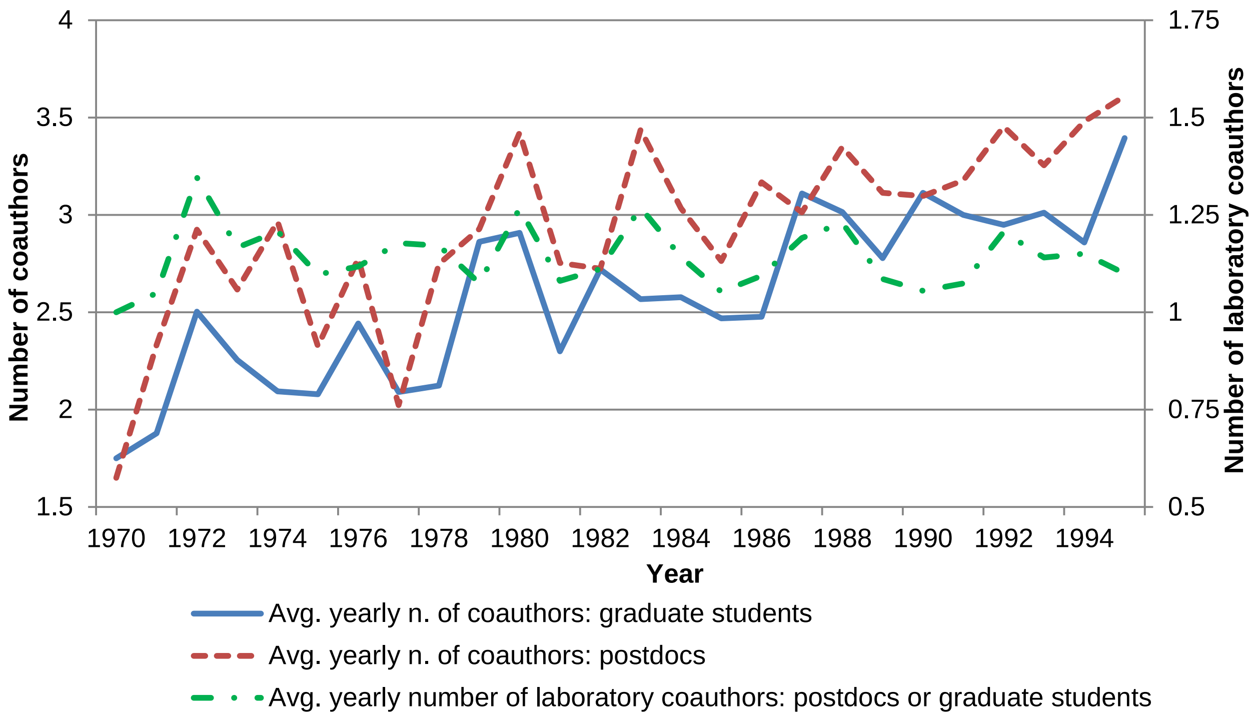

Figure 13: Share of foreign students over time

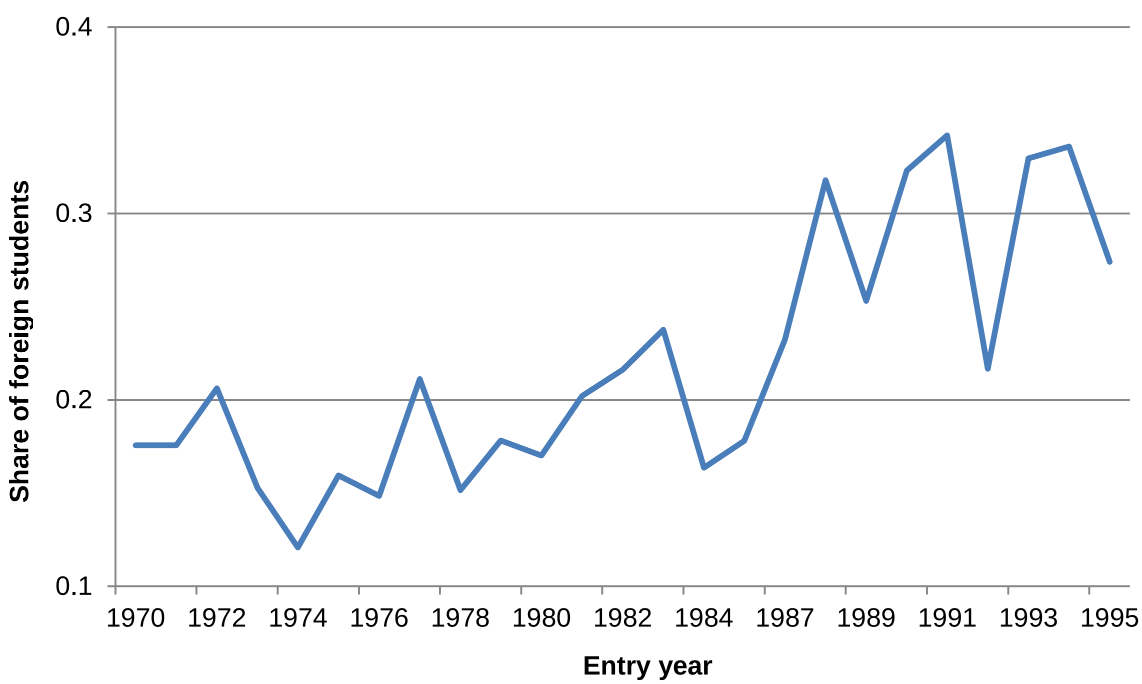

Note: We restrict the sample to foreign students with Asian last names, and to foreign students with French and Italian first *and* last names.