The Race between Academia and Industry for AI Researchers*

Francesca Miserocchi† Savannah Noray‡ Alice Wu§

First Draft: August 2025
This Draft: March 29, 2026
Please click here for the latest version.

Abstract

The advances of artificial intelligence (AI) are built on the groundwork laid by researchers. We study the labor market competition between academia and industry for AI researchers and its consequences for public knowledge production. Using data on 150,000 computer science researchers, we document a major reallocation of AI talent toward top technology firms between 2005 and 2020. Publications at AI conferences predict transitions to top firms more strongly than to academia. Exploiting acceptance decisions at a leading AI conference, we compare accepted authors with similar rejected authors and find that a publication increases the probability of moving to a top firm by 2–6 percentage points in the next 1-3 years. Sorting to top firms is stronger for male researchers, whereas female students and postdocs are more likely to get tenure-track positions following a publication. Researchers who move to top firms subsequently publish fewer papers, resulting in approximately 1,000 fewer AI papers and 2,000 fewer papers in other computer science areas per year in the public domain.

*We thank Anna Le Sun, Jonathan Tong, and Alejandro De La Torre for outstanding research assistance, and Antonio Coran, David Card, Omeed Maghzian, Peter McCrory, Bilge Mutlu, Gabrielle Pepin, Anna Salomons, Chris Taber, seminar audience at NBER Labor (Fall 2025), Anthropic, UW Madison INTEGRATE (CS), Wisconsin Business School, and Midwest SOLE for many helpful comments. This research benefited from a grant from the Lab for Economic Applications and Policy at Harvard Economics.

†Department of Economics, Tilburg University. F.Miserocchi@tilburguniversity.edu
‡Reimagining the Economy, Harvard Kennedy School. savannahnoray@g.harvard.edu
§Department of Economics, University of Wisconsin - Madison. hwu526@wisc.edu

1 Introduction

Technological progress relies heavily on academic research. Traditionally, most research works are produced in academia, but in one of the most important areas today — artificial intelligence (AI) — over 45 percent of the publications at academic conferences involve researchers from industry, driven by scientists hired by the largest tech firms.[1] 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. Given the different incentives in academia versus the private sector (Aghion, Dewatripont, and Stein 2008), the allocation of these researchers across employers could influence the private versus social value of AI research.

This paper focuses on the labor market competition for computer science researchers between academia and industry. We study the impacts of conference publications on the reallocation of researchers between employers, and investigate the consequences of sectoral shifts for the production of public knowledge. Leveraging a longitudinal dataset on the career trajectories of about 150,000 researchers who publish at computer science (CS) conferences, we begin by analyzing 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 are more predictive of job movements to top firms than to academia. Then, to identify the causal effects of publications, we use 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 that an accepted paper increases movement into top tech firms substantially in the next 1-3 years, especially among students and postdocs. Finally, we show that the reallocation of researchers from academia to industry has consequences for future research: workers who move to top firms do not maintain their lead in AI publications relative to workers who enter the tenure track, and they publish fewer papers in other CS areas.

Building on publication data collected from 89 major CS conferences in Wu (2025), we match individual authors with professional profiles collected by Revelio Labs by name and affiliation. We track the education and job histories of 150,000 researchers who published at least one CS paper between 2000 and 2023 and reported at least a college degree.[2] We begin by showing trends in the initial job placement of these authors. At the PhD level, the share working in academia three years after graduation declined from around 61 percent to 47 percent between 2005 and 2020.

[1] The share of publications with industry researchers is computed from the publication and author affiliation data from Scopus. We plot the trend in Appendix Figure A.1.

[2] The supply of computer scientists has been increasing over time. Every year about 4,000 individuals graduate with a PhD degree in Computer Science or Electrical Engineering in the U.S. alone——about three times the number of economics PhD graduates (Appendix Figure A.2).

Meanwhile, the share of PhD graduates joining one of the top tech firms, defined as the six big tech firms {Alphabet, Meta, Amazon, Apple, Microsoft, IBM}, doubled from 5 percent to 11 percent. These opposing trends highlight a substantial reallocation of CS researchers from academia toward industry.

How do industry and academic employers identify research talent? We show that publications – a signal of underlying research ability – play an important role in predicting the initial sorting of workers, including their placement into firms. 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 0.2 percentage points (p.p.) and increases the chance of working for a top firm by 0.5 p.p., which is a 7 percent increase relative to the mean. The correlation between AI publications and employment at top firms is the strongest among PhD graduates, consistent with the sorting patterns among CS PhD graduates in Wu (2025). In contrast, AI papers are less predictive of academic placement than papers published in other CS areas. After workers enter the labor market, AI publications produced on the job continue to predict transitions into top firms from both academia and nontop firms in industry, with effects persisting for at least three years after a researcher's first AI publication.

These results suggest that top tech firms attract researchers who publish at AI conferences and are able to hire from both industry and academia. The predictive effects could be driven by a mix of supply-push and demand-pull factors. Workers with stronger preferences for industry jobs could disproportionately select themselves to AI conferences. On the demand side, employers use publications to identify high-ability computer scientists. An ideal experiment for estimating the signal values of publications to employers would randomly publish papers by authors with similar skills and preferences. We proxy for this experiment by using data from one of the leading AI conferences, the International Conference on Learning Representations (ICLR), which publicly releases all submitted papers, referee reports, and editor decisions.

There is substantial uncertainty about whether a paper is accepted (and therefore published) by ICLR, especially for papers that receive similar referee ratings around the margin of acceptance. We use a difference-in-differences design to estimate the causal impacts of a publication on the career outcomes of researchers, such as moving to a top firm or getting a tenure-track job.

We focus on eighteen thousand authors who submit to ICLR for the first time between 2017 and 2020. On average, accepted authors are different from rejected authors: they submit papers of higher quality (according to referee ratings), are slightly younger, more likely to be white, and more likely to be in academia before submitting to ICLR. Given the lack of balance in pre-treatment characteristics, we follow Abadie (2005) and assume parallel career trajectories of accepted and rejected authors conditional on observables, which include referee ratings, demographics, and job history before the event. Under this identifying assumption, we use reweighting and matching

techniques to compare accepted authors with rejected authors who are observably similar and, importantly, submit similar-quality papers. In particular, we use inverse propensity tilting to achieve covariate balancing, which is shown to be doubly robust in staggered adoptions in Słoczyński et al. 2025.

In the original sample without reweighting, the acceptance by ICLR is estimated to increase movements into top firms from other places by 2.4 p.p. in the next three years after the conference. By reweighting rejected authors to achieve covariate balancing, we eliminate differential pre-trend in sorting to top firms and estimate that the publication increases mobility to top firms by 3 p.p. one year after the conference, and by 4.5-6.7 p.p. in the following two years.

We further examine the impacts separately by the original employer of the authors in the year before the ICLR conference. The largest estimated effect of an ICLR publication on sorting to top firms is observed among students and postdocs, suggesting that top firms use AI publications to identify research talent early in their career – before they are likely to reach peak productivity (Jones, Reedy, and Weinberg 2014). For this group, we also find a 10 p.p. increase in moving to an ICLR sponsor within 3 years. Corporate sponsorship represents an investment in networking and recruitment, indicating strong demand for research talent. The increase in mobility to sponsors is again driven by top firms, which sponsor the conference almost every year.

The ICLR quasi-experiment also allows us to investigate gender differences in career trajectories following a publication. An ICLR publication significantly increases the probability that female students secure tenure-track jobs in academia but does not increase their employment in top firms. For male students, the pattern reverses: acceptance increases mobility to top firms, with no corresponding boost into tenure-track academia. The gender differences in sorting have implications for the gender wage gap among computer scientists and on the direction of research, both of which we aim to investigate further.

Lastly, we study the consequences of the growing competition between academia and industry for top AI researchers. Using the Scopus-Revelio linked job and publication data, we track the publication trajectories of computer scientists who move to top firms versus tenure-track academic positions. Before moving, researchers who will join top firms exhibit steeper increases in publication output, especially in AI-related work, consistent with either firms selecting highly productive AI researchers or scholars increasing AI output in anticipation of recruitment. After the move, publication activity diverges sharply: output among authors moving to tenure-track jobs stabilizes and resumes growth, while that of researchers joining top firms falls and remains well below pre-move levels. A staggered difference-in-differences analysis confirms these patterns: five years after the move, publications decline by roughly 57 percent, AI-related publications by about 46 percent, and non-AI publications by approximately 65 percent among authors moving to top firms relative to researchers entering tenure-track positions.

Our findings have implications for the quantity and quality of research in the public domain. Based on the difference-in-differences analysis, a back-of-the-envelope calculation suggests that 1,050 AI papers (and 1,978 papers in other CS areas) are missing per year due to researchers opting for top firms over academic positions. We interpret this estimate with caution: if researchers increase output before moving in anticipation of top firm recruitment, we overstate the true decline; if top-firm movers would have continued on a steeper trajectory absent the transition, we understate it. We further show that top-firm movers are 10 percentage points less likely to produce any publicly cited work across all CS areas. The reallocation also reshapes the direction of high-impact public output (measured as the top decile of forward citations): while high-impact AI publications increase among top-firm movers, high-impact research in other CS areas declines.

This paper contributes to several strands of the literature. First, we contribute to the literature that examines the career progression of scientists and researchers, especially the role of publications in driving promotions and sorting between employers. Publications are arguably the most important measure of research ability. Research focusing 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. Our evidence on sorting to top firms is consistent with recent findings based on U.S. Census data that the average earnings of top 1% AI researchers have grown significantly since the 2010s (Akcigit, Chikis, Dinlersoz, and Goldschlag 2026).

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 also junior researchers, such as students and postdocs. AI publications are more predictive of initial job placement in top firms than of placement in academia among PhD graduates. This finding is robust in the quasi-experiment based on ICLR. This result indicates that the competition for talent begins at the very start of researchers' careers, thus potentially amplifying its impact on research production.

More broadly, our results also speak to the literature on industry–academia research trade-offs (Aghion, Dewatripont, and Stein 2008). Liang et al. (2024) finds that in the context of AI,

academic researchers are more likely to produce novel work, while industry researchers produce work with a higher impact (as measured by citations). These findings are echoed in Färber and Tampakis (2023), which also finds that AI research produced by industry coauthors receives more online attention. Jurowetzki et al. (2021) finds that researchers who move to industry experience only 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 documents trends in the reallocation of talent from academia to industry and describes the role of publications in predicting initial placement and subsequent job-to-job mobility. Section 4 presents our empirical design of using submissions to ICLR to estimate the causal impacts of publications on career outcomes. Section 5 examines the consequences of moving to top firms versus academia on the quantity and quality of publications. We conclude by discussing some immediate next steps for this project.

2 Data

We collected data on the career trajectories of computer science researchers who published at 89 CS conferences.[3] Our data has three main sources: (1) Scopus for CS publications, (2) Revelio 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.[4] We collected conference proceedings on Scopus, a large-scale publication database produced by Elsevier.[5] 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 them have at least one publication at an AI-related conference.[6] Figure A.3 shows the trends in topics of AI research in the past 25 years. Most AI publications in our data

[3] We include all conferences that are used to rank computer science departments by CSRankings (Berger 2017). The conferences listed on CSRankings were developed in consultation with faculty across a range of institutions, including via community surveys.

[4] See the list of publication venues that are used to rank computer science departments: `https://csrankings.org/#/index?all&us` (Berger 2017).

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

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

predate the rise of large language models (LLM), and focus more broadly on machine learning, neural networks, and applied areas such as natural language processing and computer vision.

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 150,000 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 build 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 and otherwise is her main employer as reported in her job history. We refer to individuals in this dataset as either "CS authors" or "CS researchers." 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.[7] ICLR is also one of the few major conferences that makes public both their 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, 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 to trackindividual research productivity before and after an ICLR

[7]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).

submission.

3 Descriptive Analysis: Job Placements and The Role of Publications

3.1 Placement of CS Researchers in Industry versus Academia

Using data spanning the past 20 years, we document that an increasing fraction of new graduates have sorted into industry jobs – specifically positions at top technology firms – rather than pursuing academic careers. We begin by documenting the trend within the least restrictive sample, which includes all CS authors who report at least a college degree on their professional profiles. Figure 1 (a) plots the share of graduates employed in industry (left y-axis) and in academic positions (right y-axis) three years after their highest degree completion against the year in which that degree was obtained.[8] The share of CS researchers whose initial placement is a job in industry rises from roughly 50 percent to 59 percent between 2005 and 2020; the share in academic positions displays a commensurate decline over the same period.

Academic positions primarily consist of postdoctoral and tenure-track positions and, therefore, require a PhD. Panel (b) shows that among PhDs, the share who work in industry three years post graduation increases from 39 percent in 2005 to 53 percent in 2020. Relatively speaking, this increase in industry employment is larger than the increase among bachelor or master graduates.

Panels (c) and (d) of Figure 1 examine these trends while focusing on the roles that are considered the most prestigious or competitive within each sector, for the full sample and the PhD subsample, respectively. In academia, we focus on tenure-track jobs, whereas in industry, we focus on positions at the top tech firms.[9] In the 2005 cohort, roughly 5 percent of PhD researchers were employed by top firms three years post graduation, compared to 26 percent on tenure track. Since then, placements at top firms have risen to about 11 percent for the 2020 cohort, while placements in tenure-track positions have declined to 15 percent. These opposing trends underscore a major shift of CS research talent from academia to industry, with PhD researchers becoming almost as likely to take a role in a top tech firm as they are on tenure track.[10]

[8] We focus on the third year after degree completion, when researchers are still at an early stage of their careers but are more likely to hold a stable position than immediately after graduation (e.g., PhD researchers may enter into temporary postdoctoral positions after graduating).

[9] Throughout the paper, we define top firms as the six largest tech firms: {Alphabet, Meta, Amazon, Apple, Microsoft, IBM}. As shown in Appendix Figure A.3, research on large language models did not appear until 2022. We plan to consider a more dynamic definition of top firms.

[10] An analogous – but muted – pattern emerges for less prestigious positions (Appendix Figure A.4), where we find employment in non-tenure-track academic positions declining (especially from 2011 onward) while non-top firm positions make up an increasingly larger share of employment over the period.

A natural question is whether the shift toward employment in top tech firms is driven by graduates with high research potential or instead by those with lower research potential or a weaker interest in academic careers. To explore this, we use several different proxies for research potential. Figure 2 depicts the share of top firms vs tenure-track positions, separately for students whose PhD institution is a top 50 CS department and students from lower-ranked programs. The substitution from academia to industry is evident for both types of graduates and, if anything, is a slightly sharper trend for top graduates. These trends replicate when using publications before graduation as proxies. Appendix Figure A.5 Panel (a) examines this pattern when defining top graduates as those with at least one AI publication in the year of (or those preceding) graduation. And, to account for the fact that the "bar" for who is considered a top graduate may rise over time, Appendix Figure A.5 Panel (b) plots these trends when defining top graduates as those in the top quartile of the number of pre-PhD AI publications.

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. In fact, publications are often mentioned as preferred qualifications in research scientist jobs in industry (see Appendix Figure A.6 for example).

3.2.1 Initial Sorting

To understand the role of publications in the initial sorting of workers between industry and academia, we regress characteristics of the first job after a person's highest degree on variables capturing her publication history:

$$y_i = \beta_0 + \beta_1 \text{ AI-Pubs}_i + \beta_2 \text{ Oth-Pubs}_i + X_i'\Theta + u_i \tag{1}$$

The outcome y_i represents the sector or position characteristic associated with person i's first job in the year after obtaining her highest degree (bachelor/master/PhD). We count the number of publications at AI conferences before graduation year (AI-Pubs_i), and the number of publications at other CS conferences before graduation (Oth-Pubs_i). AI publications are measured by conference proceedings at 24 conferences in the areas of machine learning, computer vision, NLP, and information retrieval on the web. We also control for observable characteristics such as gender, race, cohort, and school fixed effects.

We find that publications at AI conferences are strongly predictive of being placed in industry. Every additional AI paper before graduation predicts a 0.2 p.p. increase in the probability of working in industry (column 1 of Table 1) and a 0.5 p.p. increase in employment specifically by a

top firm (column 2), which is roughly a 7% increase relative to the mean. When we estimate the regression separately by highest degree, we find the strongest relationship between AI publications and placements at top firms for PhD graduates. Every AI paper before PhD increases the probability of initial placement at top firms by 0.7 p.p., a 9% increase relative to the mean (Appendix Figure A.7(a)).

In particular, AI papers are positively correlated with becoming a research scientist (column 3 of Table 1), consistent with the notion that these jobs typically explicitly require research skills. In contrast, such publications are negatively correlated with becoming an engineer (column 4), a role that does not emphasize research.

The positive selection into industry by researchers is a primarily an AI-specific phenomenon. The more traditional pattern of research ability predicting academic employment emerges when focusing on other CS papers outside of AI. Such publications are positively related to getting a job in academia, as shown in columns 5-7 of Table 1. This pattern is echoed when focusing on PhD graduates in particular, for whom academic placements are most relevant. PhD graduates with AI publications are less likely to be placed on the tenure track, while graduates with publications in other CS areas are more likely to stay in academia (Appendix Figure A.7 and Table A.1).

3.2.2 Job Mobility Post Graduation

Do publications after graduation matter for sorting in the labor market between academia and industry, and between nontop and top firms? We answer this question by estimating OLS regressions of year-to-year job transitions (after graduation from each person's highest degree) on indicators for publications in the current year:

$$y_{i,t+1} = b_0 + b_1 \,\text{Any-AI-Pub}_{it} + b_2 \,\text{Any-Oth-Pub}_{it} + X_{it}'\theta + \varepsilon_{it} \tag{2}$$

Any-AI-Pub_{it} and Any-Oth-Pub_{it} indicate at least one publication by worker i in year t in an AI-related or non-AI related conference, respectively. 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).[11] The outcomes $y_{i,t+1}$ are defined based on employment in the next year, including indicators for whether the worker is employed by industry, by top firms in particular, or by academia in $t+1$. We estimate the regression separately depending on the worker's employer in year t, taking into account that workers outside top firms may move at a different rate from workers at top firms, or in academia.

Table 2 reports the results. Echoing earlier findings, we find that publications play a role in transitions to between academia and industry. Academics who publish a paper in AI or other

[11] See Table A.2 for the full set of controls.

areas experience a 1.1 p.p. or almost 18% increase in the probability of moving to industry.[12] For workers already employed by industry (nontop or top firms), having either type of publication is negatively correlated with staying in industry (columns 1 and 4) and positively correlated with moving to academia (columns 3 and 6).

Moreover, publishing any AI paper strongly predicts moving to a top firm from nontop firms or academia (columns 2 and 8). Among workers at nontop firms, any AI publication this year is estimated to increase the probability of moving to top firms by 1.5 p.p., doubling the likelihood that workers move up the job ladder within industry in a given year (column 2). This relationship is nearly three times as large as the transitions into academia from nontop firms (column 3). Any AI paper also predicts a 0.8 p.p. increase in movements from academia to top firms, a 118% increase relative to the mean (column 8).Figure A.7(c)-(d) breaks down these 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.

Having established that AI publications predict one's placement or mobility into top tech firms, we estimate a descriptive event study around AI publications. We focus on the first AI publication after (or in the year of) graduation and estimate a two-way fixed-effect event study regression:

$$y_{it} = \beta_0 + \sum_{\forall l \neq -1} \beta_l D_{it}^{(l)} + \delta_t + \alpha_i + \varepsilon_{it} \tag{3}$$

where i indexes person and t indexes calendar year. Letting τ_i denote the year in which person i publishes her first AI paper after graduation, we define the event dummy $D_{it}^{(l)} = 1$ if it has been l years since the first AI paper, where $l \in [-23, 23]$. We include individual and year fixed effects (α_i and δ_t), and standard errors are robust and clustered at the person level. Only individuals who publish at least one AI paper after graduating are included in the sample.

Panel (a) of Figure A.8 shows the event study estimates. In the year of the first AI publication ($l = 0$), authors are 4.1 p.p. more likely to work in a top firm relative to the year before. The increase could be driven by (1) workers who move into top firms and publish in the same year, (2) workers who publish first and then get poached by top firms, and (3) incumbents who publish and are retained by top firms. Given that we do not find a correlation between publications and retention at top firms based on year-to-year transitions (column 5 of Table 2), we estimate the same event study with the sample restriction that authors who are not employed by top firms at -1, in order to focus on the impacts of the publication on moving to top firms from *other* places. Panel (b) of Figure A.8 shows the estimates under this sample restriction. The similar increases in

[12] Workers in academia include postdocs or non-tenure-track researchers. We carefully differentiate between them in the causal analysis among authors who submit to ICLR in Section 4.

employment by top firms suggest that most of the effects are driven by movers into top firms rather than stayers.

We then compare the estimates $(\beta_1 - \beta_0)$ to ensure that we are focusing on workers who first publish and are subsequently poached. One year after the publication, we find an additional 1.8 p.p. increase in the employment at top firms that represents an inflow of publishing authors to top firms. This estimate is very close to the 1.5 p.p. predictive effect of any AI publication for moving to top firms from nontop firms shown earlier (column 2 of Table 2). The probability of working at top firms continues to rise for two years and stabilizes around 7 p.p. three to five years after the publication.

Taken together, the results in this section suggest that top tech firms (increasingly) demand and compete for AI research talent. However, a researcher's publication record may reflect unobserved worker characteristics – such as preferences for landing a job at a top tech firm – that also drive sorting into top tech firms. To isolate the role of firms demand for AI talent in driving these patterns, we now turn to the next section, where we provide evidence that AI publications have a causal effect on moving to top firms.

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 difference-in-differences design, in which the control group comprises rejected authors who are observably similar as accepted ones. This design allows us to identify the extent to which sorting into top tech firms is driven by the demand for AI researchers rather than supply-side factors such as preferences. We also analyze the gender differences in the impacts of a publication on obtaining a tenure-track job versus a job at top firms.

4.1 Empirical Strategy

We observe the numeric ratings and comments by referees, based on which editors decide whether a paper can be accepted into the conference and hence published at the ICLR conference proceedings. 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.[13] More than 80% of the papers with an average rating above 6.5 were accepted by the editor, and most

[13]The rating scale changed in 2020, but the margin of acceptance remained the same at 6. We take this change into account when estimating the weights in the covariate balancing and construction of propensity score of acceptance described later in this section.

