The golden leap: gender differences in the Matthew effect*

Jennifer Doleac† Erin Hengel‡ Myra Mohnen§

July 2025

Abstract

In this paper, we explore whether the Matthew effect can mitigate under-recognition of women's high quality research by indirectly drawing attention to it. To identify the effect, we exploit the timing of an author's first publication in a top-five economics journal to estimate the impact it has on citations to his previously published work. We find that male and female co-authors on the same top-five paper receive significantly more citations to their earlier research around the timing of their first top-five paper, but the effect is much larger for the female co-author than it is for her male co-author. Empirical tests informed by a model of the decision to cite suggest that gender differences are driven by poor awareness of women's earlier research; among citations entirely motivated by strategic considerations—*e.g.*, a desire to cite well-known authors—gender differences disappear. These results suggest that status shocks shine a light on under-recognised work by women without granting them prestige out of proportion to their accomplishments. They also emphasise that gender citation gaps can dramatically change over the life cycle of a paper.

*We are especially grateful to Hayoung Jo and Shyrley P. from Upwork for excellent research assistance. We also appreciate helpful comments from Thomas Breda, Anja Prummer and seminar/conference participants at the Yosemite Retreat, Société canadienne de science économique Canada, DRUID Université Côte d'Azur, SInnoPSis/TWIN4MERIT University of Cyprus, Career Structures in Economics Workshop, Royal Holloway, Freie Universität Berlin, Brunel University of London and the Gender Gaps in Academia Workshop Centre d'Economie de la Sorbonne.

†Arnold Ventures, jdoleac@arnoldventures.org.

‡Brunel University of London, erin.hengel@gmail.com.

§University of Ottawa, mmohnen@uottawa.ca.

1 Introduction

In general, women's research is under-recognised relative to men's. Men are less likely than women to cite female-authored articles (Dion, Sumner, and Mitchell, 2018; Dworkin *et al.*, 2020; Ferber, 1986; Ferber, 1988; Koffi, 2025; Teich *et al.*, 2022). Referees often evaluate papers by men as higher quality compared to equivalent—and even identical—papers by women (Bikard, Fernandez-Mateo, and Mogra, 2025; Card *et al.*, 2020; Krawczyk and Smyk, 2016). Grant reviewers score women's applications lower than they score men's (Royal Statistical Society, 2025), and tenure committees discount their contributions to collaborative work (Sarsons *et al.*, 2021).

In this paper, we study whether the Matthew effect can mitigate under-recognition of women's high quality research by indirectly drawing attention to it. According to Merton (1968, p. 62) "a scientific contribution will have greater visibility in the community of scientists when it is introduced by a scientist of higher rank than when it is introduced by one who has not yet made his mark." This "visibility Matthew effect" is often criticised for penalising researchers who have "not yet made their marks"; however, it may also serve a positive function by reducing gender disparities in the recognition awarded to those who *have* made their marks.

To investigate, we compile a database of the publication histories and career trajectories of academics who plausibly experience a boost in attention that leads to a Matthew effect. To proxy for this shock, we use the dates authors first published in a "top-five" economics journal. Among economists, there is almost unanimous agreement that the five best economics journals are the *American Economic Review* (*AER*), *Econometrica* (*ECA*), *Journal of Political Economy* (*JPE*), *Quarterly Journal of Economics* (*QJE*) and *Review of Economic Studies* (*REStud*). Publications in these journals are assumed to be high quality and are consequently greatly valued—indeed, the average economist would sacrifice half a thumb to publish in the *AER* (Attema, Brouwer, and van Exel, 2014)![1] Authors accomplishing this feat therefore probably enjoy a subsequent boost in visibility among economists and other academics.

To identify the Matthew effect, we adopt an event study approach that estimates citation counts to authors' previously published research around the date their first top-five paper was published (the event), controlling for top-five paper fixed effects.[2] We focus on citations to authors' previously published papers as their quality and contributions are unaffected by the timing of the event. We control for top-five paper fixed effects to assess gender differences in recognition between co-authors who experience identical visibility shocks.

Our results suggest that male and female co-authors on the same top-five paper receive significantly more citations to their earlier work around the event date. Assuming citations just after this date would have evolved similarly to citations accrued just before it, this pattern suggests that publishing in a top-five journal brings attention—and consequently citations—to previously published work. According to our estimates, this visibility Matthew effect is worth, on average, 1–3 citations a year, per paper.

Furthermore, the visibility effect is much larger for the female co-author than it is for her male co-author on the same top-five paper. Before jointly publishing in a top-five journal, women's previously published research is cited less than early research by their future male co-authors. Afterwards, it is cited significantly more.

Our results hold in the samples of papers published in economics and non-economics journals, among

[1] As noted by Heckman and Moktan (2020, p. 419), "Faculty meetings about hiring, promotion, tenure, and prize committee discussions assess candidates by the number of [top-five] articles they have published or have in the pipeline and the rapidity with which they were generated."

