

Dissertation Paths: Advisors and Students in the Economics Research Production Function

Joshua Angrist, Marc Diederichs

JANUARY 2025

Dissertation Paths: Advisors and Students in the Economics Research Production Function*

Joshua Angrist[†] Marc Diederichs[‡]

December 2024

Abstract

Elite economics PhD programs aim to train graduate students for a lifetime of academic research. This paper asks how advising affects graduate students' post-PhD research productivity. Advising is highly concentrated: at the eight highly-selective schools in our study, a minority of advisors do most of the advising work. We quantify advisor attributes such as an advisor's own research output and aspects of the advising relationship like coauthoring and research field affinity that might contribute to student research success. Students advised by research-active, prolific advisors tend to publish more, while coauthoring has no effect. Student-advisor research affinity also predicts student success. But a school-level aggregate production function provides much weaker evidence of causal effects, suggesting that successful advisors attract students likely to succeed—without necessarily boosting their students' chances of success. Evidence for causal effects is strongest for a measure of advisors' own research output. Aggregate student research output appears to scale linearly with graduate student enrollment, with no evidence of negative class-size effects. An analysis of gender differences in research output shows male and female graduate students to be equally productive in the first few years post-PhD, but female productivity peaks sooner than male productivity.

*We thank our exceptionally committed and hard-working research assistants, Alison Fang, Merilin Martinson, Isabel Munoz, and Tianyuan (Margaret) Zheng for two years of painstaking work. This essay is dedicated to the memory of Oberlin College Professor Hirschel Kasper, a pioneer in the study of economics education and undergraduate advisor to Josh Angrist. Kasper's report on economics education was published as Kasper et al. (1991).

[†]MIT and NBER. Email: angrist@mit.edu

[‡]ROCKWOOL Foundation Berlin and University of Passau. Email: marc.diederichs@uni-passau.de

1 Introduction

The most selective doctoral programs in economics promise to teach their students to write and publish journal articles reporting on their research. Although many economics PhD students land in non-research positions in consulting, finance, or government, elite program curriculum is research-oriented. This can be seen in the programs' lengthy reading lists, demanding research exercises such as second- and third-year papers, and the time students are expected to invest in their job market papers and thesis chapters. Stratospherically—and increasingly—selective economics PhD programs target bright and ambitious students who appear committed to and well-prepared for a career of academic economics research.¹

These scholarly aspirations notwithstanding, half of elite economics PhDs, graduates of Harvard, MIT, Stanford and the like, publish next to nothing in the 6 years following degree completion, while only 5-10% publish more than a paper or two (Conley and Önder, 2014).² Elite schools employ stellar faculty with lengthy, influential research careers. Surprisingly, the Conley and Önder (2014) data suggest that graduates of good-but-not elite programs (classified according to widely-used rankings), like CMU, Rochester and San Diego, publish about as well as do graduates from Harvard, MIT, and Stanford. Why do so few highly-selected elite program graduates follow the path to research success taken by their extraordinarily successful advisors? What aspects of economics advisee training might be changed or enhanced so as to boost graduate student success and total research output?

These questions motivate our study of the economics PhD education production function at elite universities. The principal production inputs in this function are the faculty who teach and advise graduate students, along with aspects of the advising relationship that faculty and students develop together. We aim to measure features of the advising relationship and to link these features with students' research success. Graduate education has many features; we focus here on those most obviously tied to research. Specifically, we look at advisors' own research success and aspects of an advisor's advising history such as the number of past advisees and the scholarship of former students. We also consider measures of research affinity such as whether an advisee's doctoral thesis cites advisor work and the extent of coauthoring between advisors and their students. Notably, we ignore teaching and graduate classes. This reflects our view that students in the elite programs we study are exposed to broadly similar levels of coursework. Advising relationships, by contrast, are highly idiosyncratic even within programs.

¹The MIT Department of Economics, long a flagship for economics graduate education, admitted 38 of 794 PhD applicants in 2014 and just 27 of 895 PhD applicants in 2024. MIT Economics' current 3% acceptance rate is below that of even the most selective Ivy League colleges.

²These statistics are for 1986-2000 PhD graduates of 30 top departments.

Our analysis starts with an ambitious data collection effort linking doctoral dissertations with advisor characteristics and measures of recent graduates' research productivity for cohorts of PhD students completing an economics degree since 1994. We focus on eight elite schools, some of which have multiple programs training economics research PhDs. These schools are, in alphabetical order, Berkeley, Chicago, Harvard, MIT, Northwestern, Princeton, Stanford, and Yale. Our linked sample contains roughly 8,000 graduates from economics departments and economics-adjacent programs like those in business and public policy schools granting economics-related PhDs. By the early 2000s, the graduates in our sample account for roughly 20% of research articles published in the 137 or so most cited journals indexed by the AEA's bibliographic database Econlit (their advisors likewise account for around 20%). The students of interest authored roughly half of articles published in top-6 economics journals in 2020.³

Our investigation of the economics graduate education production function is related to earlier analyses of economics graduate education and to studies of research training in other fields. Waldinger (2010), for instance, shows that Nazi Germany's expulsion of Jewish and politically unreliable mathematics faculty degraded PhD student success in affected departments. Corsini, Pezzoni and Visentin (2022) examines the effect of advisor characteristics on STEM PhD student success, focusing on advisor gender as a causal factor. Gaule and Piacentini (2018) and Neumark and Gardecki (1998) investigate the importance of gender matching between advisors and graduate students in the sciences and economics, respectively. Following a broader analysis of student research output, we look briefly at gender effects and interactions in the PhD student research production function as well.

Also related, Hilmer and Hilmer (2009, 2011) attempt to disentangle the effects of advisor scholarship from program prestige in a sample of economics PhDs. García-Suaza, Otero and Winkelmann (2020) does something similar using recently-available data from the [Repec](#) economics research paper repository. Building on this work, our analysis examines the role of advisor networks and considers estimates with and without control for school fixed effects. Natural sciences PhD students appear to benefit from advising and mentoring by highly visible and productive superstar scholars (Li et al., 2019). At the same time, an analysis of life sciences students suggests PhD graduates are more successful when their dissertation research synthesizes work from areas an advisor or mentor's field (Liénard et al., 2018). We likewise aim to assess the importance of research affinity for economics PhDs.

Our analysis is distinguished from related earlier work by our large recent sample and by

³Journal lists come from Angrist et al. (2020), which identifies the journals most cited by the *American Economic Review*. Top-6 journals include the usual top 5 plus the *Review of Economics and Statistics*. For a sample graduating 1987-92, Collins, Cox and Stango (2000) likewise report a preponderance of top program graduates' research publications in 36 highly-cited and top-5 economics journals.

the scope of the explanatory variables considered as inputs in the economics PhD research production function.⁴ Importantly, we also tackle the problem of selection bias and systematic sorting in estimates of advisor effects on student success. Successful advisors, however defined, likely attract successful students. If so, the relationship between an advisor’s past students’ success and a new advisee’s research productivity need not be causal. A school-level analysis mitigates this by asking what happens to average success at say, Princeton and Berkeley, when these schools employ more or fewer prolific advisors. The resulting estimates show surprisingly little evidence of superstar advisor effects at the school level. A similar analysis considers the effect of program size, asking whether larger and therefore less selective programs face declining returns to scale in research success. Results here suggest that the best way to increase a program’s total research impact, defined as the number of high-quality publications its graduates produce, is to increase the number of PhD students the program graduates.

The next section sketches aspects of our data set construction, with details covered in an accompanying data appendix. Section 3 presents a descriptive overview of the research output generated by the PhD students in our sample. This section also looks briefly at gender gaps and gender matching in the advising relationship. Section 4 discusses estimates of the relationship between advisor attributes—advisor research prominence, past advisee numbers, and former student success—and current advisee publication success. This section also includes an examination of the role played by advisor-advisee coauthoring and research topic affinity in determining student publication rates. Section 5 looks at aggregate research productivity by school in a grouped data instrumental variables (IV) setup. This section also shows that aggregate school-level PhD student research output scales linearly with graduate program enrollment. Section 6 summarizes our findings and concludes.

2 Data and Descriptive Statistics

2.1 Data, Definitions, and Descriptive Statistics

Our database starts with economics dissertations listed in the *Journal of Economic Literature* and indexed in the ProQuest Dissertations and Theses Global database (formerly Dissertation Abstracts), augmented with information from individual schools and a few other sources detailed in the data appendix. We used ProQuest’s “research field” classifier to identify economics dissertations. The initial sample includes nearly 10,000 PhD graduates who

⁴Buchmueller, Dominitz and Hansen (1999) reports regression estimates of the effect of pre-graduation publication on later research output computed using a non-representative survey sample of two older cohorts.

completed their doctoral degrees between 1989 and 2023. Much of our analysis is limited to graduates from 1994-2017, allowing 6 years of post-PhD follow-up for the most recent cohort. The 1994 start was chosen in view of what appears to be reduced coverage of graduates from some schools in earlier years.

Schools and Departments

ProQuest identifies new degree-granting institutions (which we refer to as schools) more reliably than graduates' departments within schools. We were able to classify many students into departments, however, using a combination of ProQuest thesis PDFs and information supplied by economics departments at six of our eight schools. Much of our analysis combines (and sometimes distinguishes) two groups: identifiable economics department graduates and graduates completing economics-related theses while earning degrees in economics-adjacent departments and programs. We refer to the combined sample as containing graduates of *economics and related programs*. Students appearing in department-provided spreadsheets are classified as graduating from an economics department. Otherwise, classification into departments within schools relies on machine-reading of thesis PDFs. The data appendix details this classification process.

As can be seen in the right panel of Figure 1, economics and related-program PhD cohort size ranges from a high of over 40 for Berkeley in some years to 20 or fewer for Princeton and Yale. Average cohort sizes by school for sample periods underpinning our analysis appear in Table 1. The addition of graduates from economics-related programs increases economics and related cohort size much more for some schools than for others. Specifically, economics-adjacent programs produce a larger share of graduating cohorts from Berkeley, Chicago, Harvard, MIT, and Stanford than cohorts graduating from Northwestern, Princeton, and Yale. Economics programs at the super elites—Harvard and MIT—trend downwards throughout the sample period.

Publication Data

The sample of graduates is augmented with information on publications drawn from Econlit, an AEA database that indexes nearly 2000 economics-related journals. We focus on a shorter list of 137 journals classified as the most cited by articles published in trunk journals for economics and related disciplines plus a few other influential journals. This “Deep Impact” (DI) list comes from an analysis reported in Angrist et al. (2020). The economics trunk journal is the *American Economic Review*. We also look at publications in top-6 (T6) economics journals, defined as the usual top-5 plus the *Review of Economics and Statistics*,

which was once seen as roughly comparable to the top-5.⁵

Research activity—defined by publications—varies considerable by school and over time. Column 3 in Table 1 shows that roughly 37-56% of 1994-2017 graduates from economics departments published at least one DI paper in the six years following degree receipt. As can be seen in column 6, research activity rates are lower for graduates of economics-related programs.

3 Research Output

The advisors and graduates in our data account for a substantial slice of academic economics journal output, especially at more selective and more widely-read outlets. Figure 2 documents this by plotting advisor and student publication shares for various journal tiers. In the extensive set of all EconLit journals, graduates from economics and related programs and their advisors each account for roughly 5% of articles on average over time, with the share authored by the former rising and the share authored by the latter falling through the sample period. Publication shares are markedly higher in more influential outlets. Advisors and advisees contribute over 20% each of DI journal articles in peak years. Moreover, advisors' and advisees' research shares each peak at around half of publications in T6 outlets. In contrast with the share of T6 publications by advisees, which rises near monotonically, the share of T6 publications by advisors peaks in the early 2000s.⁶ These patterns underscore the outsized impact elite PhD program graduates have on academic economics research.

Research activity slows early in most graduates' careers, a trend reflected in the annual activity profiles plotted in Figure 3. Specifically, the figure plots the proportion of research-graduates with one or more publications in each year before and after graduation. Among graduates of economics and related programs, activity rates increase in the four years after degree completion, peaking six years after at about 8%, 25%, and 34% for T6, DI, and all EconLit journals. Activity declines thereafter, a trend also documented in Brogaard, Engelberg and Van Wesep (2018). Perhaps surprisingly, annual activity profiles of (identifiable) economics program graduates, shown in the right-hand panel of the figure, are only a little above those for the broader sample that includes economics department and related program graduates.

A comparison of activity profiles by school, presented in Figure 4, shows profiles that align with widely-held views of program prestige. Activity rates for Berkeley and Chicago

⁵The list used here adds new AEA journals not on the DI list and a relatively new Econometric Society journal, Quantitative Economics.

⁶In this figure, papers coauthored by advisors and advisees contributed to the shares of each.

graduates peak lower than do activity rates at other schools. MIT and Harvard graduates are consistently among the most active. But Northwestern and Princeton graduates' activity rates mostly track and sometimes exceed those of the Cambridge schools, while Stanford and Yale graduates are mostly in the middle. It's noteworthy that 15 years post-PhD, activity gaps by *alma mater* are much diminished.

Many graduates fail to produce a single publication, a finding reported earlier in Conley and Önder (2014). Figure 5 shows something similar for our graduates. Among graduates of economics and related programs, fewer than 60% ever place a paper in a DI journal; only around 45% place two. These long-run success rates are only a little higher among economics department graduates. T6 publications are rare: only about a quarter of graduates place a paper in a top-6 outlet and only around 15% manage two in T6. The likelihood of a T6 publication flattens a little sooner than does the odds of placing a paper in DI journals. T6 success goes higher for economics graduates, though still fails to clear 30%.

Cumulative publication rates also vary markedly across schools. As can be seen in Figure 6, cumulative DI publication rates approach 70% for Princeton graduates, while leveling off under 50% for Chicago graduates. Princeton and Northwestern graduates' research output tracks that of Cambridge alumni initially, eventually pulling ahead. The right side of this figure shows that MIT graduates are the most likely to place in T6 journals, with Princeton graduates a close second in the long run. Berkeley and Chicago graduates are least likely to place in a T6 outlet, a gap that persists. But the general picture is remarkably similar. Most graduates make their mark, if any, in the first 6 years after graduation. Perhaps not coincidentally, in their 7th year of academic employment, many academics are granted a lifetime employment contract based in large part on their initial scholarship.

In the 35 years covered by our sample, elite PhD programs have grown more selective and economics doctoral student funding has grown more generous. Have more selective admissions and higher graduate-education spending yielded higher research output? Cross-cohort trends in research success are captured by a Poisson regression of DI and T6 publication counts on cohort and school effects that can be written:

$$E[R_{isc}|c(i), s(i)] = e^{\delta_{s(i)} + \gamma_{c(i)}}, \quad (1)$$

where dependent variable R_{isc} is the year $c + 1$ to $c + 6$ research output of graduate i from school $s(i)$ in cohort $c(i)$, parameters δ_s are school effects and parameters γ_c are cohort effects. Cohort effects in expression (1) capture differences in research output relative to the reference year, controlling for cross-cohort changes in the graduate distribution over schools.⁷

⁷Appendix Figure A1 plots research activity by cohort and school.

