The Race between Academia and Industry for AI Researchers*

Francesca Miserocchi† Savannah Noray‡ Alice Wu§

August 2025

Preliminary and incomplete. Please do not circulate.

Abstract

The advances of artificial intelligence (AI) are built on the substantial groundwork laid by researchers. In this paper, we study the labor market competition between academia and industry for AI researchers. Consistent with academia and industry increasingly rewarding the same talent, we find that over the past decade, publications in AI-related conferences are more predictive of employment in industry vis-a-vis academia. We also provide causal evidence via a matched difference-in-differences design, exploiting data from a leading AI conference that publicly releases submissions and referee reports. On average, a publication increases the chance of moving to a top firm by 2-4 percentage points in the next 1-3 years. Poaching by top firms is stronger for male academics, whereas female students and postdocs are more likely to get tenure-track positions following a successful publication. Researchers who move to industry are subsequently less productive in publishing than academics, suggesting that an increasing concentration of talent in industry could compromise the efficiency of research production.

*We are grateful to Antonio Coran for many helpful discussions, and to Anna Le Sun for outstanding assistance with the publications data. This research was funded by the Lab for Economic Applications and Policy (Harvard).

†Department of Economics, Tilburg University. F.Miserocchi@tilburguniversity.edu

‡Reimagining the Economy Lab, Harvard University.savannahnoray@g.harvard.edu

§Department of Economics, University of Wisconsin - Madison.hwu526@wisc.edu

1 Introduction

The advances of technology rely heavily on academic research. Traditionally, most research works are produced in academia, but in one of the most important areas today — artificial intelligence (AI), almost half of the publications at academic conferences involve researchers from industry, driven by scientists hired by the largest tech firms. Despite the large and growing literature on the economic consequences of AI (Acemoglu 2024; Furman and Seamans 2019), less is known about the researchers who produce the original ideas that advance these technologies. The allocation of these researchers across employers could influence the private versus social value of AI research given the different incentives in academia versus the private sector (Aghion, Dewatripont, and Stein 2008).

This paper aims to examine the labor market competition for AI researchers between academia and industry, study how researchers are rewarded for their publications, and investigate the consequences of the research talent allocation on the direction and quality of research. By building a longitudinal dataset on the career trajectories of over 270,000 researchers who publish at computer science (CS) conferences, we analyze the initial sorting of graduates into different types of employers, and the role of publications in driving this sorting and subsequent job movements. We find that AI publications have become more predictive of movements to top firms than to academia. We then corroborate this finding by leveraging data from a leading AI conference that has publicly released all submissions and referee reports since 2017. Using a difference-in-differences design, we find an accepted paper increases sorting into top firms, especially among students and postdocs. The reallocation of researchers from academia to industry has implications on future research: movers to top firms publish fewer papers afterwards.[1]

Starting with publications data from 89 major CS conferences, we matched individual authors with professional profiles from Revelio, and we constructed the education and job history for 276,600 CS authors. We begin by showing trends in the job placement of these authors within three years of obtaining their highest degrees. The fraction of new graduates joining academia within three years of graduation went down from 55 percent to 45 percent between 2005 and 2020. Meanwhile, the share of graduates joining one of the top tech firms, defined as the six big tech firms {Alphabet, Meta, Amazon, Apple, Microsoft, IBM}, increased from 9 percent to 15 percent, a 67% increase within a 15-year window. These opposing trends highlight a substantial reallocation of CS research talent from academia toward industry. These trends are driven by graduates who are from top schools or are positively selected in research productivity prior to graduation, suggesting that top tech firms successfully recruit research talent in the beginning of their career.

How do industry and academic employers identify research talent? We show that publications

[1] We will continue to examine these consequences (on the direction and quality of research) as this project develops.

play an important role in the initial sorting of workers between industry and academia. Publications carry information about the research ability of authors. We estimate the relationship between publications before graduation and the characteristics of a person's first primary job post highest degree, conditional on observable characteristics such as demographics and school and cohort fixed effects. Publications at AI conferences are strongly predictive of being placed in industry. Every additional AI paper before graduation increases the chance of working in industry by 0.4 percentage points (henceforth pp), and increases the chance of working for a top firm by 0.6 pp, which is an over 8% increase relative to the mean. The correlation between AI publications and getting hired by top firms is the strongest among PhD graduates, consistent with the sorting patterns among CS PhDs in Wu (2025). In contrast, AI papers are less predictive of being placed in academia relative to papers published in other CS areas, suggesting this poaching is specific to AI research. In fact, conditional on the overall number of papers, publishing in AI conferences is negatively predictive of working in academia for PhDs. We further show that AI publications continue to matter for subsequent sorting in the labor market, predictive of upward mobility from other firms in the industry (nontop) to top firms, and of movements from academia to industry or to top firms in particular.

These results suggest that top tech firms attract researchers who publish at AI conferences and are able to hire from both industry and academia. However, publications may simply proxy for underlying ability and need not, by themselves, lead to a job at top firms. To address this endogeneity concern, we exploit the International Conference on Learning Representations (ICLR), one of the leading AI conferences that publicly releases all submitted papers, referee reports, and editor decisions. The availability of rejected papers allows us to use a difference-in-differences design to estimate the causal impacts of a publication on the career trajectories of researchers.

To compare authors who submit papers of similar quality, we use both the numeric ratings by referees and the text of the referee comments to match authors on accepted papers with authors on rejected ones. In addition, we also match authors by their original jobs (industry/academia/still in school) before submitting to ICLR for the first time, and the predicted propensity score of having an accepted paper based on demographics and publication history. The combination of coarsened exact matching and propensity score matching allows us to match 4,401 authors whose first submissions to ICLR were accepted with 943 rejected authors (drawn with replacement). On average the matched authors are more balanced in the quality of their submissions and observable characteristics. They also show parallel trends in employment before submitting to the conference.

Under the identifying assumption that the matched accepted and rejected authors would have followed parallel career trajectories, we use a difference-in-differences design to estimate the impacts of an accepted publication on career outcomes, such as sorting to top firms or getting a tenure-track job. Among authors who are not employed by top tech firms one year before submitting to

ICLR, the acceptance by the conferences (resulting in a publication) increases the movements into top firms by 3 pp on average, within three years after the conference. Event study estimates show that the increase in sorting into top firms among accepted authors kicks in one year after the ICLR conference, and continues to rise from 2 pp to 4 pp in the next two years. We further estimate separate event studies based on the original employer/affiliation of the authors in the year before the ICLR conference. Relative to authors employed in academia, the estimated effects of an accepted paper on movement to top firms are larger for authors who are already in industry, and for students and postdocs who are yet to settle in their first jobs.

Given the policy focus on gender differences in STEM, we examine heterogeneity by gender in our main result—that paper acceptance increases mobility to top firms. We find that acceptance raises the probability that female students secure tenure-track jobs in academia but does not increase their moves to top firms. For male students, the pattern reverses: acceptance increases mobility to top firms, with no corresponding boost into tenure-track academia. These divergent responses may reflect many forces (e.g., diversity initiatives in CS academia or gender differences in preferences over sector-specific amenities). Regardless, if men and women differ in the type or volume of research they produce, a gender-skewed shift toward industry could be a mechanism for reshaping the composition and quantity of public AI research. We plan to explore how this heterogeneity mediates the consequences of this race in future work.

Lastly, we study the consequences of the growing competition between academia and industry for top AI researchers. Using event-study analyses of computer science researchers with linked publication and job-history data, we compare those who transition to top tech firms with those who move into tenure-track academic positions. Prior to the move, both groups display similar upward publication trends. Following the move, research output declines modestly but rebounds for those entering academia, while it falls sharply and persistently for those joining top firms. Building on this descriptive evidence, we will further examine how such transitions affect research quality, measured by citations, and research direction, testing whether moves to industry shift effort away from topics such as algorithmic bias toward more applied domains. To strengthen identification, we will compare industry movers with a matched control group of similar researchers who remain outside of top firms.

This paper contributes to several strands of the literature. First, we contribute to the literature that examines the career progression of scientists or researchers, and especially the role of publications in driving promotions and sorting between employers. Publications are arguably the most important measure of research ability. Research with a focus on the economics profession has shown the value of publications in getting tenured (Sarsons 2017), and being awarded as a fellow of the Econometrics Society (Card, DellaVigna, Funk, and Iriberri 2022). We expand this area of research by looking at a richer set of career outcomes in both academia and industry. Job

placements in industry have become more common in computer science and economics; yet there has been little research on whether and how much industry employers value academic publications. We fill in this gap by showing that in the field of computer science, not only does industry value publications, but top firms are more successful in recruiting researchers publishing in AI than academia.

Second, in the context of computer science, we contribute to a growing area of research on the brain drain of professors in AI fields from academia. Although professors may continue to produce high-quality research when collaborating with industry (Yue 2024), the poaching of AI professors by tech firms has negative spillover on students, reducing students' startup formation (Gofman and Jin 2022). We document that the race between academia and industry not only affects senior researchers but junior researchers such as students and postdocs. AI publications are more predictive of an initial job placement in top firms than of a placement in academia among PhD graduates. This finding is robust in the quasi-experiment based on a leading AI conference. These findings suggest that the race starts as early as the beginning of a person's career, and could affect the longer-term research production if talented researchers are concentrated at top firms from the beginning of their careers.

Through ongoing work that investigates the impacts of sorting to top firms vs. academia on future research, we also aim to speak to the literature on industry–academia research trade-offs (Aghion, Dewatripont, and Stein 2008). Jurowetzki et al. (2021) finds that researchers who move to industry only experience a temporary increase in citations. In the long term, both the direction and the quality could differ between research produced in academia and industry. Privately funded research may also lead to inefficiently low creative destruction, as evidenced by high-tech entrepreneurship arising from universities (Babina, He, Howell, Perlman, and Staudt 2020).

Section 2 describes the data we use in our analysis. Section 3 provides descriptive evidence on the relationship between job placement or mobility and paper publications. Section 4 presents our difference-in-difference analysis on submissions to ICLR. We analyze how researchers' output changes after a move to industry in Section 5. We conclude by discussing future directions for this project.

2 Data

We collected data on the career trajectories of computer science researchers who published at 89 CS conferences.[2] Our data has three main sources: (1) Scopus for CS publications, (2) Revelio

[2]We include all conferences that are used to rank computer science departments by CSRankings. The conferences listed on CSRankings were developed in consultation with faculty across a range of institutions, including via community surveys.

individual profiles for education and job history, and (3) OpenReview for submissions to the International Conference on Learning Representations (ICLR), a top AI conference that makes public referee reports for accepted and rejected papers.

Publications. The most prestigious publication venues in computer science are often conferences rather than journals. We collected conference proceedings on Scopus, a large-scale publication database produced by Elsevier.[3] Each paper comes with a complete list of authors and their affiliations, which indicate the academic affiliation or the employer of an author at the time of publication. Between 2000 and 2023, there are almost 900,000 unique authors (full names) across the globe, and 36.3% of whom have at least one publication at an AI-related conference.[4]

Education and Job Histories. We used individual-level data collected by Revelio Labs to obtain education and employment histories self-reported on the largest online professional network. We matched computer science authors with individual profiles by full name and affiliations, and ranked potential matches based on the number of affiliations in common. We established over 276,600 one-to-one matches between CS authors on Scopus and Revelio individual profiles, and focused on roughly 171,000 matches that report at least one post-secondary degree, and at least one job on their profiles. 56% of the matched individuals report a PhD, and 31% report a master's as the person's highest degree. About 45% of the CS authors report living or having worked in the US, and the fraction with a PhD degree is 61%. Outside the US, Germany, India, the UK, China, and Canada each account for 3-6% of the matched CS authors.

We built a person$\times$ year panel for matched authors across the globe. The panel is balanced: for each person, it includes all years between max{2000, year of graduation} and 2023. The primary affiliation at person$\times$ year level is a person's university if she is still a student, or her main employer as reported in her job history. Section 3 will present our findings on the initial job placement of CS authors, and the predictive relationship between publications and career outcomes in industry or academia based on this panel (henceforth "Scopus-Revelio").

International Conference on Learning Representations (ICLR) Based on the h5-index on Google Scholar, ICLR is ranked as the second most impactful conference in artificial intelligence (below NeurIPS and above ICML). ICLR is also one of the few major conferences that makes public accepted and rejected submissions, using a double-blind, open peer-review process on the OpenReview platform. This allows us to observe detailed review materials, including referee scores,

[3] We are especially grateful to Anna Sun for her help with the Scopus Search API.

[4] AI-related conferences include 24 conferences in AI, machine learning, computer vision, natural language processing, and information retrieval on the web.

written reports, and editorial decisions. We scraped such information on the universe of ICLR submissions between 2017 and 2023 on OpenReview.

Similar to "Scopus-Revelio", we also matched ICLR authors with their education and job histories in the Revelio data by full name and affiliations. We established over 20,000 one-to-one matches between ICLR authors and Revelio individual profiles. 91% of the matches report at least one job on their online professional profiles. We further merged the ICLR-Revelio sample with publications from Scopus and Web of Science to keep track of individual research productivity before and after an ICLR submission.