[2] We estimate the Matthew effect brought about by publishing a high impact paper. We do not capture the *specific* impact of publishing in a top-five journal as opposed to another journal, as we do not adjust for a counterfactual in which a top-five worthy paper is published elsewhere. (See Appendix F.1 for evidence suggesting that a smaller Matthew effect is present when authors publish in non-top journals.)

authors who solo-author their first top-five paper, controlling for top-five papers' primary and secondary *JEL* codes and number of co-authors, cited papers' journal impact factors, authors' institutions' ranks when their first top-five papers were published and their pre-top-five paper counts. Not only is the visibility Matthew effect always present, but it is consistently larger for women than it is for men.

To investigate the mechanisms driving our results, we construct a model of the decision to cite the previously published articles of authors who eventually publish a high impact paper. According to our model, a Matthew effect and its gender difference are present in this context because of: (i) increases in the information that potential citing papers have about the existence and relevance of an author's pre-top-five research; (ii) decreases in the threshold of relevance they apply when deciding whether to cite this research, a behaviour known "strategic citation" (see, *e.g.*, Rubin and Rubin, 2021); and (iii) increases in the value to others of publishing papers related to this research, a phenomenon we label "strategic publication".

As we show in Lemma 3, we can identify the combined impact of strategic citation and strategic publication by restricting the sample of citing papers to those whose authors were already aware of top-five authors' previous research before they published in a top-five journal. The intuition here is simple. Suppose a potential citing paper was aware of an author's early work *before* he published for the first time in a top-five journal; however, it only cited this work *afterwards*. Because the citing paper was always aware of the top-five author's early work, the decision to cite was not motivated by better information. Instead, it was sparked by a decline in the threshold of relevance it applied to this research ("strategic citation") or an increase in the value it attributed to writing and publishing papers related to it ("strategic publication"). Either way, the citing paper responded strategically to the top-five author's first top-five publication by citing his previous research when it otherwise would not have.

To test our theory, we identify two "treatment" environments—*i.e.*, environments where citing papers are more likely to be aware of authors' previous work before they publish in a top-five journal—and compare them to their relevant counterfactual environments. The treatment environments are: (i) the subset of previously published papers that the author also cited in his first top-five paper; and (ii) the subset of top-five authors who also released their first top-five paper as an NBER working paper. The counterfactual environments are: (i) the subset of previously published papers that were *not* cited by the author in his first top-five paper; and (ii) the subset of authors who did *not* release their first top-five paper as an NBER working paper.

As we argue in Section 4.2, citing papers in treatment environments should be more aware of authors' pre-top-five work before their first top-five publications, compared to their respective counterfactual environments. Economists almost always widely present and release pre-print versions of their future top-five research; as a consequence, most other economists with a paper related to this research will be aware of it—as well as the papers that it cites—several years before it is actually published. Similarly, working papers released in the NBER series receive more downloads and abstract views than their published versions (Lusher, Yang, and Carrell, 2023); thus, authors are probably exposed to more attention when they release their future top-five papers as NBER working papers than when they officially publish them in top-five journals.

In all treatment and counterfactual environments, we find that publishing for the first time in a top-five journal boosts citations to an authors' previously published work. Because citing papers in the treated environments are probably aware of this work before the top-five article was published, the bump is evidence that they respond to its acceptance by strategically citing earlier work by the author and/or strategically publishing papers related to it.

However, we find no evidence that citing authors' strategic behaviour differs by top-five author gender. Gender differences in the Matthew effect are essentially zero in treatment environments, where citing papers are always familiar with authors' pre-top-five work, so the effect is caused by strategic behaviour, according to our theory. In contrast, we document an obvious gender difference in counterfactual environments, where the Matthew effect could also be caused by better awareness as citing papers were not already familiar with authors' pre-top-five work. Assuming gender differences in citing papers' strategic behaviour is unrelated to selection into the treated (vs. counterfactual) environment, our combined results suggest that the Matthew effect increases awareness of under-recognised research by women; however, when awareness of women's research is already high, it confers them no advantage relative to comparable men.

We emphasis once again that gender differences in citing papers' strategic behaviour must be unrelated to selection into treatment in order to come to this conclusion. Nevertheless, available evidence strongly supports this assumption. First, gender differences in the characteristics of citing, cited and top-five papers are similar in treatment and counterfactual environments. Second, we find no evidence that citing papers face additional incentives to behave strategically in response to a woman (vs. a man) publishing for the first time in a top-five journal. Although there is a noticeable bump in the probability that authors are promoted to professor and named journal editors around the publication date of their first top-five papers, gender differences on both dimensions are negligible. There are also no gender differences in the impact publishing a top-five article has on the number and dollar amount of grants awarded to authors, their employers' reputations or the number of seminar invitations they receive.

Finally, our theory predicts that gender differences in the Matthew effect should: (i) cease to exist in estimates using authors' *second* top-five publication date as the event date; (ii) disappear among job-market stars; and (iii) re-emerge when the event date is the date a future top-five article was released in the NBER working paper series.[3] Our empirical evidence supports all three predictions: there are no gender differences in the Matthew effect around the publication dates of authors' second top-five papers and among stars on the economics job market; in contrast, they re-emerge around the dates authors release their future top-five papers in the NBER series.