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.9). This uncertainty suggests that the editorial decisions among papers that receive similar ratings could be arbitrary, depending crucially on the assignment of editor.[14] The importance of editors' own evaluation is not unique to ICLR. NeurIPS, the highest-ranked AI conference, conducted a consistency experiment in 2014 and found that 60% of accepted papers were rejected if they were reviewed a second time by a different editor (Cortes and Lawrence 2021).

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 obtaining a tenure-track job. The underlying assumption is that, absent the ICLR decision, accepted and rejected authors would have followed parallel career trajectories. Focusing on the first time an author submits to ICLR, we consider the following difference-in-differences regression:

$$Y_{it} = b_0 + b_1 Post_{it} + b_2 Accepted_i + b_3 Post_{it} \times Accepted_i + \delta_t + \alpha_i + \nu_{it} \tag{4}$$

where $Post_{it} = 1[t > \tau_i]$ indicates years after the author i submits to an ICLR conference 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 person fixed effects α_i to account for time-invariant unobserved heterogeneity between individuals. The coefficient of interest is b_3, which represents the average treatment effect of an accepted paper on the treated (ATT) under the parallel trend assumption. We also estimate a dynamic difference-in-differences or event study model:

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

$D_{it}^{(l)} = 1[t - \tau_i = l]$ 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 (4) can be viewed as an average of $\{\gamma_l : l > 0\}$, pooling the post periods.

[14] Given that the editors' opinions play an important role in the final decision, the numeric ratings by referees alone do not generate a well-defined running variable in the regression discontinuity design. Figure A.9 shows a smooth "S" curve between acceptance and mean ratings. Fuzzy RD around the threshold (mean=6) does not yield a significant first-stage.

4.1.1 Identifying Assumptions

The identification of ATT in the equations (4) and (5) hinges on the parallel trend assumption that, given any periods $s \neq t$, $E[Y_{is}(0) - Y_{it}(0)|Accepted_i = 1] = E[Y_{is}(0) - Y_{it}(0)|Accepted_i = 0]$, in which $Y_{it}(0)$ represents the counterfactual career outcome of an author absent the ICLR acceptance. As pointed out in Abadie 2005, this assumption may be implausible if pre-treatment characteristics that are thought to be associated with the career trajectories of authors are unbalanced between the accepted and rejected authors.

Table A.3 provides a comparison between accepted and rejected authors who submit to the ICLR for the first time. On average, the submissions by accepted authors have a mode of rating above 6 (the margin of acceptance), whereas the papers by rejected authors have a mode around 4. The minimum rating received by a paper represents the most negative review. Accepted authors on average get a minimum around 5.5 while rejected authors get a minimum around 3.3. The two groups also look different in demographics and educational background: accepted authors are 0.35 years younger, more likely to be male, white, and US-based, and more likely to report a PhD degree on their Revelio profiles. Almost 14% of accepted authors are employed by a top firm in the year before the ICLR conference, 6% higher than the rate among rejected authors.

We are most interested in estimating the impacts of an ICLR publication on sorting to top firms from other places. Among authors who are not employed at top firms in the year prior to their first ICLR conference, accepted and rejected authors show different pre-trends (Appendix Figure A.10).[15] Accepted authors are more likely to have worked for top firms two to five years prior to the ICLR event. The gap only becomes zero by the sample restriction at -1. The lack of balance in pre-treatment characteristics in Table A.3 and the pre-trend in Figure A.10 suggest we have to act with caution when implementing the difference-in-differences design.

To address this concern, we make alternative assumptions as in Abadie (2005), and adopt a two-step strategy.

Assumption 1 (Parallel Trends Conditional on Observables) *Let X_i represent pre-treatment observable characteristics of the author and the quality of the paper. Given any periods $s \neq t$,*

$$E[Y_{is}(0) - Y_{it}(0)|Accepted_i = 1, X_i] = E[Y_{is}(0) - Y_{it}(0)|Accepted_i = 0, X_i] \tag{6}$$

Assumption 2 (Common Support) *Given pre-treatment observable characteristics X_i,*

$$Pr(Accepted_i = 1 \,|\, X_i) < 1 \tag{7}$$

[15] The sample is restricted to authors who are not employed by top firms at -1, so there is no difference between accepted and rejected authors at -1. Figure A.10 shows that accepted authors are 1 p.p. more likely to be employed by top firms at -3 than (unweighted) rejected authors, suggesting that they may have been more positively selected or recognized by top employers even before the result of the ICLR event is known.

Under Assumptions 1 and 2, we can identify ATT by estimating regressions (4) and (5) if the non-parallel career trajectories are caused by the observed compositional differences between accepted and rejected authors. We include in X_i the mode and the minimum rating received by the author's submission, the number of authors on the paper, calendar year group of the conference, demographic characteristics such as age, gender, and education level of an author, publication history and job history between [-3,-1] years relative to the ICLR conference (see Table 3 for the complete list). Other than the referee ratings and team size, every control in X_i is either time-invariant or a pre-treatment characteristic. Controlling for the mode and the minimum referee ratings allows us to compare authors who submit papers of similar quality.

4.1.2 Weighting and Matching Procedures

Our two-step strategy involves first estimating a propensity score of acceptance based on the observable characteristics of the authors and the quality of their papers, then estimating difference-in-differences models with inverse propensity weights.

First, following Rosenbaum and Rubin (1983), Hirano, Imbens, and Ridder (2003), and Imbens (2004), we fit a logistic regression of paper acceptance on X_i, record the propensity score $\hat{p}_i$ from maximum likelihood estimation, and assign weight $\frac{\hat{p}_i}{1-\hat{p}_i}$ to rejected authors and 1 to accepted authors. The inverse propensity weighting (IPW) allows us to compare accepted authors with rejected authors who have a similar likelihood of having their papers accepted on average, and then we interpret the difference-in-difference estimates as ATT under the identifying assumptions.

Second, we adopt inverse propensity tilting (IPT) in Graham, De Xavier Pinto, and Egel (2012) and estimate the propensity scores by balancing the covariates between accepted and rejected authors:

$$E[X_i] = E[\frac{1-Accepted_i}{1-p(X_i'\theta)} \times X_i]. \tag{8}$$

in which the propensity score is logit $p(X_i'\theta) = Pr(Accepted_i|X_i) = \frac{exp(X_i'\theta)}{1+exp(X_i'\theta)}$. We estimate the parameters θ by solving the moment conditions (8) and assign weight $\frac{p(X_i'\hat{\theta})}{1-p(X_i'\hat{\theta})}$ on each rejected author. Given the IPT weights, the means of pre-treatment X among reweighted rejected authors equal the means among treated or accepted authors, allowing us to make causal comparisons between groups with identical mean observables. Słoczyński, Uysal, and Wooldridge (2025) shows that in difference-in-differences settings, the doubly robust estimator proposed by Sant'Anna and Zhao (2020) is numerically equivalent to the IPW estimator of Abadie (2005) with IPT weights. The covariate balancing estimator with a logistic propensity score model is also equivalent to entropy balancing proposed by Hainmueller (2012) (Zhao and Percival 2017; Tan 2019).

Given our interest in the impact of a publication on sorting to top firms from other places, we focus on authors who were not employed by top firms in the year before their first ICLR conference.[16] We estimate the IPT weights by solving (8) under this sample restriction. Table 3 shows the differences between accepted and rejected authors, before and after reweighting. Relative to the full sample shown in Table A.3, we drop 15% of the accepted authors and 10% of the rejected authors, but the (unweighted) mean characteristics remain largely unchanged as shown in the first 2 columns of Table 3.[17] The mode of referee rating for an accepted author is 6.31, which is 2.31 points higher than the rating received by a rejected author. After the reweighting via IPT, rejected authors also on average have a mode rating of 6.31, which is by construction, given that the mode is included in X and we solve the moment equations in (8). The only variable summarized in Table 3 that is not included in X is mean rating. The reweighted average among rejected authors is 6.16, which is 0.38 points lower than the average among accepted authors. This reweighted difference is much smaller than the 2-point gap without reweighting. Since X includes indicators for whether a worker is employed by top firms -3, -2, -1 years relative to her first ICLR conference, Figure A.10 shows that after reweighting, there is no differential pre-trend in employment by top firms during the window $[-3,-1]$.

In addition to the propensity score weighting (likelihood or covariate balancing), we also apply coarsened exact matching to select a matched control group for accepted authors (Iacus et al. 2012; Jäger and Heining 2022; Sarsons 2017). We match authors based on the quality of the papers, demographic characteristics, and publication and job histories before they submit to the ICLR for the first time. The matching procedure is detailed in Appendix A1. If we draw without replacement, we can match 1,263 accepted authors to 1,263 rejected authors, who look more similar to each other than in the original sample (without reweighting). However, since authors are first put in the bins based on observable characteristics, the covariates are not perfectly balanced as in the IPT procedure. Relative to reweighting techniques, establishing 1-1 matches also drops a large number of accepted and rejected authors, reducing the power in our analysis. We use the matching method mainly as a robustness check.

The reweighting and matching methods above allow us to compare accepted authors with rejected authors who submit papers of similar quality and are observably similar. In the second step, we estimate the difference-in-differences model (4) and the event study (5) in the reweighted or matched sample. Our preferred approach is to use IPT weights that achieve balance in covariates

[16] ICLR conferences are typically held in April or May. The submission deadline is often in Q4 (Oct/Nov) of the previous year. We are most interested in the causal impact of publishing an ICLR paper on moving to top firms from other places, and therefore exclude authors who were at top firms at -1, and those who moved to top firms before the conference even happened.

[17] Under the sample restriction that authors were not employed by a top firm at -1 (the year before the person's first ICLR submission), the variable "Employed by Top at -1" always equals 0 as shown in Table 3.

(unlike likelihood-based IPW) and do not throw away any observation (unlike matching).[18] We present estimates from every method below to show robustness, but we will focus more on the results with IPT weights, which we interpret as ATT under identifying assumptions 1 and 2.

4.2 Impacts of Acceptance on Moving to Top Firms

Motivated by our finding in Section 3 that publications are strongly correlated with sorting to top firms, we estimate the difference-in-differences models for the outcome of moving to top firms, which could either indicate the hiring of researchers from academia or poaching from other firms in the industry. Table 4 presents estimates of the difference-in-differences regression (4) without reweighting. The estimation sample comprises authors who are not at a top firm right before submitting to ICLR ($l = -1$). Accepted authors are 2.4 p.p. more likely to move to top firms following the ICLR conference than rejected authors, which is 54% higher than the mobility rate among rejected authors. We then estimate (4) separately based on the author's employment in the year before their first ICLR conference: industry (nontop firms), academia, or students/postdocs. Among students and postdocs, an ICLR publication increases the probability that they work at top firms by 3.1 p.p., suggesting that early publication success is important for getting a good job in the industry.

We also present the difference-in-differences estimates in the reweighted sample, where the rejected authors are (after the reweighting) comparable to accepted authors and therefore are more likely to be on parallel trajectories. Table 5 shows that the acceptance of a person's initial submission increases the probability of moving to a top firm by 4.4 p.p., driven by authors from nontop firms in the industry (4.2 p.p. in column 2), and students and postdocs (5.7 p.p. in column 4). Before reweighting, rejected authors are more likely to come from industry and less likely to be students/postdocs prior to the event (Table 3). The IPT procedure assigns a lower weight to rejected authors from industry who are not comparable to accepted authors and a higher weight to rejected students.

Moving on to the event study regression (5), we find in the unweighted sample that an accepted paper increases mobility to top firms by 2 p.p. after the conference (red squares in Figure 3). Given the concerns about the parallel trend assumption discussed in Section 4.1, we estimate the same model on the reweighted samples in which each rejected author is weighted by the $\frac{\hat{p}_i}{1-\hat{p}_i}$. The propensity score is estimated by maximum likelihood or by solving the moment equations (8). Under reweighting by likelihood-based propensity scores (IPW), the estimates (blue triangles in Figure 3) are slightly larger than the estimates without weighting.

[18] Table A.3 compares accepted authors with rejected authors who are reweighted by likelihood-based propensity scores. IPW also improves balance in covariates but in contrast with IPT, it doesn't ensure the means would be matched exactly.

Under IPT reweighting that achieves covariate balancing, an accepted paper is estimated to increase the mobility rate to top firms from other places by 2.9 p.p. in the year after the conference, and by 4.5-6.7 p.p. two to three years after (shown in green circles in Figure 3. In comparison with reweighted samples, matching accepted authors with rejected authors without replacement drops a large number of authors. The estimates from the matched sample (yellow X in Figure 3) are noisier, but similar in magnitude to the reweighted estimates.

Taken together, we find a robust 2-3 p.p. increase in mobility to top firms from other employers one year after the conference across methods shown in Figure 3. The short-term estimate is similar to the predictive effect of an AI publication on moving to top firms from our descriptive analysis that covers 24 AI conferences (column 2 of Table 2 and Figure A.8). By reweighting or matching, we find a stronger increase in mobility to top firms 2-3 years later, which suggests that an AI publication has persistent career benefits for computer scientists.

Our analysis below shows separate estimates by the origins of authors before submitting to ICLR. Given the importance of comparing the accepted authors with rejected authors who are more likely to be on parallel career trajectories absent the ICLR decision, we will present results with IPT reweighting. But the findings are robust across methods.

4.2.1 Heterogeneity by the Origin of Authors

The strongest effect of publication on movements to top firms comes from students and postdocs early in their careers, according to the estimates of regression (4) in Tables 4-5. We estimate the dynamic difference-in-differences model (5) in the IPT-reweighted sample, separately by the origin of authors at -1 (the year before each person's first ICLR conference). Figure 4 shows that having a student or postdoc's initial submission accepted boosts her chance of getting a job at top firms in the same year as the ICLR conference by 1.8 percentage points. Since we have excluded authors who get hired by top firms at 0 before the ICLR conference takes place (in May), this immediate effect is driven by students who move to top firms *after* the conference. In the following three years, the estimated effects of publication on being employed by top firms continue to increase from 4.3 p.p. to 9.4 p.p., suggesting a persistent effect of early publication success on sorting to top firms.

The effects for authors who are employed by academia (not students/postdocs) remain statistically significant at around 2 p.p. in the three years following the ICLR conference. We interpret the effects among academics as evidence that top firms are poaching productive researchers from academia, not just among early career researchers but also among experienced faculty. The estimates for authors from nontop firms in the industry are positive but noisier than those from the other groups. In three years, accepted authors from nontop firms are 5 p.p. more likely to be

working at a top firm than rejected authors with similar job histories before submitting to ICLR.[19]

We test whether the positive effect of the ICLR publication on upward mobility to top firms is driven by reporting bias. Accepted authors may be more likely to update their résumé given an additional publication, and therefore show a higher employment rate by top firms in the data. We replace the outcome in the event study regression (5) with an indicator for reporting any full-time job on LinkedIn each year. Figure A.13 shows that accepted authors are no more likely to report a full-time job than rejected authors after the conference, and the result is robust across methods. We therefore argue that the positive effects of a publication on sorting to top firms are not driven by reporting differences.

4.3 Impacts on Moving to Corporate Sponsors of ICLR

CS conferences provide employers with an opportunity to meet and recruit talent. We observe corporate sponsors of ICLR that pay for varying levels of exposure at the conference, ranging from USD 5,000 to 100,000. Sponsors have a dedicated booth at the conference and often host recruiting events. The decision to sponsor an ICLR conference indicates a strong demand for AI talent. Figure A.14 shows that Alphabet (Google), Amazon, and Facebook sponsored ICLR every year between 2017 and 2025. The other three top firms – Microsoft, Apple, and IBM – also sponsored the conference at least 5 times. A potential explanation for the short-term increase in moving to top firms (Figure 3) is that accepted authors network with recruiters and incumbent workers at top firms at the conference.

Given the full lists of ICLR sponsors, we test whether authors with an accepted paper at ICLR are more likely to move to a corporate sponsor following the conference, and whether the effect differs for top firms versus other sponsors. Figure 5 shows our event-study estimates for the outcome being employed by an ICLR sponsor in event year ($l = 0$). Within one year, accepted authors are 1.7 p.p. more likely to be employed by an ICLR sponsor that is present at the conference, according to the estimates across all methods. Similar to Figure 3, we find a stronger increase in employment among ICLR sponsors 2-3 years after the conference in the reweighted sample, in which we compare accepted authors with similar rejected authors.

Since top firms are almost always present at the ICLR conference, the longer-term increase in employment by an ICLR sponsor could be driven by sorting to top firms rather than a direct effect of networking at the conference. Indeed, in the full sample we do not see a significant increase in employment by ICLR sponsors outside top firms (Appendix Figure A.15). This contrast suggests

[19]The estimates in the unweighted sample are shown in Appendix Figure A.11. Among authors from nontop firms, accepted authors are 2-4 p.p. significantly more likely to have worked for one of the top firms at some point four or three years before the ICLR conference. The pre-trend disappears when we include job history as covariates in the moment equations (8) for IPT. The matched difference-in-differences in Figure A.12 also shows a 5 p.p. increase in employment by top firms among accepted authors initially employed in nontop firms.

the estimated impacts of the publication on moving to ICLR sponsors – firms with a strong demand for AI talent – are also driven by the top firms in the industry.

Figure A.16 shows the origin-specific estimates of the acceptance on employment by an ICLR sponsor. It is even more striking than the estimates in Figure 4 that a publication has the strongest impact on students and postdocs. The effects for incumbent workers in industry or academia are close to zero. Focusing on authors who are students/postdocs, we show in Figure 6 that a publication increases employment by non-top ICLR sponsors by about 1 p.p. within one year of the conference. Although the overall effect on employment by non-top ICLR sponsors is small (Appendix Figure A.15), there remains a short-term impact of networking at conference on the mobility of early-career authors. This finding provides evidence that recruiting young researchers is a key benefit of sponsoring CS conferences.

To investigate whether networking at the conference is the main mechanism driving the results for young researchers, we provide a placebo exercise using ICLR sponsors in the year before or after the event. Appendix Figure A.17 clearly shows zero increase in employment at sponsors in other years. Since top firms sponsor almost every ICLR conference, the placebo estimates rely on year-to-year variation from non-top sponsors, which receive a smaller inflow of publishing authors even if they show up in the same year, as shown in Figure 6.

4.4 Gender Differences

Both industry and academia have publicized their commitments to reducing the under-representation of women in STEM.[20] What does the industry–academia race imply for gender differences among CS 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.

Given the large impacts of an ICLR acceptance on the sorting of students and postdocs into top firms, we focus on the gender differences within this group of early career researchers. We estimate the event study regression (5) (with IPT weights) separately for 728 female authors versus 2,354 male authors who are students/postdocs in the year before their first ICLR conference. Male authors, comprising 76% of students/postdocs, are 1.8 p.p. more likely to work full-time for a top firm immediately after acceptance into ICLR, and are 9.4 p.p. more likely to work for top firms 3 years after (Figure 7a). These estimates are closer to the estimates for all students/postdocs shown in Figure 4. An accepted paper at ICLR also increases the probability that female students/postdocs get hired by top firms by 2.8 p.p. immediately after the conference. However, this effect is not persistent. We find dynamic effects noisily fluctuating between -3 p.p. and 2 p.p. in the next three

[20] Following the MeToo movement, tech companies invested in Diversity, Equity, and Inclusion (DEI). to recruit underrepresented minorities into their workforce. Our data cover years up to 2023, before most tech firms halted their DEI initiatives.

years for female authors. Taken together, the large impact of an ICLR publication on employment by top firms in the next three years applies primarily to male students and postdocs.

Are the academic returns to a publication different by gender? We use the position titles in individual job histories to identify tenure-track jobs, and we provide difference-in-differences estimates of the impact of ICLR acceptance on sorting to such jobs via regression (5) with IPT reweighting. Relative to rejected authors, female accepted authors are 4 p.p. significantly more likely to get a tenure-track job in the year of the ICLR conference and are 7 p.p. more likely to be on tenure-track 1-2 years after the conference (Figure 7b). These estimates suggest that a successful AI publication benefits the academic placement of female students and postdocs at least in the short term. In contrast, we do not find a positive impact of the acceptance by ICLR on the academic placement of male students/postdocs. Given that the earliest year in the estimation sample is 2012, this finding could reflect the commitment of academic employers to increase gender diversity, which also appears in other disciplines such as economics (e.g., Card et al. 2022).

These patterns hold in the original sample without reweighting (Figure A.18). Female accepted authors are 7 p.p. more likely to be on tenure-track than female rejected authors in three years post event, whereas male accepted authors are 3 p.p. more likely to work for top tech firms than their counterparts. These findings suggest that an ICLR publication facilitates transitions into top firms for men and into tenure-track academia for women. Industry roles tend to pay more on average than academic jobs. The divergence in sectoral returns to a publication could widen gender earnings gaps among computer scientists.

In summary, we exploit acceptance decisions at ICLR to estimate the causal impacts of publications on career outcomes. We find a significant and persistent increase in movements to top firms following the publication. The increase in movements to ICLR sponsors – firms with strong demand for research talent – is also driven by top firms. Heterogeneity analyses by origin of authors, and by gender, suggest that publications have the largest impacts on early-career researchers.

5 Consequences for the Direction and Quality of Public Research

Having established that a race for AI talent is underway and that top technology firms have gained the upper hand in recruiting top AI researchers from traditionally academic conferences, we now examine what happens to research output when talent reallocates. We study three dimensions: the quantity of public research, its quality as measured by citations, and the direction of knowledge production across computer science areas. We conduct descriptive event-study analyses of research output around moves to top firms versus tenure-track academic positions, using the full sample of computer science researchers with linked job histories. The sample includes all authors in the Scopus–Revelio dataset who made a first observed career transition either to a top technology firm

or to a tenure-track position. We compare these two groups in an event window of five years before and after the transition, using a two-way fixed effects regression of the form:

$$y_{it} = \beta_0 + \sum_{\ell \neq -1} \beta_\ell D_{it}^{(\ell)} \times Treated_i + \delta_t + \alpha_i + \varepsilon_{it} \tag{9}$$

where $Treated_i$ equals one for top-firm movers, $D_{it}^{(\ell)}$ are event-time dummies indicating years relative to the transition, and α_i and δ_t are author and year fixed effects. Standard errors are clustered at the author level. The control group of tenure-track movers reduces the concern of selection into job mobility relative to a stayer control group, since both groups reveal a willingness to change affiliations. We restrict to a balanced panel, requiring each author to be observed for the full five-year window on both sides of the event.