Increased program selectivity and spending notwithstanding, estimated cohort effects suggest graduate research productivity has changed little since the 1989 cohort. This is apparent in Panel A in Figure 7, which plots estimates of γ_c for DI and T6 publications; 1989 is the reference cohort with an effect of zero. For DI publications, estimates hover around zero (meaning no difference from 1989). For T6 publications, cohort effects range from mostly negative before the early 2000s to mostly small and positive after 2007. This modest increase may reflect longer degree completion times, allowing early-career scholars to learn how to clear ever-higher bars for T6 acceptance.⁸ Panel B of the figure reports estimates computed using the sample of identifiable economics department graduates only. Estimated cohort effects for economics graduates differ little from the estimates for the broader sample plotted in Panel A.

3.1 Gender Gaps

We conclude this descriptive graduate student research productivity with a brief look at gender gaps in research output⁹. As can be seen in Figure 8, fewer advisors than students are female, a pattern that likely reflects the fact that the advising load is concentrated among successful academic researchers, fewer of whom are female. At the same time, the share of advisors who are female has climbed steadily from around 9% in 1989 to roughly 22% in 2023. Although more volatile than advisor share female, the share of PhD graduates who are female has also trended upwards, and now exceeds 30 %.

Research activity profiles by time since degree, plotted by gender in Figure 9 show male and female graduates to be similarly active in the first few years post-PhD. By year five, however, female DI activity rates have crested while male DI activity rates continue to climb. This leads to a gap in activity rates that begins to close only 15 years after graduation. Consistent with much lower T6 activity, gender gaps in T6 publication rates are smaller than those for DI, and appear to close sooner than those for DI publications. As a proportion of corresponding male activity rates, however, the largest gaps (observed between post-PhD years 5-10) are similar for DI and T6 publications.¹⁰

Gender gaps in activity rates appear to be unrelated to advisor gender. Figure 10 plots activity profiles similar to those plotted in Figure 9, separately by advisor as well as advisee

⁸Card and DellaVigna (2013) document declining acceptance rates at top-5 journals from 1990-2012, while also showing a roughly proportional rise in number of authors per paper published. Ellison (2002) documents increasing review times at T6 journals.

⁹Gender differences in academic publishing have also been studied by Card et al. (2020), who analyze gender differences in the peer review process at four economic journals; Sarsons et al. (2021), who examine gender differences in credit attribution for group work; and Iaria, Schwarz and Waldinger (2024), who highlight lower publication and hiring rates for female scientists.

¹⁰Proportional gaps are 0.22 for DI and 0.21 for T6.

gender. Activity rates of female and male students evolve similarly regardless of advisor gender. In particular, activity rates of female students with at least one female advisor fall relative to activity rates for men who were advised by either men or women. At the same time, the gender gap in DI publication rates is sharply lower for recent cohorts, a pattern documented in the left panel of Figure 11. The gender gap in T6 activity, by contrast, is reasonably stable across cohorts.

Persistent gender gaps in research success and changes in gender composition across cohorts lead us to control for student gender in regression models that aim to predict student research success. Advisor gender and student-advisor gender matching appear less important. We focus, therefore, on advisor characteristic that seem directly related to research and advising success.¹¹

4 The Advising Relationship

The distribution of graduate advising is highly skewed, with a minority of advisors advising a large share of graduates. This is documented in Figure 12, the left panel of which shows the histogram of number of advisees for the 2499 advisors in our data who advised at least one student in the sample of 1989-2023 graduates from economics and related programs and who were affiliated with one of our eight schools.¹² Roughly a quarter of advisors have only one advisee in the relevant cohorts, while the busiest 15% of advisors advised 20 or more PhD students. The Lorenz curve shown in the right panel of the figure highlights this concentration further. The least prolific 50% of advisors (in terms of advising load) advised fewer than 10% of graduates (indicated by a red line at 0.5 on the x-axis). At the other end of the distribution, the 10% of advisors who are busiest account for roughly half of the advising relationships in our sample.

4.1 Advisor Attributes and Student Success

Advisors vary greatly both in their own research output and in the number of students they advise. Do the most successful faculty researchers advise the most successful PhD students? Is past advising success—both in terms of numbers and in terms of student publications—a

¹¹Appendix Figures A2 and A3 are versions of Figures 9 and 10 for publication counts; these show gender gaps similar to those in activity rates.

¹²Advisors necessarily have at least 1 advisee who graduated from sample schools, but need not be *affiliated* with a sample school or have advised an economics+related program graduate. Advisor affiliations are taken from advisor publications in EconLit; for 1171 with no EconLit publications, this information is missing. And many of the roughly 4800 advisors in our data have an affiliation outside our eight-school sample.

predictor of future advisee performance? We construct three sorts of variables to quantify these aspects of the advising relationship.

Advisor research productivity for a cohort graduating in year c is characterized here by an advisor’s publication record in the five years preceding c . Specifically, *advisor research*, denoted AR_i , counts advisors’ DI publications for each 5-year window ($c - 5$ to $c - 1$), averaged over advisors for each student for up to six advisors (roughly 80% of students have 3 or fewer advisors). Let A_{ic} denote the set of advisors, indexed by j , who advised PhD graduate i finishing in cohort c . Then AR_i is defined by:

$$AR_i \equiv \frac{1}{|A_{ic}|} \sum_{\{j \in A_{ic}\}} R_{jc} , \quad (2)$$

where $|A_{ic}|$ is the number of advisors to i and R_{jc} is the j th advisor’s DI publication count from $c - 5$ to $c - 1$. Similarly, our measure of *advising load* averages the number of a cohort- c graduate’s advisors’ advisees finishing in $c - 5$ to $c - 1$, over up to six advisors for a given student. This quantity, denoted, AL_i , can be written

$$AL_i \equiv \frac{1}{|A_{ic}|} \sum_{\{j \in A_{ic}\}} L_{jc} , \quad (3)$$

where L_{jc} counts advisor j ’s advisees in the relevant cohorts. Note that, for a given graduate, both AR_i and AL_i are averages over the advisor team.

Finally, *past student success* averages the number of DI publications by an advising team’s past students. For cohort- c graduate i advised by advisors indexed by $j \in A_{ic}$, this variable averages DI publications by these advisors’ advisees in the six years since the advisees graduated, looking at advisor- j advisees who finished in years $c - 5$ to $c - 1$. Specifically, let $S_j^{(i)c}$ denote the set of students indexed by $t \neq i$, advised by j who graduated in cohorts $d \in [c - 5, c - 1]$. Past student success, denoted, PS_i , is defined as

$$PS_i \equiv \frac{1}{|A_{ic}|} \sum_{\{j \in A_{ic}\}} \frac{1}{|S_j^{(i)c}|} \sum_{\{t \in S_j^{(i)c}\}} R_{td} , \quad (4)$$

where R_{td} counts DI publications by student t who graduated in cohort d in years $d + 1$ to $d + 6$.

Like research success, advising success is rare. We therefore look at upper-tail measures of success as well as advisor averages. For student i who graduated in cohort c , three *super* advising dummy variables indicate whether at least one of i ’s advisors was in the top 10% of advisors among those who advised students who graduated from $c - 5$ to $c - 1$. For example,

the *super advising load* dummy indicates graduates in cohort c with at least one advisor in the top 10% of advisors measured by number of advisees graduating in $c - 1$ to $c - 5$. This amounts to having one advisor $j \in A_{ic}$ with L_{jc} in the top 10% of the distribution of L_{jc} for the relevant set of cohorts as a group. Super advisor research and super past student success are defined similarly. It's enough to have one super advisor for a student-level super dummy to switch on. A parallel set of three *dupes* variables indicate graduates with at least one advisor among the 5% most prolific when ranked by advisor research, advising load, and past student success.¹³

Figure 13 tracks the distribution of super advisors over time by school. Consistent with the traditional view of Harvard as employing many highly successful senior scholars, Panel A in this figure shows Harvard advisors enjoy standout research success, while Northwestern and Yale advisors are the least successful.¹⁴ Interestingly, beginning in the early 2000s, Berkeley moves ahead with the largest number of super advisors in terms of advising load. Perhaps not coincidentally, this increase follows legendary advisor David Card's move to Berkeley. Not only is Card a super advisor, he attracted top younger advisors to the Berkeley faculty. The bottom panel of Figure 13 offers an interesting counterpoint to the top two panels: student research success shows less dispersion across schools than do advising load and advisor research success.

A student-level regression connects graduates' post-PhD research success with advisor attributes, controlling for cohort and school effects. Let R_{isc} denote publications by student i graduating from school $s(i)$ in cohort $c(i)$, including zeros. Specifically, the impact of advisor attributes on R_{isc} is estimated using the following regression model:

$$R_{isc} = \alpha'W_i + \tau'D_i + \beta N_i + \gamma_{c(i)} + \delta_{s(i)} + \varepsilon_{isc}, \quad (5)$$

where D_i is a vector of one or more of the advisor variables defined above, N_i is the number of i 's advisors, and $\delta_{s(i)}$ and $\gamma_{c(i)}$ are the relevant school and cohort effects, respectively. Control for team size is motivated by the fact that, for a given graduate, average advisor productivity, advising loads, and past student success are diluted when advising teams are larger. At the same time, the probability of having at least one advisor with upper-tail values

¹³Advisors are defined as super using the universe of economics-related PhDs, not limited to advisors of PhDs from economics departments and related programs. The super dummy for past student success is coded as follows: for each advisor, advisees graduating $d \in [c - 5, c - 1]$ are identified. DI publications for this group in years $d + 1$ to $d + 6$ are summed and divided by the number of advisees graduating in this period. This quantity (the inner summation in expression (4)) gives an advisor's average past student success looking at their advisees from the last five years. A student is advised by a super advisor if one of their advisors is among the top 10% of advisors when ranked by average past student success.

¹⁴Harvard advising outcomes are also affected by school size, with substantial economics-adjacent graduate enrollment at Harvard's Kennedy School, Business School, and Graduate School of Education.

of these variables increases with the size of the advising team. Vector W_i in equation (5) contains a set of student controls that includes a female dummy, a dummy for PhDs with gender unclassified, a dummy for graduates with no readable thesis PDF, and a dummy for identifiable economics department graduates (and a constant). Vector τ contains the coefficients of primary interest. We also report estimates of a Poisson analog of equation (5), where estimates of τ give the percentage change in publications attributable to D_i .

In models entering advisor attributes one at a time, graduates advised by advisors who publish more, advisors with more advisees, and advisors with more successful former students have greater publication success. The first column of Panel A in Table 2 shows, for instance, that graduates see 0.12 more DI publications, on average, when advising team research increases by one. This is a gain of roughly 7% (the corresponding Poisson coefficient appears in column 5). As can be seen in column 3, advisor research success appears to boost advisees' T6 publications much less, though effects on advisees' T6 publications in percentage terms (reported in column 7) are larger and approach 9%. The estimates in the second row of Panel A show that graduates advised by super researchers generate more publications. Super-research effects are 0.69 and 0.21 for DI and T6 levels, respectively, and 0.48 and 0.68 in percent.¹⁵

The estimates in Panel B of Table 2 suggest that an advising team's advising load is (mostly) a weaker predictor of student research success than the team's own research record. At the same time, patterns here are similar to those in Panel A, with consistently positive effects: prolific advisors (in terms of numbers of past students) have advisees that see greater post-graduation research success. As with the estimated advisor-research effects in Panel A, effects on student output in levels are larger for DI than T6, while effects in percentage terms are larger than the corresponding levels coefficients for T6. Super advising-loaded advisors have substantially more successful students than do advisors with fewer advisees.

Among the three continuous advisor attributes examined in Table 2, the estimates in Panel C of the table show that average past student success is the best predictor of current student success. Increasing an advisor team's average past student DI publications by 1, for instance, is associated with 0.43 more current student DI publications, a 26% increase. Increasing the average of past student DI publications is predicted to yield a 36% increase in current student T6 publications. Again, coefficients on super dummies indicating the strongest 10% of advisors based on their advisees' past success are even larger.

Multivariate models that capture effects of advisor research, advising load, and past student success jointly suggest advisor research and past student success are more important

¹⁵García-Suaza, Otero and Winkelmann (2020); Hilmer and Hilmer (2009) likewise find a strong association between advisor research productivity and economics advisees' post-PhD publishing success.

drivers of current student success than advising load. This can be seen in Table 3, which reports OLS estimates in columns 1-4 and Poisson estimates in columns 5-8, separately for DI and T6 student publications. In particular, the estimates in columns 1, 3, 5, and 7 show advising load effects close to zero and not significantly different from zero. The corresponding estimates for advisor research and past student success are smaller than those in Table 2, though still substantial and significantly different from zero.

Even-numbered columns in Table 3 report results from models that add duper dummies indicating graduates advised by at least one advisor in the upper 5% of the relevant advisor attributes distribution. In these specifications, super and duper advising dummies are mutually exclusive, so super dummies indicate advisors with characteristics in percentiles 6-10. Estimated super advising effects remain large and relatively precisely estimated for all 3 variables, though they're smaller than the corresponding one-at-a-time estimates in Table 2. The multivariate super advising load coefficient is the most diminished from Table 2. Duper advising coefficients are larger than the corresponding super advising coefficients, suggesting that student success is monotone in advisor quality as proxied by advisor research, advising load, and past student success.

Multivariate models generate dampened effects of advising load on student research success. Interestingly, however, the estimates in Tables 2 and 3 show no evidence of a productivity-diminishing advising burden: advisees guided by advisors carrying a heavy load do not appear to suffer from their advisors' more finely-divided attention. Of course, this finding and others in Table 2 and 3 may reflect selection bias as well as causal effects. We return to this point following an analysis of student-driven aspects of the advising relationship.

4.2 Advisor-Advisee Coauthoring and Research Affinity

Advisor-advisee coauthoring has long been common in the natural sciences.¹⁶ Mirroring the rise of empirical economic research since the early 1990s, this sort of teamwork has grown in economics publications too (Angrist et al., 2017; Jones, 2021). Figure 14 documents increasing student-advisor and student-classmate coauthoring for cohorts graduating since 1994. Specifically, the figure tracks the share of students coauthoring with either an advisor or a classmate (defined as same-school PhDs who graduated the same year or within two years before or after), before and after degree receipt.

Unsurprisingly, coauthoring of any kind is far more common after degree receipt than before (since PhD students must publish to coauthor). For all cohorts, coauthoring with

¹⁶See, for instance, Azoulay, Graff Zivin and Wang (2010), which examines collaborations in the life sciences.

advisors is more common than coauthoring with classmates. But both sorts of coauthoring are on the rise. Among the most recent cohorts for which data appear in Figure 14, advisor-advisee coauthoring rates hit 20%, while classmate coauthoring reaches roughly 15%.

The share of PhD theses citing an advisor’s work has also trended upwards for most of the cohorts included in our samples. This trend is visible in the advisor citation rates plotted in Figure 15, albeit with considerable cross-school variation and within-school variation over time. Citations to advisors have been highest at Harvard, peaking with over 70% of early-2010s theses citing advisor work. Citations to advisors increased most at MIT, from an initial rate around 40% to a rate similar to that of Harvard PhDs by 2015. Princeton PhDs’ citation rates evolve much like those for Harvard PhDs, while Berkeley, Chicago, Northwestern, Stanford and Yale advisor-citation rates hover between 50-65% by the end of our sample period. These patterns suggest growing research affinity between advisors and advisees in general, though with persistent differences in the degree of intellectual alignment across institutions.