This paper makes several contributions to the literature. First, we contribute to a large body of research documenting evidence of the Matthew effect, *e.g.*, on research dissemination (Azoulay, Stuart, and Wang, 2014), peer review (Simcoe and Waguespack, 2011), collaboration patterns (Perc, 2014) and funding (Bol, de Vaan, and van de Rijt, 2018). Most relevant to our study, Azoulay, Stuart, and Wang (2014) estimate the impact of the Howard Hughes Medical Institute Investigators award on citations to scientists' previously published papers. Our own study complements their research by focusing on theoretically and empirically analysing the gendered impact of a similar shock, and finding that it increases citations to women's research more than men's. Moreover, we show that citations triggered by better information must be present for women to enjoy a larger Matthew effect; when the effect is motivated by status alone, its gender difference is zero. This suggests that status shocks shine a light on women's under-recognised work without granting them prestige out of proportion to their accomplishments, relative to comparable men.

Our second contribution is to an emerging literature studying agency behaviour in decisions to cite and publish (see, *e.g.*, Bornmann and Daniel, 2007; Franzoni, Scellato, and Stephan, 2011; Kuhn, Younge, and Marco, 2023; Lampe, 2012; Rubin and Rubin, 2021; Seeber *et al.*, 2019; Siler and Larivière, 2022). Like most of this research, our evidence suggests that strategic considerations play a role in the decision to cite. We extend the literature by studying how this strategic behaviour interacts with cited authors' genders.

Third, our study complements research exploring how uncertainty creates and exacerbates gender out-

[3] Potential citing papers are not necessarily aware of authors' earlier work around the NBER release date, but they should be aware of this work if the authors are recent job market "stars" (see Appendix D.5 for a more detailed discussion) and if they already published a paper in a top-five journal.

come gaps. For example, Bohren, Imas, and Rosenberg (2019) show that evaluators discriminate less as they learn more about women's skills;[4] similarly, Alexander *et al.* (2023) find that referees produce more favourable outcomes for female authors as they gain experience in the peer review process and learn more about the standards of acceptance at a particular journal. We add to this research by investigating how a shock in citing papers' information affects their propensity to cite women. Our results suggest that poor quality information drives early career citation gaps favouring men; when potential citing papers are better informed, the gender gap disappears or reverses.

Finally, we contribute to research studying gender differences in citations. Although most studies find that men are less likely than women to cite female-authored papers (Dion, Sumner, and Mitchell, 2018; Dworkin *et al.*, 2020; Ferber, 1986; Ferber, 1988; Koffi, 2025; Teich *et al.*, 2022), evidence is less clear how this impacts citations to women overall and in selected samples (see, *e.g.*, Hengel and Moon, 2023; Larivière *et al.*, 2013). Our own results provide further nuance to the debate by suggesting that poorer quality information about female-authored papers may partially explain observed citation gaps favouring men. They also emphasise that changes in information over the lifecycle of a paper can dramatically alter the gender citation gap.

This paper proceeds in the following order. Section 2 describes the data. Section 3 identifies and estimates the Matthew effect and its gender difference; Section 4 explores the mechanisms driving them. Section 5 concludes.

2 Data

We create, expand and combine several datasets to form a comprehensive database tracking research outputs, citations and career trajectories for every economist who published at least one full-length, original research article between 1986–2015 in a "top-five" economics journal.[5] From the datasets analysed in Hengel (2022) and Hengel and Moon (2023), we obtained basic bibliographic information on top-five articles and biographic data on their authors. Authors' publication histories and citations were originally obtained from Web of Science (Clarivate, 2022); to improve accuracy, we also manually verified and corrected these records. Information on career trajectories—including education, funding, awards, employment and seminar invites—were retrieved by hand-collecting and digitising CVs.

Our final database covers the entire pre-top-five publication histories of 3,897 economists; for 78 percent of them, it also includes partial data on career trajectories. Further details on sources, coverage, collection procedures and variable definitions are provided in Appendix B.

Figure 1 illustrates numerous stylised facts about authors' pre-top-five publications. Graph (A) displays annual counts of male and female authors publishing their first top-five papers. It suggests women are under-represented: at *ECA* and *JPE*, they make up 9 and 12 percent of first-time authors, respectively; at the *AER*, *QJE* and *REStud*, they are 16 percent. And while women's numbers have increased over time, so have men's; as a result, female representation grew only 4 percentage points between 2000–2015.

The percentage of first time top-five authors who are female is low in every field. Graph (B) displays women's representation by primary *JEL* code of their first top-five paper. Although female authors only exceed 25 percent in a single field, there are nevertheless noticeable differences across them: the average

[4] Bohren, Imas, and Rosenberg (2019) build a dynamic model of evaluations to differentiate between taste-based discrimination, belief-based discrimination with correct beliefs and belief-based discrimination with incorrect, biased beliefs. Their model is ideally applied to experimental settings in which evaluators are presented with identical information about men's and women's past evaluation histories. This type of experimental setting is difficult to simulate in our observational data, as we do not know what information potential citing papers have (or could be reasonably expected to have) about authors' early work, especially before the quality of that work becomes common knowledge (*e.g.*, because the author publishes another paper in a top journal).