5.1 Quantity and Quality of Research

Panel A of Figure 8 shows raw trends in the number of publications for movers to top firms versus tenure-track around the transition event. Prior to the transition, top-firm movers publish at a substantially higher rate than their tenure-track counterparts, particularly in AI, consistent with positive selection of the most productive AI researchers into industry — mirroring our findings on the role of AI publications in sorting to top firms in Sectioins 3-4. Pre-transition trends in non-AI publications are broadly parallel across the two groups. Following the transition, the trajectories diverge sharply. Total publication output declines modestly among researchers entering tenure-track positions but resumes an upward trajectory in subsequent years. In contrast, research output drops sharply and persistently among those moving to top firms and does not return to its pre-move trend. The right panel of Figure 8 reports difference-in-differences estimates from specification (9): five years after the transition, annual publications decline by approximately 57%, AI publications by 46%, and non-AI publications by approximately 65% (see Table A.5).

Based on the difference-in-differences analysis, a back-of-the-envelope calculation suggests that 1,050 AI papers (and 1,978 papers in other CS areas) are missing per year due to researchers opting for top firms over academic positions. We interpret this estimate with caution: if researchers increase output before moving in anticipation of top firm recruitment, we overstate the true decline; if top-firm movers would have continued on a steeper trajectory absent the transition, we understate it.

Do reductions in research volume correspond to reductions in research quality as well? Panel B of Figure 8 extends the analysis to research quality along two dimensions: the extensive margin of forward citations — whether a researcher's papers receive any citation within three years of publication — and a high-impact indicator equal to one if three-year forward citations fall in the top

decile of the citation distribution. On the extensive margin, top-firm movers are approximately 10 percentage points less likely to produce any publicly cited work within five years of the transition, with the decline present across all publication types (Table A.6). The high-impact indicator reveals a striking divergence: following the transition, top-firm movers become relatively more likely to produce high-impact AI publications, while becoming less likely to produce high-impact non-AI research. While the move to top firms thus reduces the overall probability of producing any cited work, high-impact public output increases in commercially valuable AI areas at the expense of the broader computer science research.

5.2 Directions of Research

Which areas of computer science lose the most public research output as top firms increasingly recruit top AI researchers? We classify papers by conference area, grouping the 29 conference areas in our data into four categories: the AI frontier (artificial intelligence, computer vision, NLP, and web and information retrieval), infrastructure behind AI (computer architecture, operating systems, databases, networks, design automation, and measurement), security, privacy and society (cybersecurity, economics and computation, computational biology, HCI, and visualization), and software and physical computing (software engineering, programming languages, robotics, embedded systems, mobile computing, and computer graphics). Figure 9 plots raw publication trends by broad area. The positive pre-transition selection of top-firm movers is concentrated almost entirely in the AI frontier: their pre-transition publication rate in this category is almost double that of tenure-track movers, consistent with our findings on AI publications driving the sorting to top firms in Sections 3-4. For infrastructure research, both groups show growing pre-transition trends reflecting an expansion of research in this direction, though top-firm movers display a steeper pre-trend in databases and networks, areas where commercial applications are most immediate, as visible in Panel B of Figure A.19.[21] For society and software categories, pre-transition trends are broadly parallel, suggesting less differential selection into these areas.

Figure A.19 shows difference-in-differences estimates for conference subareas. The decline in AI frontier publications is present across all AI subareas, though with notable heterogeneity. The largest drop occurs in core AI and computer vision, while the decline is least pronounced for NLP. Five years after the transition, researchers who move to top firms publish approximately 0.04 fewer papers per year at core AI conferences relative to tenure-track movers. Two mechanisms may explain this heterogeneous pattern. The first is a *substitution from public to proprietary research*:

[21] Databases and networking infrastructure play a critical supporting role in modern AI development. Training large language models requires storing and retrieving vast amounts of data at speed; serving predictions to millions of users in real time requires routing queries efficiently across thousands of servers. Firms that solve these infrastructure problems more efficiently can build and deploy AI systems faster and at lower cost than competitors, creating strong incentives to recruit researchers in these areas and conduct their work internally rather than publishing it openly.

top firms may reduce public output across all areas, replacing it with internal research whose results are not disclosed. The second is a *redirection of total research effort*, both public and proprietary, toward areas where top firms hold a comparative advantage over academia and that align with firms' commercial interests. We speculate that the broad-based decline across all AI subareas primarily reflects the first mechanism, consistent with evidence in Akcigit et al. (2026), which documents a near-sixfold increase in patenting among researchers who transition to industry, suggesting systematic substitution from open science toward proprietary innovation. The relatively smaller decline in NLP may instead reflect the second mechanism: training large language models require massive compute infrastructure and proprietary datasets that academic labs cannot match, giving top firms a comparative advantage in this area.[22]

The other conference subareas experiencing the largest declines in public research following transitions to top firms are databases, human-computer interaction, security, and robotics. The decline in robotics, covering autonomous systems, surgical robots, and industrial automation, is particularly consequential: as these systems are deployed in hospitals, factories, and public spaces, the shift of robotics researchers into proprietary labs reduces the supply of independent work evaluating whether they behave safely and reliably. The declines in software engineering and programming languages raise analogous concerns: these are the researchers who develop tools for testing, verifying, and debugging complex software systems, and their migration to industry weakens the academic community's capacity to independently audit the AI code increasingly embedded in consequential decisions. Finally, the decline in security research, covering cryptography, vulnerability detection, and privacy-preserving systems, is directly relevant to public welfare: as AI systems handle ever more sensitive personal data, the shift of security researchers into proprietary labs managing them reduces independent scrutiny.

6 Conclusion

We document that AI researchers are increasingly concentrating in industry – especially at top-tier firms – relative to academia. To test whether academia and industry 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 of researchers. 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. The predictive relationship could be driven by both individual preferences on the supply side and the extent to which employers use publications to screen talent.

[22]This interpretation is consistent with the fact that the most prominent NLP breakthroughs of the past decade: the transformer architecture (Vaswani et al., 2017), BERT (Devlin et al., 2018), and GPT (Brown et al., 2020) were all published openly by researchers at Google, Google Brain, and OpenAI.

To identify the signal value of publications on the job market, we leverage variation in the acceptance of a paper at a top-tier AI conference. Using a difference-in-differences design, we compare accepted authors with similar rejected authors and find that a top AI publication increases the probability of moving to top firms by 2-6 percentage points in the next 1-3 years. The impact is the largest among students and postdocs at the beginnings of their careers. The casual evidence 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 relative to top firms in industry – we examine the implications for the production of public knowledge. We track the publication trajectories of computer scientists who move to top firms versus tenure-track academic positions. Before moving, researchers who will join top firms exhibit steeper increases in publication output, especially in AI-related work. After the move, publication activity diverges sharply: output among authors moving to tenure-track jobs stabilizes and resumes growth, while that of researchers joining top firms falls and remains well below pre-move levels. Difference-in-differences estimates suggest that there are 1,050 fewer AI papers (and 1,978 fewer papers in other CS areas) being produced per year due to researchers opting for top firms over academic positions.

While these results speak to the quantity of research being published, 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 different from academics in terms of topics and application areas. Our hypothesis is that 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. Our preliminary results show movers to top firms publish more impactful papers in AI. Immediate next steps include exploring whether the gap in citations is driven by breakthrough papers or self-citations within firms.

References

Abadie, Alberto. Semiparametric difference-in-differences estimators. *The Review of Economic Studies*, 72(1):1–19, 01 2005.

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.

Akcigit, Ufuk, Craig A Chikis, Emin Dinlersoz, and Nathan Goldschlag. Attention (and money) is all you need: Why universities are struggling to keep ai talent. Technical report, National Bureau of Economic Research, 2026.

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.

Berger, Emery D. CSrankings, 9 2017.

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

Cortes, Corinna and Neil D. Lawrence. Inconsistency in conference peer review: Revisiting the 2014 neurips experiment, 2021.

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

Färber, Michael and Lazaros Tampakis. Analyzing the impact of companies on ai research based on publications, 2023.

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

Graham, Bryan S., Cristine Campos De Xavier Pinto, and Daniel Egel. Inverse probability tilting for moment condition models with missing data. *The Review of Economic Studies*, 79(3):1053–1079, 04 2012.

Hainmueller, Jens. Entropy balancing for causal effects: A multivariate reweighting method to produce balanced samples in observational studies. *Political Analysis*, 20(1):25–46, 2012.

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

Hirano, Keisuke, Guido W. Imbens, and Geert Ridder. Efficient estimation of average treatment effects using the estimated propensity score. *Econometrica*, 71(4):1161–1189, 2003.

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

Imbens, Guido W. Nonparametric estimation of average treatment effects under exogeneity: A review. *The Review of Economics and Statistics*, 86(1):4–29, 02 2004.

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

Jones, Benjamin F, EJ Reedy, and Bruce A Weinberg. Age and scientific genius. *The Wiley handbook of genius*, pages 422–450, 2014.

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.

Liang, Lizhen, Han Zhuang, James Zou, and Daniel E. Acuna. The complementary contributions of academia and industry to ai research, 2024.

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.

Sant'Anna, Pedro H.C. and Jun Zhao. Doubly robust difference-in-differences estimators. *Journal of Econometrics*, 219(1):101–122, 2020.

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.

Słoczyński, Tymon, S. Derya Uysal, and Jeffrey M. Wooldridge. Covariate balancing and the equivalence of weighting and doubly robust estimators of average treatment effects. (18147), 2025.

Tan, Z. Regularized calibrated estimation of propensity scores with model misspecification and high-dimensional data. *Biometrika*, 107(1):137–158, 12 2019.

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.

Zhao, Qingyuan and Daniel Percival. Entropy balancing is doubly robust. *Journal of Causal Inference*, 5(1):20160010, 2017.

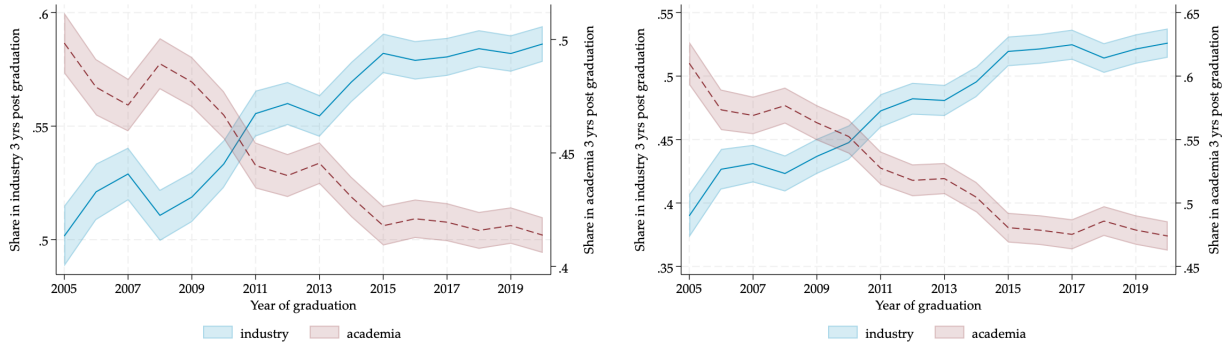
7 Main Figures and Tables

Figures

FIGURE 1. Share of early-career CS researchers in industry versus academia over time

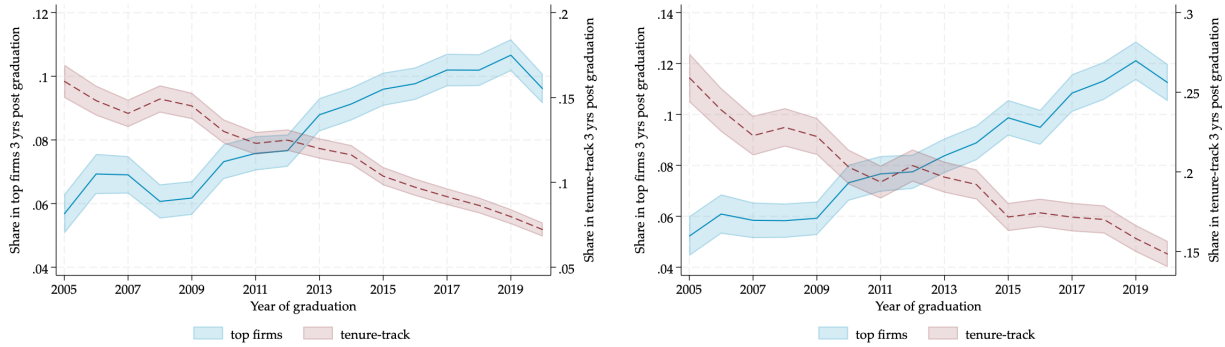
(a) All: Industry vs. Academia

(b) PhDs: Industry vs. Academia



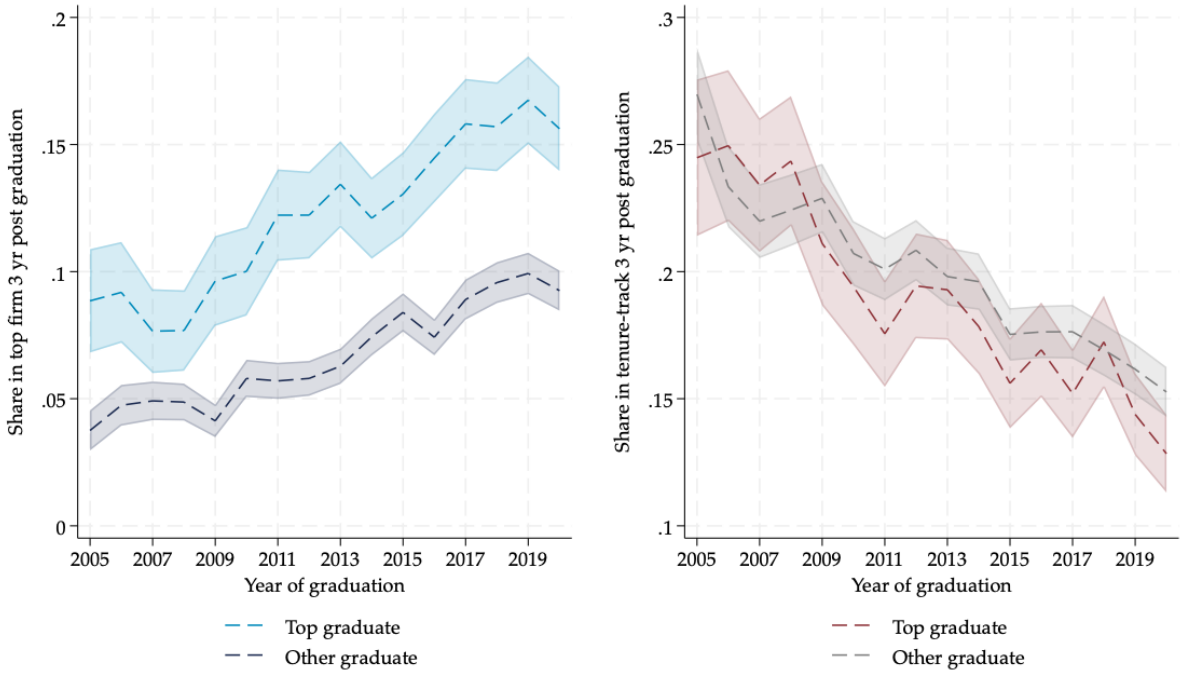
(c) All: Top Firms vs. Tenure-track

(d) PhDs: Top Firms vs. Tenure-track



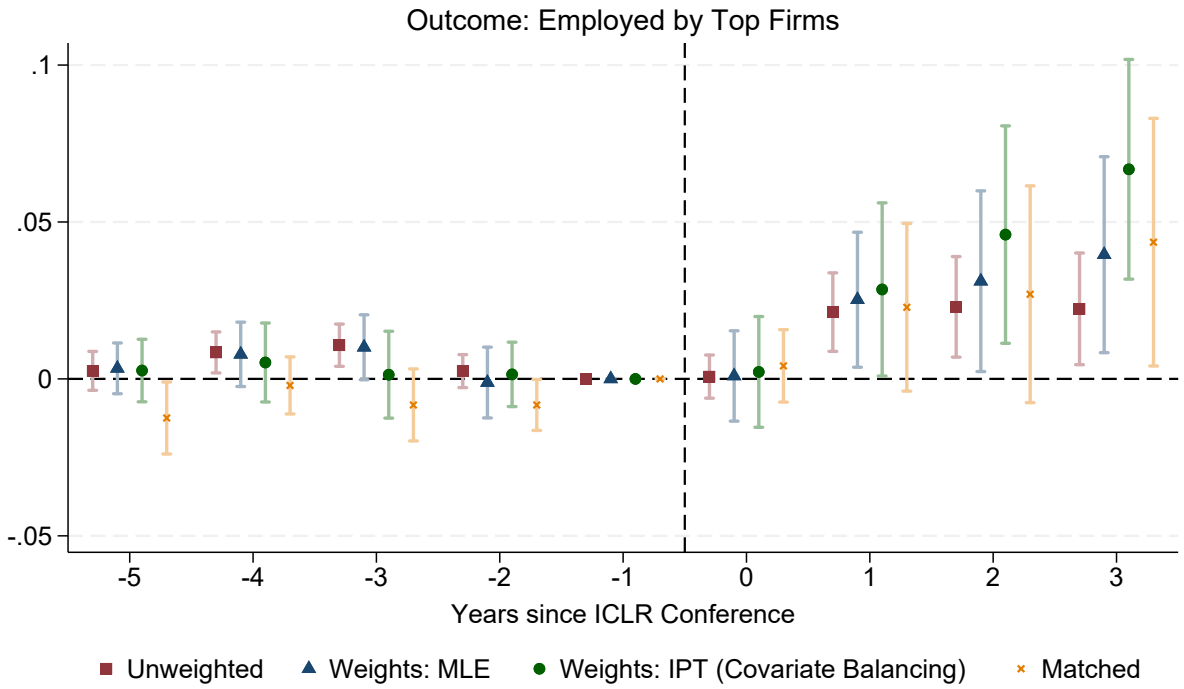
Notes: This figure shows the share of computer science researchers working in industry versus academia three years after completing their highest degree, conditional on reporting any job that year on LinkedIn. Panels (a) and (b) show the share employed in industry and academia, respectively, while panels (c) and (d) show the share employed at top technology firms and in tenure-track academic positions. Panels (a) and (c) include all graduates (bachelor's, master's, or PhD), whereas panels (b) and (d) restrict the sample to PhD holders, for whom tenure-track positions are more relevant.

FIGURE 2. Share of early-career researchers in top firms versus tenure-track: heterogeneity by PhD institution



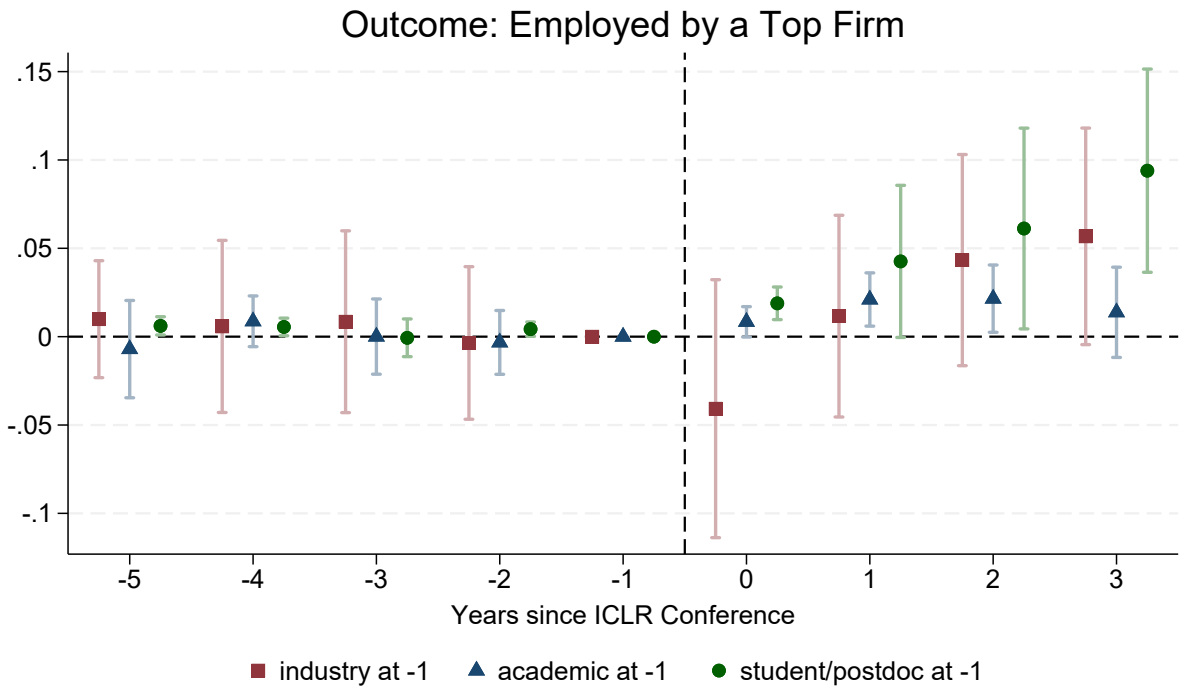
Notes: This figure shows the share of CS researchers working in top firms versus being on tenure track, 3 years after PhD. Top graduates refer to PhD graduates from top 50 CS departments, based on CSRankings, QS 2025, and THE 2025 in AI/CS. Other graduates obtain their PhDs outside the top 50.