Our measure of thesis citations to advisors is imperfect and relies in large part on algorithms that read thesis PDFs.¹⁷ Advisor affiliation is also approximate in our data, relying on author affiliations attached to publications in EconLit. Still, the view that the resulting citation rates are informative is supported by Figure 16. This figure tracks the number of students advised and student citations to advisors for advisors that change affiliations. For 367 fixed-super advisors (defined as an advisor whose advising load falls in the upper decile of the advisee distribution for any cohort), the figure counts advisees at the former affiliation and new affiliation separately. A transitioning advisor’s advising load at their previous affiliation falls sharply at the time of a move, plateauing close to zero four years later. At the same time, the number of advisees advised by transitioning advisors at the new institution increases rapidly. A similar pattern appears in the number of graduates that cite transitioning advisors’ work: their number decreases at a transitioning advisor’s previous institution while it increases at the new one (considering all students at origin and destination schools, not just graduates advised by transitioning advisors). These patterns suggest advisees citations capture a decline in intellectual influence at an advisor’s old institution and rising influence at the new one.

When included in a regression model like equation (5), variables that indicate coauthoring with advisors and classmates before or in graduation years are unrelated to PhD students’ immediate post-graduation publication success. Estimated advisor-coauthoring effects on DI publications, reported in the first four columns of Table 4, are small and not significantly

¹⁷Thesis citations are not indexed by the Web of Science or EconLit. We therefore search thesis PDFs for all article titles of advisors’ publication lists as detailed in the data appendix.

different from zero. Classmate-coauthoring coefficient estimates, reported in columns 2 and 4 of the table, are likewise not significantly different from zero, though estimated much less precisely than the corresponding advisor-coauthoring effects.

Graduates whose thesis cites an advisor, by contrast, are estimated to see 0.22 more DI publications and 0.07 more T6 publications, estimated gains that are significantly different from zero. The Poisson estimates in columns 6 and 8 of Table 4 show that these gains amount to publication increases of 16% and 27%, respectively. The strongest predictor of post-PhD publication success is a dummy variable indicating graduates with any pre-PhD publication; this is a control necessitated by the fact that our definition of coauthoring requires publication and precocious pre-degree publishers are more likely to publish after leaving the nest.

As can be seen in the first three rows of Table 4, estimated advisor effects generated by the extended version of equation (5) with coauthoring and affinity variables are similar to those in the first three rows of Table 3, with significant estimates of coefficients on advisor research and past student success. Replacing continuous advisor attributes with dummies indicating *fixed-duper* advisors generates the estimates in Table 5. In this model, fixed-duper dummies indicate advisors with a given attribute (e.g., number of advisees) ever coded as duper. This specification is motivated by a view of exceptional advisor performance as a time-invariant attribute. The resulting estimates are broadly consistent with those in Tables 3 and 4, suggesting that the most important advisor attributes are advisor research and past student success. In contrast with the estimates in Table 4, however, estimated fixed-duper advising load coefficients are larger—around 0.1—and significantly different from zero for graduates’ T6 publications. It’s noteworthy, however, that regardless of specification, student success is hard to predict: for all models reported in Tables 3-5, R^2 s are no bigger than .14 for DI publications and below .1 for T6.

5 The Aggregate Research Production Function

5.1 Aggregation IV

The strong relationship between advising features and student research success documented in Tables 3-5 suggests factors like an advisor’s research record and research affinity contribute to their advisees’ success. But estimates of this relationship may also reflect selection bias. Advisors whose students have done well in the past, for instance, may attract students who are most likely to succeed in the future. A school-level analysis mitigates this sort of bias by asking what happens to overall average success among graduates of, say, Princeton and

Berkeley, when a prolific advisor moves from the former to the latter.

As before, let $c(i) \in \{1994 \dots 2017\}$ encode graduation cohort for PhD i and let $s(i)$ encode which of the eight schools i attended. Suppose average potential research publication outcomes when $D_i = 0$ can be described by the conditional expectation function (CEF):

$$E[R_{isc}(0)|c(i), s(i); W_i] = \tilde{\alpha}'W_i + \tilde{\gamma}_{c(i)} + \tilde{\delta}_{s(i)}. \quad (6)$$

This says that, conditional on cohort and school, the CEF of potential research outcomes at a reference level of advising inputs is assumed to be an additive function of cohort and school effects, possibly with adjustment for W_i . Suppose also that the causal effects of advising features vector D_i (augmented to include coauthoring and research affinity variables) and advisor team size, N_i , on post-PhD research are constant and given by $\tilde{\tau}$ and $\tilde{\beta}$, respectively. These causal effects and the coefficients in (6), denoted by Greek letters with a $\tilde{}$, need not coincide with coefficients in an OLS estimand defined by a regression of R_{isc} on advising features and controls, as in equation (5).

Restriction (6) and this constant effects assumption imply the following conditional moment restriction:

$$E[R_{isc} - \tilde{\tau}'D_i - \tilde{\beta}N_i - \tilde{\alpha}'W_i - \tilde{\gamma}_{c(i)} - \tilde{\delta}_{s(i)}|c(i), s(i); W_i] = 0. \quad (7)$$

In other words, conditional on cohort and school effects, variation in the productivity of graduates by cohort and school is explained by variation in the mean of right-hand-side variables by cohort and school. Without individual controls, W_i (in this case, dummies for female graduates, unclassified gender, economics department graduates and missing thesis PDFs), this moment restriction generates a grouped model that regresses average research productivity by cohort and school on average advisor research, the share of graduates coauthoring with advisors, and so on. With individual covariates, these regressors are adjusted for compositional changes due to a changing mix of gender, economics department, and PDF availability for the graduates in our data.

Importantly, two-stage least squares (2SLS) estimates based on restriction (7) are not confounded by *within-school* sorting of students to advisors. Suppose, for instance, that within departments, the best students seek out advisors with strong research records, but advisor research success generates no payoff in terms of advisee research success. In this scenario, estimates like those in Table 4 are likely to be positive: the fact that productive students seek out productive advisors engenders positive omitted variables bias. Yet, such within-school sorting leaves average student success by cohort and school unchanged, making the aggregate student research production function a better guide to causal advisor effects

than equation (5) estimated using data on individual graduates. Because the aggregate model controls for cohort and school effects, restriction (7) identifies causal effects by exploiting cross-cohort *changes* in average advising features within schools. These changes are due both to advisor transitions between schools and evolving within-department changes in incumbent advisor features. Elite departments compete to attract top scholars and prolific advisors, hoping (among other hiring goals) for an immediate boost in advising horsepower. As a suggestive exploration of the role played by advisor transitions, Figure 17 plots average research output for graduates at schools losing and receiving transitioning fixed-duper advisors as determined by their past student success. This figure is constructed in a manner similar to that used to construct Figure 16. We opt here for fixed-duper rather than fixed-super advisor transitions because the former generate larger student-level estimates in Table 5. In contrast with Figure 16, which shows advising load and citation changes consistent with transitioning advisors having an impact on affected schools, the pattern of student success traced in Figure 17 shows little evidence of gains in student research success at departments bolstered by the arrival of a fixed-duper advisor.

In principle, restriction (7) justifies a 2SLS estimator using a full set of 192 cohort-by-school dummies as instruments. Many of these instruments are weak, however, in the sense that they generate noisy first-stage estimates based on only a few students in each cohort for some schools. The resulting 2SLS estimates are therefore likely to be biased, and misleadingly similar to the corresponding OLS estimates. This motivates an IV strategy using 49 3-year-cohort-by-school dummies as instruments instead of the 192 dummies generated by single-year cohorts. The individual covariate vector, W_i , appears in both the first and second stages in this 2SLS setup. Table 6 reports 2SLS estimates in which all right-hand-side variables other than W_i are instrumented as well as estimates from models in which only advisor research, advising load, and past student success are instrumented, treating other features like coauthoring and research affinity as covariates.

Consistent with Figure 17, 2SLS estimates using 49 dummy instruments provide less evidence that advising features matter than do the corresponding OLS estimates. 2SLS estimates with all features instrumented, reported in the first four columns of Table 6, generates statistically significant estimates only for the impact of advisor research. Statistically significant 2SLS estimates of the impact of advisor research, on the order of 0.16 for graduate DI publications and 0.047 for graduate T6 publications, exceed the corresponding OLS estimates in columns 1-4 of Table 4 but are less precise. 2SLS estimates instrumenting only the three advisor attributes generate large, precisely-estimated pre-grad publication effects similar those in Table 4.

Estimates of coauthoring and advisor-citation effects in Table 6 are so imprecise they

should be seen as uninformative, though large advisor-coauthoring coefficients are marginally significant for graduates' T6 publications in column 3. On balance, this table does little to bolster claims for a strong relationship between advisor attributes and student success. Most notably, the coefficient on advising load is a reasonably-precisely estimated statistical zero, while the coefficient on past student success is precise enough that some effects of the magnitude seen in Table 4 would be marginally significantly different from zero. The strongest evidence for causal effects emerges for advisors' own research output, highlighting the critical role of advisor research success in shaping their advisees post-PhD research outcomes. Factors that seem to matter in Tables 4 and 5 fail to generate consistent or precisely-estimated changes in average PhD student research success.

5.2 Research Returns to Scale in Cohort Size

Motivated by the trend towards rising economics PhD program selectivity and falling cohort size at the super-elites, we conclude with an analysis of cohort size effects on student research success. In particular, we're interested in whether the scale of economics PhD research production function is constrained by decreasing returns. Roughly speaking, cohort size can be thought of as graduate program class size. Perhaps reduced cohort size enhances students' post-graduation prospects by increasing resources and facilitating faculty-student mentoring. Larger cohorts, by contrast, may produce a critical mass of students in the classroom and more stimulating interactions between students.

The relationship between cohort size and research productivity is quantified here using a regression of school-by-cohort aggregate graduate publication output up to six years after graduation on cohort size. This regression fits the following school-by-cohort CEF:

$$T_{sc} = \gamma_c + \delta_s + \kappa_1 n_{sc} + \kappa_2 n_{sc}^2 + \nu_s c, \quad (8)$$

where n_{sc} is the size of cohort c at school s and the dependent variable, T_{sc} , sums DI or T6 publications in years $c + 1$ to $c + 6$ by cohort- c graduates from programs at s . Decreasing returns to scale in graduate student research are evinced by a negative estimate of κ_2 , the coefficient on n_{sc}^2 .

Estimates of a linear version of equation 8 using the full sample, reported in columns 1-2 and 5-6 of Panel A in Table 7, indicate that DI publications increase by 1.3 – 1.5 per student, while T6 publications increase by 0.3 – 0.4. These estimates are close to the mean post-publication statistics in Appendix Table A1. Estimates are similar when computed with and without school effects. Results using the sample of economics department graduates only,

reported in Panel B, are also similar to those for the full sample.¹⁸

Estimates of κ_2 in models that include n_{sc}^2 suggest the graduate research production function is remarkably linear. Models with and without school effects, for the full sample and for economics department graduates only, generate estimates of κ_2 that are small and not significantly different from zero. Inclusion of a quadratic term makes estimates of linear terms less precise, and nonlinear models for T6 publications are unstable and sensitive to the inclusion of school effects. Although the estimates for T6 output are sensitive to the inclusion of school effects, the results are reasonably consistent across specifications and samples.

The estimated marginal effect of cohort size in equation (8) is $\hat{\kappa}_1 + 2\hat{\kappa}_2 n_{sc}$. The estimated change in returns to scale as a function of cohort size is therefore $2\hat{\kappa}_2$. Estimates of κ_2 for DI publications in the full sample imply a modest effect-change gradient of $-.0066(2 \times -0.0033)$ when estimated without school effects. This flips to equally-modest-but positive 0.0042 when school effects are added. Inclusion of school effects flips the sign of the other small estimated gradients elsewhere in the table as well: the implied change in returns to scale as enrollment grows are therefore neither large nor consistently negative.

On balance, the estimates in Table 7 suggest that departments looking to increase social impact through academic economics research can do so cheaply and quickly by admitting more students. This possible free lunch likely partly reflects the fact that advising is so concentrated. The schools in our sample appear to have plenty of advising slack.

6 Summary and Conclusions: Selection, Training, and Luck on the Path to PhD

Top economics departments attract exceptional students and invest substantial resources in training these students for successful research careers. The most important graduate education resource is the time and attention of PhD advisors. With the help of uniquely rich and comprehensive data linking economics PhD students with their advisors, our analysis quantifies the relationship between key features of advisors and the advising relationship and the post-degree research success of economics graduate students.

Key descriptive facts emerging from our analysis include the high concentration of advising among a minority of advisors and the limited research success seen by the median graduate. It's also noteworthy that, even as elite programs have grown costlier and more selective, graduate research success has remained reasonably flat across cohorts. Other descriptive findings include the fact that research performance differs little between identifiable

¹⁸Cohort size in this case is the number of identifiable economics department graduates.

economics department graduates and their peers from economics-related programs. We've also shown that, after a brief warm-up period characterized by gender parity in research output, female graduates publish fewer papers than do male graduates. But the gender gap in research output, which is unrelated to advisor gender or advisor-advisee gender matching, may now be closing.

Which factors increase the likelihood of post-PhD research success? Multivariate models that predict graduate research success as a function of advisor and thesis characteristics yield robust positive effects of advisor research and the advising team's past student success. Surprisingly, however, a relationship between an advising team's advising load and student research success emerges only for dummy variables indicating the most prolific advisors. We've also seen that while precocious pre-PhD student publishers publish more papers post-PhD, research output appears to be unrelated to coauthoring with advisors or classmates. On the other hand, PhDs who cite an advisor in their thesis—a measure of student-advisor research topic affinity—tend to see greater research success post-PhD.

The possibility of advisor-advisee sorting and the attendant selection bias motivates 2SLS analysis using dummies for cohort-by-school to instrument advising features. 2SLS estimation uncovers only weak evidence of causal effects of advising features on PhD student research success. In particular, at the level of cohorts and schools, only average advisor research predicts research success for the average graduate. 2SLS estimates of effects of past student success and research affinity are uniformly positive but mostly not significantly different from zero.