[5] We define full-length, original research articles as any non-errata/corrigenda/editorial article published with an abstract, excluding *Papers & Proceedings* issues of the *American Economic Review*. See Appendix B for further details.

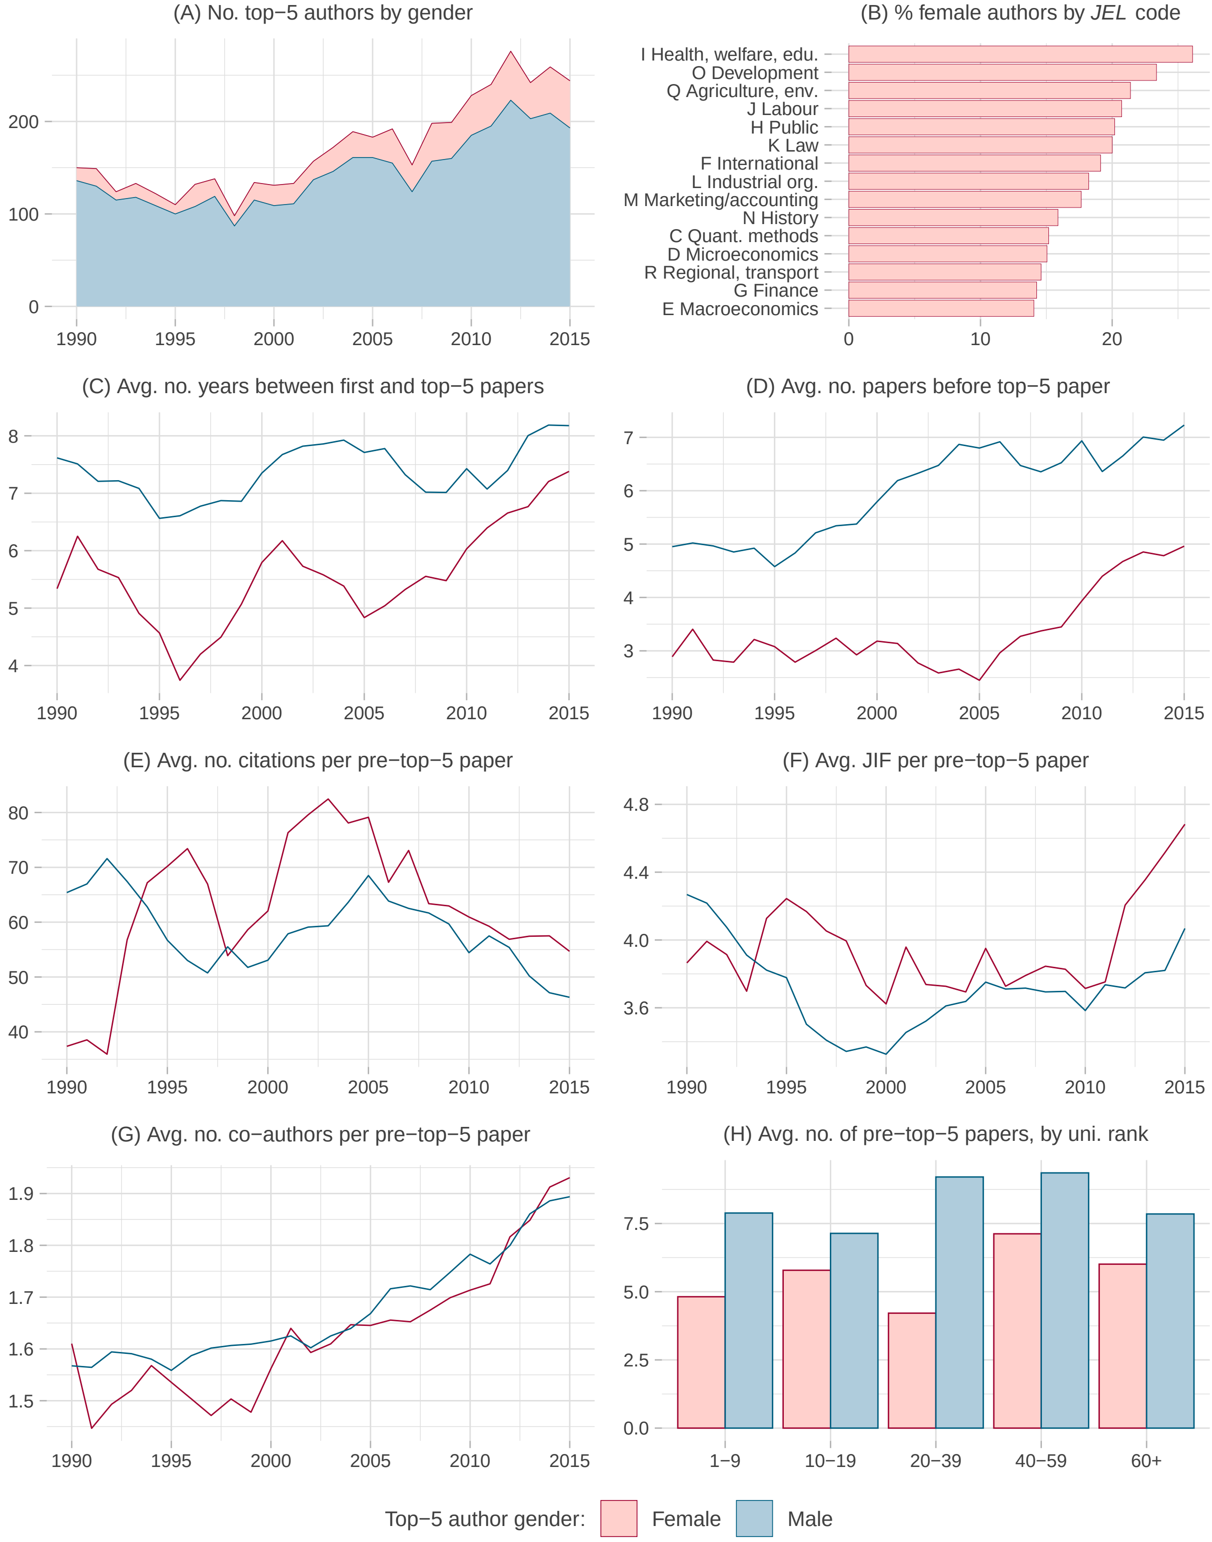

Figure 1: Characteristics of pre-top-five publications

Note. Graph (A) displays the total number of female and male authors published for the first time in a top-five journal each year. Graph (B) is the average number of years between authors' first publication and their first publication in a top-five journal. In graph (C) we plot authors' average number of papers before they publish their first top-five paper. Graphs (D) and (E) are the average number of citations and journal impact factors per pre-top-five paper. Graph (F) shows the average number of co-authors per pre-top-five paper. Graph (G) is the average number of pre-top-five papers by authors' institutional rank (see Appendix B). Graph (H) is the average number of pre-top-five papers by primary *JEL* code. Figures in graphs (B), (C), (D), (E) and (F) are five-year moving averages.

percentage of female authors is lowest in *JEL* codes E (macroeconomics and monetary economics), G (financial economics), R (urban, rural, regional, real estate and transportation economics) and D (microeconomics) and highest in I (health, education and welfare), O (economic development, innovation, technological change and growth) and Q (agriculture and natural resource economics and environmental and ecological economics).[6]

Graphs (C) and (D) plot counts of years and papers between authors' first publications and their first top-five papers, respectively. On average, men publish in top-five journals 8 years after their first papers were published. Women publish their first top-five papers 1.4 years sooner; they also produce 2.4 fewer pre-top-five papers. More recently, both gender gaps have narrowed, suggesting that the career trajectories of prominent female economists may be starting to resemble the career trajectories of comparable men.