3 Descriptive Analysis: Job Placements and The Role of Publications

3.1 Job Placement of CS Graduates in Industry versus Academia

We document trends in job placement of computer science graduates in academia and industry over the past 15 years. An increasing fraction of new graduates has joined top technology firms as researchers rather than pursuing academic careers.[5] Figure 1 plots the share of graduates employed at top firms (left y-axis) and in tenure-track academic positions (right y-axis) within three years of degree completion, by graduation cohort. In 2005, roughly 9 percent of graduates were employed by top firms within three years, compared to 55 percent in academia. Since then, top industry placements have risen to about 15 percent in 2020 (a 67 percent increase), while academic placements have declined to 45 percent (an 18 percent decrease). These opposing trends highlight a substantial reallocation of CS research talent from academia toward industry.

A natural question is whether the shift toward employment in top tech firms is driven by graduates with high research potential—suggesting increased competition between academia and industry for research talent—or instead by graduates with lower research potential or weaker interest in academic careers, as was more typical in earlier cohorts. We proxy research potential using two measures: (i) graduation from a top school, and (ii) having at least one publication prior to graduation (results using the latter proxy are similar and reported in Figure A.1 in the Appendix). Figure 2 replicates the earlier trends separately for graduates from top schools and from other institutions. Both the rise in top-firm placements (left panel) and the decline in academic placements (right panel) are present across all graduates, but are especially pronounced among those from top schools—consistent with increasing competition between academia and industry for top research talent.

[5] Top firms are defined as {Alphabet, Meta, Amazon, Apple, Microsoft, IBM} – the six big tech firms (before the more recent rise of Nvidia) throughout the paper.

3.2 Publications Predict Initial Placement and Job Mobility

How do industry and academic employers identify research talent? Conference publications carry information about the research ability of authors, and are often mentioned as preferred qualifications in research scientist jobs in industry (see Appendix Figure A.2 for example). To understand the role of publications in the initial sorting of workers between industry and academia, we estimate a regression of first job post highest degree on publications before graduation as follows:

$$y_i = \beta_0 + \beta_1 \text{Pub}_i + \beta_2 \text{AI-Pub}_i + \beta_3 \text{Top AI-Pub}_i + X_i'\Theta + u_i \quad (1)$$

We estimate (1) for a series of outcomes (y_i's) that characterize a person's job in the year after obtaining her highest degree (bachelor/master/PhD), including whether the job is in industry or academia. Among those placed in industry, we further ask whether the person is employed by one of the top tech firms, and whether she works as a research scientist or engineer according to her job title. For graduates who start working in academia, we look at whether she gets a postdoc or tenure-track position. We count the number of CS publications before graduation year (Pub_i), and the number of publications in AI conferences (AI-Pub$_i$). Top AI-Pub$_i$ represents the number of publications at top three AI conferences. [6] We also control for observable characteristics such as gender, race, school and cohort fixed effects.

Table 1 shows the OLS estimates for the main coefficients of interest: $(\beta_1, \beta_2, \beta_3)$ on the publication counts before graduation. Publications at AI conferences are strongly predictive of being placed in industry.[7] Every additional AI paper before graduation increases the chance of working in industry by 0.4 pp (column 1), and increases the chance of working for a top firm by 0.6 pp (column 2), which is an over 8% increase relative to the mean. AI papers are also positively correlated with becoming a research scientist (column 3), which often describes publishing at CS conferences as part of the job (Figure A.2). In contrast, AI papers before graduation are negatively correlated with becoming an engineer (column 4), a role that does not emphasize research.

The overall number of publications before graduation is positively related to getting a job in academia, as shown in columns 5-7. AI papers are less predictive of an academic placement relative to papers published in other CS areas. An additional top AI paper predicts a 0.79 pp higher chance of getting an academic job, but it does not strongly predict working as a postdoc or a tenure-track professor. Most academic jobs presumably require a PhD degree. We estimate

[6] We used conference ranking on Google Scholar based on h-index: https://scholar.google.com/citations?view_op=top_venues&hl=en&vq=eng_artificialintelligence. The three most impactful conferences in AI are: Conference on Neural Information Processing Systems (NeurIPS), International Conference on Learning Representations (ICLR), and International Conference on Machine Learning (ICML).

[7] We coded 24 conferences on AI, machine learning, computer vision, NLP, information retrieval on the web as AI conferences.

the same regressions (1) by the level of the highest degree, and show the estimates in Figure 3. Among PhD graduates, each additional publication (in any area) before graduation predicts a 0.2 pp increase in the probability of being on the tenure track, which is a 2.2% increase relative to the mean. Any AI paper that is not published at the top three venues predicts a 0.3 pp lower chance of a tenure-track job (see column 7 of Table A.2).

Not surprisingly, for authors who do not report a PhD degree, publications before graduation matter little for getting a tenure-track job. Figure 3(a) shows that the strong relationship between AI papers and getting hired by top firms in industry is also driven by PhD graduates, who are more likely to take research positions in industry.

Do publications after graduation matter for sorting in the labor market? We answer this by estimating OLS regressions of job mobility (after graduation from each person's highest degree) on indicators for publications at person×year level:

$$y_{it} = b_0 + b_1 \,\text{Pub}_{it} + b_2 \,\text{AI-Pub}_{it} + b_3 \,\text{Top AI-Pub}_{it} + X_{it}'\theta + \varepsilon_i \tag{2}$$

$(\text{Pub}_{it}, \text{AI-Pub}_{it}, \text{Top AI-Pub}_{it})$ are indicators for any publication by worker i in calendar year t, any publication at an AI-related conference, and any publication at a top AI conference. In addition to demographics, the controls X_{it} also include a cubic polynomial in experience since graduation, position characteristics, and a Heckit correction for reporting any job (Heckman 1974).[8] We compare the employer next year, $j(i,t+1)$, with the current employer $j(i,t)$, and examine four outcomes: (1) whether $j(i,t+1)$ is in industry, (2) whether she moves to a different employer, (3) whether $j(i,t+1)$ is a top tech firm, and (4) whether $j(i,t+1)$ is in academia. We estimate the OLS regression (2) separately based on the current employer, taking into account that workers at top firms may sort differently from those at other firms in the industry, and from those in academia.

The estimates of (b_1, b_2, b_3) on indicators of publications are presented in Table 2. For workers already employed by industry (nontop or top firms), having any publication is negatively correlated with staying in industry (column 1 or 5) but positively correlated with moving to academia (column 4 or 8). Academics who publish a paper also experience a 0.1 pp or almost 20% increase in movement to industry. Workers in academia include postdocs or non-tenure-track researchers. We will carefully differentiate between them in the causal analysis of ICLR in Section 4.

Publishing any AI paper strongly predicts moving to a top firm from other "nontop" firms or academia, and marginally predicts staying within top. On average, only 1.3% of workers at nontop firms in industry move to a top firm in one year. Publishing any AI paper is estimated to increase this upward mobility by 0.9 pp, which is a 72% increase relative to the mean (see column 3). In comparison with column 4 — moving to academia, AI papers are more predictive of moving to

[8] See Table A.3 for the full set of controls.

top firms than of moving to academia. Any AI paper also predicts a 0.5 pp increase in movements from academia to top firms, a 47% increase. Figure 4 breaks down the estimates by the level of the highest degree reported. AI papers are positively correlated with movements to top firms for workers without a PhD degree as well.

These results suggest that top tech firms attract researchers who publish at AI conferences and are able to hire from both industry and academia. A potential explanation is that top firms are better at screening candidates. Publications are correlated with a worker's ability but do not lead to getting a job at top firms. To address this endogeneity concern, in the next section we will present causal evidence that publications increase movements into top firms from one of the leading AI conferences, ICLR.

4 Causal Impacts of Publications on Career Trajectories

The International Conference on Learning Representations (ICLR), one of the top three AI conferences, has publicly released all submissions, referee reports, and editorial decisions since 2017. We exploit this "open review" feature to estimate the causal impacts of AI publications on the career outcomes of researchers, using a matched difference-in-differences design. This section presents our quasi-experimental design and estimates of an accepted paper (publication) on sorting into top tech firms from academia and industry. We also analyze the gender differences in the impacts of a publication on getting a tenure-track job versus a job at top firms.

4.1 (Matched) Difference-in-Differences Design for ICLR Authors

4.1.1 Evaluation of Paper Submissions

The open-review process of ICLR allows us to observe the numeric ratings by referees, and the exact content of the referee reports, based on which editors decide whether a paper can be accepted (published at ICLR).[9] On average, each submission gets 3-4 referees, who rate the paper on a discrete scale from 1 to 10, with 5 suggesting marginally reject and 6 suggesting marginally accept. More than 80% of the papers with an average rating above 6.5 were accepted by the editor, and most papers with an average rating below 5 were rejected. There is substantial uncertainty, however, for submissions with average ratings in the middle (Appendix Figure A.3).

Referee reports may contain additional information about the quality of a submission that cannot be captured by the numeric rating. Since editors would also take into account such information, we converted the text of each referee report to a numeric 1536-by-1 vector (embeddings), using

[9]There are papers accepted for poster sessions but not oral presentations. In this version we count all accepted papers as publications. We will look into whether the papers with a presentation slot are more valuable to workers.

large language models developed by OpenAI. We then estimated a Lasso-logit model of whether the paper is accepted by the editor on embeddings, and obtained the average predicted probability of acceptance across referee comments on each paper. On average the rejected submissions have a predicted probability of acceptance around 0.15, while accepted submissions have a predicted probability around 0.68. If we restrict to papers that are rated between 5.5 and 6.5 on average (around 6 = marginally accept), the gap in the predicted scores between accepted and rejected papers remains as large as 0.41. This comparison suggests there is substantial information about the quality of a paper conveyed in the text of referee comments, which could not be absorbed by numeric ratings.

Table 3 compares rejected and accepted authors in the full sample, and shows large gaps in submission quality, as measured by numeric ratings and the text-based likelihood of acceptance. As a result, comparisons between unlikely acceptances and well-above-threshold acceptances risk confounding treatment effects of a publication with other factors, such as underlying author ability or submission quality. To address this concern, we use a matching procedure to identify authors who are observably similar in the quality of their first submissions to ICLR, and in their origins (academic affiliations or employers) before the submissions.

4.1.2 Matching Procedure

We match each author on an accepted paper with a comparable author on a rejected paper. Our matching procedure is a combination of coarsened exact matching (Iacus et al. 2012; Sarsons 2022, Jäger and Heining (2022)) and propensity score matching (Rosenbaum and Rubin 1983; Abadie and Imbens 2016).

To have a well-defined event at the person level, we focus on the first ICLR submission of each author, and refer to the calendar year of the corresponding ICLR conference as the event year henceforth.[10] We match authors based on the quality of their submissions, the propensity score of acceptance conditional on observable characteristics, and their original jobs in the two years before their first submissions. Specifically, we consider three types of "jobs": employed by industry, employed by academia, or still in school.

First, we group the authors based on (1) the year of ICLR conference, (2) terciles of the mean rating, the minimum rating, and the predicted probability of acceptance on their submissions, and (3) the type of job in the year prior to submission. There are 290 mutually exclusive groups that have at least one accepted author and at least one rejected author.

Within each of these 290 groups, for every accepted author, we sort the potential matches (rejected authors) by (1) whether both are at nontop firms or both at top firms one year before the

[10] ICLR conferences are typically held in April or May. The submission deadline is often in Q4 (Oct/Nov) of the previous year.

event year, (2) whether both are in academia or both in industry two years before the event year, (3) whether both are at nontop firms or both at top firms two years before the event year, and (4) the difference in the propensity score of being an accepted author. We fit a logistic regression of acceptance on an author's demographics, job characteristics, and publication history before their first submission to ICLR, and use the fitted probability as the person's propensity score of success. (4) — the difference in p-score between each pair, as a continuous measure, helps break ties if there are multiple rejected authors from the same type of employers in the two years before their first submissions. The complete set of matching variables are summarized in Appendix Table A.4.

We then select the highest-ranked rejected author as the best matched unit to each accepted author, allowing a rejected author to be matched with multiple accepted authors. There are 4,401 accepted authors, who are matched to 943 rejected authors. Table 3 shows that the gap in the mean rating of papers between accepted and rejected authors shrinks from 2.085 in the full sample to 0.377 in the matched sample. The mean rating received by matched rejected authors is 6.089, above the margin of acceptance. In addition, the gap in the probability of acceptance based on the text of referee reports also shrinks from 0.517 to 0.146. Although the difference in the matched sample remains statistically significant, the rejected authors on average have a 0.503 predicted probability of acceptance, representing a control group that is much more similar to accepted papers than the vast majority of rejected manuscripts. With an average predicted score around 0.5, whether a paper is accepted by an editor is as good as random — anecdotally, computer scientists often complain about the arbitrariness of editors.[11]

4.1.3 Difference-in-Differences

We use a difference-in-differences design to estimate the impacts of an accepted publication on career outcomes, such as sorting to top firms or getting a tenure-track job. The underlying assumption is that, absent the ICLR decision, accepted and rejected authors would have followed parallel career trajectories.