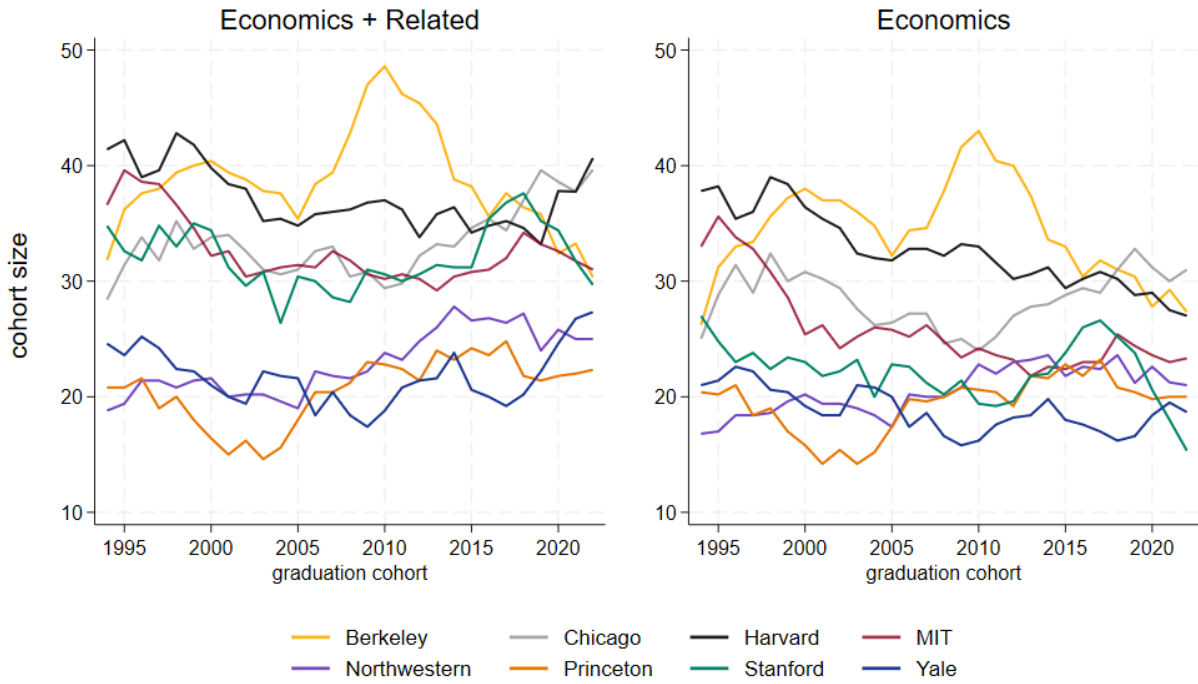
We've also explored the nature of returns to scale in graduate economics education. Our analysis suggests that aggregate student research output scales roughly linearly with graduate economics enrollment. This finding challenges conventional wisdom regarding the importance of small, highly selective cohorts in economics graduate education. On the margin, graduates in larger cohorts publish about as well as graduates in smaller cohorts.

A broader lesson suggested by our findings is that research success is hard for elite schools to engineer or even predict. In this, academic economics is like professional sports: coaches at all levels struggle to identify and nurture talent; uncertainty regarding outcomes and performance is pervasive. Even among American Division I college basketball players, who are well-trained and necessarily play a very good game, few ultimately play for a living. As in high-level sports, economics PhD students and their advisors alike should benefit from a clear-eyed view of the winding and uncertain path to research success.

Exhibits

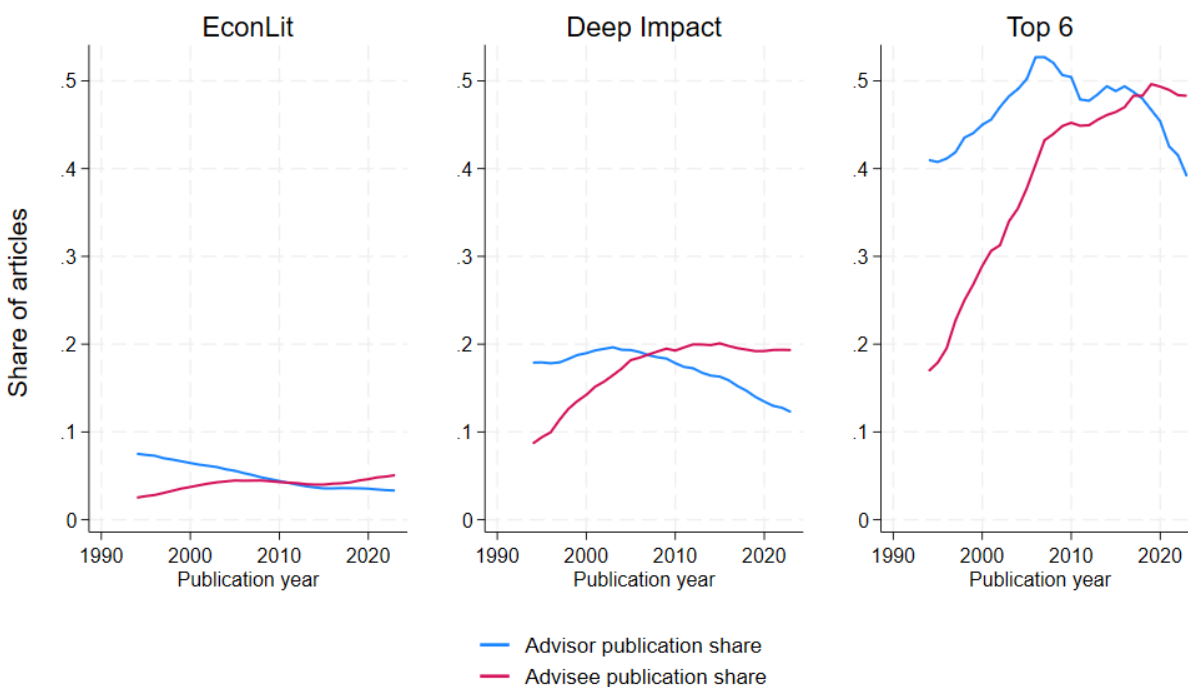
Figures

Figure 1. Economics and Related Program Cohort Sizes



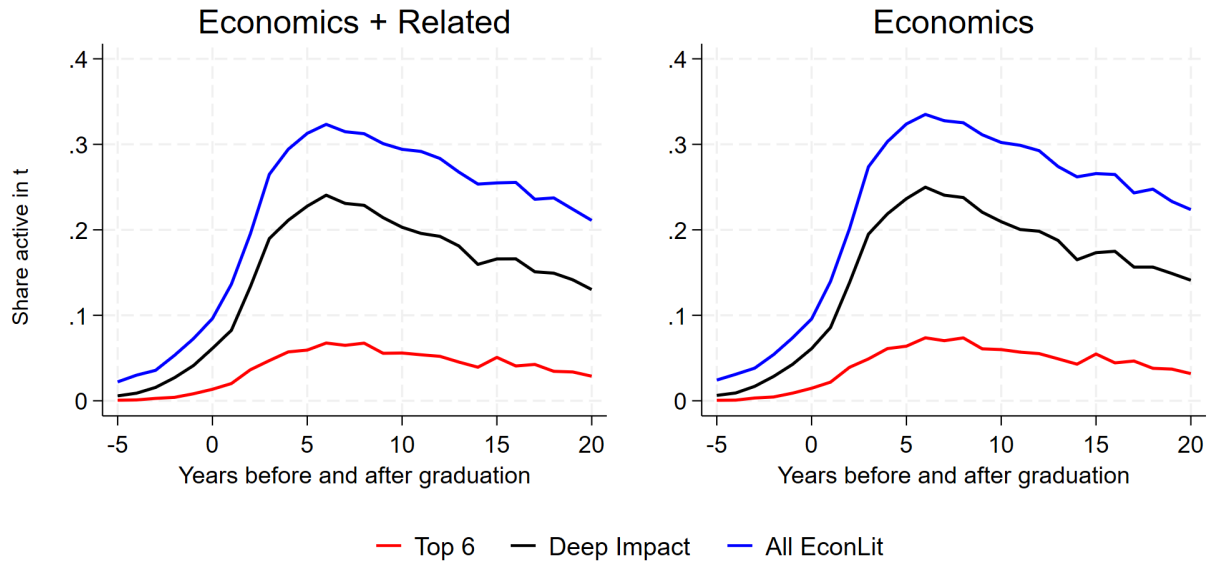
Notes: The left panel shows graduation cohort size for students identified as earning either economics department degrees or related department or program degrees with affiliation determined as described in the data appendix. The right panel shows graduation cohort size for students identified as economics department graduates. Related departments and programs include Finance, Management, Business, Accounting, Marketing, and Operations Research. A few students with no department indicated on thesis cover pages or for whom no thesis was available for download are included in the economics + related sample if one of their advisors advised at least one student identified as an economics department graduate. The figure plots five-year moving averages starting in 1994, with the first years smoothed using data back to 1989.

Figure 2. Advisor and Advisee Publication Shares



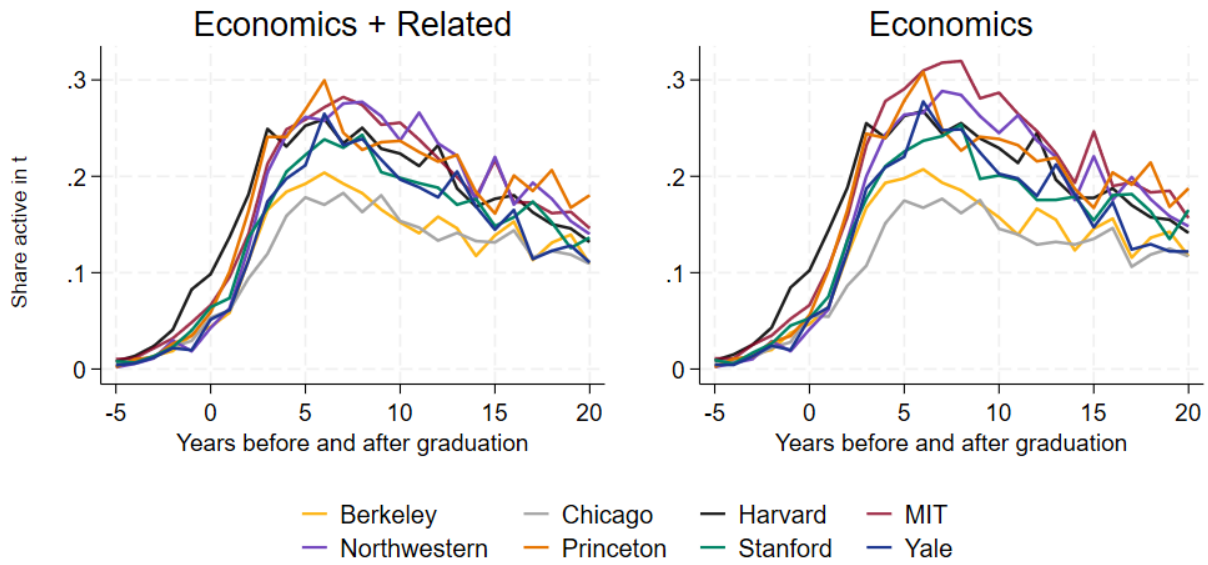
Notes: This figure plots a five-year moving average of the yearly share of publications authored or coauthored by economics and related program graduates earning degrees at one of the eight institutions in our sample, along with the share of publications attributable to their advisors. The first panel shows the share of advisor and student publications in all of the roughly 2000 journals indexed by EconLit. The second panel shows the share of advisor and student publications in relatively well-cited “Deep Impact” journals, classified by Angrist et al. (2020). The third panel shows the share of publications in top-6 economics journals. The data appendix lists the journals included in the second and third panels.

Figure 3. Annual Activity Profiles: EconLit, Deep Impact, and Top 6



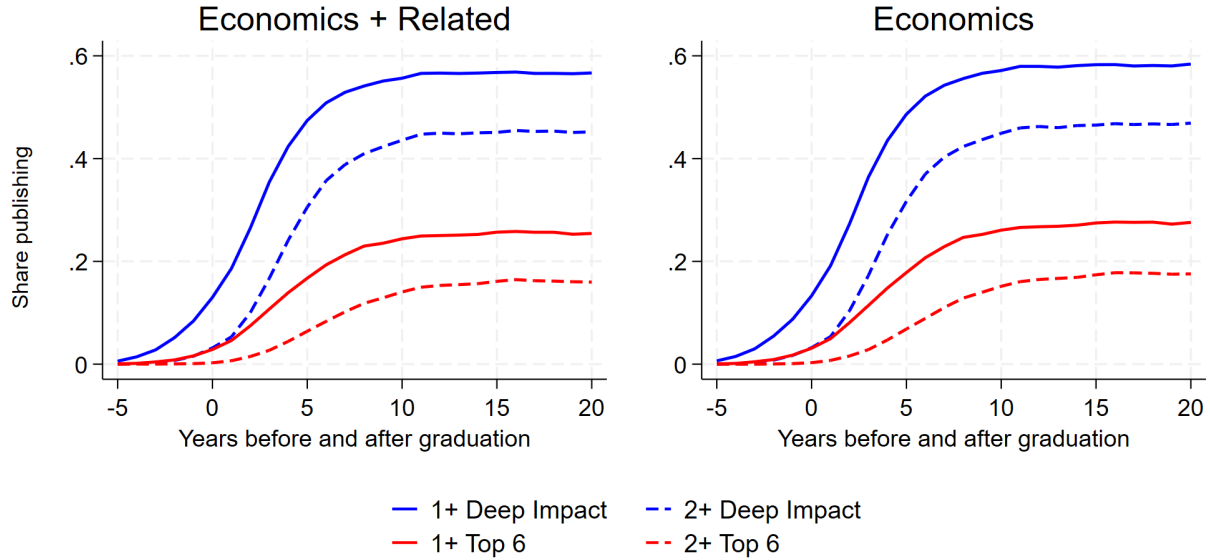
Notes: This figure plots the share of students that have at least one publication in year $t - c$, where c is graduation year, for three nested journal lists. Data are for economics and related program graduates on the left and for identifiable economics department graduates on the right. Data for 1994-2017 graduates.

Figure 4. Annual Activity Profiles by School



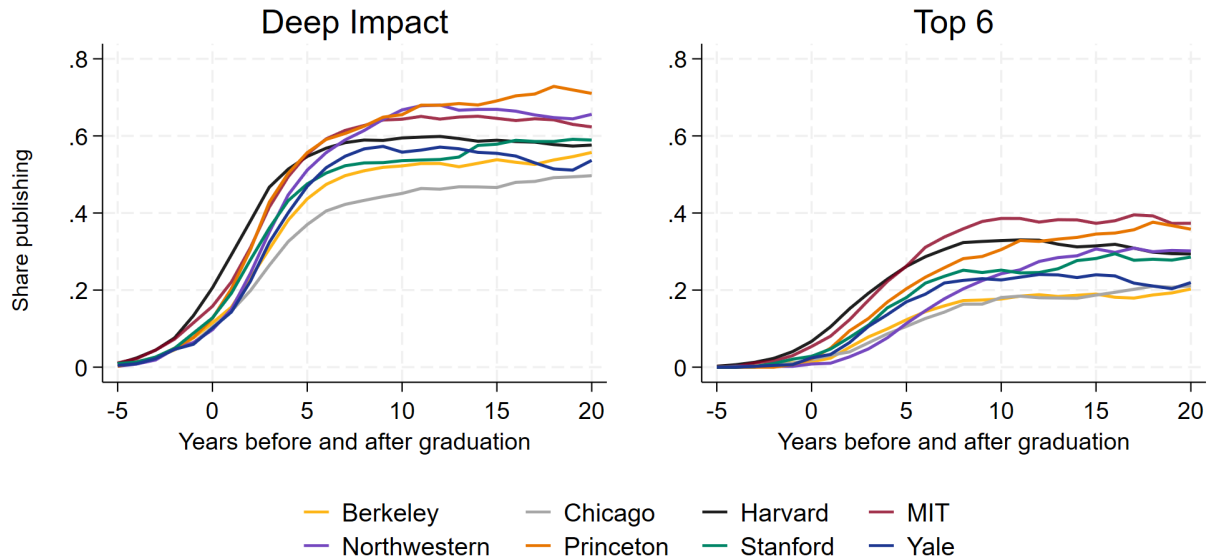
Notes: This figure plots the share of students that have at least one Deep Impact publication in year $t - c$, where c is graduation year, separately by school. Data are for economics and related program graduates on the left and for identifiable economics department graduates on the right. Data for 1994-2017 graduates.