Although women's pre-top-five output is smaller, it is cited more and published in more impactful journals. In graph (E) we plot citations per paper. On average, women's and men's pre-top-five papers are cited 60 and 56 times, respectively. Graph (F) displays journal impact factors (JIF) per paper; men's papers are published in journals with an average impact factor 0.3 points below the average impact factor for women's papers.

In graphs (G) and (H) we present average co-author counts and paper counts by authors' institutional rank .[7] Co-author counts have risen over time: in the 1990s, papers were co-authored by 1.6 people, on average; by 2015 that figure had risen to 1.9. There are also noticeable differences in publication counts across institutional rank—*e.g.*, authors at lower ranked institutions publish more papers before their first top-five paper—although we observe no meaningful gender differences.

Figure 2 describes the career characteristics of top-five authors themselves. Graphs (A) and (B) plot the percentage of professors who are women (by employer rank) and men's and women's average employer rank, respectively.[8] The share of women professors has steadily increased with time (see also Lundberg and Stearns, 2019): in 1990, 4–5 percent of professors were women; by 2015, their share had grown to 12 percent. And while top-five authors of both genders have gravitated to higher ranked institutions, the trend is especially pronounced for women: in the early 1990s, men's employers ranked 16 positions ahead of women's; by 2001, women were at higher ranked institutions than men.

The percentage of female authors among journal editors is also increasing and the average number of seminars given by men and women has converged. In graph (C), we plot the percentage of women among top-five authors who edit any journal, a top-20 journal and a top-five journal.[9] In 1990, only 5 percent of editors—and just 3 percent of top-five editors—were women. By 2015, women were 13 percent of all editors and 15 percent of editors at top journals. According to graph (D), both men and women give more seminars today than they did in the past, but while men gave more talks in 1990, women gave more in 2015.