We investigate the trends in employment by top firms among accepted and rejected authors. Appendix Figure A.4(a) shows that accepted authors are more likely to be employed by top firms than rejected authors before the ICLR conference, and the slope for accepted authors (red dots) is steeper than that for rejected ones (blue dots), suggesting stronger sorting to top firms even before the ICLR conference. The matched sample, in contrast, shrinks the gap between accepted and rejected authors, and the pre-trends appear more parallel than in the full sample. Figure A.4(b) further restricts to authors who are not employed by top firms at -1, the sample we use to estimate the regressions of sorting to top firms from other employers. Accepted authors are more likely to

[11] See for example the complaints from Yann LeCun, a "Godfather of AI" who worked as the chief AI scientist at Meta in recent years: http://yann.lecun.com/ex/pamphlets/publishing-models.html.

have been working for top firms more than one year before the ICLR events. The gap between accepted and rejected authors shrinks at -2 in the full sample, and becomes 0 by definition at -1. In the matched sample, the gap remains about the same between [-4,-2], which could be driven by our inclusion of same type of origin as a matching variable.[12]

Under the identifying assumption that conditional on the matching variables, accepted and rejected authors display parallel trends, we estimate the following difference-in-differences regression:

$$y_{it} = b_0 + b_1 Post_{it} + b_2 Accepted_i + b_3 Post_{it} \times Accepted_i + \delta_t + \alpha_{m(i)} + \nu_{it} \tag{3}$$

where $Post_{it} = 1[t > \tau_i]$ indicates years after the author i submits to ICLR for the first time (τ_i is the event year for i). $Accepted_i$ indicates if her first submission is accepted. We also control for calendar year fixed effects to absorb common trends, and $\alpha_{m(i)}$ fixed effects for a matched pair.[13] The coefficient of interest is b_3, which represents the diff-in-diff estimate of the impact of an accepted paper on the career outcome (y_{it}).

We further estimate a dynamic difference-in-differences or event study model:

$$y_{it} = \beta_0 + \sum_{l=-5, l\neq -1}^{3} \beta_l D_{it}^{(l)} + \zeta Accepted_i + \sum_{l=-5, l\neq -1}^{3} \gamma_l D_{it}^{(l)} \times Accepted_i + \delta_t + \alpha_{m(i)} + \varepsilon_{it} \tag{4}$$

$D_{it}^{(l)} = 1[t - \tau_i = l]$ is an event dummy that indicates l years relative to the event. The coefficients of interest are $\{\gamma_l\}$, which represent diff-in-diff estimates for an accepted paper on career outcomes at l. The coefficient b_3 in regression (3) can be viewed as an average of of $\{\gamma_l : l > 0\}$, pooling the post periods.

4.2 Impacts of Publications on Sorting to Top Firms

We estimate both model (3) and model (4) in the event window $l \in [-5, 3]$. We present the estimates for the outcome — employed by a top firm, which could indicate poaching by top firms from academia or poaching from other firms in the industry among authors who are not at a top firm right before submitting to ICLR ($l = -1$).

Table 4 presents estimates of the difference-in-differences regression (3). In the matched sample, accepted authors on average are 0.4 pp more likely to be employed by a top firm than

[12] Figure A.4(b) also shows that the rejected authors in the matched sample do not display an upward trend in sorting to top after the event as in the full sample. We will investigate if the matching procedure introduces some unexpected selection in the control group that would bias our estimates.

[13] If author i's paper is accepted, the fixed effect $\alpha_{m(i)}$ is turned on for i and her matched rejected author. Otherwise, $\alpha_{m(i)}$ is turned on for i and the accepted author she is matched to.

rejected counterparts before their first ICLR submissions. This difference goes up by $\hat{b}_3 = 3.0$ pp on average, as shown by the estimates in column (1) for everyone who is not at a top firm at -1. We then estimate (3) separately based on where the authors are in the year before their first ICLR conference: industry (nontop firms), academia, or students/postdocs. For all three groups, the diff-in-diff estimates for b_3 are positive, ranging from 1.8 pp in academia, to over 3 pp in industry and also among students and postdocs — the group at the stage of initial sorting.

Moving on to the event study (4), Figure 5(a) shows that among authors not at top firms at -1, an accepted ICLR paper increases mobility to top firms by 2-pp in the year after the ICLR conference, and by 4-5 pp two to three years after. Breaking it down by the origin of authors, Figure 5(b) shows the effects on movement to top firms within one year are higher for workers in the industry and for students/postdocs. The effects for authors who are employed by academia, despite being smaller in the beginning, remain statistically significant three years after the ICLR conference. In the matched sample, authors from nontop firms in the industry show the strongest sorting to top firms following an acceptance by ICLR.

One concern for the matching before estimating difference-in-differences model is that the inclusion of origins as matching variable, similar to the use of pre-treatment outcomes in matching, could introduce unobserved selection of the control group and yield biased diff-in-diff estimates (e.g., Ham and Miratrix 2024). To address the concern about matching, we also estimate the same set of models in the full sample, including accepted and rejected authors for whom we can construct a job history around the ICLR events. Instead of controlling for match fixed effects ($\alpha_{m(i)}$), we control for person fixed effects. An accepted ICLR paper is estimated to increase the movements into a top firm by 1.8 pp in the full sample (column 1 of Table A.5), which is lower than our estimate in the matched sample (Table 4). Notably, the effect of an accepted ICLR paper is no longer significant for incumbent employees in the industry (nontop firms - column 2) or incumbent academics (column 3). What remains robust is the estimate among students or postdocs who are yet to find their first primary employer. Column 4 of Table A.5 shows that students with an accepted ICLR paper are 2.3 pp more likely to be employed by a top firm after the decision is known, in comparison with the change for students with a rejected author. This positive effect among students is smaller but not significantly different from the estimate in the matched sample (column 4 of Table 4).

The event-study estimates of (4) in the full sample show a similar picture. Pooling authors from academia or industry prior to their ICLR submissions, part (a) of Figure A.5 shows that an accepted paper on average increases sorting to top firms by 2 pp. Part (b) shows separate estimates by the origin of authors at -1. The estimates for academics and students/postdocs remain positive and statistically significant around 2 pp post event, whereas the estimates for authors from nontop firms in the industry appear smaller and noisier.

Authors who were employed by top firms at -1 have been excluded from the estimation samples above. Figure 6 fills this gap by looking at retention among workers who are employed by top firms at -1. For both women and men, on average, an ICLR publication increases retention at top firms.

Going forward, we will look into potential selection issues from the matching procedure, which appear to affect authors from industry more than academics. So far the results are robust for academics, and especially for students or postdocs. Using a difference-in-differences design, we find that having an accepted paper at ICLR, one of the top AI conferences, significantly increases the chance that young researchers who are still in school get hired by top firms.

4.3 Heterogeneity in Impacts by Gender

Both industry and academia have publicized their commitments to reduce the under-representation of women in STEM.[14] What does the industry–academia race imply for gender differences among computer-science researchers? While a full answer is beyond the scope of this paper, we examine whether the labor market return to an ICLR publication differs by gender.

Figure 7(a) estimates the event study regression (4) separately by gender for researchers who were in academia (e.g., professors) in the year prior to the event. Following an ICLR publication, men are more likely to move into top firms, whereas the estimated effect for women is small and statistically indistinguishable from zero. This suggests that the mobility into top firms from academia in 5(b) is driven primarily by men.

Given that a publication does not seem to have returns to women in industry, does it have any academic returns? Figure 7(b) repeats the analysis for researchers who were students prior to the event. An ICLR publication increases the likelihood that women enter tenure-track positions, while there is no detectable effect for men. Given that the earliest year in the estimation sample is 2012, this finding could reflect the commitment of academia to increase gender diversity, which also appears in other disciplines like economics (e.g., Card, DellaVigna, Funk, and Iriberri 2022).

These patterns hold in the full, unmatched sample (Figure A.6). Taken together, the estimates suggest that an ICLR publication facilitates transitions into top firms for men and into tenure-track academia for women. Mechanisms may operate on either the supply or demand side. On the supply side, women may prefer the stability and schedule control offered by academic careers (e.g., Goldin and Katz 2011; Noray 2025). On the demand side, male-dominated industry environments may feature evaluation practices or biases that negatively impact women. Regardless of the mechanism, because industry roles tend to pay more on average than academic jobs, this divergence in sectoral returns to a publication could exacerbate gender earnings gaps among computer scientists.

[14] Following the MeToo movement, tech companies expanded DEI efforts to recruit underrepresented minorities in their workforce. Our data cover years up to 2023, before the conclusion of DEI in most tech firms since 2024.

5 Implications of the Academia-Industry Race on Future Research

Next, we study the consequences of the growing competition between academia and industry for top AI researchers. As a first step, we conduct event-study analyses of research output, using the full sample of computer science researchers with linked job histories. The sample includes all researchers in the Scopus-Revelio sample who moved either to top firms or to tenure-track positions; in cases of multiple moves, we consider the first move for each author. Figure 8 plots publication trajectories around these career transitions. Prior to the move, both the number of publications and the number of AI publications rise on similar trends for researchers who later transition to academia or to top firms. Following the move, publication output declines modestly among those entering tenure-track positions but resumes an upward trajectory in subsequent years. In contrast, output drops sharply and persistently among those moving to top firms, both in overall publications and in AI-specific work.

In the next stage of our analysis, we will examine the impact of transitioning into top firms versus tenure-track positions on two additional dimensions: research influence (measured by citations) and research direction. Because industry researchers may align their work with firm priorities, we will test whether transitions to top firms shift focus away from topics such as algorithmic bias toward more applied areas. To move closer to causal identification, we will conduct event-study analyses comparing researchers who transition to top firms with a matched control group of similar researchers who remain outside of industry, using coarsened exact matching to construct the matched control group.

6 Conclusion

In this paper, we document that AI researchers are increasingly concentrating in industry – especially at top-tier firms – relative to academia. To test whether the two sectors compete for the same talent, we study how publications — a canonical signal of research ability — are related to initial job placements and subsequent job mobility. We find that AI publications are strongly predictive of getting a job at top tech firms, but less so for getting a job in academia. We provide supportive causal evidence by leveraging variation in the acceptance of a paper at a top-tier AI conference. Using a matched difference-in-differences design, we find publications raise moves to top firms among academics, including faculty and students/postdocs who have yet to kick start their career. This pattern is consistent with industry placing high returns on research ability that was historically rewarded mainly in academia.

Having established that a "race for AI talent" is underway – and that academia is losing – we examine the implications for the research landscape. After researchers move from academia to industry, their publication output declines. Because we cannot observe proprietary work, total research output is unknown. However, this evidence suggests that increased industry poaching reduces the volume of *public* AI research.

While these results speak to the quantity of AI research, we hope to also explore various measures of the quality or content of AI research in future work. For example, we will begin by documenting whether AI researchers in industry produce (public) research that is categorically different from academics. Our hypothesis is that AI topics that serve a public-interest (e.g., bias mitigation and algorithmic fairness) or that have long-term or uncertain commercial gains (e.g., theory behind black-box models) are at risk of being under-produced as a consequence of this race. On the other hand, topics that require large amounts of computing power or data (e.g., large language models) may be more readily produced in industry, which typically has more resources to accommodate this type of research. Moreover, we will examine how citations – a measure of research influence – change with a transition into industry from academia.

References

Abadie, Alberto and Guido W. Imbens. Matching on the estimated propensity score. *Econometrica*, 84(2):781–807, 2016.

Acemoglu, Daron. The simple macroeconomics of ai. Working Paper 32487, National Bureau of Economic Research, 2024.

Aghion, Philippe, Mathias Dewatripont, and Jeremy C Stein. Academic freedom, private-sector focus, and the process of innovation. *The RAND Journal of Economics*, 39(3):617–635, 2008.

Babina, Tania, Alex Xi He, Sabrina T Howell, Elisabeth Ruth Perlman, and Joseph Staudt. The color of money: Federal vs. industry funding of university research. Working Paper 28160, National Bureau of Economic Research, 2020.

Card, David, Stefano DellaVigna, Patricia Funk, and Nagore Iriberri. Gender differences in peer recognition by economists. *Econometrica*, 90(5):1937–1971, 2022.

Furman, Jason and Robert Seamans. Ai and the economy. *Innovation Policy and the Economy*, 19:161–191, 2019.

Gofman, Michael and Zhao Jin. Artificial intelligence, human capital, and innovation. *Journal of Finance, Forthcoming*, 2022.

Goldin, Claudia and Lawrence F. Katz. The cost of workplace flexibility for high-powered professionals. *The Annals of the American Academy of Political and Social Science*, 638:45–67, 2011.

Ham, Dae Woong and Luke Miratrix. Benefits and costs of matching prior to a difference in differences analysis when parallel trends does not hold. 2024.

Heckman, James. Shadow prices, market wages, and labor supply. *Econometrica*, 42(4):679–694, 1974.

Iacus, Stefano M., Gary King, and Giuseppe Porro. Causal inference without balance checking: Coarsened exact matching. *Political Analysis*, 20(1):1–24, 2012.

Jäger, Simon and Jörg Heining. How substitutable are workers? evidence from worker deaths. Technical report, National Bureau of Economic Research, 2022.

Jurowetzki, Roman, Daniel Hain, Juan Mateos-Garcia, and Konstantinos Stathoulopoulos. The privatization of ai research(-ers): Causes and potential consequences – from university-industry interaction to public research brain-drain?, 2021.

Noray, Savannah. Can women "have it all"? the trade-off between social tasks and workplace flexibility. 2025. Working Paper.