FIGURE 3. Employment by Top Firms among Authors From Other Places



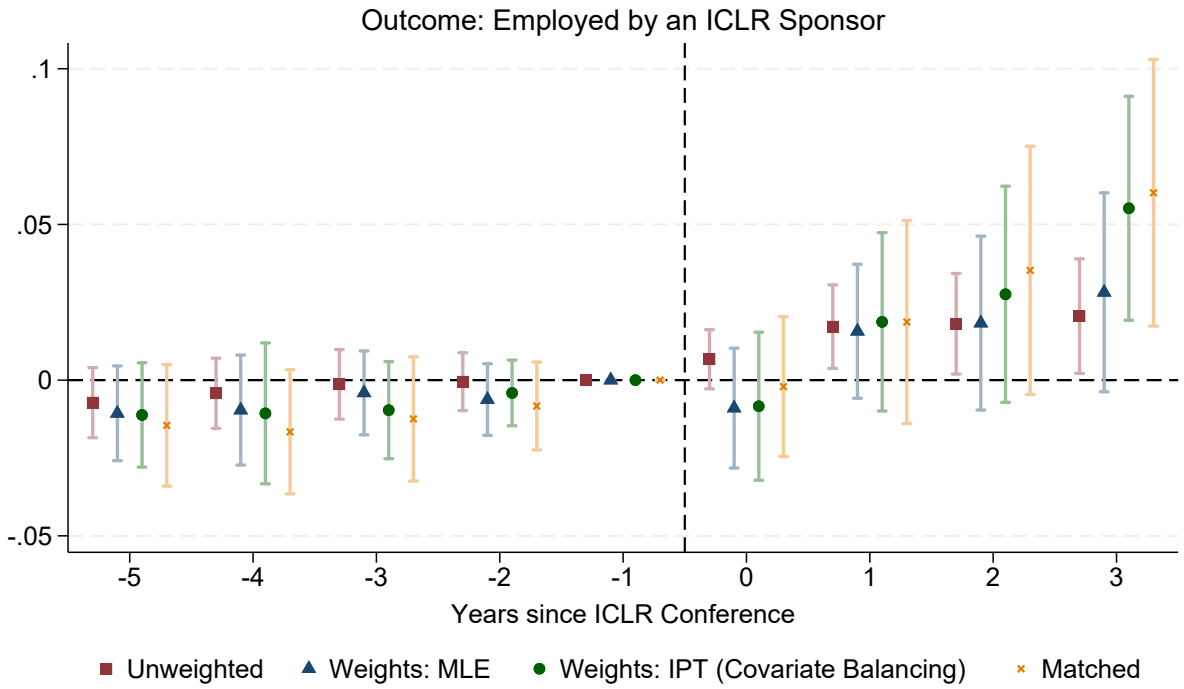
Notes: This figure shows the difference-in-differences estimates of the impacts of an ICLR publication on moving to top firms from employers outside the top firms, corresponding to $\{\gamma_l\}$ in model (5). The estimation sample comprises ICLR-Revelio matched authors who submitted to an ICLR conference for the first time between 2017 and 2020, and who were not employed by top firms at -1, the year before their first ICLR conference (event). Standard errors are robust and clustered at the person level. The outcome variable indicates whether the worker is employed by a top firm in the current year. We show the estimates across methods. "Unweighted" refers to the estimates from the full sample of accepted and rejected authors. "Weights: MLE" or "IPT" refers to the estimates after we weight each rejected author by $\frac{\hat{p}_i}{1-\hat{p}_i}$ (accpeted authors are weighted by 1). The propensity score $\hat{p}_i$ is estimated by maximum likelihood under "Weights: MLE", and by method of moments under "Weights: IPT" to achieve covariate balancing. "Matched" refers to estimates among 1-1 matches between accepted and rejected authors. The methods are discussed in detail in Section 4.1.2.

FIGURE 4. Employment by Top Firms: Separate Estimates by Origin Outside the Top



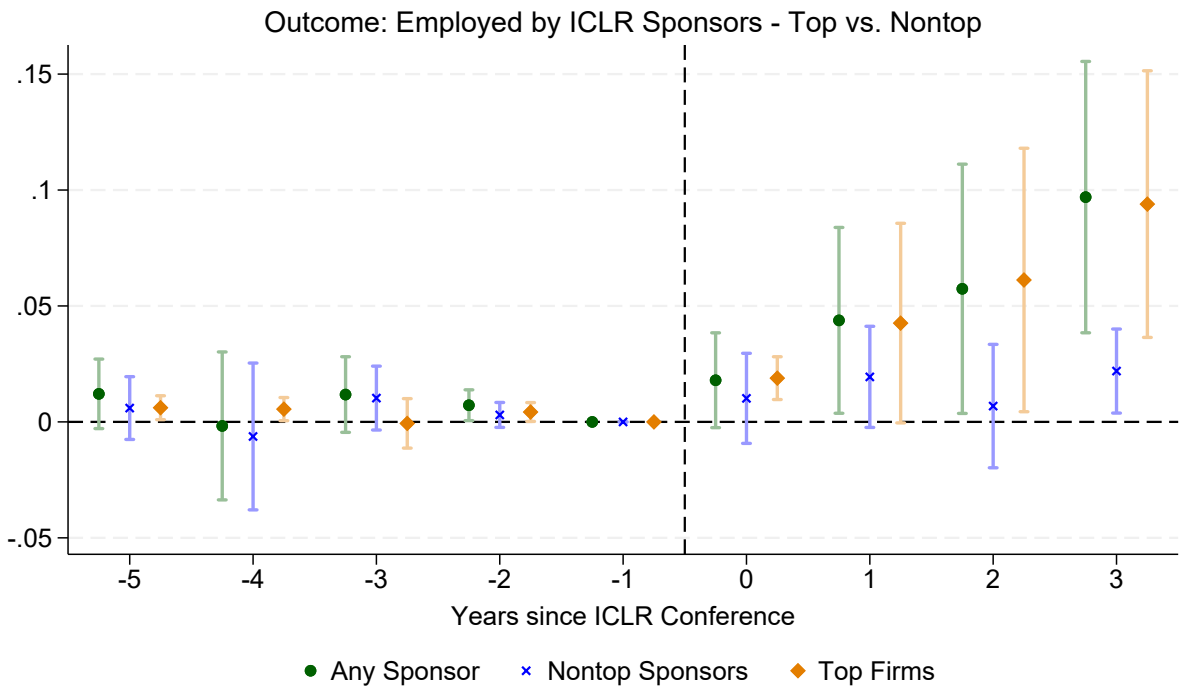
Notes: This figure shows the difference-in-differences estimates of the impacts of an ICLR publication on moving to top firms from different types of employers outside the top firms. Model (5) is estimated on authors who are outside the top firms at -1, separately by whether the author was employed by nontop firms in industry at -1 (red square), employed by academia (blue triangle), or were students or postdocs (green circles). Rejected authors are reweighted by $\frac{\hat{p}_i}{1-\hat{p}_i}$, where the propensity score $\hat{p}_i$ is estimated by GMM (8) to achieve covariate balancing. Standard errors are robust and clustered at the person level.

FIGURE 5. ICLR Acceptance on Employment by Corporate Sponsosrs



Notes: This figure shows the difference-in-differences estimates of the impacts of an ICLR publication on employment by ICLR sponsors, corresponding to $\{\gamma_l\}$ in model (5). The estimation sample comprises ICLR-Revelio matched authors who submitted to an ICLR conference for the first time between 2017 and 2020, and who were not employed by top firms at -1, the year before their first ICLR conference (event). Standard errors are robust and clustered at the person level. The outcome variable indicates whether the worker is employed by a company that sponsors the ICLR conference in the event year. "Unweighted" refers to the estimates from the full sample of accepted and rejected authors. "Weights: MLE" or "IPT" refers to the estimates after we weight each rejected author by $\frac{\hat{p}_i}{1-\hat{p}_i}$ (accpeted authors are weighted by 1). The propensity score $\hat{p}_i$ is estimated by maximum likelihood under "Weights: MLE", and by method of moments under "Weights: IPT" to achieve covariate balancing. "Matched" refers to estimates among 1-1 matches between accepted and rejected authors. The methods are discussed in detail in Section 4.1.2.

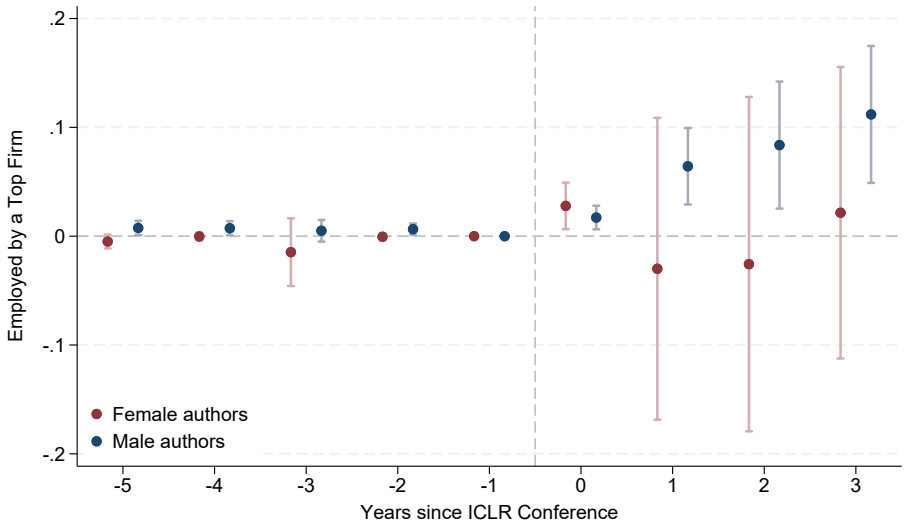
FIGURE 6. Employed by Conference Sponsors - Top versus Nontop Firms



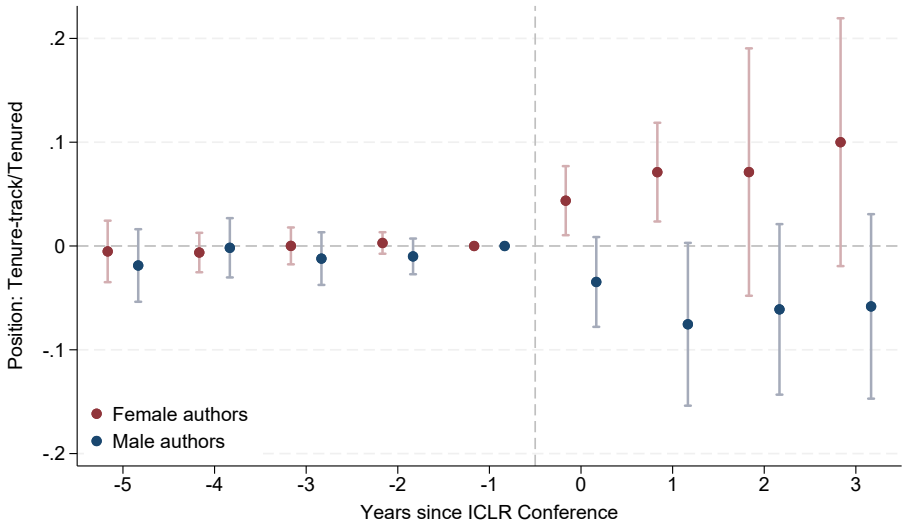
Notes: This figure shows the difference-in-differences estimates of the impacts of an ICLR publication on employment by ICLR sponsors, by sponsors that are not top firms (nontop sponsors), and by top firms for comparison. Corporate sponsors refer to firms that pay for advertising and recruitment at ICLR in the event year (0 = each author's first ICLR conference). The estimation sample comprises authors who are students or postdocs at -1. Rejected authors are reweighted via inverse propensity titlting. Standard errors are robust and clustered at the person level. The estimated impacts of ICLR acceptance on employment by sponsors (green circles) are very close to the estimates for moving to top firms (orange diamonds), which are repeated from Figure 4 for comparison. Top firms almost sponsor every year. The effects on moving to sponsors that are not top firms (blue cross) are smaller in comparison.

FIGURE 7. Gender Differences in Employment Outcomes Among Students and Postdocs

(a) Top Firms



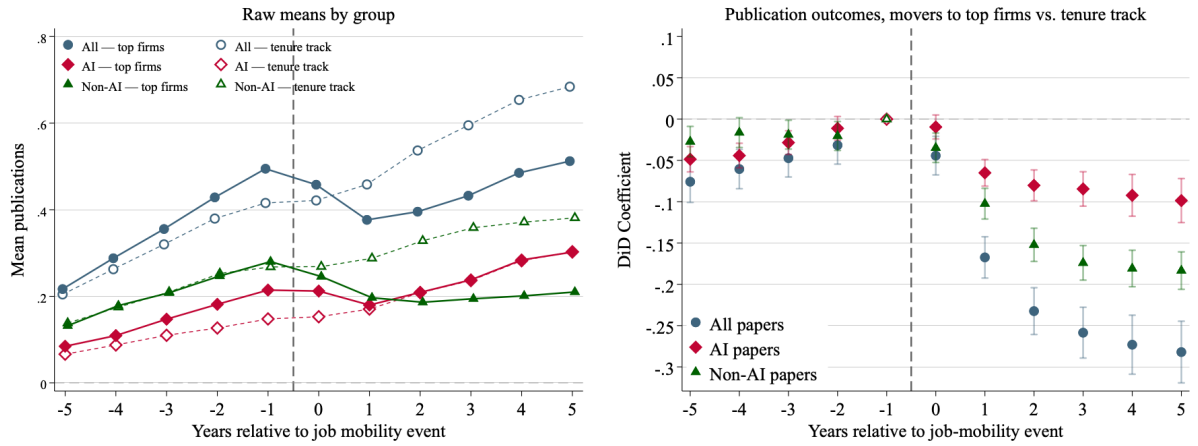
(b) Tenure-track Jobs



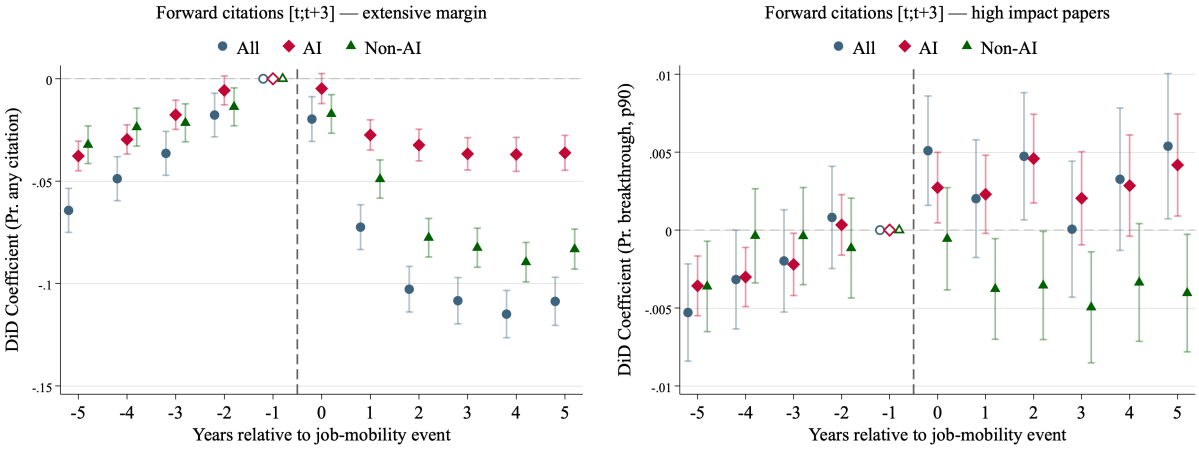
Notes: This figure shows the estimated impacts of an ICLR publication on employment at (a) top firms, and (b) tenure-track/tenured positions, separately by gender of the authors. We estimate the dynamic difference-in-differences model (5) separately for female and male authors who are students/postdocs at -1 (the year prior to the event). Rejected authors are reweighted by $\frac{\hat{p}_i}{1-\hat{p}_i}$, where the propensity score $\hat{p}_i$ is estimated by GMM (8) to achieve covariate balancing. Standard errors are robust and clustered at the person level.

FIGURE 8. Quantity and Quality of Research - Movers to Top Firms vs. Tenure Track

Panel A: Publication productivity

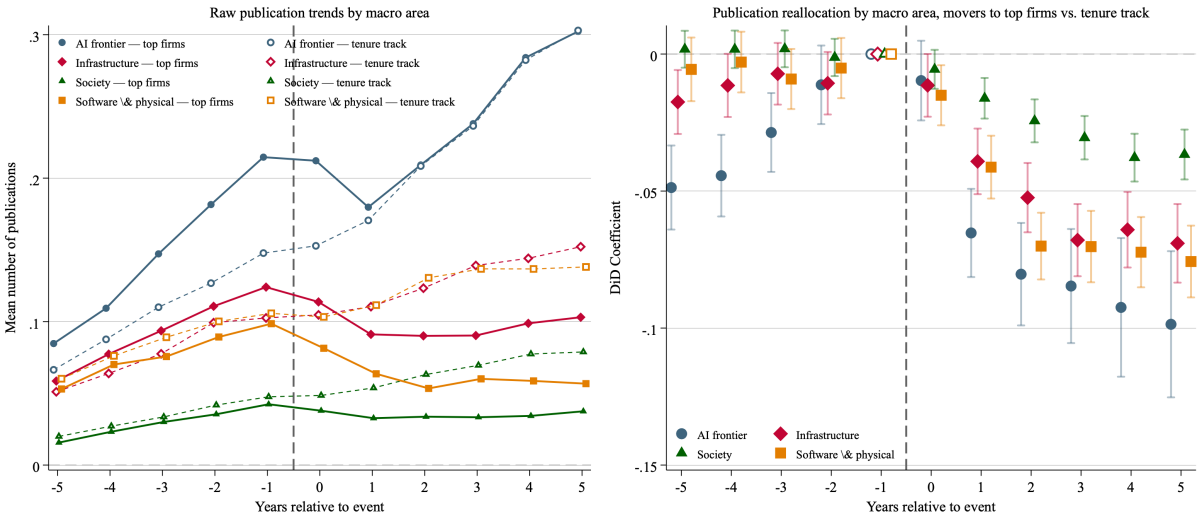


Panel B: Forward citations [t, t+3]



Notes: Both panels compare researchers whose first observed move is to a top technology firm against those whose first observed move is to a tenure-track academic position. All difference-in-differences estimates are from two-way fixed effects regressions including author and year fixed effects, with standard errors clustered at the author level. The sample is restricted to authors observed in a balanced window of five years before and after the career transition. *Panel A* — the left subfigure plots raw mean publication counts by year relative to the event, for top-firm vs tenure-track movers. The right subfigure reports DiD estimates for four publication outcomes: total publications, publications at AI conferences, top AI papers (a subset restricted to top-tier AI venues), and non-AI papers (publications at other computer science conferences). *Panel B* — the left subfigure reports DiD estimates for the extensive margin of forward citations, defined as an indicator equal to one if a researcher's papers published in a given year receive at least one citation within three years. The right subfigure reports estimates for high-impact research, defined as an indicator equal to one if forward citations fall in the top decile of the citation distribution, computed conditional on receiving at least one citation.

FIGURE 9. Number of publications by broad area, movers to top firms vs. tenure track



Notes: The figure plots mean annual publication counts by year relative to the career transition event, separately for researchers whose first observed move is to a top-six technology firm (solid lines, filled markers) and those whose first observed move is to a tenure-track academic position (dashed lines, hollow markers). Each color corresponds to one of four broad areas. *AI frontier* (blue circles) aggregates publications in artificial intelligence (AAAI, IJCAI), computer vision (CVPR, ICCV, ECCV), natural language processing (ACL, EMNLP, NAACL), and web and information retrieval (SIGIR, WWW). *Infrastructure* (red diamonds) aggregates computer architecture (ISCA, MICRO, ASPLOS), operating systems (OSDI, SOSP), high-performance computing (SC, HPDC), databases (SIGMOD, VLDB), computer networks (SIGCOMM, NSDI), design automation (DAC, ICCAD), and measurement and performance analysis (SIGMETRICS, IMC). *Security, privacy and society* (green triangles) aggregates computer security (CCS, IEEE S&P, USENIX Security), economics and computation (EC, WINE), computational biology (ISMB, RECOMB), human-computer interaction (CHI, UIST), and visualization (VIS, VR). *Software and physical computing* (orange squares) aggregates software engineering (ICSE, FSE), programming languages (PLDI, POPL), robotics (ICRA, IROS), embedded and real-time systems (EMSOFT, RTSS), mobile computing (MobiCom, MobiSys), and computer graphics (SIGGRAPH). The sample is restricted to authors observed in a balanced window of five years before and after the transition event. The dashed vertical line marks the transition year ($t = 0$).

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
AI papers	.001638	.005313	.007901	-.006523	-.0005383	.001164	-.0001045
	(.0005854)	(.0006793)	(.0007935)	(.0007865)	(.0005795)	(.0004056)	(.0003302)
Other Papers	-.001432	-.000968	-.002891	.0003503	.002499	.0009579	.002169
	(.0004486)	(.0003914)	(.000406)	(.000434)	(.0004597)	(.0003235)	(.0003082)
N	148,080	148,080	148,080	148,080	148,080	148,080	148,080
Adj. R2	.0491	.0535	.02351	.04566	.0493	.02853	.05392
Mean of y	.5435	.07686	.07703	.2703	.3435	.1123	.07862
AI Papers|y=1	.5316	1.128	1.159	.3668	.4793	.6172	.5264
Other Papers	.9464	1.456	1.166	.8694	.9164	1.076	1.127

Notes: This table shows the OLS estimates of regression (1) of the initial job placement on publications before graduation, for 148,080 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 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. Job Mobility After Graduation (Year-to-Year Transitions between Employers)

	At Nontop Firms			At Top Firms			In Academia		
	(1) Industry	(2) Top	(3) Academia	(4) Industry	(5) Top	(6) Academia	(7) Industry	(8) Top	(9) Academia
Any AI Paper	-.02011	.01534	.005264	-.003424	.000722	.003424	.003937	.007566	-.008393
	(.001341)	(.0009414)	(.0008524)	(.001373)	(.002473)	(.0009959)	(.001213)	(.0005738)	(.001684)
Any Other Paper	-.007237	.002425	.002431	-.005195	-.004703	.005306	.006731	.001513	-.01021
	(.001025)	(.0006295)	(.0006874)	(.00141)	(.002597)	(.001093)	(.001014)	(.0003839)	(.001384)
N	651,567	651,567	651,567	108,822	108,822	108,822	469,132	469,132	469,132
Adj. R2	.04756	.01143	.02758	.01844	.01361	.0335	.05582	.009738	.3093
Mean	.9567	.01478	.01731	.9856	.9366	.005587	.05902	.006422	.8005

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). The estimation sample is at person×year level, including for each person the years between graduation from her highest degree and 2023. There are three outcomes that represent the employer in the next year (denoted by $j(i,t+1)$): (1) whether next year the person is employed in industry, (2) whether she is employed by top firms, and (3) whether she is employed by academia as opposed to industry. We estimated separate regressions for those who are currently employed by nontop firms (denoted by $j(i,t) \in$ Nontop), by top firms, or by academia. Standard errors are robust.