Women also win more grants, although men still surpass them in the amount awarded. Graphs (E) and (F) plot average numbers and sizes (in thousand USD) of grants awarded to top-five authors.[10] Initially, women were awarded fewer grants than men; however, that trend reversed in 2003. By 2015, women were winning noticeably more grants, although the dollar amounts awarded to them are generally smaller.

[6] We omit *JEL* codes A, B, P and Z due to small numbers of top-five articles.

[7] In graph (G), institution refers to the author's highest ranked institutions among all institutions listed on his first top-five paper. Institutions are annually ranked (in descending order) by the number of top-five articles affiliated to them, smoothed over a five-year period (see Appendix B for further details).

[8] To limit the impact of outlier observations and impose ascending order, we additionally apply the log transformation multiplied by -1 to the sample in graph (B).

[9] We define top-20 journals as the economics and finance journals ranked 1–15 in Combes and Linnemer (2010), plus the *Journal of the European Economic Association* and the four *American Economic Journals*.

[10] Grant amounts were converted to 2024 USD using exchange rate data from the World Bank and the U.S. Consumer Price Index.

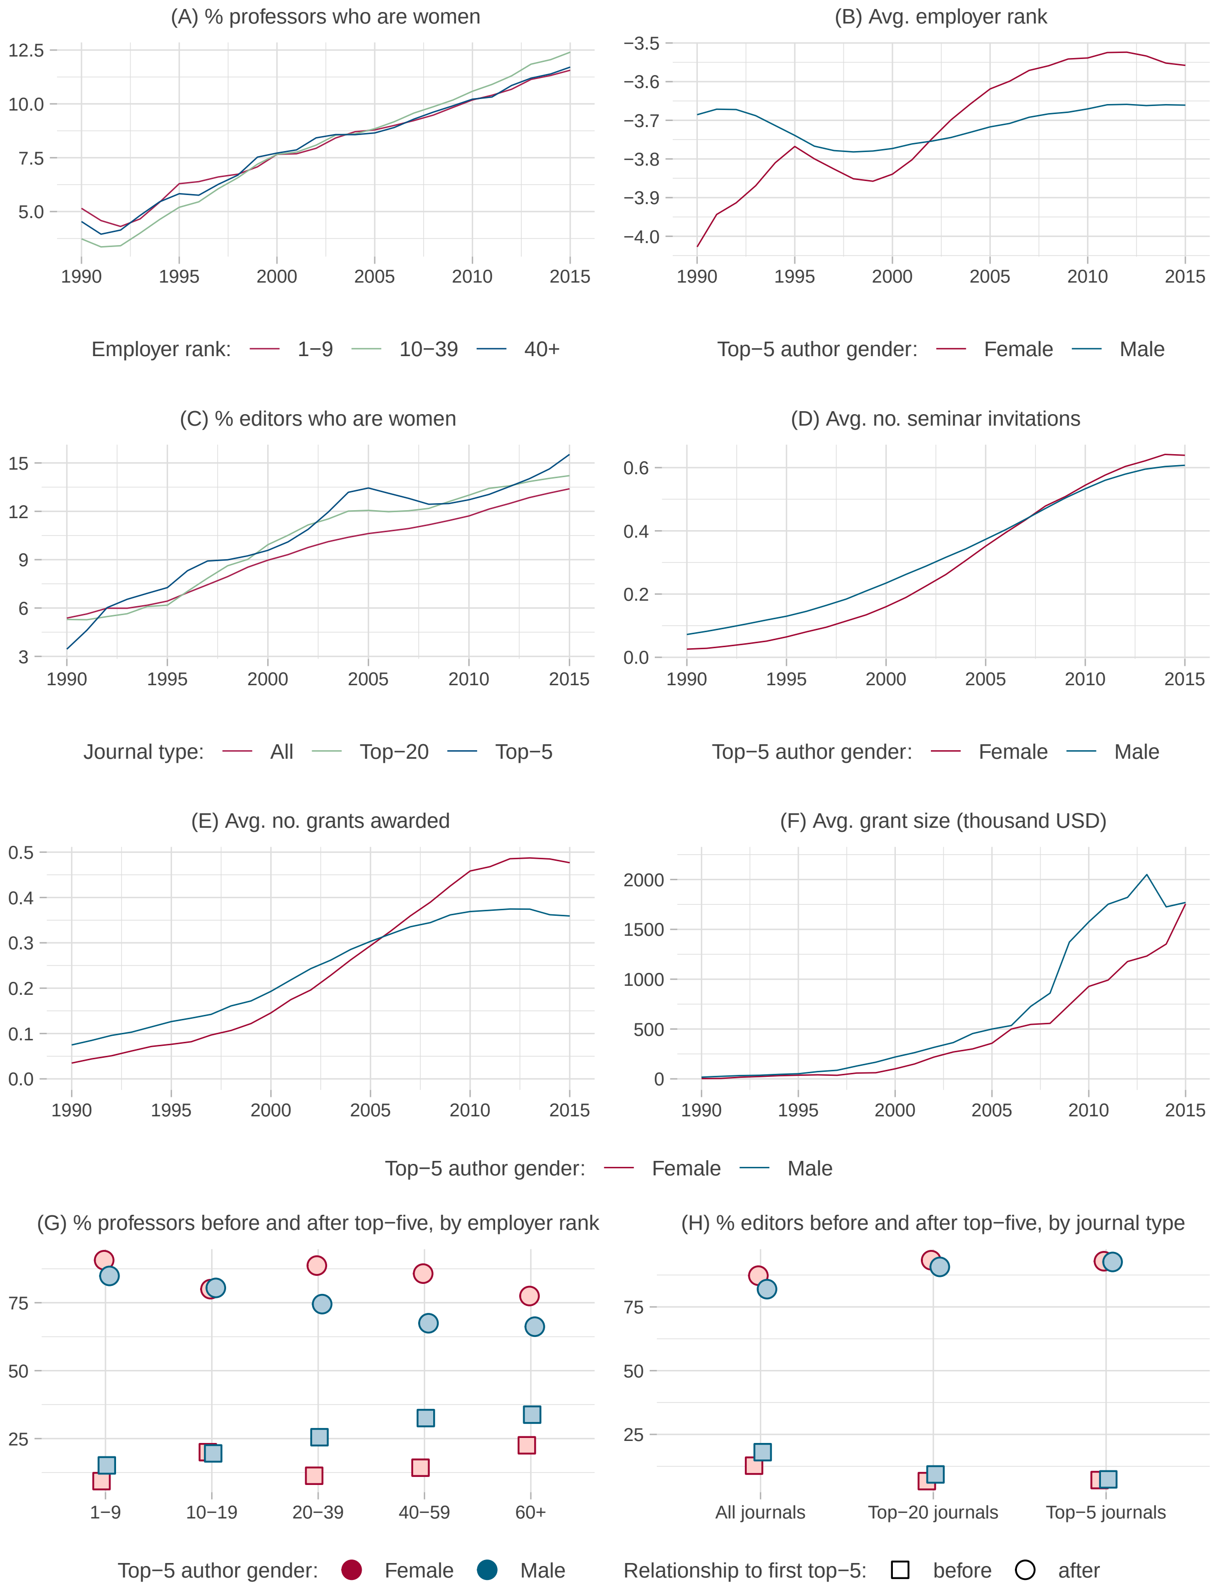

Figure 2: Characteristics of top-five authors

Note. Graph (A) displays the percentage of top-five authors who are professors and women. Graph (B) plots employers' average institutional rank. In graph (C), we show the percentage of editors who are women, and in graph (D) we plot authors' average number of seminar invitations. Graphs (E) and (F) show the average number and size (in thousand USD) of grants awarded; graphs (G) and (H) plot the percentage of authors promoted to professor and named editor before and after their first top-five paper was published. Figures in graphs (A)–(F) are five-year moving averages.

Graphs (G) and (H) plot the percentage of each gender's authors who were professors and editors, respectively, before and after they publish their first top-five papers. Men are more likely to have been promoted to professor and edited a journal beforehand; women are more likely to achieve both milestones afterwards. Interestingly, both professorship gaps widen as employer rank declines, suggesting that women at lower ranked institutions are especially unlikely (relative to men) to earn their promotions before they've published in a top-five journal.[11]