Figure 5. Cumulative Publication Profiles



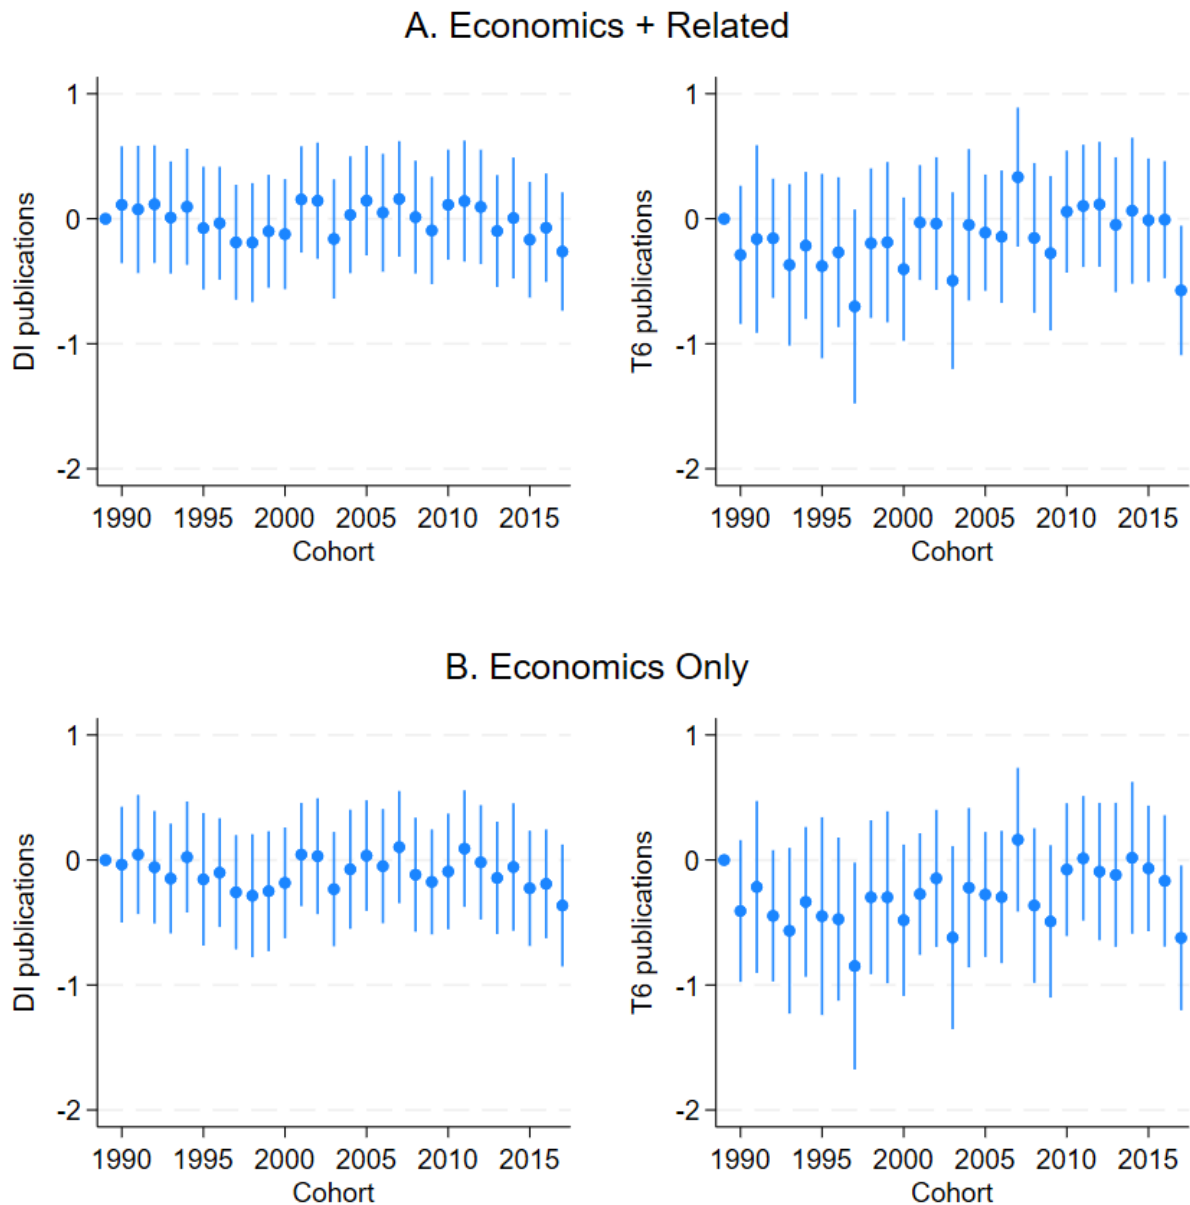
Notes: This figure plots the share ever-publishing as of year $t - c$, where c is graduation year, for publication types and levels indicated in the legend. Data are for economics and related program graduates on the left and for identifiable economics department graduates on the right. Data for 1994-2017 graduates.

Figure 6. Cumulative Publication Profiles (1+ Pubs), by School



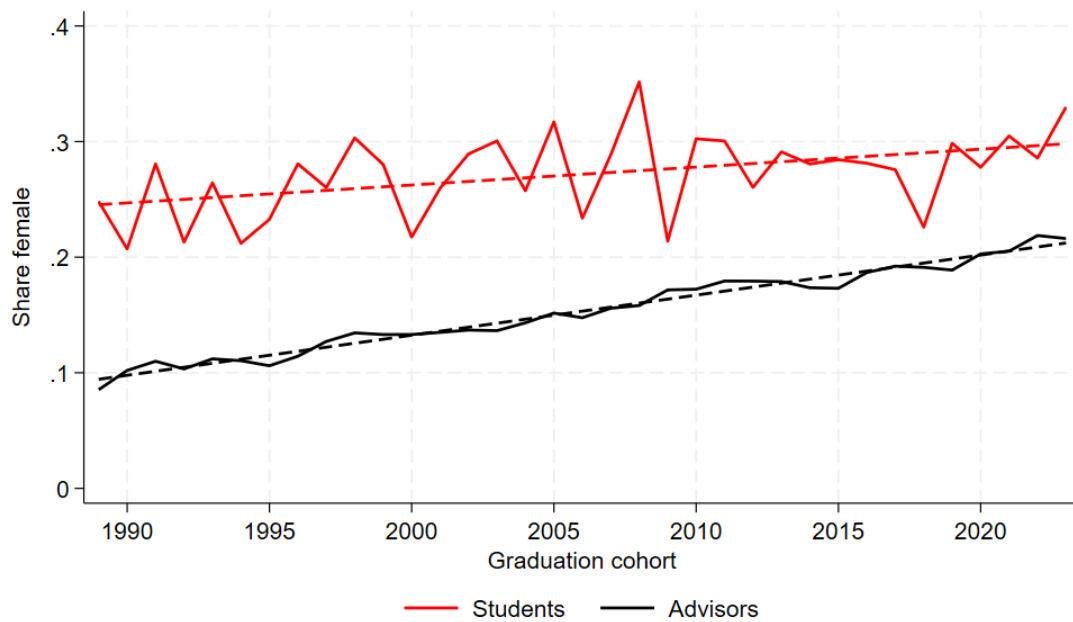
Notes: This figure plots the share publishing (1+ pubs) in Deep Impact journals (left) and Top 6 journals (right) as of year $t - c$, where c is graduation year, separately by school. Data for 1994-2017 economics+related program graduates.

Figure 7. Cohort Effects in Activity Rates



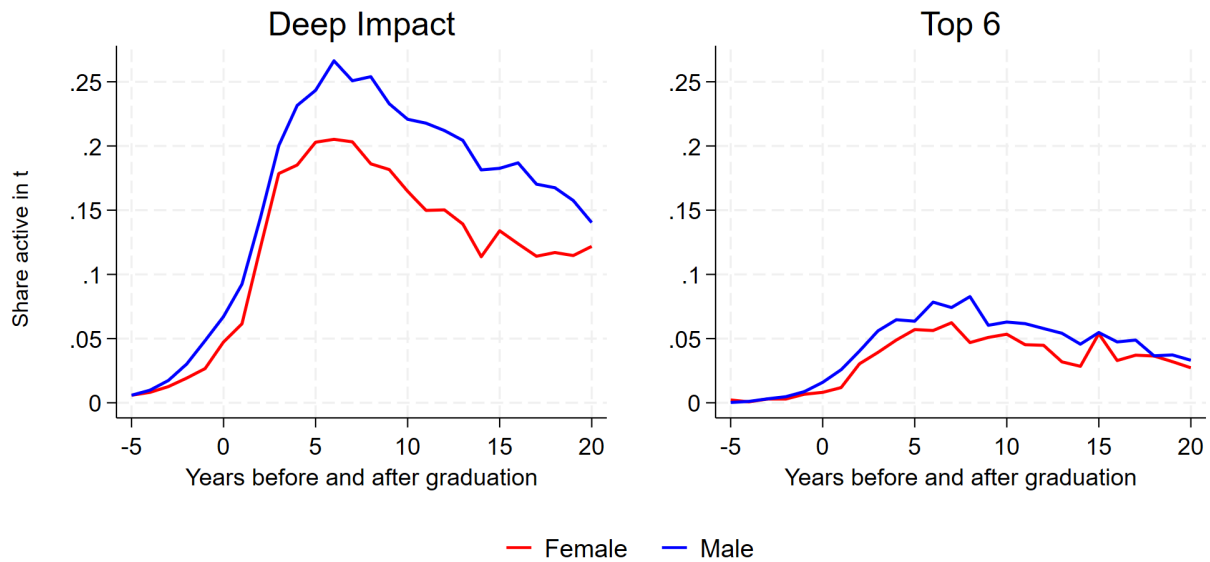
Notes: This figure reports estimated cohort effects from a student-level Poisson regression of Deep Impact and Top 6 publication counts for years $c+1$ to $c+6$ on cohort and school effects. The reference year is 1989 and the sample starts with the 1989 cohort. Results for Deep Impact are on the left; results for Top 6 are on the right. Estimates in Panel A use the sample of economics and related program graduates (dependent variable means are 1.53 for Deep Impact publication counts and 0.34 for Top 6 publication counts). Panel B reports estimates for identifiable economics department graduates only (dependent variable means are 1.59 for Deep Impact publication counts and 0.36 for Top 6 publications). Standard errors are clustered on school-by-year.

Figure 8. Graduate and Advisor Share Female



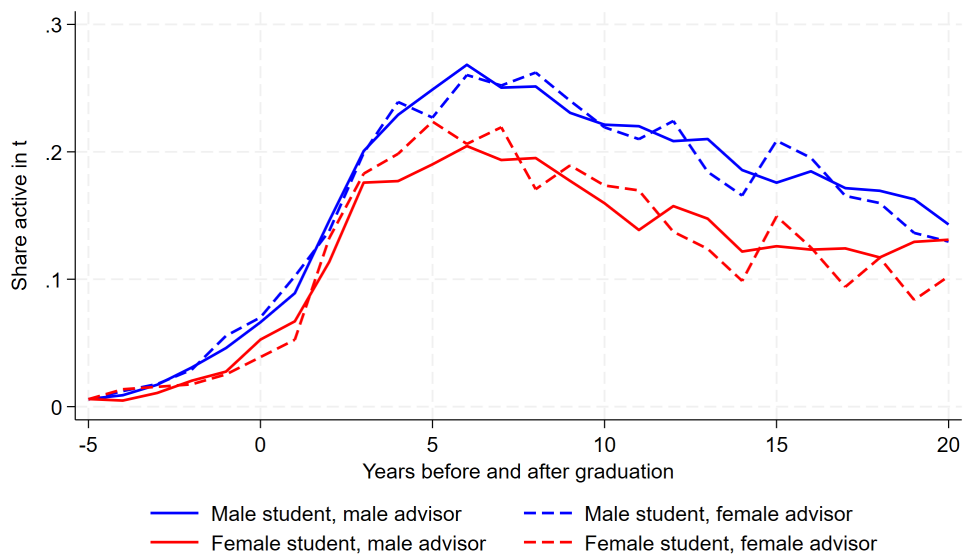
Notes: The red line plots the graduate share female by graduation cohort; the black line plots the share female among advisors. Gender is coded from first-name gender frequencies in Social Security records, as described in the data appendix. Advisor data is limited to years of active advising, defined as the period between an advisor's first and last advisee graduation year. Data are for the economics + related sample of 1989-2023 graduates and their advisors.

Figure 9. Gender Gaps in Graduate Research Annual Activity Profiles



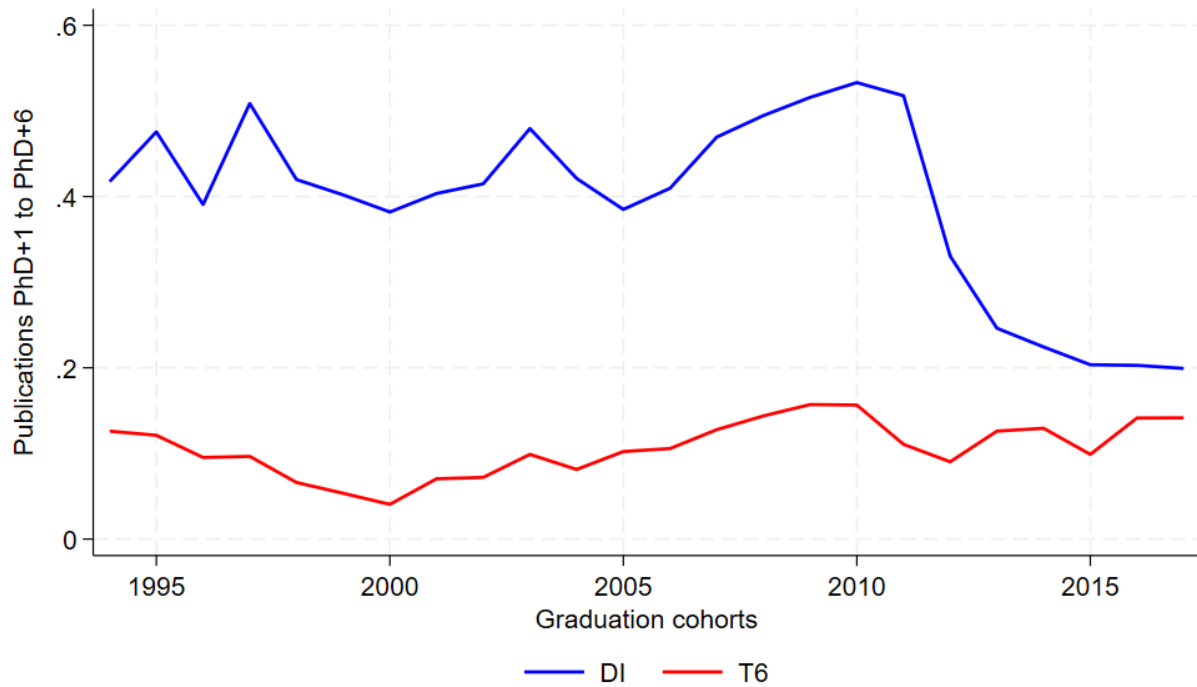
Notes: See notes to Figure 8. This figure plots the share of male and female students with at least one publication in the Deep Impact (left panel) or Top 6 (right panel) journal list in year $t - c$, where c is graduation year. Appendix Figure A2 plots publication counts in the same format. Data are for the economics and related sample of 1994-2017 graduates.