TABLE 3. Balance Table: Accepted vs. Rejected Authors in the Original Sample and in the Reweighted Sample (IPT)

	Original Sample			Reweighted Sample (IPT)	
Variable	Accepted	Rejected	Diff.	Rejected	Diff.
Mode of Referee Rating	6.306	3.997	2.309***	6.306	0.000
	(1.207)	(1.423)	(0.020)	(0.786)	(0.085)
Min Rating	5.463	3.273	2.190***	5.463	0.000
	(1.242)	(1.237)	(0.020)	(0.935)	(0.053)
Mean Rating	6.539	4.433	2.105***	6.159	0.380***
	(0.726)	(1.089)	(0.014)	(0.602)	(0.065)
Num. Authors	5.158	4.686	0.472***	5.158	-0.000
	(2.409)	(2.077)	(0.037)	(2.378)	(0.131)
Age	30.571	30.924	-0.353***	30.571	-0.000
	(5.426)	(5.603)	(0.088)	(5.197)	(0.367)
Female	0.244	0.255	-0.011*	0.244	-0.000
	(0.430)	(0.436)	(0.007)	(0.430)	(0.023)
Foreign	0.490	0.550	-0.060***	0.490	0.000
	(0.500)	(0.498)	(0.008)	(0.500)	(0.035)
White	0.375	0.349	0.026***	0.375	0.000
	(0.484)	(0.477)	(0.008)	(0.484)	(0.032)
Asian	0.580	0.602	-0.022***	0.580	0.000
	(0.494)	(0.490)	(0.008)	(0.494)	(0.033)
Minority (Not White/Asian)	0.045	0.050	-0.004	0.045	-0.000
	(0.208)	(0.217)	(0.003)	(0.208)	(0.009)
Any PhD	0.482	0.458	0.024***	0.482	0.000
	(0.500)	(0.498)	(0.008)	(0.500)	(0.035)
Any Master	0.630	0.620	0.010	0.630	0.000
	(0.483)	(0.485)	(0.008)	(0.483)	(0.029)
Employed by Industry at -1	0.217	0.233	-0.016**	0.217	-0.000
	(0.412)	(0.423)	(0.007)	(0.412)	(0.025)
Employed by Top at -1	0.000	0.000	0.000	0.000	0.000
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
Employed by Academia at -1	0.184	0.205	-0.021***	0.184	-0.000
	(0.388)	(0.404)	(0.006)	(0.388)	(0.025)
Student/Postdoc at -1	0.490	0.446	0.044***	0.490	0.000
	(0.500)	(0.497)	(0.008)	(0.500)	(0.035)
Observations	5,826	11,963	17,789	11,963	17,789

Notes: This table displays average paper and author characteristics for rejected versus accepted authors in both the full sample (Columns 1-2) and the reweighted sample (Columns 4-5). In the original or unweighted sample, there are 5,826 accepted authors and 11,963 rejected authors who are not employed by top firms in the year prior to submitting to an ICLR conference for the first time. In the reweighted sample, accepted authors are weighted by 1, while rejected authors are weighted by $\omega_i^{IPT} = \frac{\hat{p}_i}{1-\hat{p}_i}$. The propensity score $\hat{p}_i$ of acceptance based on observable characteristics are estimated by solving the moment equations (8). As shown in column 5 - the weighted average among rejected authors equal the average among accepted authors, except for the variable: mean rating, which is not included in the estimation of propensity scores. Standard errors for the differences between accepted and rejected authors are shown in parentheses.

TABLE 4. 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	0.0455	0.0396	0.0116	0.0646
	(0.0049)	(0.0112)	(0.0048)	(0.0074)
Post event=1 × Accepted	0.0244	0.0079	0.0126	0.0307
	(0.0071)	(0.0146)	(0.0075)	(0.0107)
Constant	0.0165	0.0225	0.0032	0.0195
	(0.0015)	(0.0033)	(0.0017)	(0.0023)
Observations	50994.0000	11439.0000	11817.0000	27738.0000
Adj. R^2	0.2175	0.2192	0.1941	0.2353

Notes: This table shows the estimates of difference-in-differences regression (4) in the original sample without reweighting. We focus on events ≤ 2020 so the panel at person-year level is balanced during the event window [-5,3] years relative to the first ICLR conference of an author. The sample is limited to authors who are not employed by a top firm in the year prior to the event. Column (1) pools everyone in this sample, whereas columns (2)-(4) estimate the regression (4) separately by the type of origin — employer of an author in the year before the ICLR conference. Standard errors are robust and clustered at the author level.

TABLE 5. Diff-in-diff estimates for Employment by Top Firms (Reweighted)

	Employed by a Top Firm			
	(1) Pooled	(2) From Industry	(3) From Academia	(4) From Student/Postdoc
Post event=1	0.0353	-0.0246	0.0119	0.0597
	(0.0146)	(0.0248)	(0.0050)	(0.0252)
Post event=1 × Accepted	0.0439	0.0417	0.0170	0.0573
	(0.0145)	(0.0293)	(0.0095)	(0.0231)
Constant	0.0147	0.0430	0.0043	0.0101
	(0.0036)	(0.0074)	(0.0022)	(0.0059)
Observations	50994.0000	11439.0000	11817.0000	27738.0000
Adj. R^2	0.2137	0.2367	0.1781	0.2283

Notes: This table shows the estimates of difference-in-differences regression (4) in the reweighted sample where accepted authors are weighted by 1, while rejected authors are weighted by $\omega_i^{IPT} = \frac{\hat{p}_i}{1-\hat{p}_i}$. The propensity score $\hat{p}_i$ of acceptance based on observable characteristics are estimated by solving the moment equations (8). We focus on events ≤ 2020 so the panel at person-year level is balanced during the event window [-5,3] years relative to the first ICLR conference of an author. Column (1) include every author who is not employed by a top firm at -1, whereas columns (2)-(4) estimate the regression (4) separately by the type of origin — employer of an author in the year before the ICLR conference. Standard errors are robust and clustered at the author level.

A Appendix

A.1 A1. 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 blocking on the propensity score (Rosenbaum and Rubin 1983).

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.[23] 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.[24] 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 the 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 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 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. If we allow a rejected author to be matched with multiple accepted authors, we could match 4,424 accepted authors to 942 rejected authors. If we draw without replacement (consistent

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

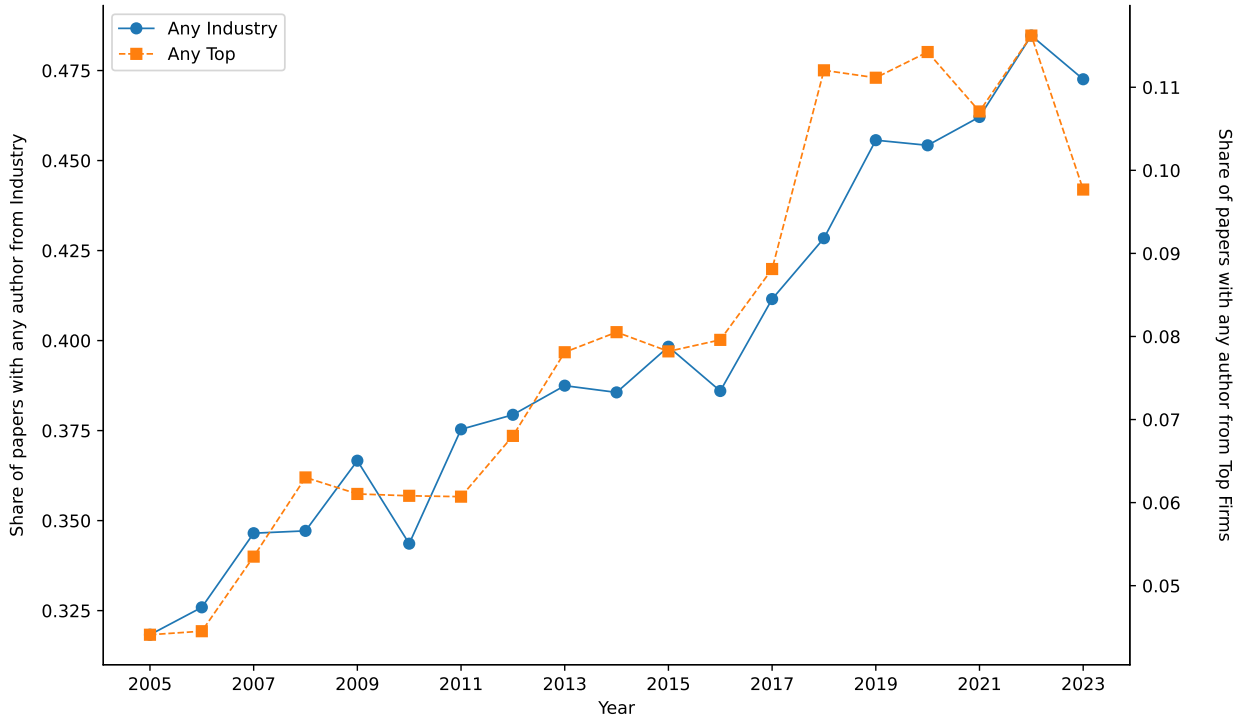
[24] FTo measure the quality of papers, we use both the numeric ratings and the content of the referee comments. We converted the text of each referee report to a numeric 1536-by-1 vector (embeddings), using large language models developed by OpenAI. We then estimated a Lasso-logit model of whether the 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.

with Jäger and Heining 2022 and Sarsons 2022), we could match 1,263 accepted authors to 1,263 rejected authors. Despite the fact we have to drop 70% of the accepted authors to establish a 1-1 match without replacement, the main event study result is very close to what we have seen in the full sample or the reweighted sample (see Figure 3).

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 (comparing Table A.3 with Table A.4). 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.

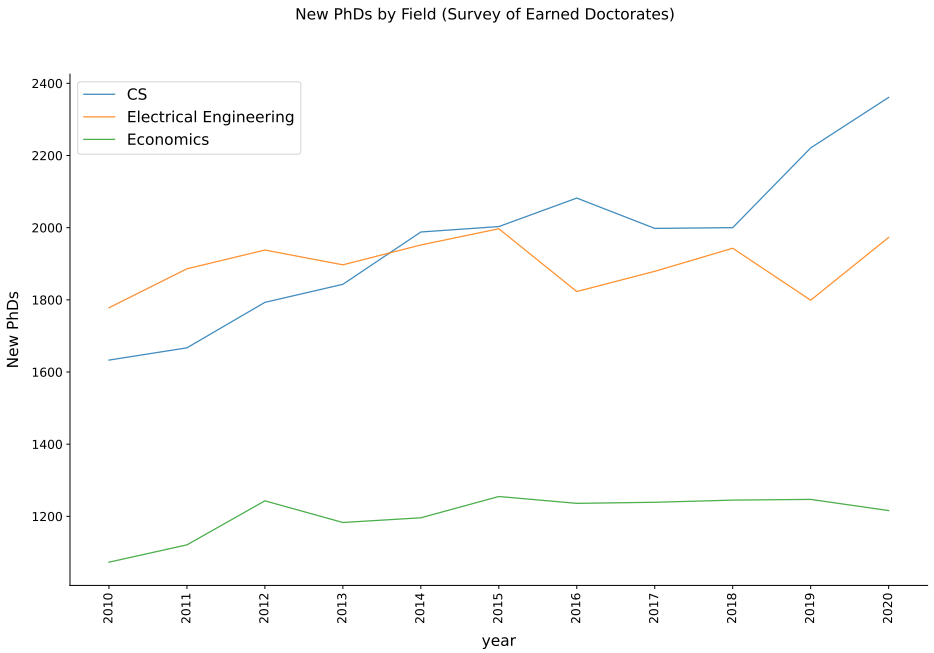
A2. Appendix Figures

FIGURE A.1. Trend in Industry Involvement in CS Research



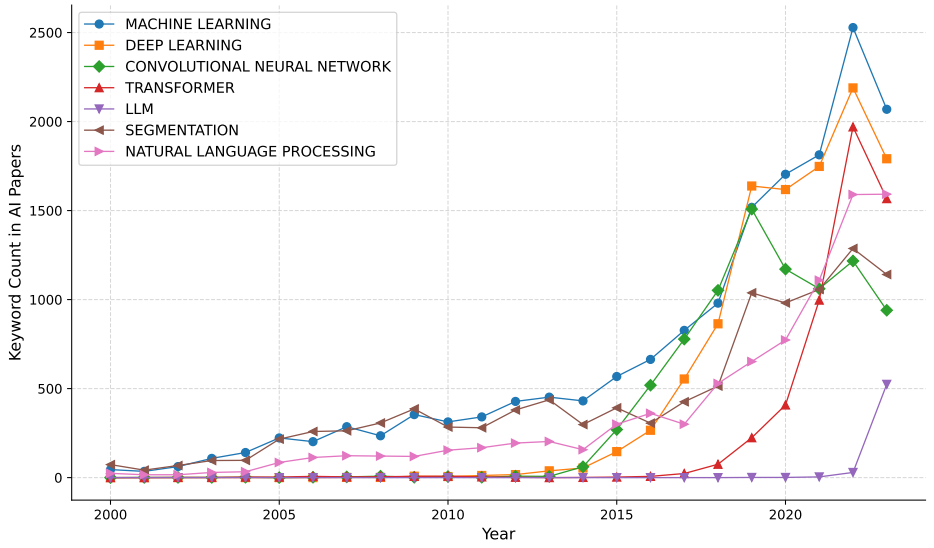
Note: This figure shows the trend in the share of conference proceedings with at least one author from industry (left y-axis), and the share with at least one author from top tech firms (right y-axis). We use records of CS conference proceedings on Scopus (see Section 2).

FIGURE A.2. Trend in CS/EE/Econ PhD Graduates in the U.S.



Note: This figure comes from Wu (2025). The PhD cohort size is measured based on the Survey of Earned Doctorates by the National Science Foundation.

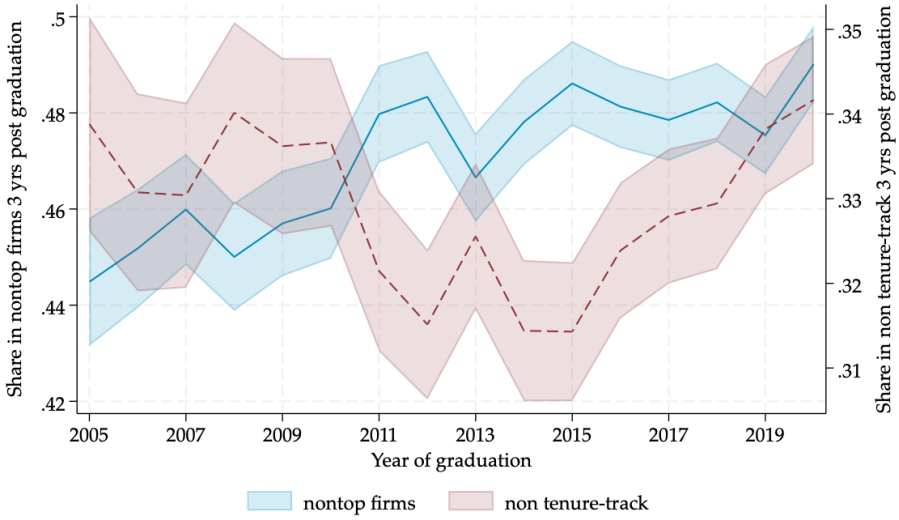
FIGURE A.3. Trend in AI Research by Keyword



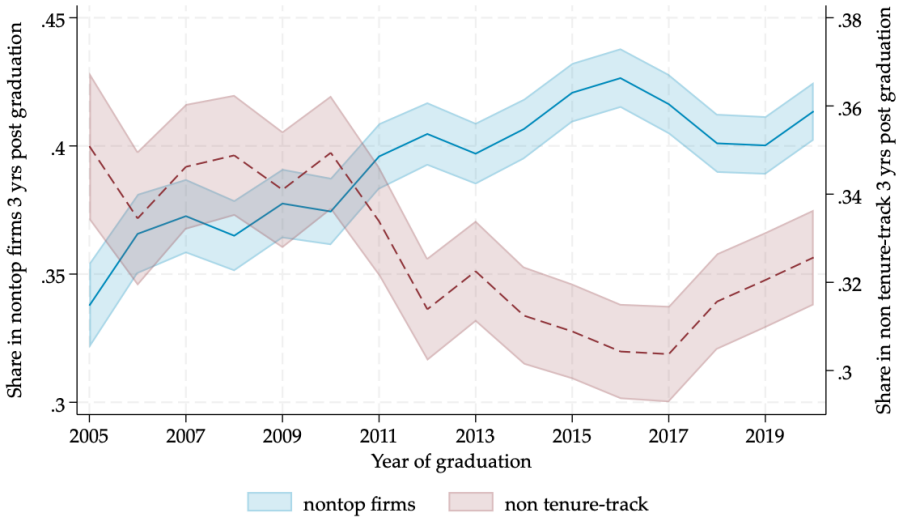
Note: This figure shows the count of top 5 AI keywords plus "transfomers" and "large language models" ("LLM") that show up in the titles or abstracts of publications at 24 AI conferences by year of publication. "LLM" did not show up until 2022. Most AI publications in our sample period concern machine learning, neural networks, and applications in computer vision or natural language processing.

FIGURE A.4. Share of early career CS researchers in non-top firms vs non-tenure-track academic positions over time

(a) All Graduates



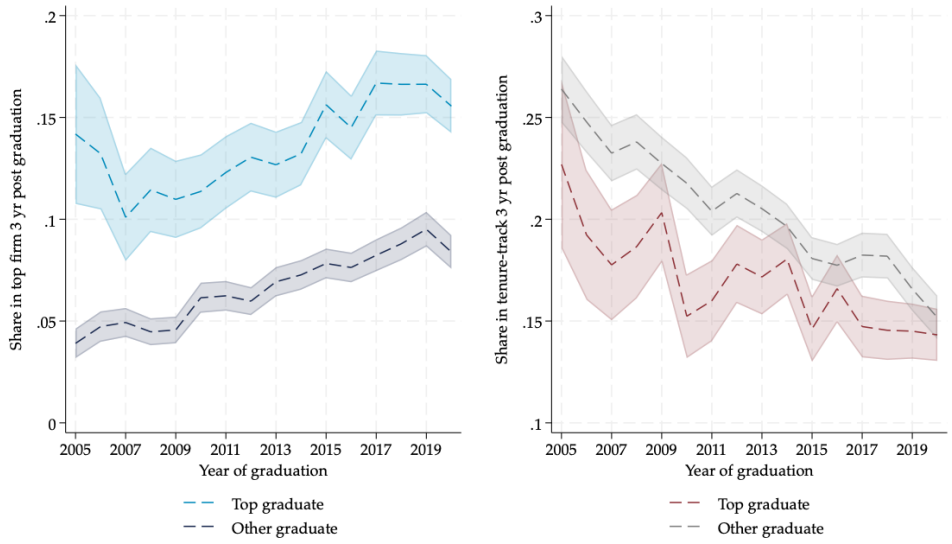
(b) PhD Graduates



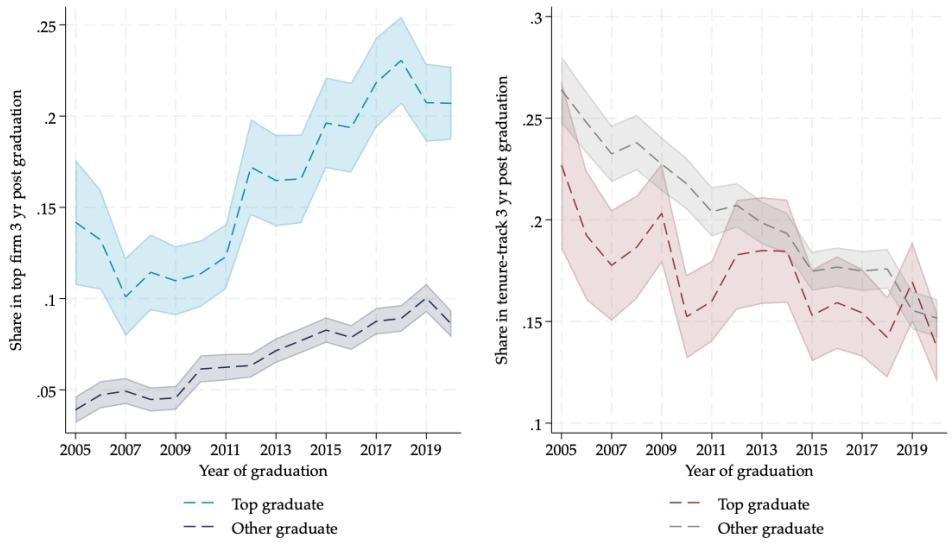
Note: This figure replicates Figure 1 (c) and (d), but instead depicting the share of graduates in non-top firms vs. non-tenure-track academia positions three years after graduation.