3 The visibility Matthew effect

3.1 Conceptual framework

Equation (1) measures how citations to paper j change after its author i experiences the initial success of publishing the high impact top-five paper k at time $t = 0$:

$$\Delta_{ijk}^{\text{Matthew}} = \mathbb{E}\left[\text{citations}_{j0} - \text{citations}_{j-1} \,\middle|\, \text{citations}'_{j0} = \text{citations}_{j-1}\right], \tag{1}$$

where j is another paper by i published before k, citations_{jt} counts the number of times j is cited at time t and $\text{citations}'_{jt}$ the number of times it would have been cited in the absence of k.

Equation (1) assumes that citations to j would not have changed between times $t = -1$ and $t = 0$ had k not been published. Or in other words, any "bump" in citations at $t = 0$ is thanks to k's publication. Conditional on this assumption, $\Delta_{jk}^{\text{Matthew}} > 0$ implies that publishing a high impact paper increases citations to one's previously published work. We interpret this as evidence of a "visibility Matthew effect"—*i.e.*, "that a scientific contribution will have greater visibility in the community of scientists when it is introduced by a scientist of higher rank than when it is introduced by one who has not yet made his mark" (Merton, 1968, p. 59).[12]

The visibility Matthew effect could be caused by a variety of factors. The first is a change in information about j at $t = 0$ relative to $t = -1$. For example, imperfectly informed individuals may have underestimated the relevance of j—or been unaware of its existence—before i published k; as a result, j is under-cited at time $t = -1$. Alternatively, publicity from k may have affected the research preferences of others, *e.g.*, by encouraging new PhD students to work on topics studied by j. Relatedly, publishing k likely enhances i's reputation and influence, and this may invite strategic citations—*e.g.*, after k is published, i may be more often asked to referee papers and write tenure letters, thereby incentivising others to favourably cite j.