Rosenbaum, Paul R. and Donald B. Rubin. The central role of the propensity score in observational studies for causal effects. *Biometrika*, 70(1):41–55, 1983.

Sarsons, Heather. Recognition for group work: Gender differences in academia. *American Economic Review*, 107(5):141–45, May 2017.

Sarsons, Heather. Interpreting signals in the labor market: Evidence from medical referrals. 2022.

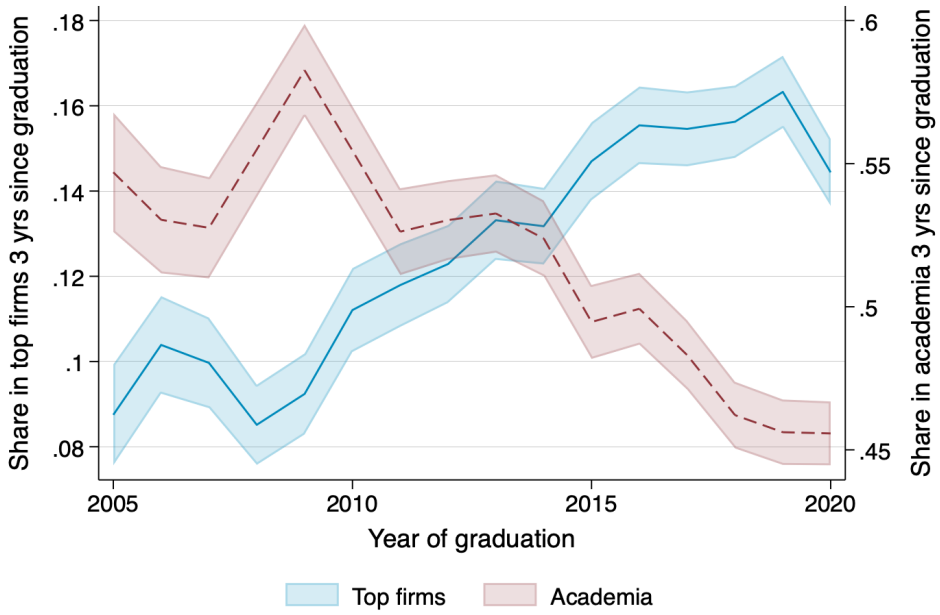
Wu, Alice H. Reveal or conceal? employer learning in the labor market for computer scientists. 2025.

Yue, Daniel. I, google: Estimating the impact of corporate involvement on ai research. 2024.

7 Main Figures and Tables

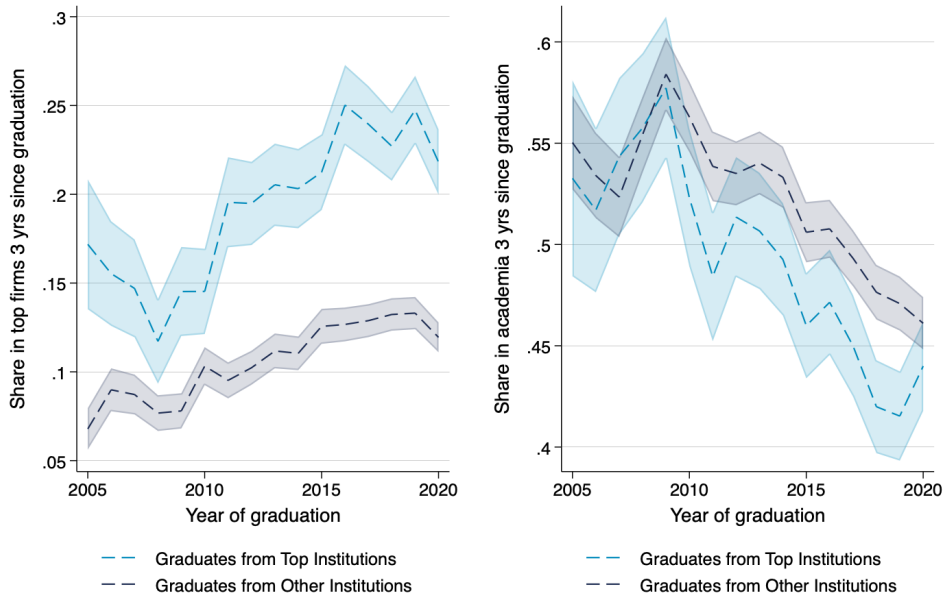
Figures

FIGURE 1. Share of AI researchers in industry vs academia over time



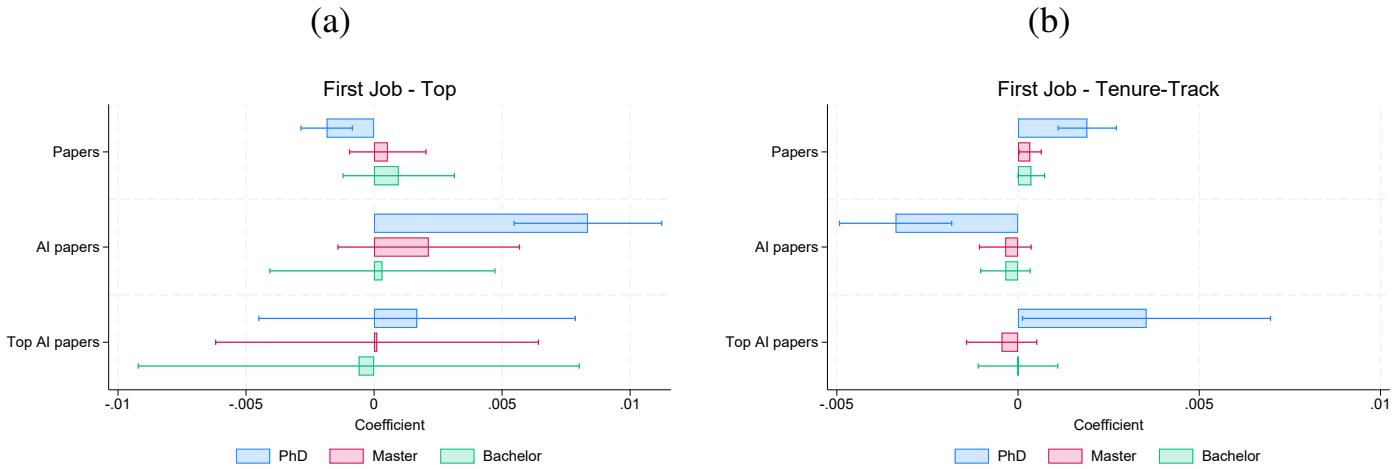
Notes: Data from around 60k AI researchers for whom we have publication data matched with job history from Revelio (conditional on nonmissing degree information, and graduation from 2015 to 2020). Share of AI researchers working in top tech firms and tenure-track academic jobs within 3 years of graduation, where graduation is defined using each person's highest degree.

FIGURE 2. Top AI researchers in industry vs academia over time



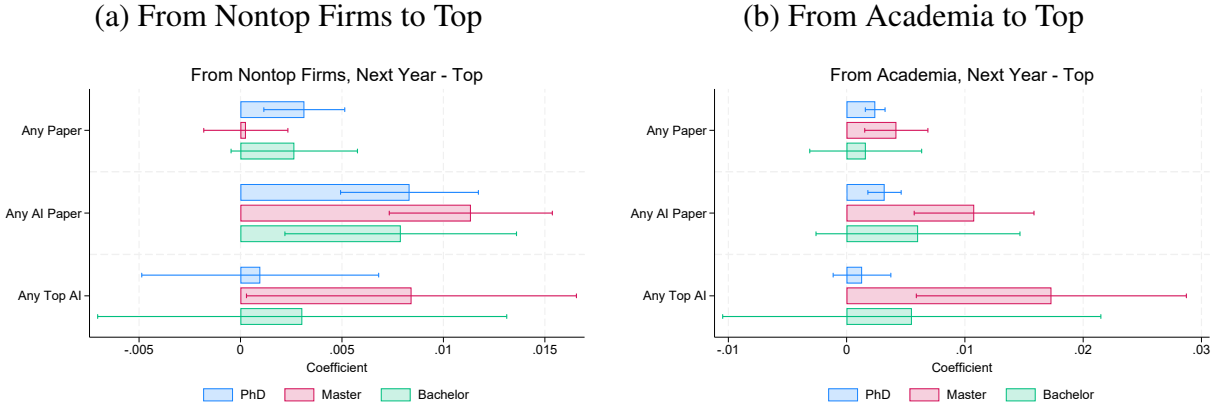
Notes: Data from around 60k AI researchers for whom we have publication data matched with job history from Revelio (data on graduation non-missing, and graduation from 2015 to 2020). Share of AI researchers working in top firms and in academia within 3 years of graduation, by top school/non top school. Top school=1 if school of graduation is one of top 50 schools based on CSRankings, QS 2025, and THE 2025 in AI/CS.

FIGURE 3. First Job in the Year After a Person's Highest Degree



Notes: This figure shows the OLS estimates of regression (1) of first job on publications before graduation, separately by the level of highest degree. We showed the estimates for two outcomes: (a) first job at top firms, (b) first job on the tenure track, which correspond to columns 2 and 7 of Table 1, respectively, for the pooled sample with authors with a college degree or above.

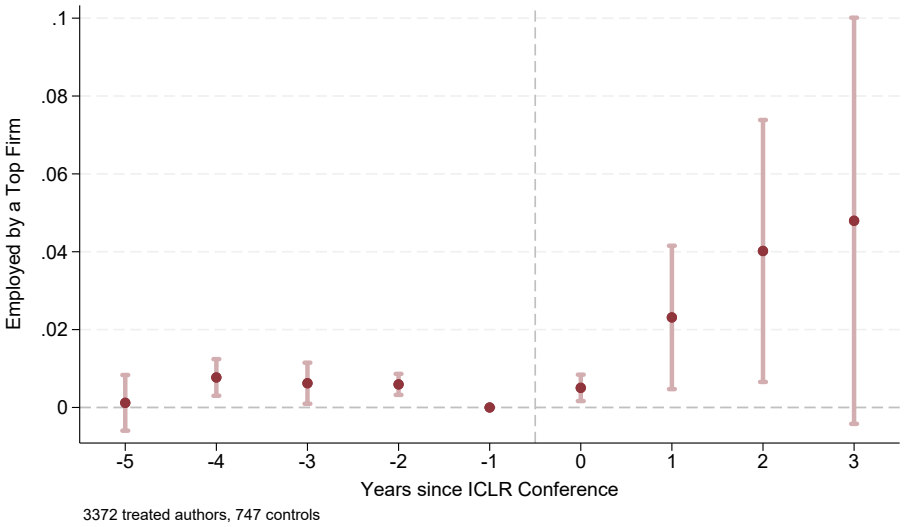
FIGURE 4. Job Mobility After a Person's Highest Degree



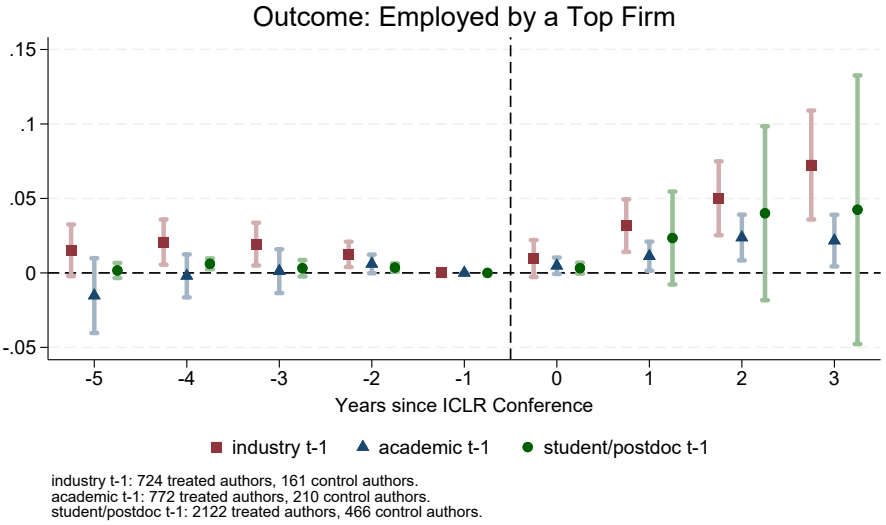
Notes: This figure shows the OLS estimates of regression (2) of job mobility on indicators for publications, separately by the level of highest degree. We showed the estimates for the outcome of moving to top firms, either from (a) other firms ("nontop") in the industry, or (b) from academia. The regressions correspond to columns 3 and 11 of Table 2, respectively, for the pooled sample with authors with a college degree or above. Standard errors are robust.