FIGURE A.5. Share of early career CS researchers in top firms vs tenure-track positions over time: heterogeneity by alternative measures of research potential

(a) Top graduate defined as any pre-PhD AI publication



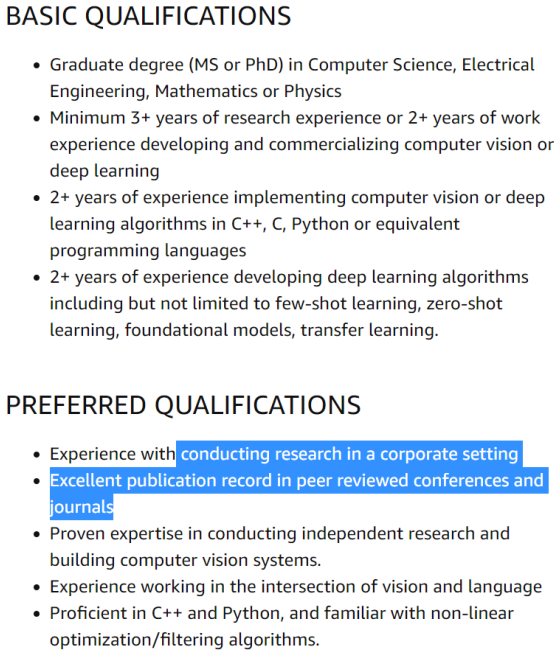
(b) Top graduate defined as top quartile in pre-PhD AI publications



Notes: This figure replicates Figure 2 to show the share of PhD graduates entering top firms (left) vs. academia (right), separately by whether they have any publication before graduation (panel (a)) or by whether they are in the top quartile of number of AI publications in their graduation cohort (panel (b)). Pre-PhD publications refer to publications that occur leading up to, or in the year of, graduation.

FIGURE A.6. 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:

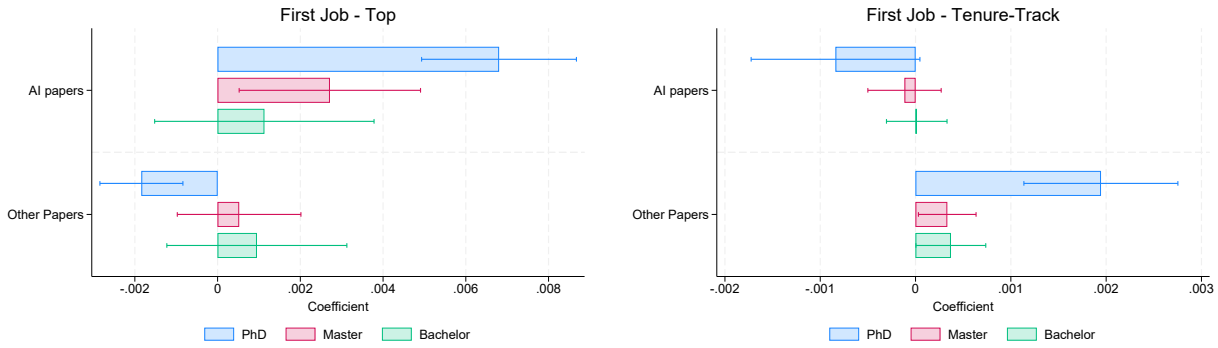
- Relevant work experience, including full time industry experience or as a re
- Strong publication record
- Ability to design and execute on research agenda.

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

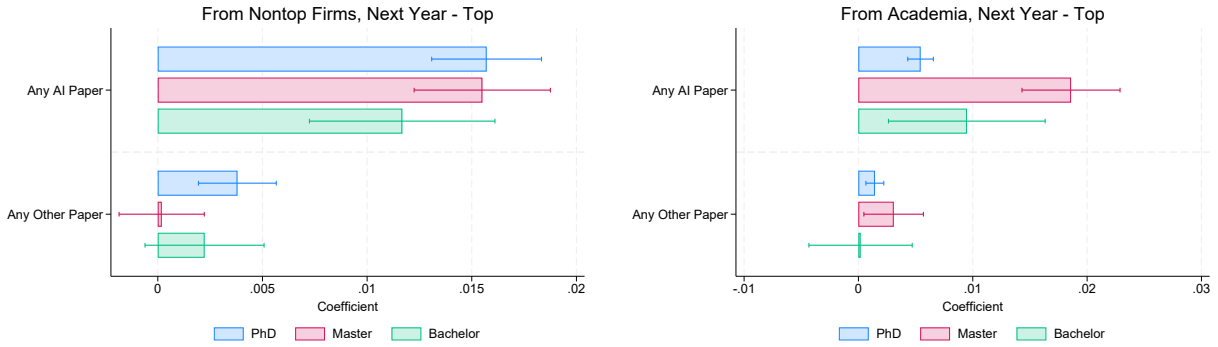
(a) First Job - Top

(b) First Job - Tenure-track



(c) From Nontop Firms to Top

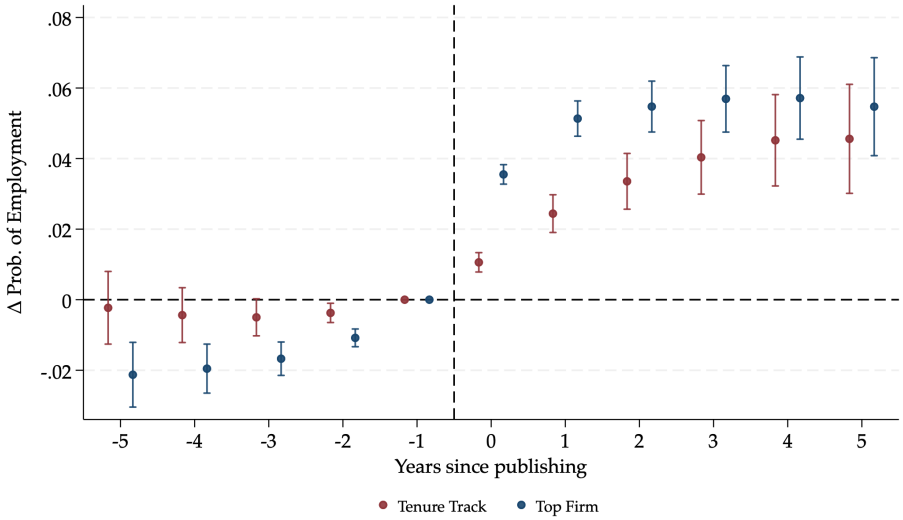
(d) From Academia to Top



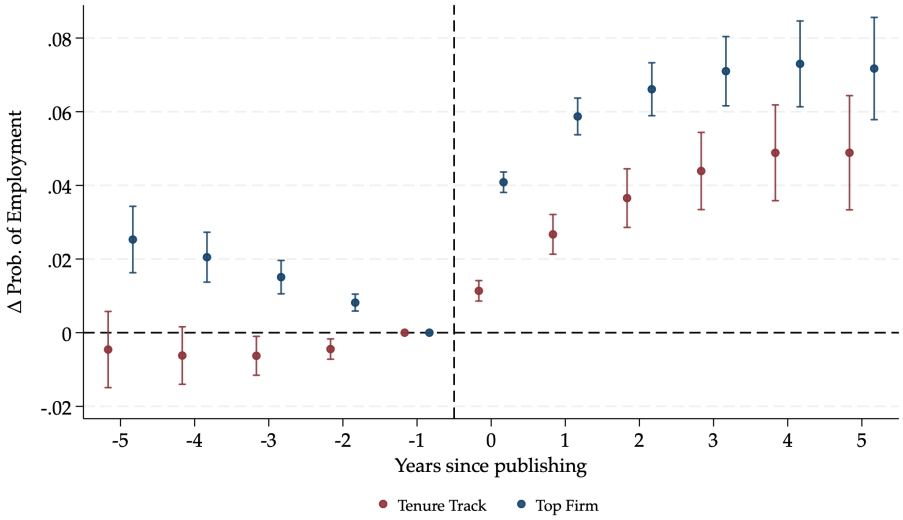
Notes: The top panel of this figure show the OLS estimates of regression (1) of a person's initial job placement on publications before graduation, separately by the level of highest degree. There are two initial placement outcomes: (a) first job at top firms, and (b) first job on the tenure track. Table 1 shows the estimates for these outcomes in the pooled sample (college or master or PhD) in columns 2 and 7, respectively. The bottom panel shows the estimates of regression (2) of (post-graduation) job mobility on indicators for publications at person$\times$year level, separately by the level of highest degree. (c) corresponds to the outcome of moving from other firms in the industry ("nontop") to top firms, and (d) shows the estimates for moving from academia to top firms. Table 2 shows the estimates for these outcomes in the pooled sample (college or master or PhD) in columns 2 and 8, respectively. Standard errors are robust.

FIGURE A.8. Event study around a person's first AI publication

(a) All AI Authors

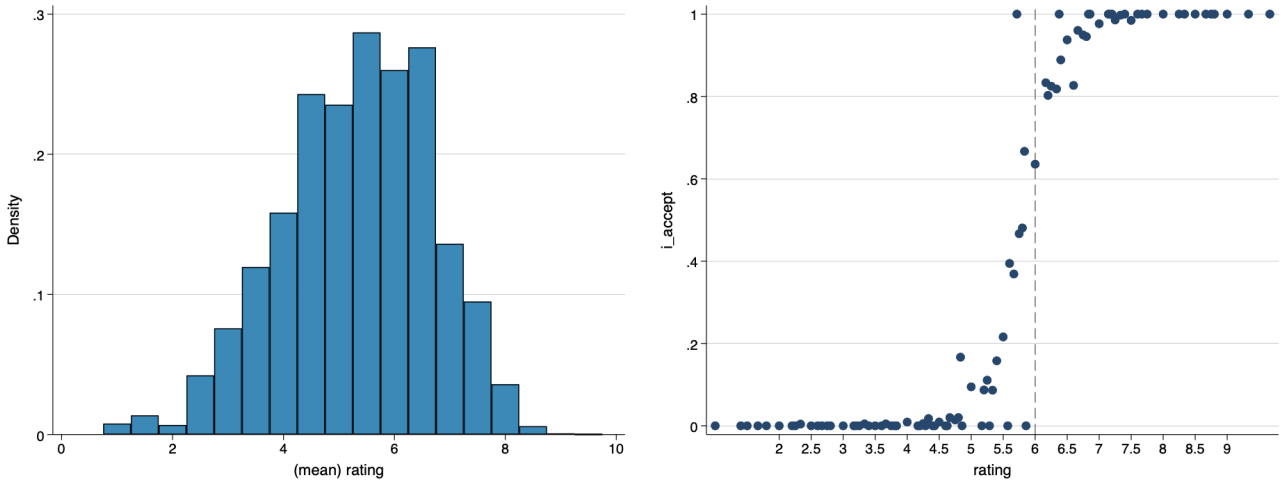


(b) Authors outside top firms at -1



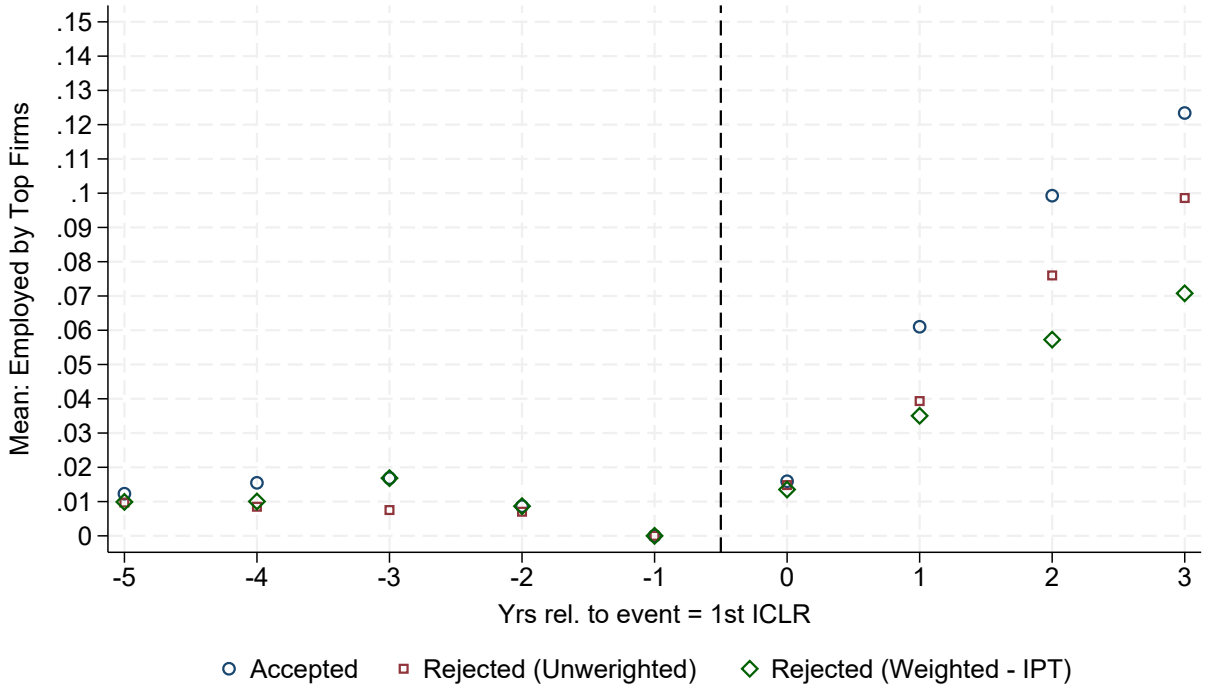
Notes: This figure presents estimates of the event study regression (3) for two outcomes around the event defined as a person's first AI publication post graduation from her highest degree. The red markers represent results where the outcome is employment in a tenure-track position, while the blue markers correspond to employment in a top firm. Panel (a) reports results for the full sample, and Panel (b) restricts to AI authors not employed at a top firm in the year prior to the event. All regressions include person and year fixed effects. Standard errors are robust and clustered at the person level. Data are from the Scopus-Revelio person-year panel.

FIGURE A.9. Referee Ratings and Acceptance by ICLR



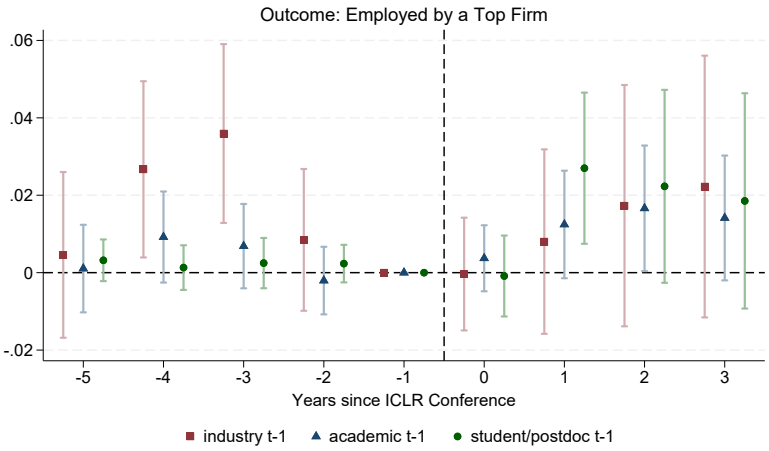
Notes: This figure shows the distribution of the average referee ratings received by ICLR submissions, and how the average rating relates to the likelihood of acceptance. Each paper is evaluated by 3–4 referees, who provide an overall assessment on a numeric scale. The scale was changed in 2020, but a score of 6 has consistently represented the "Marginally Accept". The left panel shows the distribution of average ratings per paper, while the right panel plots the relationship between the average rating and the probability of acceptance. Papers with an average rating below 5 are unlikely to be accepted, whereas most papers with an average above 7 are accepted.

FIGURE A.10. Parallel Trends in Employment by Top Firms



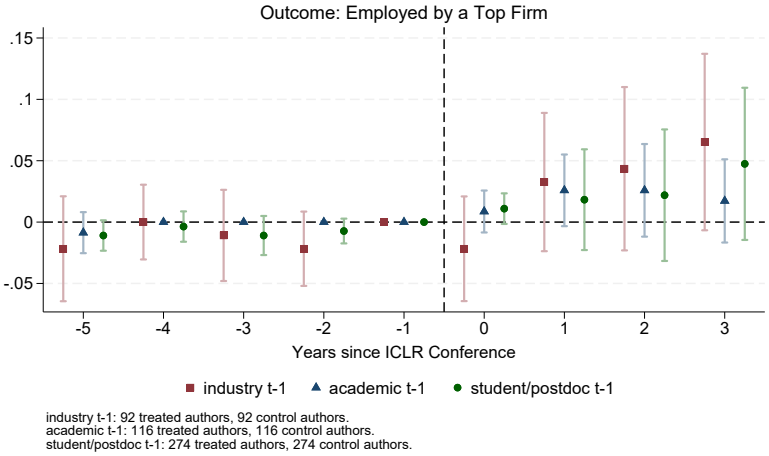
Notes: This figure shows the share of authors employed by top firms by the number of years relative to each person's ICLR event (first ICLR conference). The sample is restricted to authors who are not working at top firms one year before the ICLR conference. We show the trends in three groups: (1) accepted authors, (2) rejected authors (unweighted), and (3) rejected authors reweighted by $\hat{p}_i/(1-\hat{p}_i)$. The propensity score in (3) is estimated by solving the moment equations (8) to achieve covariate balancing. Accepted authors are weighted by 1 in the IPT-reweighted sample. See Section 4.1.2 for details.

FIGURE A.11. Employment by Top Firms among Authors from Other Places



Notes: This figure shows the estimated impacts of an ICLR publication on moving to top firms from different types of employers outside the top firms, which are represented by the coefficients $\{\gamma_l\}$ in the dynamic difference-in-differences model (5). The model is estimated on authors who are outside the top firms at -1, separately by whether the author was employed by nontop firms in industry at -1 (red), employed by academia (blue), or were students or postdocs (green). In contrast with Figure 4, the estimation sample is unweighted. Standard errors are robust and clustered at the author level.

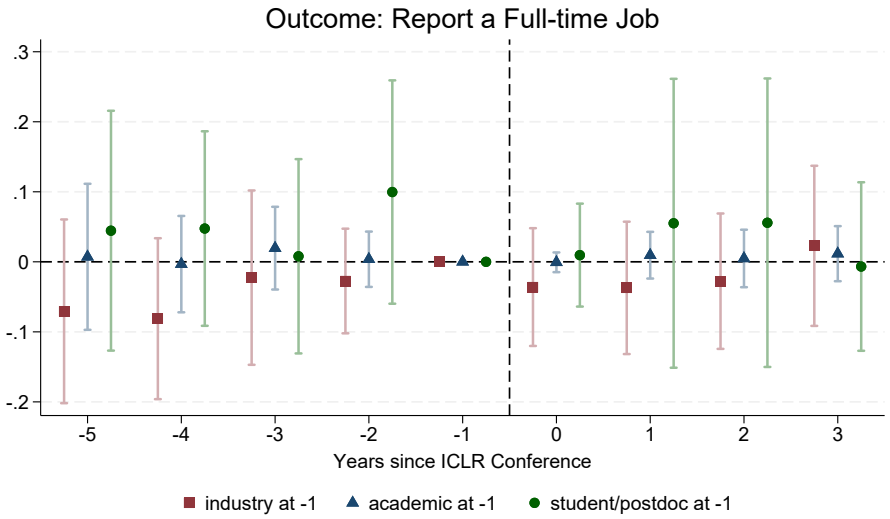
FIGURE A.12. Mobility outcome: Top | $j(i,-1) \notin$ Top (Matched Sample)



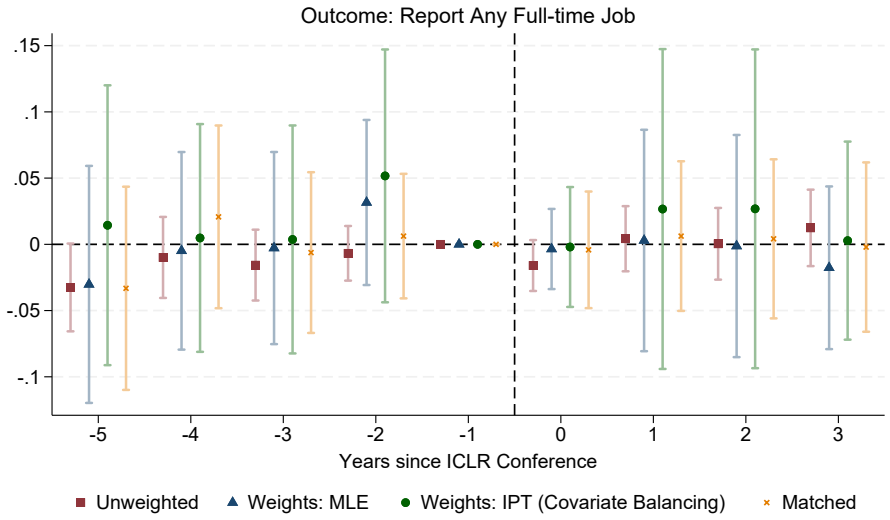
Notes: This figure shows the estimated impacts of an ICLR publication on moving to top firms from different types of employers outside the top firms, which are represented by the coefficients $\{\gamma_l\}$ in the dynamic difference-in-differences model (5). The model is estimated on authors in the matched sample who are outside the top firms at -1, separately by whether the author was employed by nontop firms in industry at -1 (red), employed by academia (blue), or were students or postdocs (green). The control/rejected authors in the matched sample are drawn without replacement. Standard errors are robust and clustered at the author level.