Figure 10. Annual Activity Profiles by Graduate Student and Advisor Gender



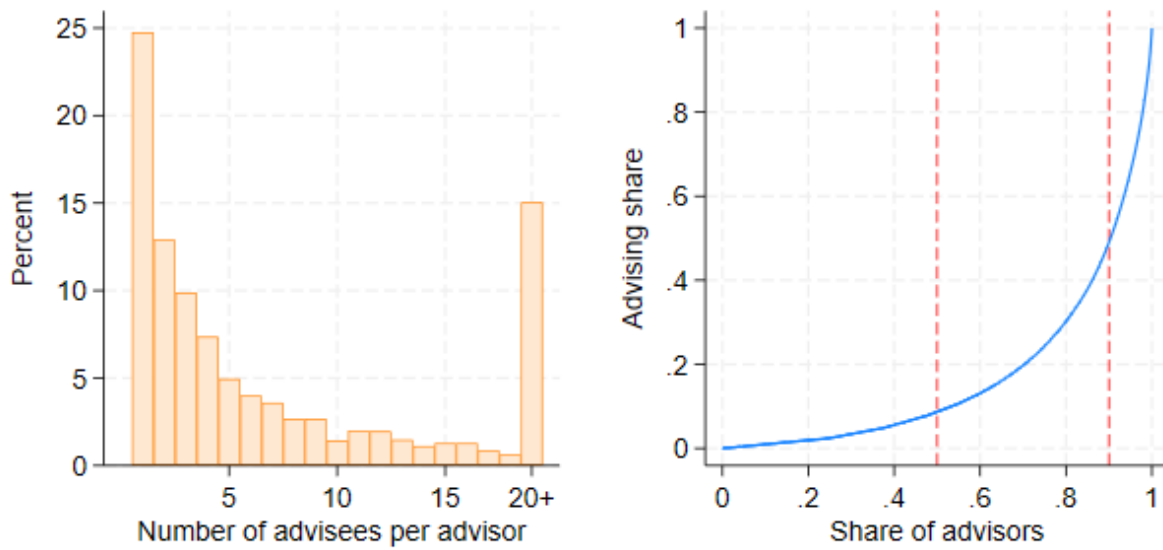
Notes: See notes to Figure 8. This figure plots the share of students with at least one publication in the Deep Impact journal list by advisee gender and advisor gender in year $t - c$, where c is graduation year. Graduates marked as having male advisors were advised by all-male advising teams. Graduates marked as having female advisors had at least one female advisor. Appendix Figure A3 plots publication counts in the same format. Data are for the economics + related sample of 1994-2017 graduates.

Figure 11. Gender Gaps in Annual Research Productivity by Cohort



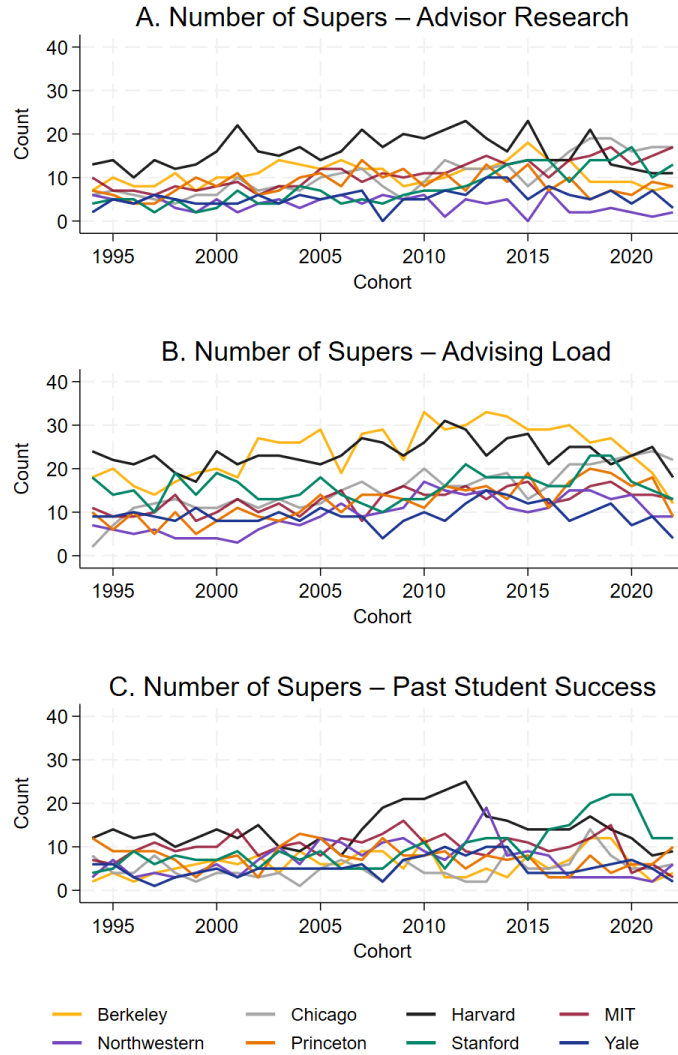
Notes: See notes to Figure 8. This figure plots gender gaps in the number of Deep Impact and Top 6 publications by cohort. For graduates in cohort c , publications are counted in years $c + 1$ to $c + 6$. The figure shows the difference in average productivity between male and female graduates, smoothed using a five-year moving average. Data are for the economics + related sample of 1994-2017 graduates.

Figure 12. The Advising Load Distribution



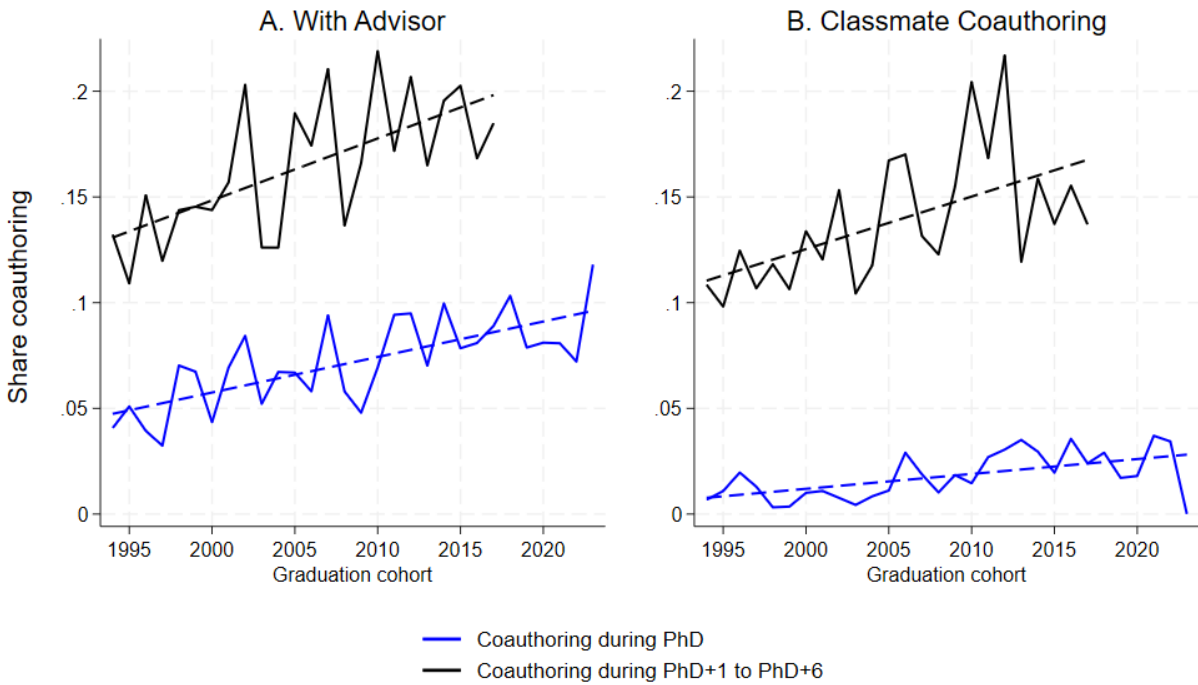
Notes: The left panel shows the histogram of number of advisees for advisors who advised an economics+related program graduate and were affiliated with one or more of the eight schools in our sample (advisors necessarily have at least 1 advisee who graduated from sample schools but need not be affiliated with a sample school or have advised an economics+related program graduate). Each advisor contributes one observation. The right panel shows the advising Lorenz curve: this orders advisors on the x-axis by number of advisees, with the cumulative share of advisees advised up to this point plotted on the y-axis. Red lines mark median and upper-decile advisors. Data are for 2499 advisors of economics + related program students who graduated 1989-2023; advisors in this sample have at least one EconLit publication from which an affiliation can be gleaned.

Figure 13. Super Advisors by Cohort and School



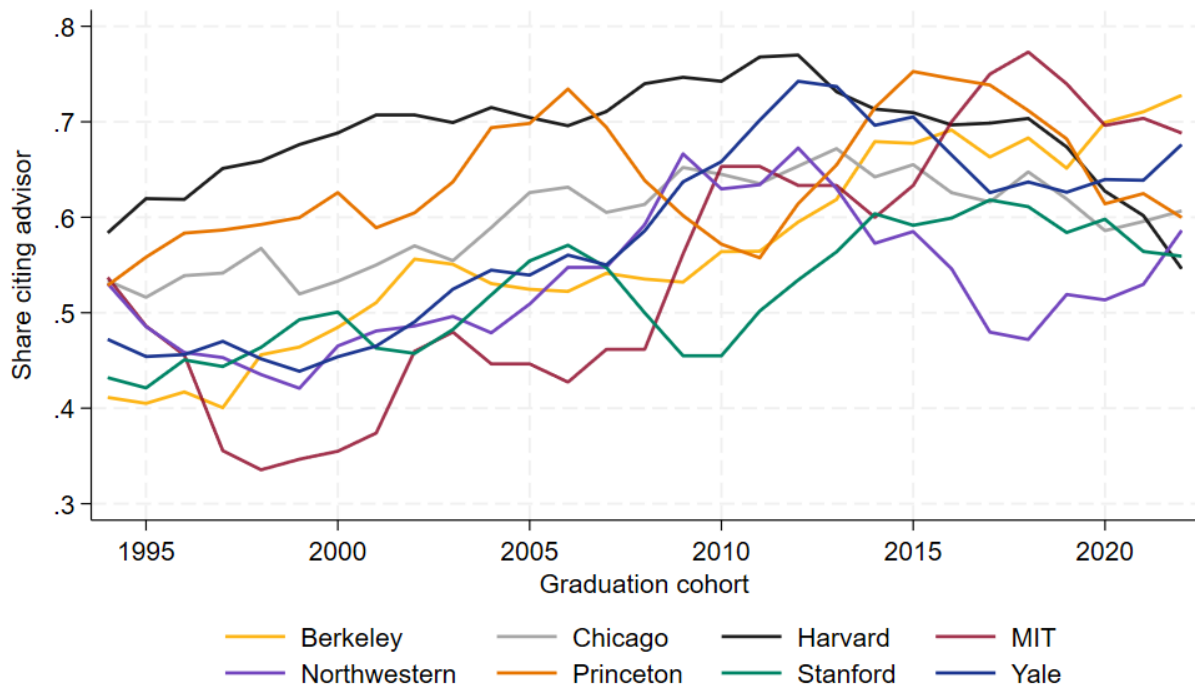
Notes: This figure plots the number of super advisors affiliated with a given school in a given year that advised at least one student in cohort *c*. Super advisors are classified by cohort using advisor attributes in the five years preceding the cohort’s graduation year. Successful-research advisors have published many Deep Impact articles; advising-loaded advisors have advised many PhDs; successful-student advisors have advised graduates who published many Deep Impact articles in the six years following their graduation. Super advisors are in the top 10% of advisors in these categories, defined separately for each distribution. Advisor distributions are for the universe of economics PhDs, not limited to graduates of economics+related programs.

Figure 14. Early Career Advisor-Advisee and Classmate Coauthoring by Cohort



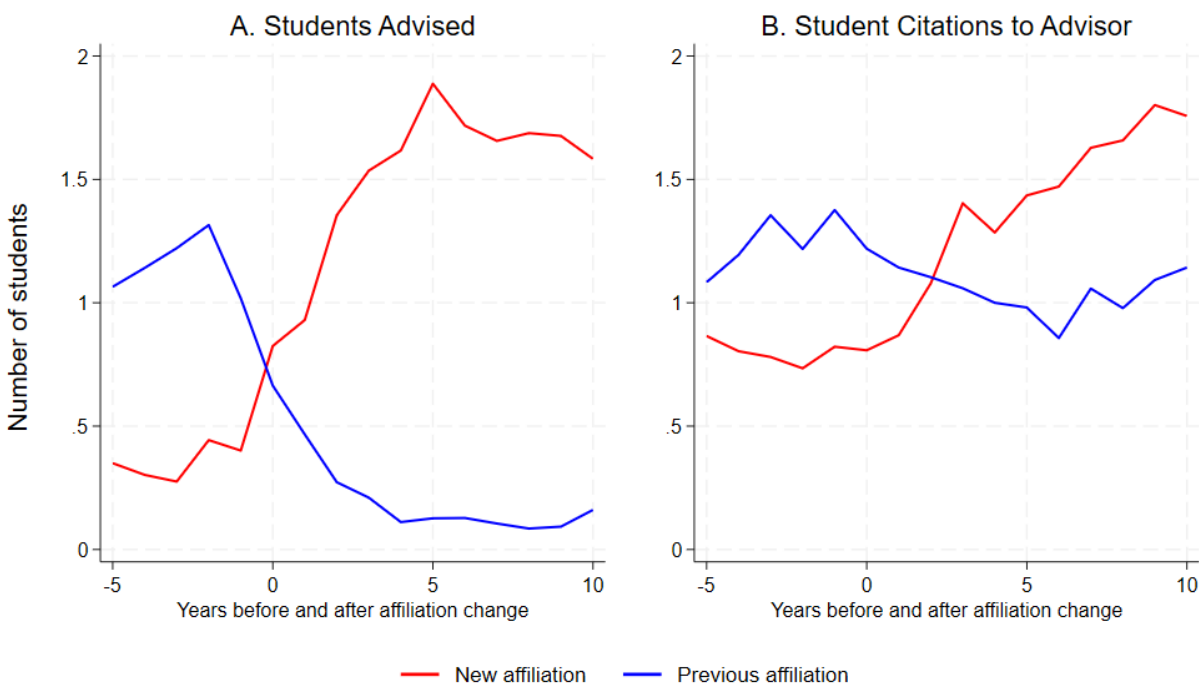
Notes: The left panel shows trends in student-advisor coauthoring; the right panel shows trends in PhD student coauthoring with classmates. Classmates are defined as students from the same school who graduated in the same year or within two years before or after. Dashed blue lines mark the share of students with a pre-graduation publication authored jointly with an advisor or classmate including in the graduation year. Dashed black line plots the share of students with at least one joint publication in the six years following graduation, excluding the graduation year. Coauthoring is determined from all publications indexed in EconLit for 1994-2023 cohorts.