FIGURE 5. Mobility outcome: Top | $j(i,-1) \notin$ Top

(a) Pooled

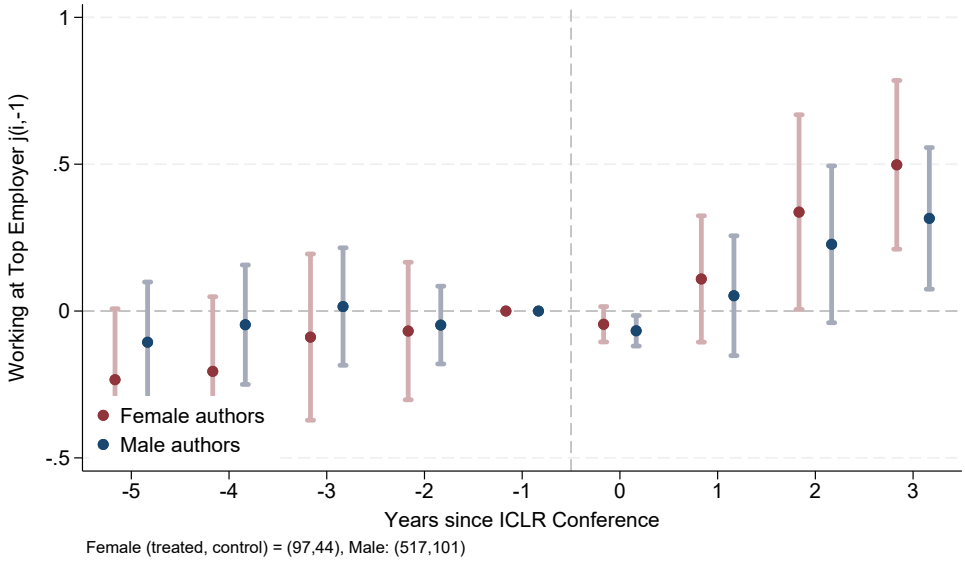


(b) By Job



Notes: This figure shows the estimates of $\{\gamma_t\}$ on the interaction between acceptance and event dummies in regression (4). Outcome variable "top" indicates whether the worker is employed full-time by a top firm in a given year. The regression is estimated on authors in the matched sample who are not employed by a top firm at -1 (the year before her first ICLR submission). The pooled sample (shown in (a)) includes 3,372 accepted authors matched to 747 rejected authors (drawn with replacement). In (b) we estimate separate regressions by the type of job at -1 (industry vs. academia vs. student/postdoc), conditional on the worker not employed by a top firm at -1. All regressions include year and fixed effects for a matched pair of authors. Standard errors are robust and clustered at the author level.

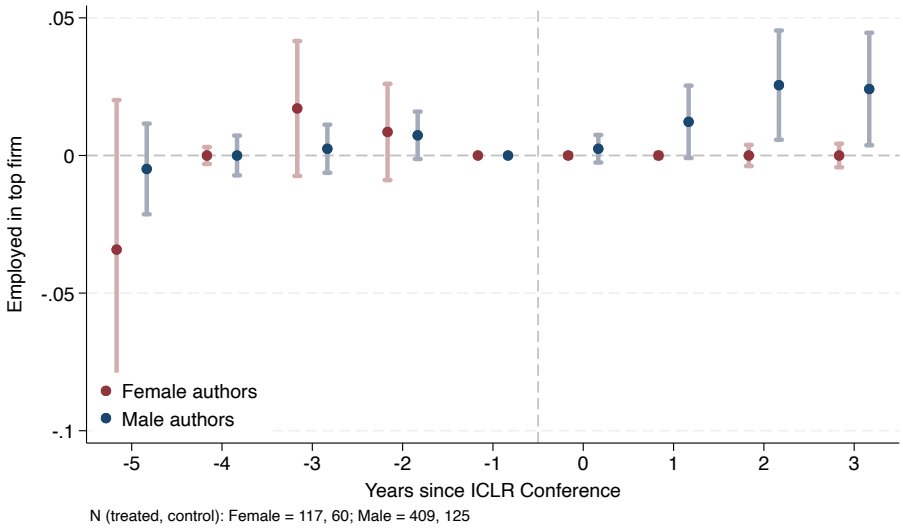
FIGURE 6. Retention at Top Firms (cond. on employment by top at -1



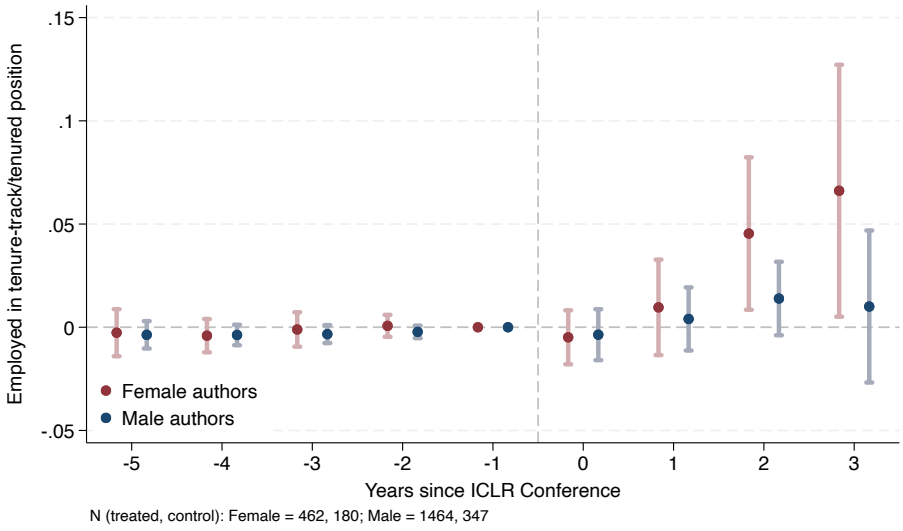
Notes: This figure shows the estimates of $\{\gamma_l\}$ on the interaction between acceptance and event dummies in regression (4). Outcome variable "top" indicates whether the worker is employed full-time by a top firm in a given year. The regression is estimated on authors in the matched sample who are employed by a top firm at -1 (the year before her first ICLR submission) — that is, authors who are excluded from the estimation sample for Figure 5. Standard errors are robust and clustered at the author level.

FIGURE 7. Gender Differences in Employment Outcomes

(a) Mobility to top firms among academics



(b) Mobility to academic position among students/postdocs

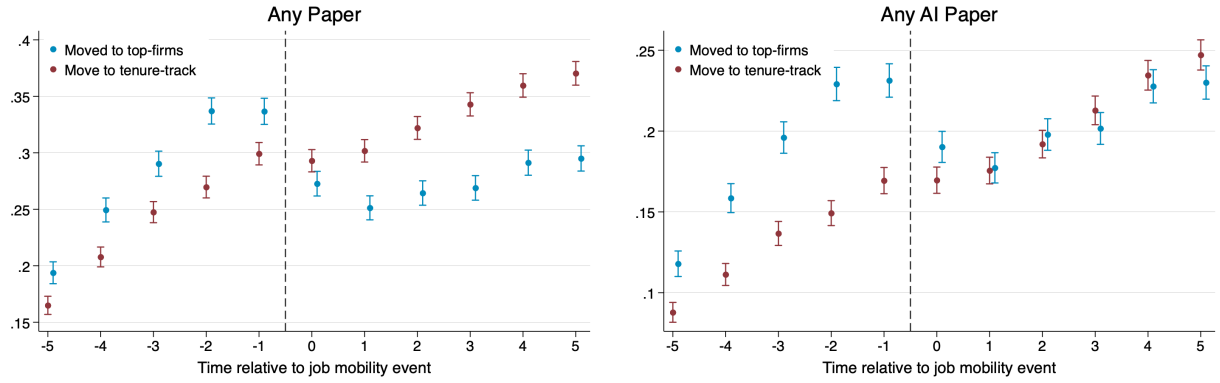


Notes: This figure shows the estimates of $\{\gamma_l\}$ on the interaction between acceptance and event dummies in regression (4). We estimate separate regressions for whether the treated/accepted author is female, allowing the matched control unit to be either sex. In Panel (a), the outcome is an indicator for working full-time at a top firm and the sample is limited to authors who are academics in -1. In Panel (b), the outcome is an indicator for employment in a tenure-track or tenured position and the sample is limited to authors who are students or postdocs in -1. All regressions include year and fixed effects for a matched pair of authors. Standard errors are robust and clustered at the author level.

FIGURE 8. Publications. The event is move to top firms vs move to tenure track job

(a) Papers

(b) AI Papers



Notes: The figure plots the average number of publications (panel (a)) and AI publications (panel (b)) for researchers moving to top firms (blue) and to tenure-track positions (red), relative to the year of the move. The sample includes all researchers in the Scopus–Revelio matched data with observed job transitions and complete event windows. For researchers with multiple moves, we consider the first move.

Tables

TABLE 1. First Job in the Year After a Person's Highest Degree

	Industry				Academia		
	(1) Industry	(2) Top	(3) Re-Scientist	(4) Engineer	(5) Academia	(6) Postdoc	(7) Tenure-Track
Papers before graduation	-0.0014	-0.0010	-0.0029	0.0004	0.0024	0.0009	0.0022
	(0.0005)	(0.0004)	(0.0004)	(0.0004)	(0.0005)	(0.0003)	(0.0003)
AI papers	0.0044	0.0062	0.0110	-0.0059	-0.0045	-0.0004	-0.0026
	(0.0010)	(0.0011)	(0.0012)	(0.0011)	(0.0010)	(0.0007)	(0.0006)
Top AI papers	-0.0072	0.0002	-0.0009	-0.0052	0.0079	0.0032	0.0016
	(0.0026)	(0.0023)	(0.0027)	(0.0031)	(0.0026)	(0.0017)	(0.0012)
N	152900	152900	152900	152900	152900	152900	152900
Adj. R2	.0494558	.0538363	.0234839	.045831	.0497791	.028532	.0546584
Mean of dep. var	.5417986	.0754349	.0763048	.2679791	.3461347	.1128515	.0793721
Mean: Papers | dep. var =1	1.436378	2.549593	2.273506	1.207888	1.341603	1.63228	1.586849
Mean: AI Papers	.5165679	1.112797	1.133282	.3583004	.4607928	.5950159	.5051088
Mean: Top AI papers	.0939873	.2320964	.220708	.060697	.0913385	.1219357	.1054713

Notes: This table shows the OLS estimates of regression (1) of first job on publications before graduation, for authors in the Scopus-Revelio matched sample — who have published a CS paper according to Scopus database, and have matched to a Revelio individual profile with self-reported education and job history. The sample is restricted to 152,900 matched authors who report at least a college degree. Each person's first job is characterized based on her primary employer in the year after obtaining her highest degree. The first four columns are outcomes related to employment in industry, whereas the next three columns are related to employment in academia.

TABLE 2. Mobility after the Highest Degree (Outcomes based on $j(i,t+1)$ next year)

	At Nontop Firms				At Top Firms				In Academia			
	(1) Industry	(2) Move	(3) Top	(4) Academia	(5) Industry	(6) Move	(7) Top	(8) Academia	(9) Industry	(10) Move	(11) Top	(12) Academia
Any Paper	-0.0123	-0.0037	0.0023	0.0036	-0.0071	0.0052	-0.0079	0.0051	0.0095	0.0143	0.0029	-0.0172
	(0.0012)	(0.0018)	(0.0007)	(0.0008)	(0.0017)	(0.0034)	(0.0033)	(0.0013)	(0.0011)	(0.0014)	(0.0004)	(0.0015)
Any AI Paper	-0.0072	0.0069	0.0094	0.0012	0.0016	-0.0034	0.0085	-0.0009	-0.0027	-0.0008	0.0046	0.0007
	(0.0018)	(0.0028)	(0.0012)	(0.0012)	(0.0023)	(0.0043)	(0.0042)	(0.0017)	(0.0016)	(0.0021)	(0.0007)	(0.0022)
Any Top AI Paper	-0.0118	0.0023	0.0036	0.0013	0.0033	-0.0033	0.0030	-0.0009	-0.0041	0.0022	0.0029	0.0021
	(0.0032)	(0.0043)	(0.0022)	(0.0019)	(0.0027)	(0.0049)	(0.0048)	(0.0020)	(0.0024)	(0.0031)	(0.0013)	(0.0032)
N	608574	603973	608574	608574	101747	101076	101747	101747	490273	486379	490273	490273
Adj. R2	.0430721	.0222411	.008767	.0256379	.0153247	.0116275	.0142735	.0279952	.0554703	.0800049	.0097742	.3067004
Mean	.9624631	.1235568	.0129713	.0149891	.9864369	.0667913	.9364109	.0047274	.0588631	.1026422	.0063291	.8020185

unit: (person, year). Conditional on school (highest degree) and year FE.

Notes: This table shows the OLS estimates of regression (2) of job mobility on indicators for publications, for authors in the Scopus-Revelio matched sample (same restrictions as noted under Table 1). We built a balanced person×year panel for matched authors, keeping all years after each person obtains her highest degree, up to 2023. There are four outcomes defined by comparing the employer next year, $j(i,t+1)$, with the current employer, $j(i,t)$, for person i in year t: (1) "industry" is 1 if $j(i,t+1)$ is in industry, (2) "move" indicates $j(i,t+1) \neq j(i,t)$, (3) "top" indicates $j(i,t+1) \in$ Top Firms, and (4) "academia" indicates $j(i,t+1) \in$ Academia. We estimated separate regressions for those who are currently at nontop firms (i.e., $j(i,t) \in$ Nontop), or top firms, or academia. Standard errors are robust.