FIGURE A.13. Reporting of Any Full-time Job (with covariate balancing)

(a) Pooled

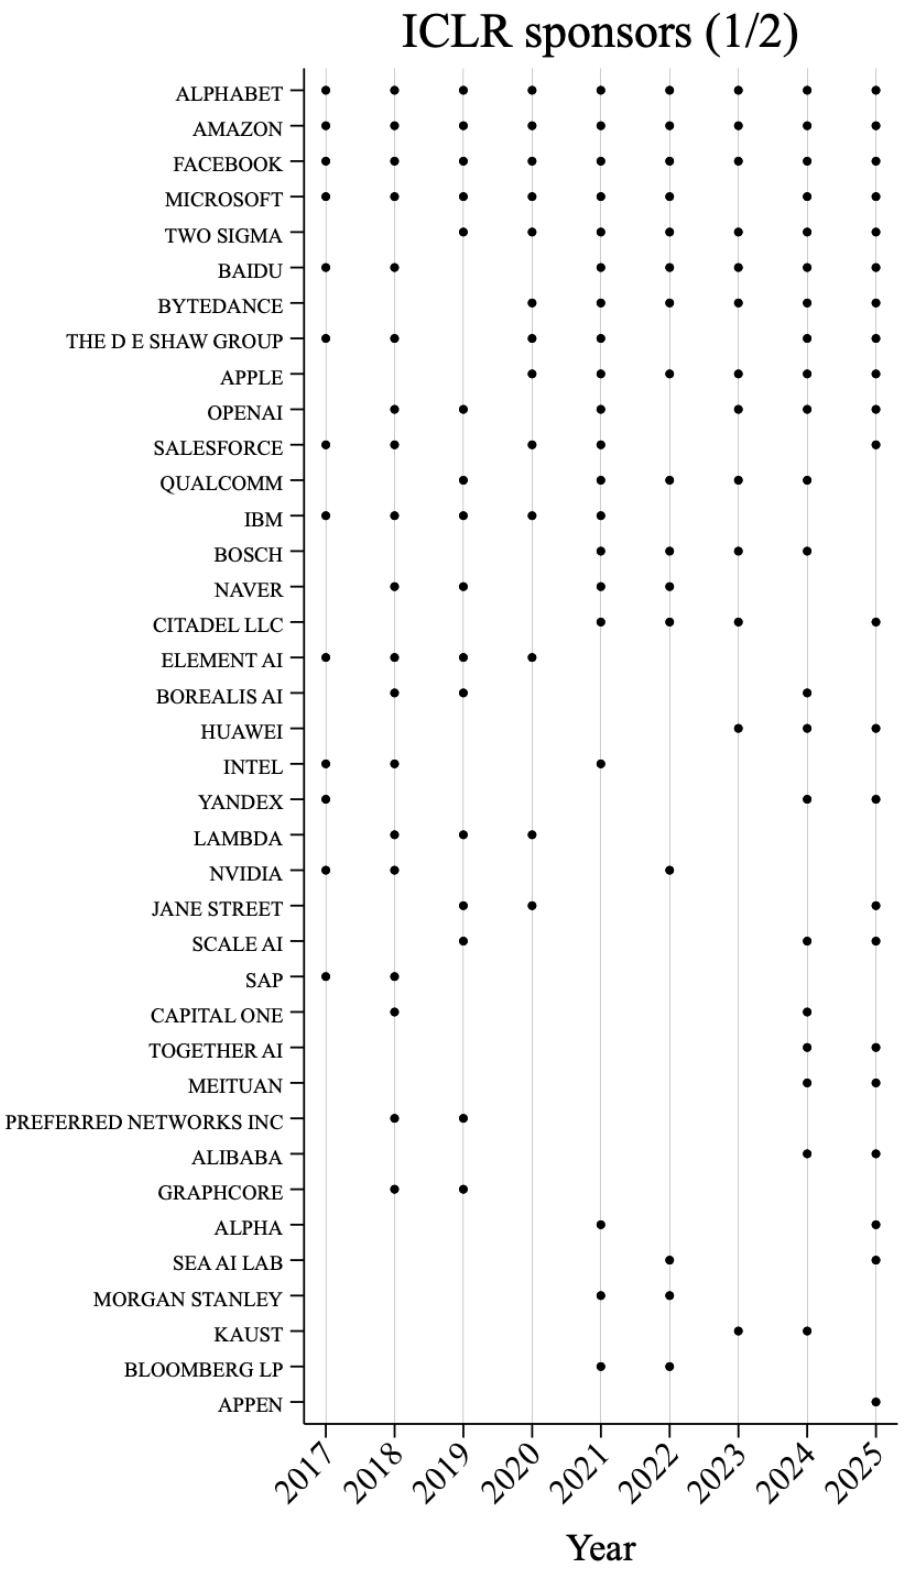


(b) By Origin at -1



Notes: This figure shows the estimated impacts of an ICLR publication on reporting any full-time job on LinkedIn (person×year level), which are represented by the coefficients $\{\gamma_l\}$ in the dynamic difference-in-differences model (5). The estimation sample comprises authors who are not employed by top firms at -1, consistent with the sample restriction in Figures 3-4. Rejected authors are reweighted by $\frac{\hat{p}_i}{1-\hat{p}_i}$, where the propensity score $\hat{p}_i$ is estimated by GMM (8) to achieve covariate balancing. Panel (a) shows the separate estimates based on the type of employers at -1: nontop firms in industry at -1 (red), employed by academia (blue), or were students or postdocs (green). Panel (b) shows the estimates by reweighting or matching method (see notes under Figure 3). Standard errors are robust and clustered at the author level.

FIGURE A.14. ICLR Sponsors by Conference Year



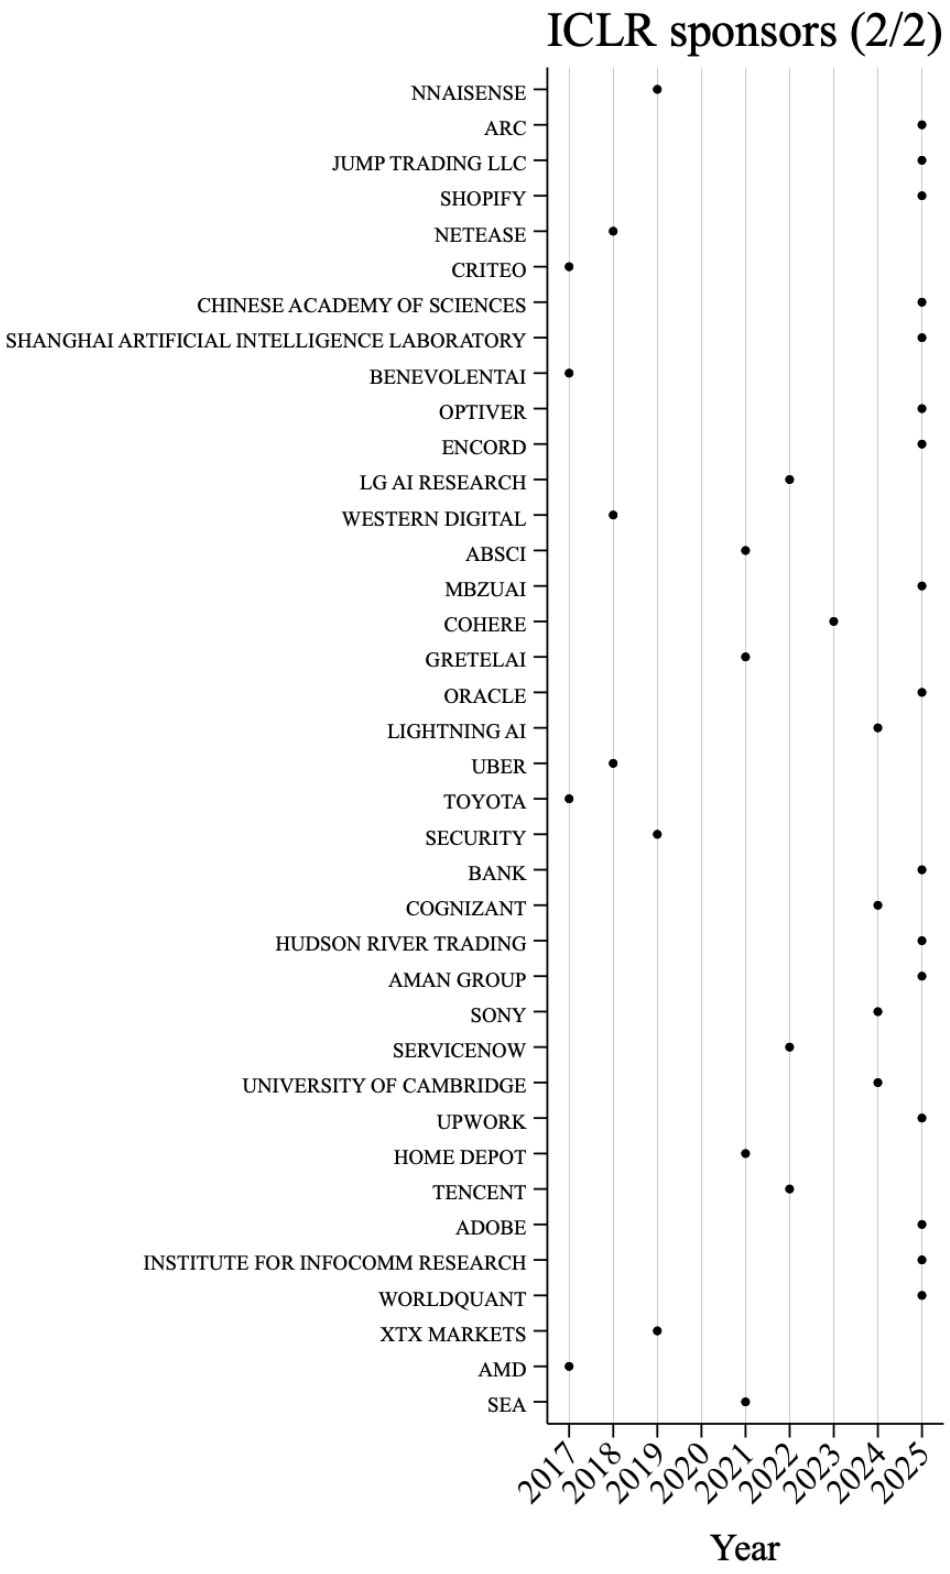

ICLR sponsors (2/2)
NNAISENSE
ARC
JUMP TRADING LLC
SHOPIFY
NETEASE
CRITEO
CHINESE ACADEMY OF SCIENCES
SHANGHAI ARTIFICIAL INTELLIGENCE LABORATORY
BENEVOLENTAI
OPTIVER
ENCORD
LG AI RESEARCH
WESTERN DIGITAL
ABSCI
MBZUAI
COHERE
GRETELAI
ORACLE
LIGHTNING AI
UBER
TOYOTA
SECURITY
BANK
COGNIZANT
HUDSON RIVER TRADING
AMAN GROUP
SONY
SERVICENOW
UNIVERSITY OF CAMBRIDGE
UPWORK
HOME DEPOT
TENCENT
ADOBE
INSTITUTE FOR INFOCOMM RESEARCH
WORLDQUANT
XTX MARKETS
AMD
SEA
2017
2018
2019
2020
2021
2022
2023
2024
2025
Year

FIGURE A.15. Employed by ICLR Sponsors that are not Top Firms

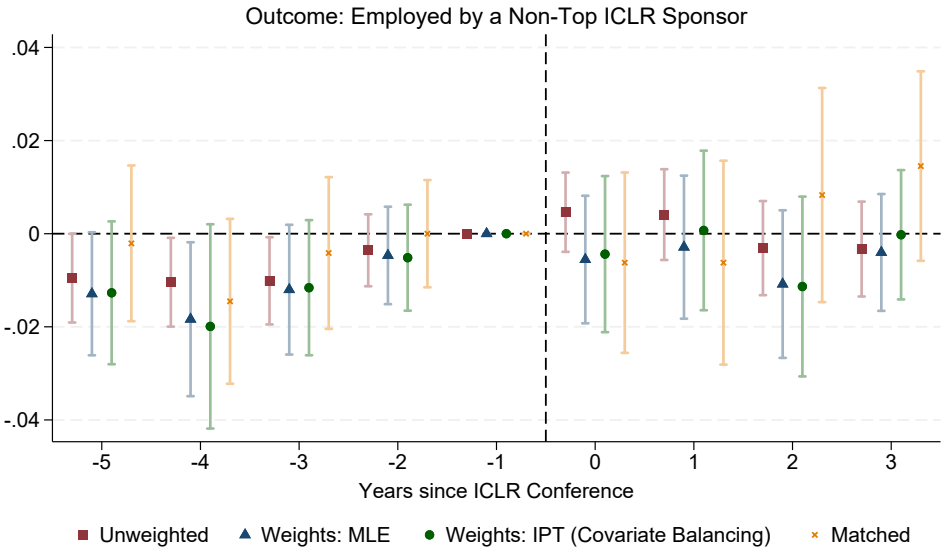
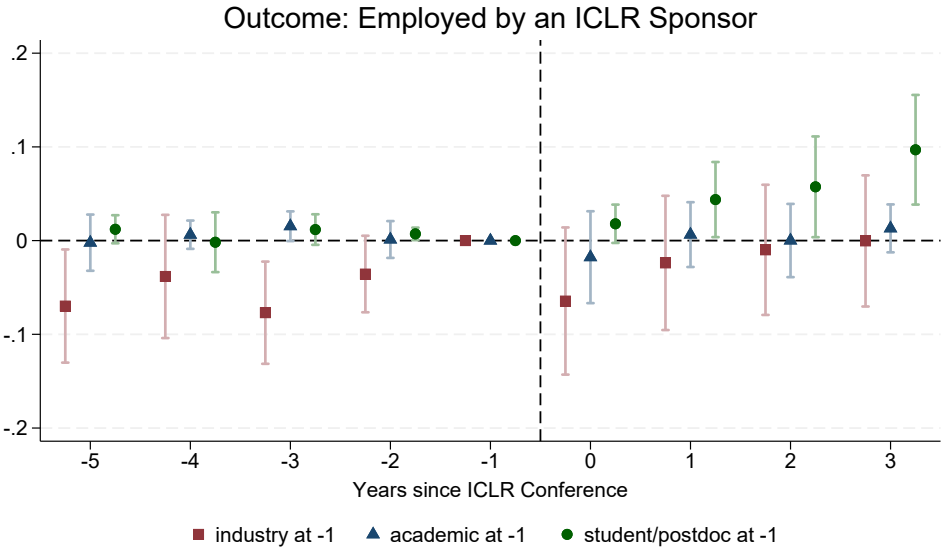


FIGURE A.16. Employed by Conference Sponsors: Separate Estimates by Origin



Notes: This figure shows the estimated impacts of an ICLR publication on being employed by an ICLR sponsor, separately by origin. The model is estimated on authors who are outside the top firms at -1, separately by whether the author was employed by nontop firms in industry at -1 (red), employed by academia (blue), or were students or postdocs (green). Rejected authors are reweighted by $\frac{\hat{p}_i}{1-\hat{p}_i}$, where the propensity score $\hat{p}_i$ is estimated by GMM (8) to achieve covariate balancing. Standard errors are robust and clustered at the author level.

FIGURE A.17. Employed by ICLR Sponsors at -1/0/1 rel. to Event Year (Weighted)

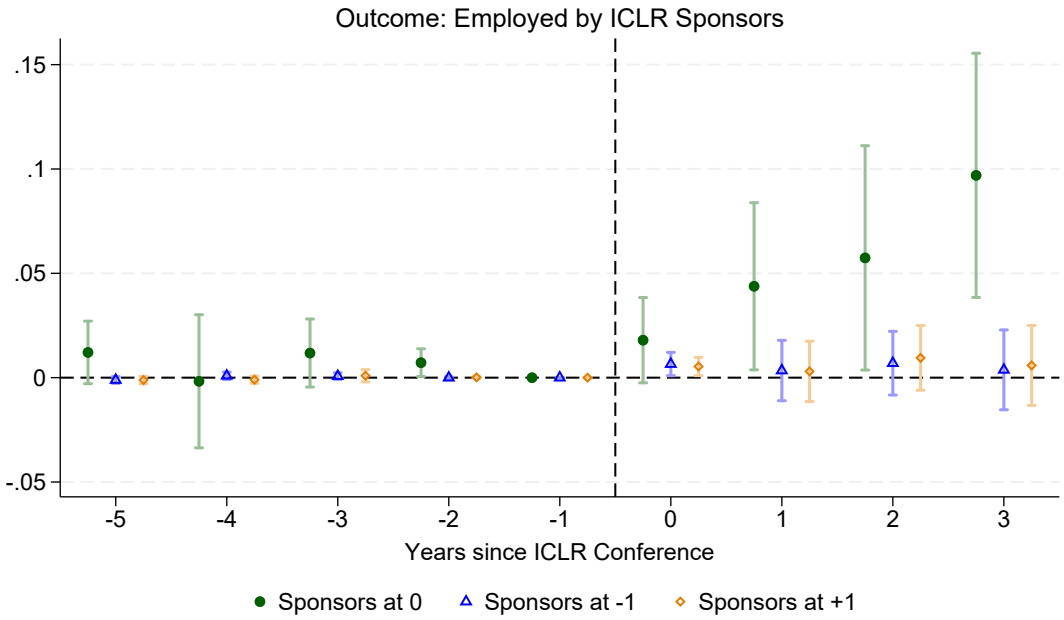
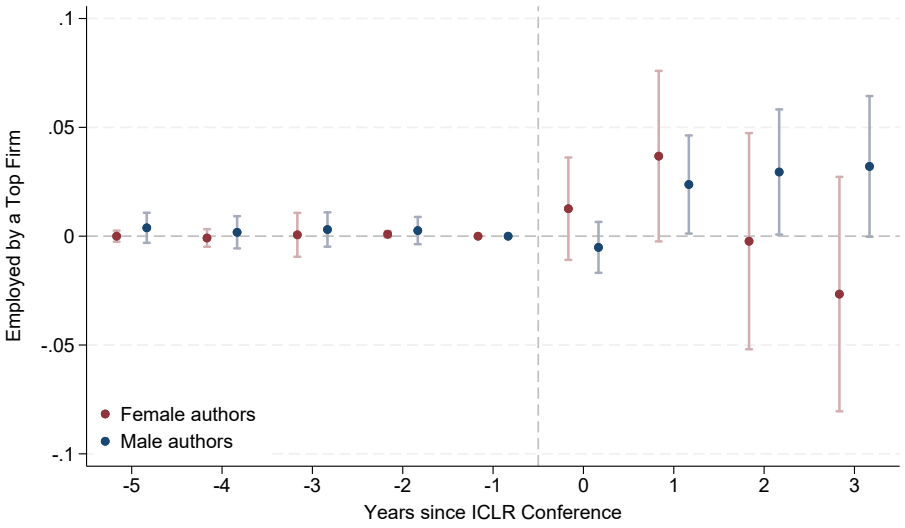
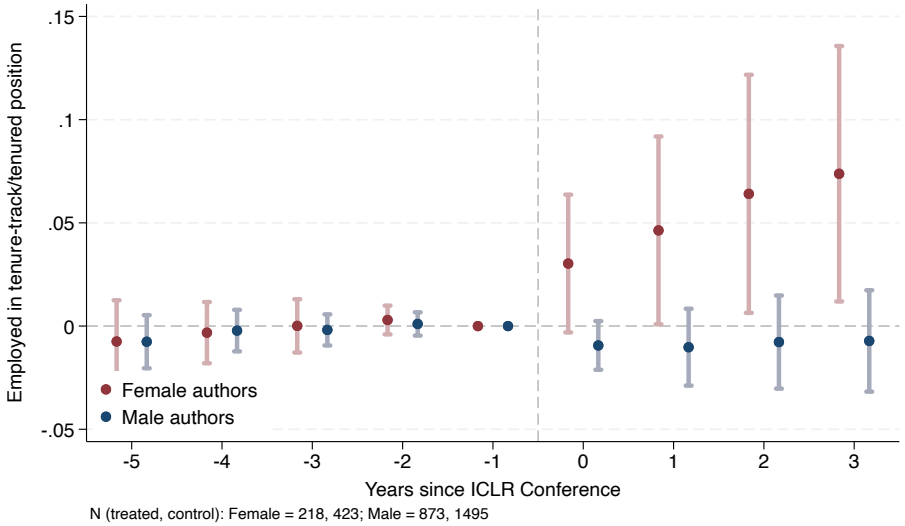


FIGURE A.18. Gender Differences in Employment Outcomes among Students/Postdocs (Unweighted)

(a) Top Firms

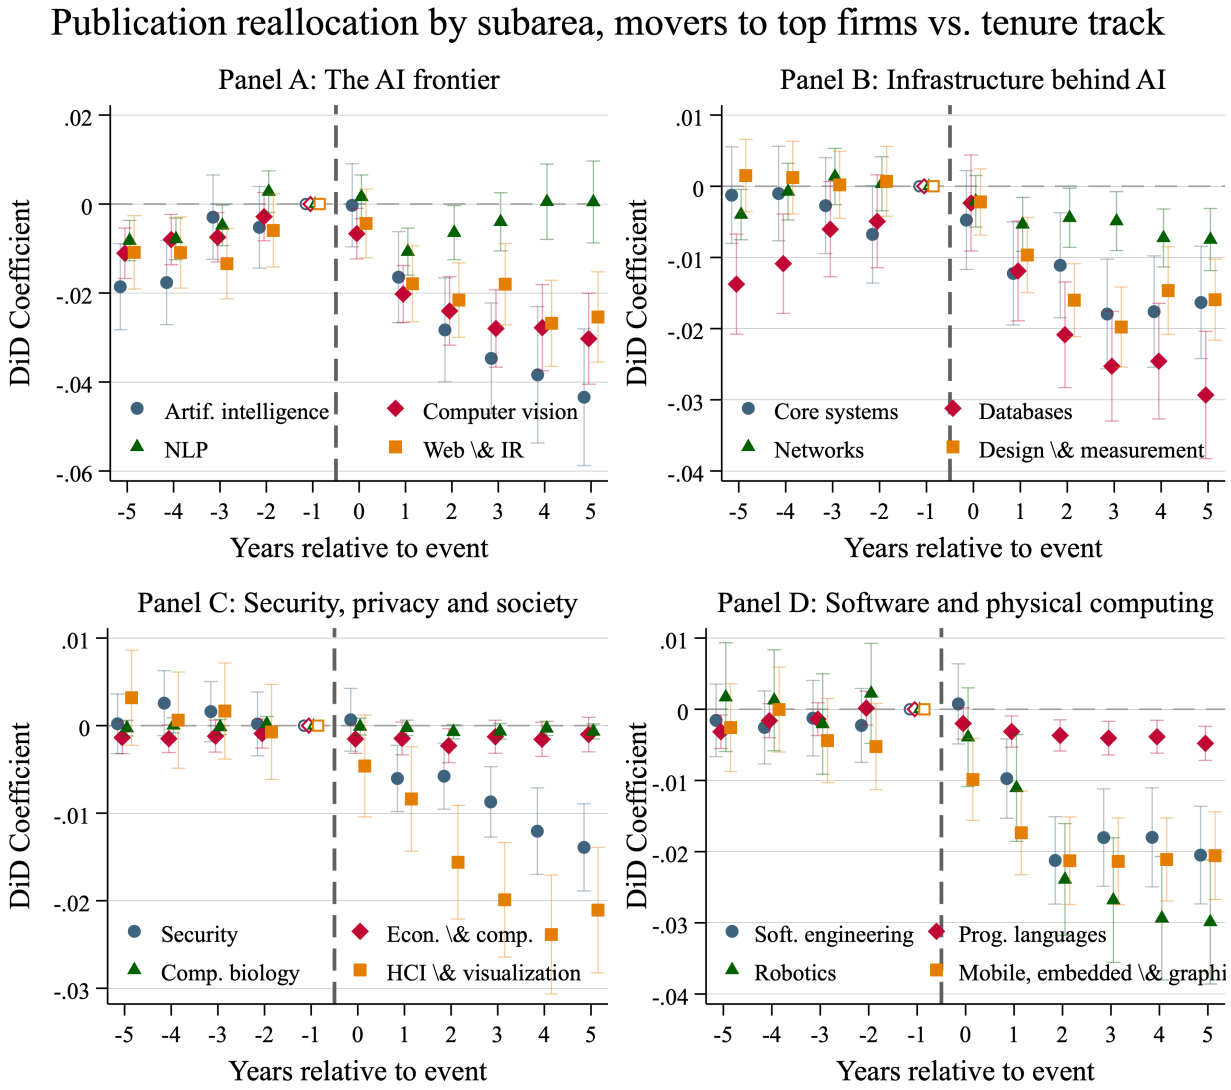


(b) Tenure-track Jobs



Notes: This figure shows the estimated impacts of an ICLR publication on employment at (a) top firms, and (b) tenure-track/tenured positions, separately by gender of the authors. We estimate the dynamic difference-in-differences model (5) separately for female and male authors who are students/postdocs at -1 (the year prior to the event). In contrast with Figure 7, the estimation sample is not weighted. Standard errors are robust and clustered at the author level.