Figure 15. Share of Theses Citing One or More Advisors by Cohort and School



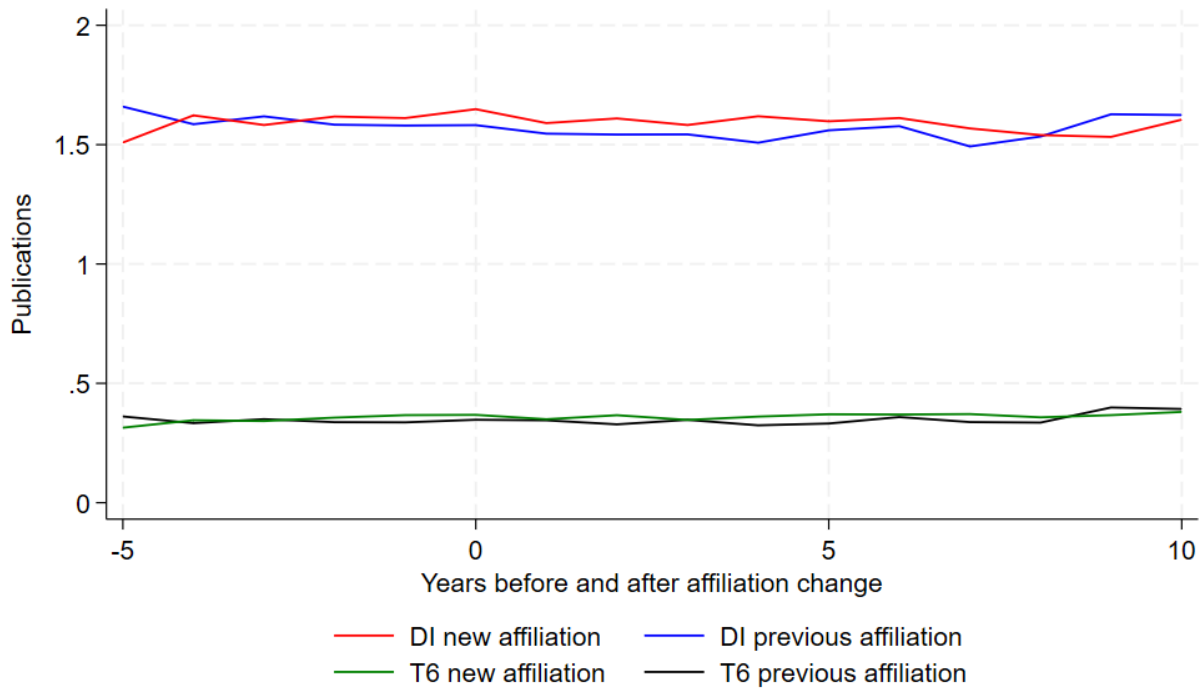
Notes: This figure plots the share of PhD theses that cite an advisor's work, computed for a sample of roughly 6000 economics + related program graduates with thesis PDFs in the 1994-2023 cohorts. Data are smoothed using a five-year moving average. Citations are identified by titles in bibliographic material; see the data appendix for details.

Figure 16. The Effect of Advisor Affiliation Changes on Graduate Students Advised and Student Citations to Advisor Works



Notes: The left panel shows the number of graduates advised by advisors who change affiliations, with separate counts for cohorts graduating at the former and new affiliation before and after a transition year. The right panel plots the number of graduates citing a transitioning advisors' work at the previous and new affiliation. The analysis includes 367 fixed-super advisors (ranked by advisee count) with at least one affiliation change, as indicated by the affiliation variable on their EconLit publications. The sample is restricted to research-active advisors whose affiliations are observed in at least one-third of their research-active years (the years between their first and last publications). Advisors in this sample advised at least one student as of an affiliation change.

Figure 17. The Effect of Advisor Affiliation Changes on Average Cohort Publications



Notes: This figure plots overall school-by-cohort averages of Deep Impact and Top 6 publication output in years $c + 1$ to $c + 6$ for graduates in cohort c . As in the previous figure, these are computed separately for schools losing or gaining a transitioning advisor, looking at cohorts completing degrees before and after a transition year. Transitioners here are 385 fixed-duper advisors (ranked by past student success) with at least one affiliation change, as indicated by the affiliation variable on their EconLit publications, limited to research-active advisors whose affiliations are observed in more than one-third of their research-active years, with at least one advisee graduating ahead of the affiliation change. Publication statistics are for all graduates from affected schools.

Tables

Table 1. Cohort Size and Activity Rates

	Economics			Related		
	Total (1)	Cohort size (2)	Active (3)	Total (4)	Cohort size (5)	Active (6)
Panel A. 1989-2023 Graduates						
Berkeley	1119	32.0		158	4.5	
Chicago	897	25.6		155	4.4	
Harvard	1140	32.6		169	4.8	
MIT	867	24.8		218	6.2	
Northwestern	675	19.3		75	2.1	
Princeton	650	18.6		39	1.1	
Stanford	761	21.7		324	9.3	
Yale	660	18.9		100	2.9	
All	6769	24.2		1238	4.4	
Panel B. 1994-2017 Graduates						
Berkeley	854	35.6	.45	103	4.3	.34
Chicago	681	28.4	.37	105	4.4	.52
Harvard	792	33.0	.53	91	3.8	.41
MIT	633	26.4	.56	151	6.3	.30
Northwestern	485	20.2	.53	54	2.3	.48
Princeton	467	19.50	.56	27	1.1	.44
Stanford	532	22.2	.47	215	9	.44
Yale	454	18.9	.50	52	2.2	.29
All	4898	25.5	.49	798	4.2	.40

Notes: Columns 1 and 4 show the number of PhDs awarded for economics subject matter from 1989-2023 and 1994-2017 for graduates of economics and related programs. Column 2 shows the average economics program cohort size; column 5 shows the average number of related-program graduates. Columns 3 and 6 in Panel B show the share of graduates who are active post-PhD, meaning they have at least one Deep Impact publication in the six years after graduation (years $c + 1$ to $c + 6$ in cohort c).

Table 2. Students of Research-Active, Prolific Advisors Do Better

	Levels				Poisson			
	DI		T6		DI		T6	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
A. Research Productivity								
Advisor Research	0.12*** (0.0099)		0.040*** (0.0044)		0.068*** (0.0057)		0.088*** (0.0100)	
Super Advisor Research		0.69*** (0.068)		0.21*** (0.025)		0.48*** (0.048)		0.68*** (0.088)
B. Advising Load								
Advising Load	0.032*** (0.0066)		0.014*** (0.0026)		0.020*** (0.0039)		0.033*** (0.0062)	
Super Advising Load		0.43*** (0.075)		0.17*** (0.028)		0.32*** (0.060)		0.69*** (0.13)
C. Successful Students								
Past Student Success	0.43*** (0.040)		0.14*** (0.016)		0.26*** (0.023)		0.36*** (0.039)	
Super Past Student Success		0.56*** (0.070)		0.18*** (0.030)		0.36*** (0.045)		0.51*** (0.086)
Dep. var. mean	1.51		0.34		1.51		0.34	

Notes: The dependent variable in columns 1, 2, 5, and 6 is student Deep Impact publications in the six years after graduation ($c + 1$ to $c + 6$ for cohort c). The dependent variable in columns 3, 4, 7, and 8 sums top-6 publications. *Research productivity* regressors average a cohort- c graduate’s advisors’ Deep Impact publications in the five years before c . *Advising load* regressors average a graduate’s advisors’ advisee counts in the five years before c . For a given graduate, the *successful student* regressor averages the number of Deep Impact publications by a graduate’s advisors’ past advisees as described by equation 4 in the text. “Super” regressors are dummies indicating whether at least one of a student’s advisors was among the 10% most prolific advisors relative to all other advisors who advised students graduating in the five years preceding the student’s graduation year. Models for each panel are run separately. All models control for school and cohort effects, graduates’ advisor team size, graduate gender, unclassified gender, missing thesis PDF, and an economics department dummy. Estimates are for the economics + related sample of 5682 students that graduated between 1994 and 2017. Standard errors are clustered on school-by-cohort (192 clusters for 8 schools \times 24 cohorts).

Table 3. Multivariate Model of Advisor Effects

	Levels				Poisson			
	DI		T6		DI		T6	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Advisor Research	0.096*** (0.012)		0.031*** (0.0049)		0.055*** (0.0068)		0.066*** (0.011)	
Advising Load	-0.0074 (0.0078)		0.0011 (0.0030)		-0.0021 (0.0052)		0.0062 (0.0079)	
Past Student Success	0.30*** (0.043)		0.091*** (0.017)		0.19*** (0.024)		0.28*** (0.039)	
Super Advisor Research		0.47*** (0.10)		0.10*** (0.038)		0.34*** (0.069)		0.40*** (0.13)
Super Advising Load		0.13 (0.085)		0.062* (0.033)		0.12* (0.066)		0.36** (0.15)
Super Past Student Success		0.28*** (0.076)		0.073** (0.032)		0.20*** (0.048)		0.24** (0.093)
Duper Advisor Research		0.57*** (0.078)		0.17*** (0.027)		0.39*** (0.054)		0.53*** (0.092)
Duper Advising Load		0.25*** (0.082)		0.12*** (0.031)		0.19*** (0.064)		0.51*** (0.14)
Duper Past Student Success		0.58*** (0.10)		0.20*** (0.042)		0.35*** (0.059)		0.50*** (0.10)
Dep. var. mean	1.51		0.34		1.51		0.34	

Notes: Dependent variables and sample are as described in the note to Table 2. Each column reports estimates from a single regression model with multiple advisor attributes on the right-hand side. Super and duper dummies indicate graduates with at least one advisor with characteristics in percentiles 6-10 (super) and 1-5 (duper). Standard errors are clustered on school-by-cohort (192 clusters). Models control for school and cohort effects, advisor team size, graduate gender, unclassified gender, missing thesis PDF, and an economics department dummy.

Table 4. Effects of Coauthoring and Research Affinity, Continuous Advisor Attributes

	Levels				Poisson			
	DI		T6		DI		T6	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Advisor Research	0.092*** (0.012)	0.086*** (0.012)	0.029*** (0.0048)	0.028*** (0.0048)	0.052*** (0.0067)	0.048*** (0.0066)	0.061*** (0.011)	0.056*** (0.011)
Advising Load	-0.0042 (0.0076)	-0.0035 (0.0075)	0.0022 (0.0029)	0.0024 (0.0029)	-0.00025 (0.0051)	0.00030 (0.0050)	0.0097 (0.0076)	0.010 (0.0074)
Past Student Success	0.29*** (0.043)	0.28*** (0.043)	0.087*** (0.016)	0.084*** (0.016)	0.19*** (0.023)	0.18*** (0.023)	0.27*** (0.038)	0.26*** (0.038)
Coauthored with Advisor (pre grad.)	-0.12 (0.24)	-0.15 (0.24)	0.16 (0.14)	0.15 (0.14)	-0.070 (0.082)	-0.087 (0.083)	0.13 (0.16)	0.099 (0.16)
Coauthored with Classmate (pre grad.)		0.56 (0.47)		0.14 (0.27)		0.13 (0.14)		0.091 (0.29)
Any Publication (pre grad.)	1.75*** (0.16)	1.68*** (0.16)	0.46*** (0.068)	0.44*** (0.073)	0.81*** (0.054)	0.79*** (0.058)	0.90*** (0.096)	0.88*** (0.10)
Cites Advisor in Thesis		0.22*** (0.067)		0.071*** (0.026)		0.16*** (0.046)		0.27*** (0.087)
Dep. var. mean	1.51		0.34		1.51		0.34	

Notes: Dependent variables and sample are as described in the note to Table 2. Models control for school and cohort effects, graduates' advisor team size, graduate gender, unclassified gender, missing thesis PDF, and an economics department dummy. Standard errors are clustered on school-by-cohort (192 clusters). Of the 5682 economics+related program graduates in the 1994-2017 cohorts, 547 published by the end of graduation year. Of those, 187 coauthored with an advisor and 58 with a classmate. *Cites Advisor in Thesis* is a dummy variable indicating PhD theses that cite an advisor's work.

Table 5. Effects of Coauthoring and Research Affinity, Fixed-Duper Advisors

	Levels				Poisson			
	DI		T6		DI		T6	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Fixed Duper Advisor Research	0.45*** (0.073)	0.41*** (0.073)	0.14*** (0.023)	0.13*** (0.023)	0.34*** (0.060)	0.31*** (0.060)	0.57*** (0.11)	0.53*** (0.11)
Fixed Duper Advising Load	0.087 (0.081)	0.084 (0.080)	0.10*** (0.025)	0.10*** (0.025)	0.11 (0.069)	0.11 (0.070)	0.62*** (0.15)	0.62*** (0.15)
Fixed Duper Past Student Success	0.87*** (0.065)	0.85*** (0.065)	0.16*** (0.017)	0.16*** (0.017)	0.90*** (0.074)	0.89*** (0.073)	0.98*** (0.12)	0.96*** (0.12)
Coauthored with Advisor (pre grad.)	-0.14 (0.24)	-0.18 (0.24)	0.16 (0.15)	0.15 (0.15)	-0.087 (0.079)	-0.11 (0.078)	0.12 (0.16)	0.087 (0.16)
Coauthored with Classmate (pre grad.)		0.63 (0.46)		0.17 (0.26)		0.17 (0.13)		0.15 (0.28)
Any Publication (pre grad.)	1.76*** (0.16)	1.68*** (0.16)	0.46*** (0.069)	0.44*** (0.073)	0.82*** (0.054)	0.79*** (0.056)	0.93*** (0.091)	0.90*** (0.096)
Cites Advisor in Thesis		0.27*** (0.065)		0.095*** (0.026)		0.18*** (0.045)		0.32*** (0.086)
Dep. var. mean	1.51		0.34		1.51		0.34	

Notes: Dependent variables and sample are as described in the note to Table 2. Models control for school and cohort effects, graduates' advisor team size, graduate gender, unclassified gender, missing thesis PDF, and an economics department dummy. Fixed-duper dummies indicate advisors ever coded duper for a given advisor attribute. Within advisors, fixed-duper dummies are cohort-invariant.