TABLE 3. Characteristics of Accepted vs. Rejected Authors

	Full Sample			Matched Sample		
	(1)	(2)	(3)	(4)	(5)	(6)
Variable	Rejected	Accepted	Diff.	Rejected	Accepted	Diff.
Mean Rating	4.441	6.526	2.085	6.089	6.466	0.377
	(1.132)	(0.729)	(0.018)	(0.585)	(0.726)	(0.014)
Min Rating	3.328	5.494	2.166	5.298	5.425	0.126
	(1.270)	(1.263)	(0.023)	(1.064)	(1.271)	(0.025)
Predicted Prob. of Acceptance	0.146	0.663	0.517	0.503	0.649	0.146
	(0.186)	(0.278)	(0.004)	(0.220)	(0.282)	(0.005)
Year of ICLR	2,020.368	2,020.114	-0.254	2,020.368	2,020.368	0.000
	(1.389)	(1.581)	(0.027)	(1.465)	(1.465)	(0.031)
Prob. Female	0.280	0.266	-0.014	0.276	0.276	0.000
	(0.312)	(0.309)	(0.006)	(0.326)	(0.311)	(0.007)

Notes: This table displays average paper and author characteristics for rejected versus accepted authors in both the full sample (Columns 1-2) and matched sample (Columns 4-5). Standard deviations are in parentheses. Columns 3 and 6 display the difference in means with the corresponding standard error in parentheses. In the full sample, 8,654 rejected and 4,624 accepted authors are represented. In the matched sample, 4,401 rejected and 4,401 accepted authors are represented. Mean and min rating refer to the set of ratings given by referee reports for the author's paper constituting the event. Predicted probability of acceptance is a score based on a Lasso-logit model leveraging embeddings based on text of referee reports, as described in the text. The probability a person is female is from the Revelio data and is based on the author's name.

TABLE 4. Matched Diff-in-diff estimates for Employment by Top Firms

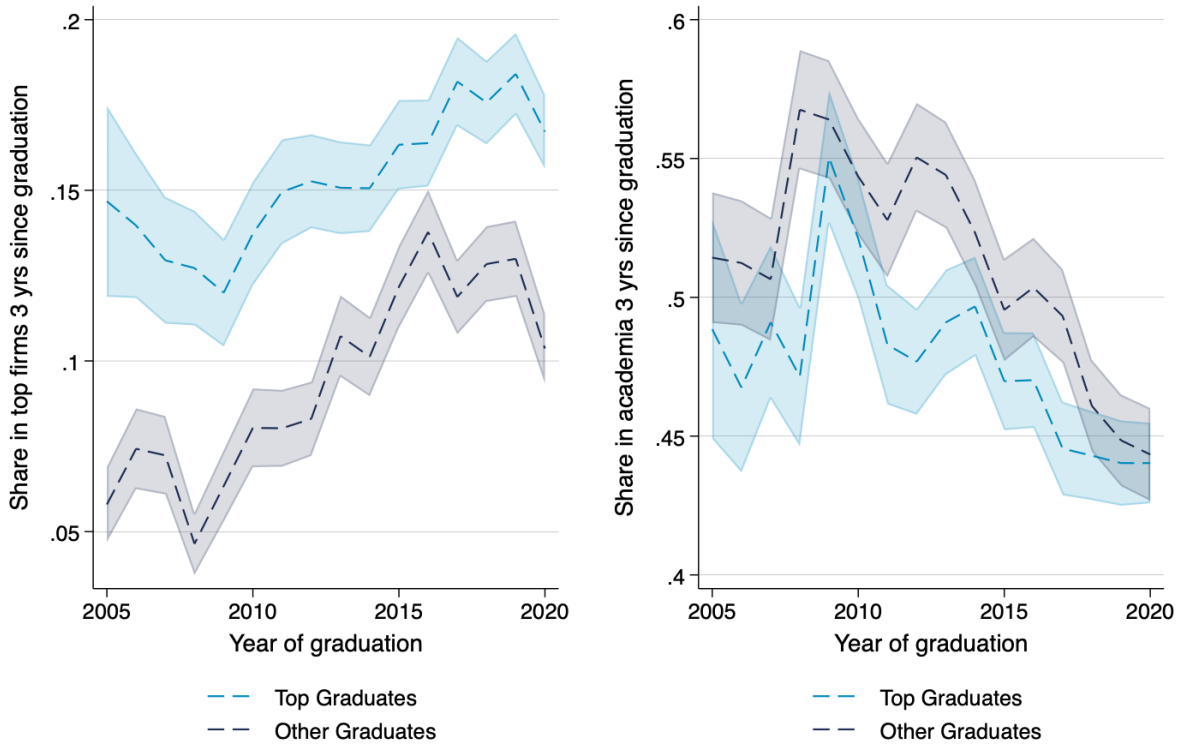
	Employed by a Top Firm			
	(1) Pooled	(2) From Industry	(3) From Academia	(4) From Student/Postdoc
Post event=1	0.0365	0.0089	0.0035	0.0537
	(0.0126)	(0.0058)	(0.0027)	(0.0218)
Accepted	0.0043	0.0129	-0.0004	0.0014
	(0.0021)	(0.0029)	(0.0025)	(0.0032)
Post event=1 $\times$ Accepted	0.0296	0.0324	0.0181	0.0311
	(0.0140)	(0.0105)	(0.0069)	(0.0238)
Constant	0.0030	0.0041	0.0030	0.0043
	(0.0017)	(0.0020)	(0.0016)	(0.0028)
Observations	55112.0000	11646.0000	10651.0000	32815.0000
Adj. R^2	0.1043	0.0988	0.1479	0.1219

Notes: This table shows the estimates of difference-in-differences regression (3) among authors in the matched sample. The estimation sample is restricted to the event window [-5,3] years relative to the first ICLR submission of each author. Standard errors are clustered at author level.

A Appendix Figures and Tables

Figures

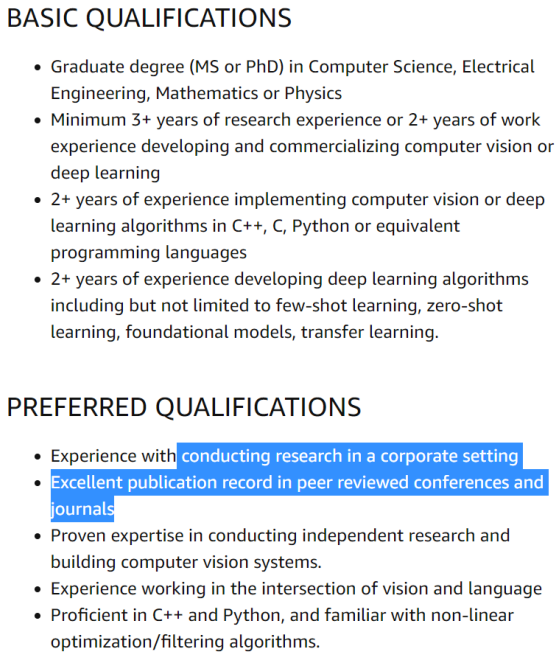
FIGURE A.1. Share of AI researchers working in top firms and in academia within 3 years of graduation, by top graduate.



Notes: Data from around 60k AI researchers for whom we have publication data matched with job history from Revelio (data on graduation non-missing, and graduation from 2015 to 2020). Top graduate is 1 if the author has at least one publication before graduation.

FIGURE A.2. Job Postings for Research Scientists

(a) Amazon Science



BASIC QUALIFICATIONS

- Graduate degree (MS or PhD) in Computer Science, Electrical Engineering, Mathematics or Physics
- Minimum 3+ years of research experience or 2+ years of work experience developing and commercializing computer vision or deep learning
- 2+ years of experience implementing computer vision or deep learning algorithms in C++, C, Python or equivalent programming languages
- 2+ years of experience developing deep learning algorithms including but not limited to few-shot learning, zero-shot learning, foundational models, transfer learning.

PREFERRED QUALIFICATIONS

- Experience with conducting research in a corporate setting
- Excellent publication record in peer reviewed conferences and journals
- Proven expertise in conducting independent research and building computer vision systems.
- Experience working in the intersection of vision and language
- Proficient in C++ and Python, and familiar with non-linear optimization/filtering algorithms.

(b) Google Research



Minimum qualifications:

- PhD in Computer Science, related technical field or equivalent practical exp
- Experience in Natural Language Understanding, Computer Vision, Machine Optimization, Data Mining or Machine Intelligence (Artificial Intelligence).
- Programming experience in C, C++, Python.
- Contributions to research communities/efforts, including publishing papers NeurIPS, ICML, ACL, CVPR).

Preferred qualifications:

- Relevant work experience, including full time industry experience or as a re
- Strong publication record
- Ability to design and execute on research agenda.

Notes: This figure shows recent postings of research scientist jobs at Amazon and Google. Both ads explicitly indicate a graduate degree in computer science as a basic qualification for this type of jobs, and list "publication records" as preferred qualifications.

FIGURE A.3. Referee ratings and publication outcomes

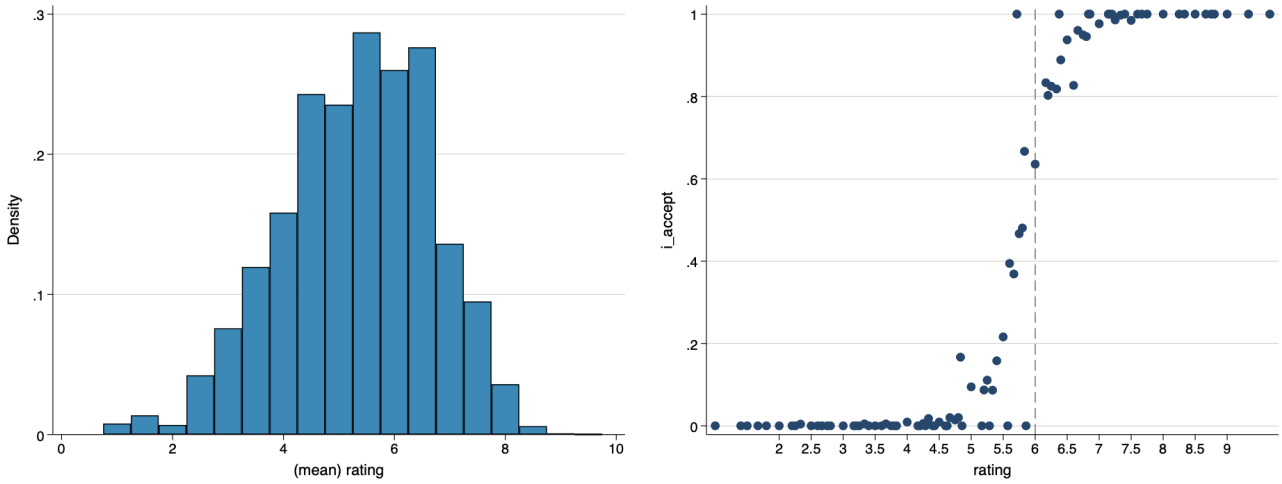
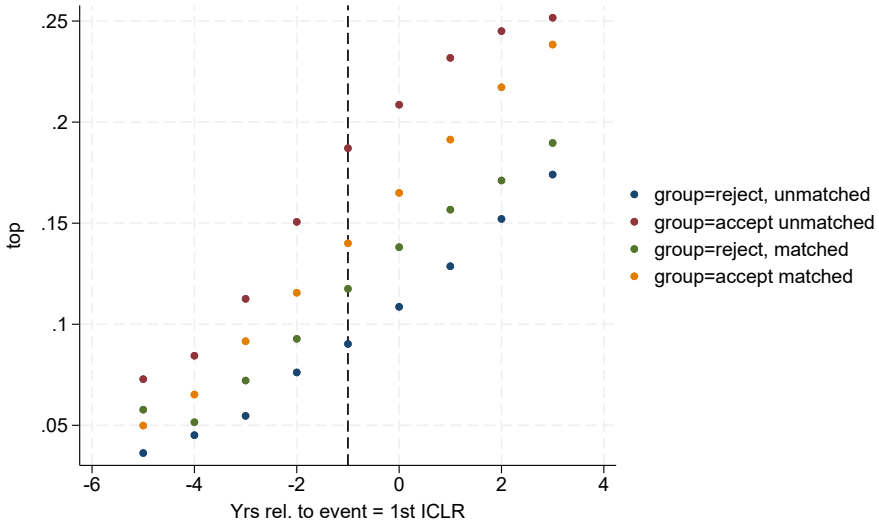