FIGURE A.19. Direction of Research - Detailed Areas



Notes: The figure reports difference-in-differences estimates of the effect of moving to a top tech firm on annual publication counts across computer science subareas, relative to researchers moving to tenure-track academic positions. Coefficients are from a two-way fixed effects regression including author and year fixed effects; standard errors are clustered at the author level. The omitted category is the year prior to the transition ($t = -1$), shown as an open marker at zero. The dashed vertical line marks the transition year. The sample is restricted to authors observed in a balanced window of five years before and after the career transition. Results are organized into four panels. *Panel A (AI frontier)* disaggregates publications into four AI subareas: artificial intelligence (AAAI, IJCAI), computer vision (CVPR, ICCV, ECCV), natural language processing (ACL, EMNLP, NAACL), and web and information retrieval (SIGIR, WWW). *Panel B (Infrastructure behind AI)* reports results for four infrastructure subgroups: core systems, combining operating systems (OSDI, SOSP), computer architecture (ISCA, MICRO), and high-performance computing (SC, HPDC); databases (SIGMOD, VLDB, ICDE); networks (SIGCOMM, NSDI); and design and measurement, combining design automation (DAC, ICCAD) and performance analysis (SIGMETRICS, IMC). *Panel C (Security, privacy and society)* covers security (CCS, IEEE S&P, USENIX Security), economics and computation (EC, WINE), computational biology (ISMB, RECOMB), and HCI and visualization, combining human-computer interaction (CHI, UIST) and visualization (VIS, VR). *Panel D (Software and physical computing)* covers software engineering (ICSE, FSE), programming languages (PLDI, POPL), robotics (ICRA, IROS), and a combined category of mobile computing (MobiCom, MobiSys), embedded systems (EMSOFT, RTSS), and computer graphics (SIGGRAPH).

A3. Appendix Tables

TABLE A.1. First Job the Year After PhD

	Industry				Academia		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	Industry	Top	Re-Scientist	Engineer	Academia	Postdoc	Tenure-Track
AI papers	.003514	.006803	.01038	-.007327	-.002858	-.0004205	-.0008396
	(.0007037)	(.0009546)	(.001125)	(.0008984)	(.0006762)	(.0004957)	(.0004518)
Other Papers	-.001238	-.001841	-.005178	.00382	.001787	-.0009024	.001945
	(.0005495)	(.0005118)	(.0005978)	(.0005072)	(.0005484)	(.0004522)	(.0004122)
female	-.08686	-.008415	-.00808	-.06833	.05895	.01367	.02279
	(.004993)	(.002558)	(.003109)	(.003668)	(.005119)	(.004202)	(.00349)
asian	.07394	.04715	.02582	.04231	-.07737	-.03455	-.02118
	(.007236)	(.003851)	(.004646)	(.005707)	(.007363)	(.005982)	(.005097)
white	.03923	.008931	.001407	.01206	-.03616	-.01287	-.009764
	(.006795)	(.003305)	(.004193)	(.005238)	(.006956)	(.005724)	(.004803)
Constant	.4194	.05281	.0949	.1813	.4964	.2176	.1375
	(.006435)	(.003163)	(.003998)	(.004991)	(.006578)	(.005415)	(.004536)
N	81,225	81,225	81,225	81,225	81,225	81,225	81,225
Adj. R2	.06471	.05337	.03759	.03466	.0514	.03646	.09152
Mean of y	.4454	.07596	.1043	.1845	.4639	.1998	.1322
AI Papersly=1	.8048	1.624	1.431	.5462	.5698	.6183	.5544
Other Papers	1.327	1.935	1.395	1.401	1.082	1.079	1.178

Notes: This table shows the OLS estimates of regression (1) of the initial job placement on publications before graduation, for 81,225 PhD 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 authors who report a PhD 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.2. Job Mobility After PhD (Year-to-Year Transitions between Employers)

	At Nontop Firms			At Top Firms			In Academia		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
	Industry	Top	Academia	Industry	Top	Academia	Industry	Top	Academia
Any AI Paper	-.01535	.01571	.009634	-.004354	.001241	.004956	.003816	.005431	-.003643
	(.001812)	(.001338)	(.001472)	(.001804)	(.003085)	(.001513)	(.001271)	(.0005714)	(.001739)
Any Other Paper	-.008466	.003801	.005805	-.007559	-.008193	.00761	.002699	.001435	-.004688
	(.001483)	(.0009481)	(.001228)	(.001916)	(.003294)	(.001673)	(.001046)	(.0003993)	(.001435)
exp1	.1362	-.02243	-.1125	.01181	-.04857	-.02447	-.1419	-.02472	.03495
	(.006847)	(.003682)	(.005703)	(.009673)	(.01677)	(.008121)	(.006291)	(.002206)	(.0085)
exp2	-.1265	.01294	.09872	-.006476	.05943	.0208	.1023	.0187	-.001895
	(.007301)	(.003995)	(.005921)	(.01062)	(.01876)	(.008362)	(.006503)	(.002235)	(.009174)
exp3	.03572	-.003155	-.02663	.0008856	-.01717	-.005826	-.02509	-.004635	.0003702
	(.002377)	(.001294)	(.001869)	(.003646)	(.006414)	(.00268)	(.002019)	(.0006814)	(.002956)
female	-.009857	-.002045	.00944	.001438	.008494	-.001931	-.01013	-.002209	.01875
	(.001728)	(.000936)	(.001309)	(.002447)	(.004214)	(.001863)	(.001447)	(.000496)	(.002114)
asian	.01306	.004462	-.003276	.008481	.00707	-.005645	.002442	.0007162	.01228
	(.00195)	(.001179)	(.001502)	(.002824)	(.005048)	(.002225)	(.001767)	(.0006135)	(.002501)
white	.006727	-.001346	-.004055	.005278	.005581	-.005461	.003446	.0006162	-.006047
	(.00172)	(.0009819)	(.001358)	(.002761)	(.004814)	(.00224)	(.001495)	(.0004876)	(.002113)
MI_pos	.03488	-.01019	-.02455	-.09279	-.03784	-.005685	-.05851	-.01596	.2944
	(.008802)	(.0035)	(.008937)	(.09731)	(.09918)	(.001636)	(.02991)	(.001792)	(.04281)
Re-Scientist	.02057	.004871	-.009188	.008411	.009634	-.006276	-.004328	-.001381	.2241
	(.001057)	(.0007166)	(.000869)	(.001811)	(.003377)	(.001482)	(.002969)	(.000996)	(.003543)
Engineer	.02766	.005303	-.01693	.01239	.02045	-.01041	.02331	-.002521	.1841
	(.0008266)	(.0005242)	(.0006526)	(.001496)	(.002832)	(.001185)	(.004313)	(.001133)	(.004938)
Manager	.02002	-.000658	-.01092	.005719	.001604	-.004766	-.009563	-.003145	.2059
	(.0008725)	(.0005923)	(.0007052)	(.00122)	(.002881)	(.0008291)	(.004384)	(.001119)	(.00528)
Postdoc	-.107	.008111	.05895	-.1054	-.1232	.1015	.008978	.001305	.1739
	(.005094)	(.002116)	(.004136)	(.01112)	(.01322)	(.01067)	(.001813)	(.0006568)	(.002417)
Inverse Mills	-.08335	.03434	-.07037	-.06354	-.09299	.009055	-.02651	.007041	-.4198
	(.02571)	(.01417)	(.01837)	(.03573)	(.05716)	(.0284)	(.02302)	(.008197)	(.03272)
Junior Faculty							-.07268	-.008893	.3203
							(.001118)	(.0003889)	(.001813)
Assoc. Prof							-.06393	-.0066	.3138
							(.00103)	(.0003223)	(.001816)
Full Faculty							-.06287	-.006615	.3117
							(.001083)	(.0003309)	(.001868)
Constant	.9158	.01439	.07335	.9767	.9405	.02443	.1307	.01467	.6872
	(.004719)	(.002559)	(.003639)	(.007066)	(.01172)	(.005824)	(.004306)	(.001513)	(.005985)
N	302,994	302,994	302,994	59,599	59,599	59,599	366,181	366,181	366,181
Adj. R2	.04672	.008104	.03283	.02645	.01417	.03772	.05506	.006569	.2017
Mean	.9565	.01519	.0271	.9847	.9395	.009161	.05036	.005336	.865

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. The estimation sample is restricted to authors who report a PhD degree. See notes under Table 2 for details on the outcomes.

TABLE A.3. Characteristics of Accepted versus Rejected Authors (Unweighted vs. IPW - likelihood)

	Original Sample			Weighted (IPW-Likelihood)	
Variable	Accepted	Rejected	Diff.	Rejected	Diff.
Mode of Referee Rating	6.313	4.007	2.305***	5.871	0.441***
	(1.211)	(1.425)	(0.019)	(0.934)	(0.036)
Min Rating	5.495	3.293	2.202***	5.116	0.379***
	(1.238)	(1.242)	(0.018)	(1.068)	(0.039)
Mean Rating	6.556	4.452	2.104***	5.841	0.715***
	(0.728)	(1.088)	(0.013)	(0.716)	(0.028)
Num. Authors	5.395	4.756	0.639***	5.324	0.072
	(2.665)	(2.123)	(0.037)	(2.455)	(0.077)
Age	30.846	31.066	-0.220***	30.995	-0.149
	(5.331)	(5.557)	(0.080)	(5.593)	(0.209)
Female	0.233	0.250	-0.016***	0.217	0.016
	(0.423)	(0.433)	(0.006)	(0.413)	(0.013)
Foreign	0.465	0.527	-0.063***	0.479	-0.014
	(0.499)	(0.499)	(0.007)	(0.500)	(0.017)
White	0.388	0.353	0.035***	0.399	-0.011
	(0.487)	(0.478)	(0.007)	(0.490)	(0.017)
Asian	0.566	0.597	-0.032***	0.564	0.002
	(0.496)	(0.490)	(0.007)	(0.496)	(0.017)
Minority (Not White/Asian)	0.046	0.050	-0.003	0.038	0.009*
	(0.210)	(0.217)	(0.003)	(0.190)	(0.005)
Any PhD	0.479	0.460	0.019**	0.494	-0.016
	(0.500)	(0.498)	(0.007)	(0.500)	(0.017)
Any Master	0.632	0.623	0.009	0.650	-0.018
	(0.482)	(0.485)	(0.007)	(0.477)	(0.015)
Employed by Industry at -1	0.322	0.295	0.028***	0.327	-0.005
	(0.467)	(0.456)	(0.007)	(0.469)	(0.016)
Employed by Top at -1	0.137	0.082	0.055***	0.148	-0.010
	(0.344)	(0.275)	(0.005)	(0.355)	(0.014)
Employed by Academia at -1	0.156	0.186	-0.030***	0.147	0.009
	(0.363)	(0.389)	(0.006)	(0.354)	(0.010)
Student/Postdoc at -1	0.427	0.414	0.013*	0.446	-0.019
	(0.495)	(0.493)	(0.007)	(0.497)	(0.017)
Observations	6,879	13,234	20,113	13,234	20,113

Notes: This table displays average paper and author characteristics of accepted versus rejected authors in the original (unweighted) sample in columns 1-3, and the reweighted sample using inverse propensity weighting in columns 4-5. The propensity score is estimated by maximum likelihood. The sample includes every author who submitted to ICLR conferences for the first time between 2017 and 2020. 13.7% of accepted authors and 8.2% of rejected authors were employed by top firms at -1, the year prior to their first ICLR conferences. Given that the main outcome of interest is moving to top firms from other places, authors employed by top at -1 are dropped from the estimation sample for the difference-in-differences models (4) and (5). We also compare accepted and rejected authors in Table 3 conditional on authors not employed by top firms at -1, and we estimated the covariate-balancing IPT weights under that restriction.

TABLE A.4. Characteristics of Accepted versus Rejected Authors (Matched Samples)

	Matched Sample (w replacement)			Matched Sample (w/o replacement)		
Variable	Accepted	Rejected	Diff.	Accepted	Rejected	Diff.
Min Rating	5.429	5.302	0.127***	4.390	4.378	0.011
	(1.270)	(1.063)	(0.025)	(1.259)	(1.194)	(0.049)
Mean Rating	6.469	6.099	0.370***	5.857	5.576	0.280***
	(0.726)	(0.589)	(0.014)	(0.672)	(0.725)	(0.028)
Predicted Prob. Accepted (Text)	0.650	0.502	0.148***	0.487	0.395	0.091***
	(0.282)	(0.220)	(0.005)	(0.301)	(0.246)	(0.011)
Num. Authors	5.273	5.104	0.169***	5.029	4.859	0.169*
	(2.560)	(2.806)	(0.057)	(2.351)	(2.261)	(0.092)
Age	30.577	31.249	-0.673***	30.778	31.059	-0.282
	(5.152)	(5.758)	(0.116)	(5.196)	(5.709)	(0.217)
Female	0.230	0.252	-0.022**	0.226	0.221	0.006
	(0.421)	(0.434)	(0.009)	(0.419)	(0.415)	(0.017)
Foreign	0.451	0.496	-0.045***	0.482	0.487	-0.005
	(0.498)	(0.500)	(0.011)	(0.500)	(0.500)	(0.020)
White	0.396	0.419	-0.022**	0.390	0.416	-0.027
	(0.489)	(0.493)	(0.010)	(0.488)	(0.493)	(0.020)
Asian	0.557	0.548	0.009	0.565	0.523	0.042**
	(0.497)	(0.498)	(0.011)	(0.496)	(0.500)	(0.020)
Minority	0.046	0.033	0.013***	0.045	0.060	-0.015*
	(0.210)	(0.178)	(0.004)	(0.208)	(0.238)	(0.009)
Employed by Industry at -1	0.308	0.308	-0.000	0.287	0.287	0.000
	(0.462)	(0.462)	(0.010)	(0.453)	(0.453)	(0.018)
Employed by Top at -1	0.127	0.092	0.034***	0.092	0.101	-0.010
	(0.333)	(0.289)	(0.007)	(0.289)	(0.302)	(0.012)
Employed by Academia at -1	0.147	0.147	0.000	0.170	0.170	-0.000
	(0.355)	(0.355)	(0.008)	(0.376)	(0.376)	(0.015)
Student/Postdoc at -1	0.476	0.476	-0.000	0.466	0.466	-0.000
	(0.499)	(0.499)	(0.011)	(0.499)	(0.499)	(0.020)
Observations	4,424	4,424	8,848	1,263	1,263	2,526

Notes: This table displays average paper and author characteristics for accepted versus in the matched sample. The first three columns summarize the matched sample in which control (rejected) authors are drawn with replacement, whereas the next three columns summarized the matched sample in which control authors are drawn without replacement. Columns 3 and 6 display the difference in means with the corresponding standard error in parentheses. 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 A.1. Drawing without replacement reduces the number of matched accepted authors from 4,424 to 1,263, but the covariates are more balanced.

TABLE A.5. Publications

	Publication volume			Non-AI breakdown		
	(1)	(2)	(3)	(4)	(5)	(6)
	All	AI	Non-AI	Society	Infrastructure	Software
Year 0 × Top firm	-0.044***	-0.010	-0.035***	-0.006	-0.011*	-0.015***
	(0.012)	(0.007)	(0.009)	(0.004)	(0.006)	(0.006)
Year +1 × Top firm	-0.167***	-0.065***	-0.102***	-0.016***	-0.039***	-0.041***
	(0.013)	(0.008)	(0.009)	(0.004)	(0.006)	(0.006)
Year +2 × Top firm	-0.232***	-0.080***	-0.152***	-0.024***	-0.052***	-0.070***
	(0.014)	(0.010)	(0.010)	(0.004)	(0.006)	(0.006)
Year +3 × Top firm	-0.259***	-0.085***	-0.174***	-0.031***	-0.068***	-0.070***
	(0.016)	(0.011)	(0.011)	(0.004)	(0.007)	(0.007)
Year +4 × Top firm	-0.273***	-0.092***	-0.181***	-0.038***	-0.064***	-0.072***
	(0.018)	(0.013)	(0.011)	(0.004)	(0.007)	(0.007)
Year +5 × Top firm	-0.282***	-0.099***	-0.183***	-0.037***	-0.069***	-0.076***
	(0.019)	(0.014)	(0.012)	(0.005)	(0.007)	(0.007)
Pre-move mean (=-1, $treated$)	0.495	0.214	0.280	0.042	0.124	0.099
Effect at year +5 (%)	-57.0	-46.0	-65.5	-86.5	-55.7	-76.8
Author FE	Yes	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Observations	333,168	333,168	333,168	333,168	333,168	333,168

Notes: This table reports difference-in-differences estimates of the effect of moving to a top technology firm on annual publication counts, relative to researchers moving to tenure-track academic positions. Each column corresponds to a separate regression of the form $y_{it} = \beta_0 + \sum_{\ell \neq -1} \beta_\ell D_{it}^{(\ell)} \times \text{Treated}_i + \delta_t + \alpha_i + \varepsilon_{it}$, where Treated_i equals one for top-firm movers and zero for tenure-track movers, and $D_{it}^{(\ell)}$ are event-time dummies relative to the career transition. The omitted category is $t = -1$. All regressions include author and year fixed effects. Standard errors are clustered at the author level and reported in parentheses. The sample is restricted to authors observed in a balanced window of five years before and after the career transition. Society aggregates cybersecurity, economics and computation, computational biology, HCI, and visualization. Infrastructure aggregates computer architecture, operating systems, high-performance computing, databases, networks, design automation, and measurement. Software aggregates software engineering, programming languages, robotics, embedded systems, mobile computing, and computer graphics. Pre-move mean is computed at event time $t = -1$ for top-firm movers. Effect at year +5 (%) is the coefficient at $t = +5$ divided by the pre-move mean, expressed as a percentage. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

TABLE A.6. Citations

	Extensive margin (any citation)			High-impact (top decile)		
	(1)	(2)	(3)	(4)	(5)	(6)
	All	AI	Non-AI	All	AI	Non-AI
Year 0 × Top firm	-0.020***	-0.005	-0.017***	0.005***	0.003**	-0.001
	(0.006)	(0.004)	(0.005)	(0.002)	(0.001)	(0.002)
Year +1 × Top firm	-0.072***	-0.027***	-0.049***	0.002	0.002*	-0.004**
	(0.006)	(0.004)	(0.005)	(0.002)	(0.001)	(0.002)
Year +2 × Top firm	-0.103***	-0.032***	-0.078***	0.005**	0.005***	-0.004**
	(0.006)	(0.004)	(0.005)	(0.002)	(0.001)	(0.002)
Year +3 × Top firm	-0.108***	-0.037***	-0.083***	0.000	0.002	-0.005***
	(0.006)	(0.004)	(0.005)	(0.002)	(0.002)	(0.002)
Year +4 × Top firm	-0.115***	-0.037***	-0.090***	0.003	0.003*	-0.003*
	(0.006)	(0.004)	(0.005)	(0.002)	(0.002)	(0.002)
Year +5 × Top firm	-0.109***	-0.036***	-0.083***	0.005**	0.004**	-0.004**
	(0.006)	(0.004)	(0.005)	(0.002)	(0.002)	(0.002)
Pre-move mean (=-1, $treated$)	0.257	0.119	0.161	0.020	0.008	0.015
Effect at year +5 (%)	-42.2	-30.5	-51.6	27.3	50.1	-26.5
Author FE	Yes	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Observations	333,168	333,168	333,168	333,168	333,168	333,168

Notes: This table reports difference-in-differences estimates of the effect of moving to a top technology firm on research citation outcomes, relative to researchers moving to tenure-track academic positions. Each column corresponds to a separate regression of the form $y_{it} = \beta_0 + \sum_{\ell \neq -1} \beta_\ell D_{it}^{(\ell)} \times \text{Treated}_i + \delta_t + \alpha_i + \varepsilon_{it}$, where Treated_i equals one for top-firm movers and zero for tenure-track movers, and $D_{it}^{(\ell)}$ are event-time dummies relative to the career transition. The omitted category is $t = -1$. Columns 1–3 report estimates for the extensive margin of forward citations, defined as an indicator equal to one if a researcher's papers published in a given year receive at least one citation within three years of publication. Columns 4–6 report estimates for the high-impact indicator, defined as an indicator equal to one if three-year forward citations fall in the top decile of the citation distribution, computed conditional on receiving at least one citation. All regressions include author and year fixed effects. Standard errors are clustered at the author level and reported in parentheses. The sample is restricted to authors observed in a balanced window of five years before and after the career transition. Pre-move mean is computed at event time $t = -1$ for top-firm movers. Effect at year +5 (%) is the coefficient at $t = +5$ divided by the pre-move mean, expressed as a percentage. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.