Table 6. 2SLS Estimates Using 49 School \times Cohort Dummies as Instruments

	All features instrumented				First 3 instrumented			
	DI		T6		DI		T6	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Advisor Research	0.16*** (0.053)	0.16*** (0.051)	0.045* (0.024)	0.047* (0.024)	0.17*** (0.052)	0.16*** (0.051)	0.052** (0.023)	0.052** (0.022)
Advising Load	-0.018 (0.026)	-0.019 (0.025)	-0.0076 (0.012)	-0.0075 (0.012)	-0.026 (0.026)	-0.027 (0.026)	-0.014 (0.011)	-0.014 (0.010)
Past Student Success	0.23 (0.19)	0.17 (0.20)	0.100 (0.069)	0.085 (0.071)	0.24 (0.18)	0.22 (0.18)	0.11* (0.064)	0.10 (0.063)
Coauthored with Advisor (pre grad.)	0.30 (2.08)	-0.075 (2.14)	1.46* (0.82)	1.34 (0.84)	-0.16 (0.25)	-0.17 (0.24)	0.14 (0.14)	0.14 (0.14)
Coauthored with Classmate (pre grad.)		1.49 (2.99)		0.98 (1.30)		0.51 (0.46)		0.12 (0.26)
Any Publication (pre grad.)	0.94 (1.16)	0.68 (1.23)	0.20 (0.47)	0.061 (0.49)	1.74*** (0.16)	1.68*** (0.16)	0.45*** (0.068)	0.44*** (0.073)
Cites Advisor in Thesis		0.72 (0.67)		0.088 (0.26)		0.14 (0.11)		0.044 (0.044)
Dep. var. mean	1.51		0.34		1.51		0.34	

Notes: The first 4 columns report 2SLS estimates computed using 49 dummies for schools \times 3-year cohorts as instruments, with all listed variables plus advisor team size instrumented. Estimates in columns 5-8 are from models instrumenting the first three advisor attributes only. Models control for school and cohort effects, graduates' advisor team size (instrumented in the first 4 columns), graduate gender, unclassified gender, missing thesis PDF, and an economics department dummy. Standard errors are clustered on school-by-cohort.

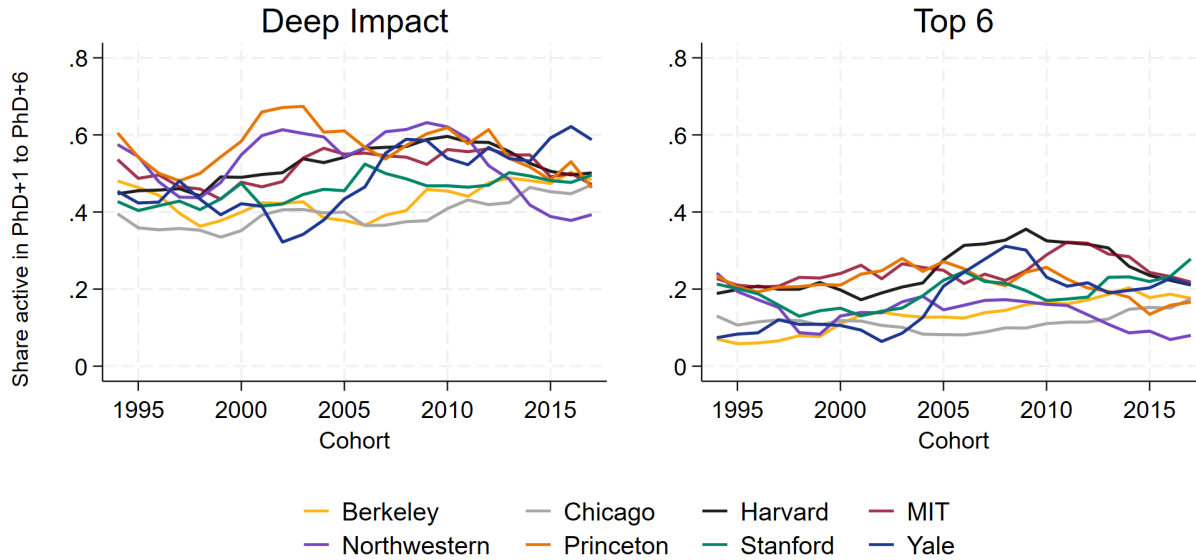
Table 7. Effects of Cohort Size on Aggregate Research Productivity

	Deep Impact				Top 6			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Panel A. Economics + Related Graduates								
Cohort Size	1.34*** (0.13)	1.53*** (0.18)	1.54** (0.66)	1.40** (0.59)	0.38*** (0.059)	0.30*** (0.074)	0.76*** (0.28)	0.24 (0.23)
Cohort Size Squared			-0.0033 (0.011)	0.0021 (0.0098)			-0.0063 (0.0047)	0.0009 (0.0039)
Panel B. Economics Graduates								
Cohort Size	1.41*** (0.16)	1.60*** (0.22)	1.68** (0.77)	1.31* (0.69)	0.39*** (0.072)	0.33*** (0.085)	0.77** (0.32)	0.14 (0.26)
Cohort Size Squared			-0.0049 (0.014)	0.0053 (0.013)			-0.0068 (0.0060)	0.0034 (0.0051)
School effects	No	Yes	No	Yes	No	Yes	No	Yes

Notes: This table reports estimates from regressions of school-by-cohort average publication output in $c + 1$ to $c + 6$ on the size of cohort c . Dependent variables sum either Deep Impact or Top 6 publications for cohorts graduating 1994-2017. The estimation sample consists of 192 school-by-year groups. Panel A reports estimates for models averaging the sample of economics+related graduates. Panel B reports estimates for averages computed using identifiable economics department graduates only. All specifications include cohort effects. Models reported in even-numbered columns include school effects. Robust standard errors appear in parentheses.

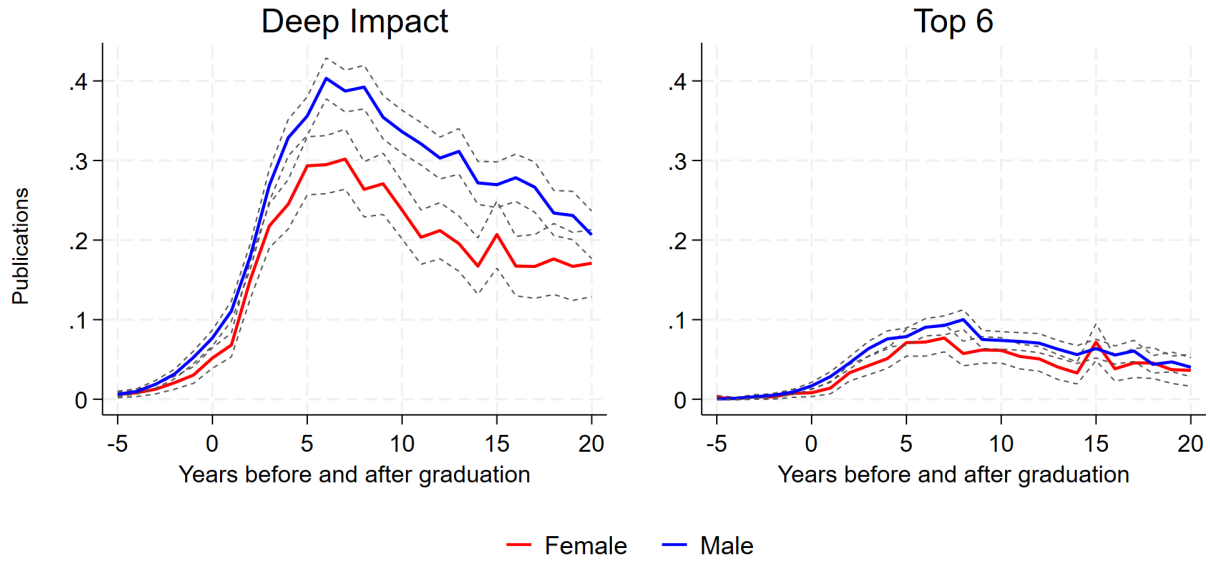
Appendix: Additional Exhibits

Figure A1. Research Activity by Cohort and School



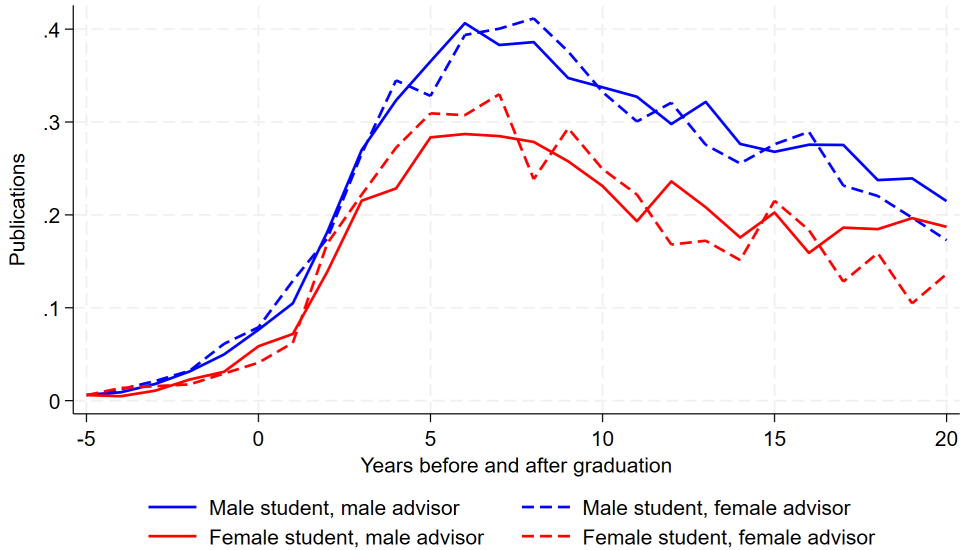
Notes: This figure plots the share of graduates with at least one publication in the first six years post-PhD ($c + 1$ to $c + 6$) for the economics + related sample. Data for 1994-2017 graduates.

Figure A2. Gender Gaps in Graduate Student Research Productivity Profiles



Notes: See notes to Figure 8. This figure plots publication counts by male and female students in year $t - c$, where c is graduation year. Data are for the economics + related sample of 1994-2017 graduates.

Figure A3. Annual Productivity Profiles by Grad Student and Advisor Gender



Notes: See notes to Figure 8. This figure plots mean publications by advisee and advisor gender in year $t - c$, where c is graduation year. Graduates marked as having male advisors were advised by all-male advising teams. Graduates marked as having female advisors had at least one female advisor. Research productivity in t is defined by the number of Deep Impact publications in that year. Economics + related sample of 1994-2017 graduates.

Table A1. Descriptive Statistics for the Sample Used in Tables 2-5

	Mean	Std. Dev.
Deep Impact publications in c+1 to c+6	1.51	2.29
Top 6 publications in c+1 to c+6	0.34	0.92
Advisor Research	5.77	3.79
Advising Load	8.27	5.84
Past Student Success	1.44	0.94
Super Advisor Research	0.54	0.50
Super Advising Load	0.77	0.42
Super Past Student Success	0.43	0.50
Duper Advisor Research	0.38	0.48
Duper Advising Load	0.59	0.49
Duper Past Student Success	0.20	0.40
Fixed-Duper Advisor Research	0.70	0.46
Fixed-Duper Advising Load	0.84	0.37
Fixed-Duper Past Student Success	0.77	0.42
Any Publication Pre-grad	0.096	0.29
Coauthored with Advisor Pre-grad	0.033	0.18
Coauthored with Classmate Pre-grad	0.010	0.10
Number of Advisors	3.08	0.73
Economics Dept Graduate	0.86	0.35
Unclassified Gender	0.14	0.34
No thesis PDF	0.069	0.25
Female	0.27	0.45
Cites Advisor in Thesis	0.60	0.49

Notes: This table reports means and standard deviations for the sample of 5682 PhD students graduating from economics and related programs between 1994-2017. The share female is calculated for 4910 students for whom gender is classified. The share citing an advisor is calculated for 5291 students with thesis PDFs.

References

- Angrist J, Azoulay P, Ellison G, Hill R, Lu SF.** 2017. “Economic Research Evolves: Fields and Styles.” *American Economic Review*, 107(5): 293–297.
- Angrist J, Azoulay P, Ellison G, Hill R, Lu SF.** 2020. “Inside Job or Deep Impact? Extramural Citations and the Influence of Economic Scholarship.” *Journal of Economic Literature*, 58(1): 3–52.
- Azoulay P, Graff Zivin JS, Wang J.** 2010. “Superstar Extinction.” *The Quarterly Journal of Economics*, 125(2): 549–589.
- Brogaard J, Engelberg J, Van Wesep E.** 2018. “Do Economists Swing for the Fences after Tenure?” *Journal of Economic Perspectives*, 32(1): 179–194.
- Buchmueller TC, Dominitz J, Hansen WL.** 1999. “Graduate Training and the Early Career Productivity of PhD Economists.” *Economics of Education Review*, 18(1): 65–77.
- Card D, DellaVigna S.** 2013. “Nine Facts About Top Journals in Economics.” *Journal of Economic literature*, 51(1): 144–161.
- Card D, DellaVigna S, Funk P, Iriberry N.** 2020. “Are Referees and Editors in Economics Gender Neutral?” *The Quarterly Journal of Economics*, 135(1): 269–327.
- Collins JT, Cox RG, Stango V.** 2000. “The Publishing of Recent Economics PhD Recipients.” *Economic Inquiry*, 38(2): 358–367.
- Conley JP, Önder AS.** 2014. “The Research Productivity of New PhDs in Economics: The Surprisingly High Non-Success of the Successful.” *Journal of Economic Perspectives*, 28(3): 205–16.
- Corsini A, Pezzoni M, Visentin F.** 2022. “What Makes a Productive PhD Student?” *Research Policy*, 104561.
- Ellison G.** 2002. “The slowdown of the Economics Publishing Process.” *Journal of political Economy*, 110(5): 947–993.
- García-Suaza A, Otero J, Winkelmann R.** 2020. “Predicting Early Career Productivity of PhD Economists: Does Advisor-Match Matter?” *Scientometrics*, 122(1): 429–449.
- Gaule P, Piacentini M.** 2018. “An Advisor Like Me? Advisor Gender and Post-Graduate Careers in Science.” *Research Policy*, 47(4): 805–813.

- Hilmer MJ, Hilmer CE.** 2009. “Fishes, Ponds, and Productivity: Student-Advisor Matching and Early Career Publishing Success for Economics PhDs.” *Economic Inquiry*, 47(2): 290–303.
- Hilmer MJ, Hilmer CE.** 2011. “Is It Where You Go or Who You Know? On the Relationship Between Students, PhD Program Quality, Dissertation Advisor Prominence, and Early Career Publishing Success.” *Economics of Education Review*, 30(5): 991–996.
- Iaria A, Schwarz C, Waldinger F.** 2024. “Gender gaps in academia: Global evidence over the twentieth century.” *Available at SSRN 4150221*.
- Jones BF.** 2021. “The Rise of Research Teams: Benefits and Costs in Economics.” *Journal of Economic Perspectives*, 35(2): 191–216.
- Kasper H, et al.** 1991. “The Education of Economists: From Undergraduate to Graduate Studies.” *Journal of Economic Literature*, 29(3): 1088–1109.
- Liénard JF, Achakulvisut T, Acuna DE, David SV.** 2018. “Intellectual Synthesis in Mentorship Determines Success in Academic Careers.” *Nature Communications*, 9(1): 1–13.
- Li W, Aste T, Caccioli F, Livan G.** 2019. “Early Coauthorship with Top Scientists Predicts Success in Academic Careers.” *Nature Communications*, 10(1): 1–9.
- Neumark D, Gardecki R.** 1998. “Women Helping Women? Role Model and Mentoring Effects on Female PhD Students in Economics.” *Journal of Human Resources*, 33(1): 220–246.
- Sarsons H, Gërkhani K, Reuben E, Schram A.** 2021. “Gender Differences in Recognition for Group Work.” *Journal of Political Economy*, 129(1): 101–147.
- 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.