FIGURE A.4. Parallel Trends in Employment by Top Firms

(a) Top



(b) Top | $j(i,-1) \neq$Top

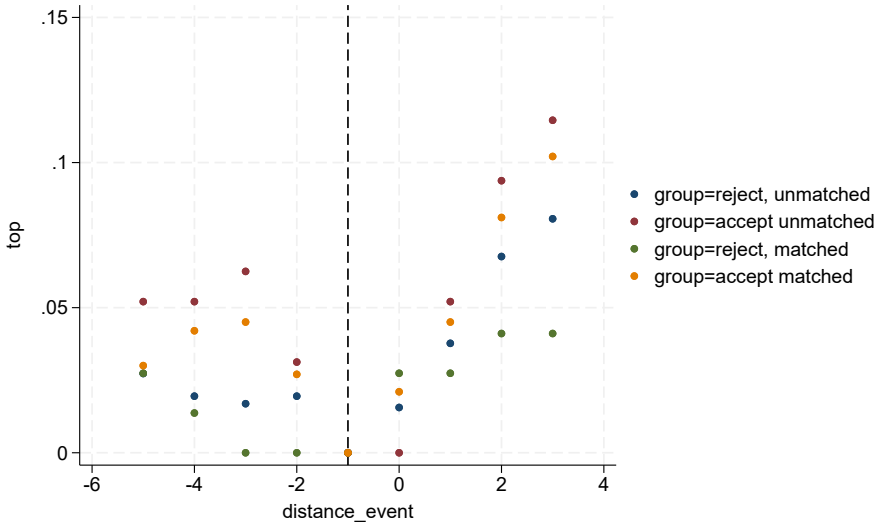
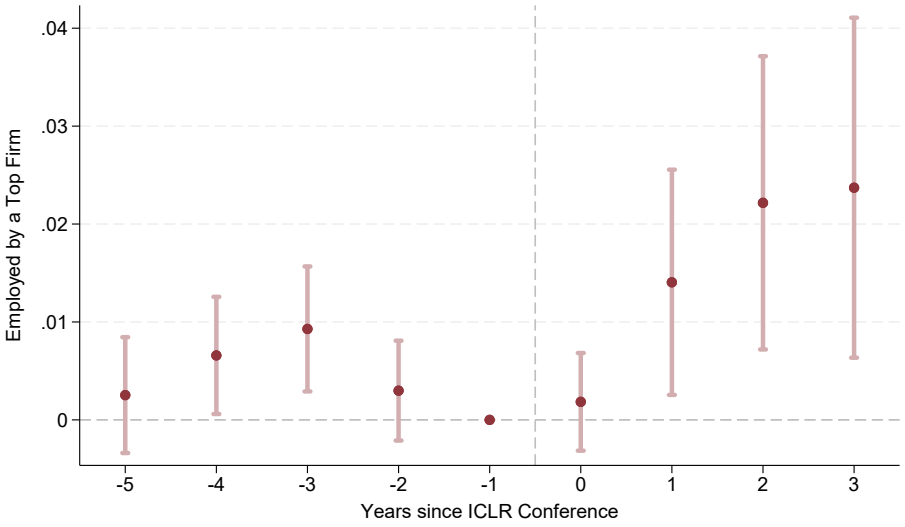
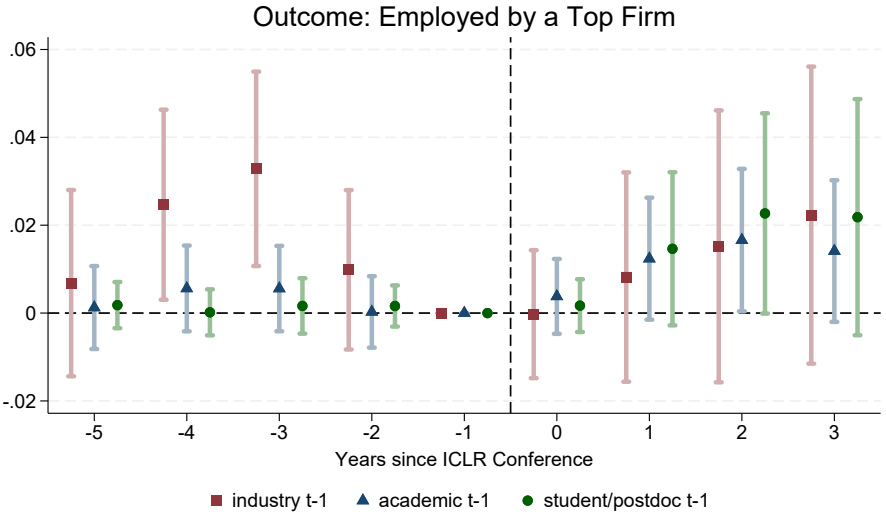


FIGURE A.5. Mobility outcome: Top | $j(i,-1) \notin$ Top (Full Sample)

(a) Full Sample -Pooled



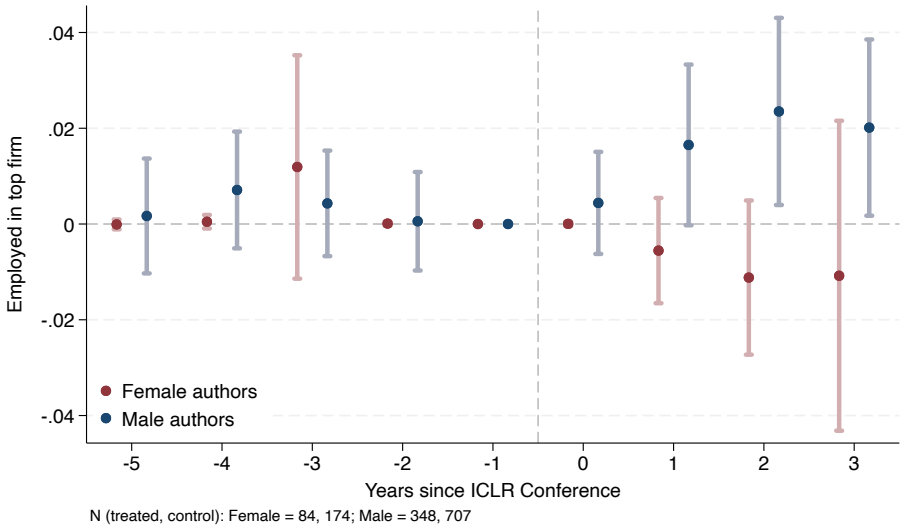
(b) Full Sample -By Job



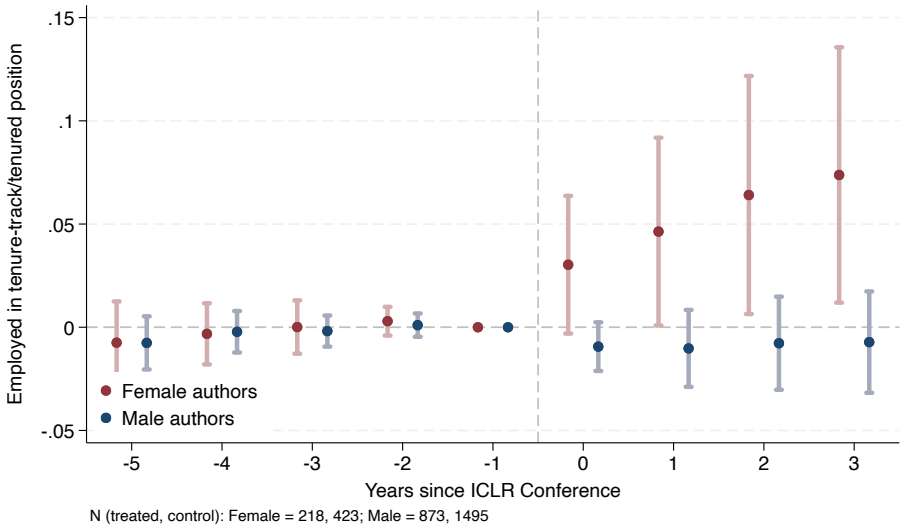
Notes: This figure shows the estimates of $\{\gamma_l\}$ on the interaction between acceptance and event dummies for the full sample (including all accepted and rejected authors with job history information). The event study specification is the same as shown in regression 4, but we control for author fixed effects instead of the fixed effects for matched pairs. Standard errors are robust and clustered at the author level.

FIGURE A.6. Gender Differences in Employment Outcomes (Full Sample)

(a) Mobility to top firms among academics



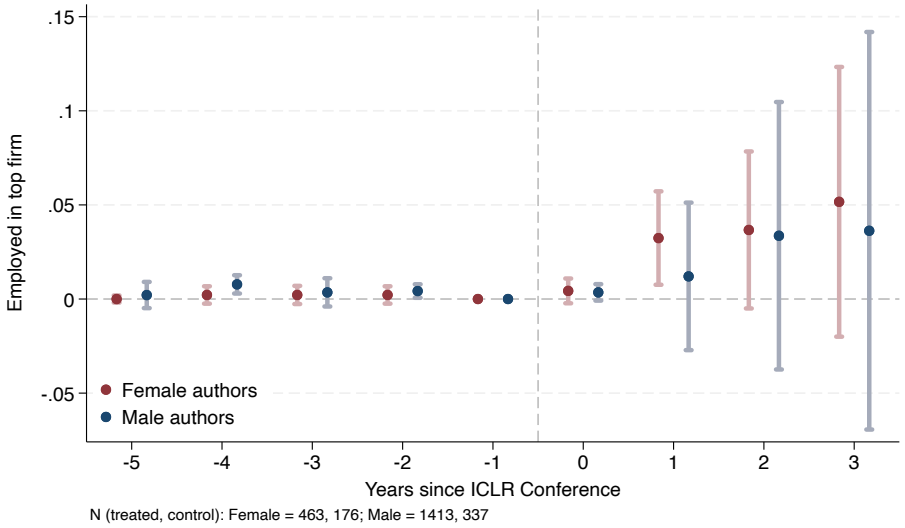
(b) Mobility to academic position among students/postdocs



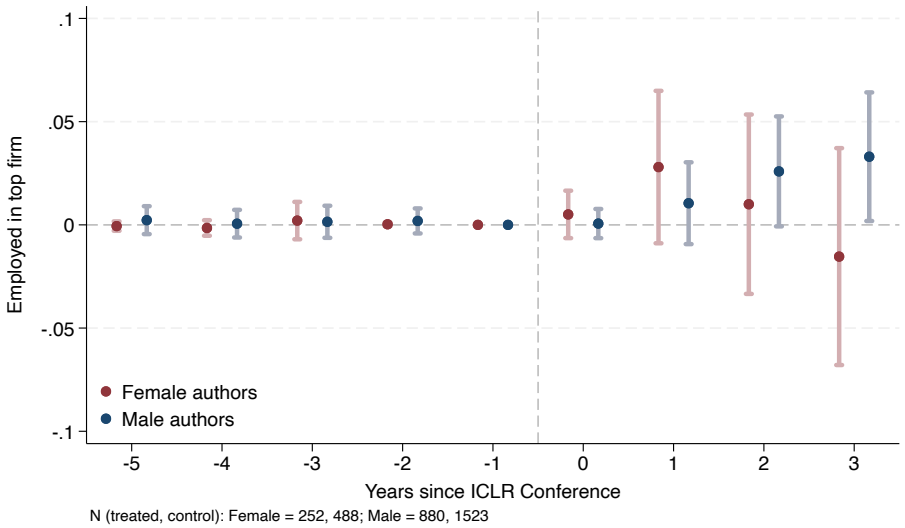
Notes: This figure shows the estimates of $\{\gamma_l\}$ on the interaction between acceptance and event dummies for the full sample (including all accepted and rejected authors with job history information). The event study specification is the same as shown in regression 4, but we control for author fixed effects instead of the fixed effects for matched pairs and estimate the regression separately by gender. Standard errors are robust and clustered at the author level. In Panel (a), the outcome is an indicator for working full-time at a top firm and the sample is limited to authors who are academics in -1. In Panel (b), the outcome is an indicator for employment in a tenure-track or tenured position and the sample is limited to authors who are students or postdocs in -1.

FIGURE A.7. Gender Difference in Mobility to Top Firms among Students/Postdocs

(a) Matched sample



(b) Full sample



Notes: This figure shows the estimates of $\{\gamma_l\}$ on the interaction between acceptance and event dummies in regression (4). Panel (a) displays the results for the matched sample and Panel (b) displays the results for the full sample. In both cases, we estimate separate regressions by gender. (Note that in the matched specification, we allow accepted female authors to match to a male or female rejected author.) The outcome is an indicator for working full-time at a top firm and the sample is limited to authors who are students/postdocs at -1. Standard errors are robust and clustered at the author level.

Tables

TABLE A.1. Summary Statistics of Authors submitting to ICLR conference (Person level)

	count	mean	sd	p50
Female	20113	0.244	0.430	0.000
Race = White	20113	0.365	0.481	0.000
Race = Asian	20113	0.586	0.492	1.000
Race = Minority	20113	0.048	0.215	0.000
Profile country outside US	20113	0.506	0.500	1.000
Highest degree = PhD	14055	0.667	0.471	1.000
Highest degree = Master	14055	0.228	0.420	0.000
Highest degree = Bachelor	14055	0.105	0.306	0.000
Yr of Highest Degree	14055	2016.834	7.023	2019.000
Num. Publications before Graduation	14055	2.597	9.272	0.000
First job in Academia	10230	0.318	0.466	0.000
First job in Industry	10230	0.543	0.498	1.000
First job in Top Firms	10230	0.156	0.363	0.000

Notes: This table provides person-level summary statistics for authors who submitted to ICLR between 2017 and 2023 and who were matched to a Revelio profile, based on name and affiliations. Out of the 20,113 ICLR-Revelio matched authors, 14,055 reported at least a college degree in their education histories. We showed the fractions with a PhD/Master/Bachelor as the highest degree in this subsample. Characteristics of first jobs are summarized in the year after a person's highest degree. Since we only focused on the window [-5,3] relative to each submission, the calendar year of the first job is out of the sample (we could expand the panel and characterize the first jobs of all 14,055 authors with a reported college degree or above).

TABLE A.2. First Job the Year After PhD

	Industry				Academia		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	Industry	Top	Re-Scientist	Engineer	Academia	Postdoc	Tenure-Track
Papers before graduation	-0.0011	-0.0018	-0.0052	0.0039	0.0017	-0.0010	0.0019
	(0.0006)	(0.0005)	(0.0006)	(0.0005)	(0.0005)	(0.0005)	(0.0004)
AI papers	0.0069	0.0084	0.0153	-0.0097	-0.0070	-0.0007	-0.0034
	(0.0013)	(0.0015)	(0.0017)	(0.0012)	(0.0013)	(0.0009)	(0.0008)
Top AI papers	-0.0130	0.0017	0.0014	-0.0086	0.0138	0.0071	0.0035
	(0.0028)	(0.0032)	(0.0038)	(0.0040)	(0.0027)	(0.0023)	(0.0017)
female	-0.0868	-0.0084	-0.0082	-0.0678	0.0586	0.0139	0.0216
	(0.0049)	(0.0025)	(0.0030)	(0.0036)	(0.0050)	(0.0041)	(0.0034)
asian	0.0747	0.0463	0.0266	0.0429	-0.0765	-0.0312	-0.0227
	(0.0071)	(0.0038)	(0.0045)	(0.0055)	(0.0072)	(0.0058)	(0.0050)
white	0.0389	0.0083	0.0020	0.0123	-0.0345	-0.0087	-0.0116
	(0.0066)	(0.0032)	(0.0040)	(0.0050)	(0.0067)	(0.0055)	(0.0047)
Constant	0.4186	0.0527	0.0934	0.1799	0.4965	0.2138	0.1401
	(0.0062)	(0.0030)	(0.0038)	(0.0048)	(0.0063)	(0.0052)	(0.0044)
N	84180	84180	84180	84180	84180	84180	84180
Adj. R2	.0650735	.0533059	.0378738	.0347614	.0518361	.0370224	.0919957
Mean	.4441435	.0746258	.1030173	.1832858	.4657995	.2	.1330245
Papers|dep. var =1	2.063309	3.495543	2.76107	1.891438	1.587845	1.63685	1.66235
AI Papers	.7787258	1.594874	1.397947	.5305593	.547831	.5959848	.5322379
Top AI papers	.1406601	.3369946	.2809041	.0872383	.1109893	.1230102	.1120736

unit: (person, year). Conditional on school and cohort FE.

TABLE A.3. Table A3: Mobility after PhD (Outcomes based on $j(i,t+1)$ next year)

	At Nontop Firms				At Top Firms				In Academia			
	(1) Industry	(2) Move	(3) Top	(4) Academia	(5) Industry	(6) Move	(7) Top	(8) Academia	(9) Industry	(10) Move	(11) Top	(12) Academia
Any Paper	-0.0087	-0.0060	0.0031	0.0075	-0.0087	0.0034	-0.0087	0.0076	0.0044	0.0099	0.0024	-0.0074
	(0.0017)	(0.0025)	(0.0010)	(0.0014)	(0.0024)	(0.0043)	(0.0043)	(0.0020)	(0.0012)	(0.0015)	(0.0004)	(0.0016)
Any AI Paper	0.0000	0.0092	0.0083	0.0006	0.0024	-0.0000	0.0086	-0.0012	0.0002	0.0005	0.0032	-0.0001
	(0.0024)	(0.0038)	(0.0017)	(0.0020)	(0.0030)	(0.0053)	(0.0051)	(0.0025)	(0.0017)	(0.0022)	(0.0007)	(0.0023)
Any Top AI Paper	-0.0098	0.0023	0.0010	0.0047	0.0037	-0.0098	0.0052	-0.0019	-0.0017	0.0030	0.0013	0.0055
	(0.0039)	(0.0059)	(0.0030)	(0.0032)	(0.0033)	(0.0059)	(0.0057)	(0.0028)	(0.0025)	(0.0033)	(0.0012)	(0.0033)
exp1	0.0778	-0.1179	-0.0074	-0.0670	0.0075	0.0239	-0.0394	-0.0178	-0.1479	-0.3203	-0.0236	0.0361
	(0.0081)	(0.0128)	(0.0045)	(0.0066)	(0.0117)	(0.0233)	(0.0224)	(0.0092)	(0.0065)	(0.0086)	(0.0022)	(0.0088)
exp2	-0.0593	0.0314	-0.0035	0.0505	0.0041	-0.0374	0.0541	0.0084	0.1056	0.2463	0.0180	-0.0087
	(0.0079)	(0.0129)	(0.0045)	(0.0063)	(0.0118)	(0.0235)	(0.0227)	(0.0087)	(0.0063)	(0.0085)	(0.0021)	(0.0090)
exp3	0.0147	0.0006	0.0019	-0.0122	-0.0033	0.0109	-0.0161	-0.0012	-0.0258	-0.0610	-0.0045	0.0026
	(0.0025)	(0.0041)	(0.0014)	(0.0019)	(0.0039)	(0.0076)	(0.0074)	(0.0027)	(0.0020)	(0.0026)	(0.0006)	(0.0029)
female	-0.0145	0.0135	-0.0016	0.0092	-0.0036	-0.0067	0.0045	0.0023	-0.0083	0.0011	-0.0023	0.0213
	(0.0020)	(0.0031)	(0.0011)	(0.0015)	(0.0029)	(0.0055)	(0.0054)	(0.0022)	(0.0016)	(0.0021)	(0.0005)	(0.0023)
asian	0.0072	0.0196	0.0043	-0.0000	0.0043	0.0042	0.0019	-0.0009	0.0053	0.0145	0.0007	0.0192
	(0.0023)	(0.0036)	(0.0014)	(0.0017)	(0.0031)	(0.0062)	(0.0061)	(0.0023)	(0.0020)	(0.0026)	(0.0007)	(0.0028)
white	0.0066	-0.0109	-0.0019	-0.0033	0.0051	-0.0047	0.0050	-0.0044	0.0035	-0.0001	0.0008	-0.0070
	(0.0016)	(0.0026)	(0.0010)	(0.0013)	(0.0027)	(0.0049)	(0.0049)	(0.0020)	(0.0014)	(0.0019)	(0.0005)	(0.0020)
MI_pos	0.0492	-0.0017	-0.0112	-0.0332	-0.0938	0.0319	-0.0521	-0.0044	-0.0584	-0.1163	-0.0159	0.2889
	(0.0085)	(0.0584)	(0.0033)	(0.0079)	(0.0972)	(0.1010)	(0.1004)	(0.0015)	(0.0299)	(0.0312)	(0.0018)	(0.0427)
Re-Scientist	0.0137	0.0252	0.0062	-0.0048	0.0074	-0.0067	0.0078	-0.0054	-0.0036	-0.0200	-0.0014	0.2195
	(0.0010)	(0.0018)	(0.0007)	(0.0008)	(0.0018)	(0.0036)	(0.0035)	(0.0014)	(0.0029)	(0.0035)	(0.0010)	(0.0035)
Engineer	0.0206	0.0257	0.0064	-0.0131	0.0111	-0.0204	0.0194	-0.0094	0.0233	-0.0031	-0.0026	0.1803
	(0.0008)	(0.0014)	(0.0005)	(0.0006)	(0.0015)	(0.0030)	(0.0029)	(0.0012)	(0.0042)	(0.0047)	(0.0011)	(0.0048)
Manager	0.0164	0.0107	0.0002	-0.0088	0.0060	0.0001	0.0013	-0.0049	-0.0090	-0.0401	-0.0029	0.2014
	(0.0008)	(0.0017)	(0.0006)	(0.0007)	(0.0012)	(0.0031)	(0.0029)	(0.0008)	(0.0043)	(0.0048)	(0.0011)	(0.0051)
Postdoc	-0.1301	0.0858	0.0107	0.0800	-0.0883	0.1163	-0.1255	0.0803	0.0096	0.0532	0.0014	0.1700
	(0.0056)	(0.0060)	(0.0023)	(0.0047)	(0.0120)	(0.0152)	(0.0153)	(0.0113)	(0.0018)	(0.0023)	(0.0006)	(0.0024)
Inverse Mills	0.0388	-0.2809	0.0079	-0.0774	0.0185	0.0071	-0.0115	-0.0526	-0.0630	-0.1664	0.0087	-0.4381
	(0.0338)	(0.0510)	(0.0195)	(0.0256)	(0.0448)	(0.0860)	(0.0835)	(0.0337)	(0.0264)	(0.0357)	(0.0092)	(0.0381)
Junior Faculty									-0.0717	-0.1006	-0.0087	0.3151
									(0.0011)	(0.0015)	(0.0004)	(0.0018)
Assoc. Prof									-0.0633	-0.0963	-0.0065	0.3078
									(0.0010)	(0.0014)	(0.0003)	(0.0018)
Full Faculty									-0.0621	-0.1004	-0.0066	0.3059
									(0.0011)	(0.0014)	(0.0003)	(0.0018)
Constant	0.9210	0.1899	0.0140	0.0584	0.9695	0.0755	0.9310	0.0282	0.1355	0.2499	0.0137	0.6947
	(0.0060)	(0.0091)	(0.0033)	(0.0047)	(0.0086)	(0.0164)	(0.0160)	(0.0067)	(0.0048)	(0.0064)	(0.0017)	(0.0068)
N	286824	285317	286824	286824	55997	55708	55997	55997	382993	380725	382993	382993
Adj. R2	.0415408	.0184909	.0065725	.0291964	.018739	.010468	.0138971	.0273261	.0548357	.0910459	.0065989	.1984217
Mean	.9623462	.1158992	.0136983	.0230385	.9863564	.0643175	.9400325	.0075718	.0502411	.0980682	.0052455	.8660445

unit: (person, year). Conditional on school (highest degree) and year FE.

TABLE A.4. Accepted vs. Rejected Authors in Matched Sample

Variable	(1) Rejected	(2) Accepted	(3) Diff.
mean_rating	6.089	6.466	0.377***
	(0.585)	(0.726)	(0.014)
min_rating	5.298	5.425	0.126***
	(1.064)	(1.271)	(0.025)
ypred	0.503	0.649	0.146***
	(0.220)	(0.282)	(0.005)
yr_event	2,020.368	2,020.368	0.000
	(1.465)	(1.465)	(0.031)
yr_event	2,020.368	2,020.368	0.000
	(1.465)	(1.465)	(0.031)
ypred_q	2.725	2.725	-0.000
	(0.504)	(0.504)	(0.011)
mean_rating_q	2.754	2.754	0.000
	(0.458)	(0.458)	(0.010)
min_rating_q	2.397	2.397	-0.000
	(0.784)	(0.784)	(0.017)
lag1_job	1.988	1.988	0.000
	(1.010)	(1.010)	(0.022)
female_prob	0.276	0.276	0.000
	(0.326)	(0.311)	(0.007)
pscore	0.355	0.370	0.015***
	(0.111)	(0.113)	(0.002)
lag2_top	0.078	0.105	0.027***
	(0.269)	(0.307)	(0.006)
lag1_top	0.093	0.127	0.034***
	(0.290)	(0.333)	(0.007)
lag0_top	0.106	0.145	0.039***
	(0.307)	(0.352)	(0.007)
lag2_acad	0.604	0.651	0.047***
	(0.489)	(0.477)	(0.010)
lag1_acad	0.621	0.621	-0.000
	(0.485)	(0.485)	(0.010)
lag0_acad	0.594	0.574	-0.020*
	(0.491)	(0.495)	(0.011)
lag2_job	1.913	2.051	0.139***
	(1.059)	(1.030)	(0.022)
lag1_job	1.988	1.988	0.000
	(1.010)	(1.010)	(0.022)
lag0_job	1.942	1.877	-0.065***
	(0.983)	(1.003)	(0.021)
Observations	4,401	4,401	8,802

Notes: This table shows the OLS estimates of regression (1) of first job on publications before graduation, for authors in the Scopus-Revelio matched sample — who have published a CS paper according to Scopus database, and have matched to a Revelio individual profile with self-reported education and job history. The sample is restricted to 152,900 matched authors who report at least a college degree. Each person's first job is characterized based on her primary employer in the year after obtaining her highest degree. The first four columns are outcomes related to employment in industry, whereas the next three columns are related to employment in academia.

TABLE A.5. Diff-in-diff estimates for Employment by Top Firms (Full Sample)

	Employed by a Top Firm			
	(1) Pooled	(2) From Industry	(3) From Academia	(4) From Student/Postdoc
Post event=1	0.0459	0.0445	0.0111	0.0618
	(0.0028)	(0.0060)	(0.0026)	(0.0043)
Post event=1 × Accepted	0.0182	0.0057	0.0087	0.0233
	(0.0048)	(0.0099)	(0.0053)	(0.0071)
Constant	0.0082	0.0155	0.0027	0.0072
	(0.0007)	(0.0014)	(0.0008)	(0.0010)
Observations	95412.0000	22700.0000	21211.0000	51501.0000
Adj. R^2	0.1842	0.2135	0.1881	0.1819

Notes: This table shows the estimates of difference-in-differences regression (3) among authors in the full sample (without matching). The estimation sample is restricted to the event window [-5,3] years relative to the first ICLR submission of each author. Standard errors are clustered at author level.