Equation (2) compares the visibility Matthew effect of men and women who experience an equivalent publicity shock from k:

$$\Delta\Delta_{ijk}^{\text{Matthew}} = \mathbb{E}\big[\Delta_{j_F k}^{\text{Matthew}} - \Delta_{j_M k}^{\text{Matthew}} \mid \text{shock}_k\big], \tag{2}$$

where shock_k accounts for gendered heterogeneity in the attention authors receive from publishing a high impact paper.

Assuming shock_k is appropriately controlled for, Equation (2) is only non-zero if there are gender differences in how publicity from k relates to characteristics of author i or her subsequently cited paper j. Examples of the latter include stereotypes against or under-promotion of female-authored work at time $t < 0$ followed by a citation correction when better information is revealed at $t = 0$. Alternatively, if

[11]In Figure B.2 (Appendix B.2), we similarly plot the percentage of each gender's authors who were awarded a grant and the amounts they were awarded. In contrast to graphs (G) and (H), we observe no obvious pattern by gender.

[12]We call this the "visibility Matthew effect" to distinguish it from competing definitions of the Matthew effect proposed by Merton (1968) and others—*e.g.*, that "centers of demonstrated scientific excellence are allocated far larger resources for investigation than centers which have yet to make their mark" (Merton, 1968, p. 62).

publishing a high impact paper enhances women's influence more than men's—*e.g.*, because women are more likely than men to become journal editors after publishing a top paper—then they may enjoy a larger increase in citations as their peers discover the career benefits of strategically citing them.

3.2 Estimation strategy

To estimate Equations (1) and (2), we adopt an event study approach that exploits the timing of authors' first top-five publications. To implement it, we regress the number of citations to authors' previously published work (j)—whose quality and contribution remains constant over the estimation period—on indicator variables for each year around the publication date of authors' first top-five papers (k) and interactions with their genders:

$$\text{citations}_{ijt} = \alpha + \delta\,\text{female}_i + \sum_{t'=-10}^{10} (\beta_t + \gamma_t \times \text{female}_i) \times \mathbf{1}[t = t'] + \zeta\,\text{shock}_k + \varepsilon_{jt}, \tag{3}$$

where female_i and $\mathbf{1}[t = t']$ are dummy variables equal to one if author i is a woman and the year is t, respectively, shock_k absorbs heterogeneity in the publicity generated by k and ε_{jt} is the error term. We estimate Equation (3) with OLS using a window spanning the 10 years before and after k's publication date at $t = 0$.

β_0 and $\beta_0 + \gamma_0$ capture men's and women's average visibility Matthew effects defined in Equation (1). As discussed in Section 3.1, they are only identified if citations to j would not have changed had k not been published at time $t = 0$. Although publishing in a top-five journal is not random—talent, persistence and work habits all play a role—the exact timing of the event arguably is. As a result, sharp changes in citations to previously published papers that coincide with $t = 0$ are plausibly orthogonal to unobserved determinants of the outcome.[13]

Assuming shock_k appropriately absorbs heterogeneity in k (*e.g.*, manuscript quality), $\gamma_0 > 0$ is evidence that female-authored j receive more attention than male-authored j conditional on their authors experiencing an identical status shock from k.[14] We primarily account for shock_k by controlling for fixed effects for k; γ_0 in this context represents differences in the visibility Matthew effect between male and female co-authors on the same top-five paper. In order to interpret γ_0 across all authors—including individuals who solo-authored their first top-five paper—we also show results that instead control for k's year and journal of publication, co-author count, citation count (asinh) and author prominence. (To proxy for the latter, we use the number of previous top-five publications by k's most prolific co-author.)

3.3 Results

Figure 3 displays β_t and $\beta_t + \gamma_t$ from estimating Equation (3) using OLS for all $t \in [-10, 10]$, where $t = 0$ corresponds to the publication date of authors' first top-five papers. Graph (A) accounts for characteristics related to publishing a top-five paper using fixed effects for each top-five article k; in graph (B), we instead control for k's co-author count, citation count (asinh), author prominence—which we proxy for using the number of previous top-five publications by k's most prolific co-author—and year and journal fixed effects.

[13] β_0 and $\beta_0 + \gamma_0$ do not adjust for visibility Matthew effects that might be present had k been published in a non-top-five journal; thus, they should only be interpreted as capturing the Matthew effect brought about by publishing a high impact paper. (See Appendix F.1 for evidence suggesting that a smaller Matthew effect is present when authors publish in non-top journals.)

[14] If this assumption fails, then gender differences in the characteristics of high impact papers create gender differences in the attention their authors subsequently receive. For example, if women's manuscripts are held to higher standards in peer review at top journals (Card *et al.*, 2020; Hengel, 2022; Hengel and Moon, 2023), then their high impact published papers are probably higher quality—and will thus generate more attention for their authors—compared to men's.

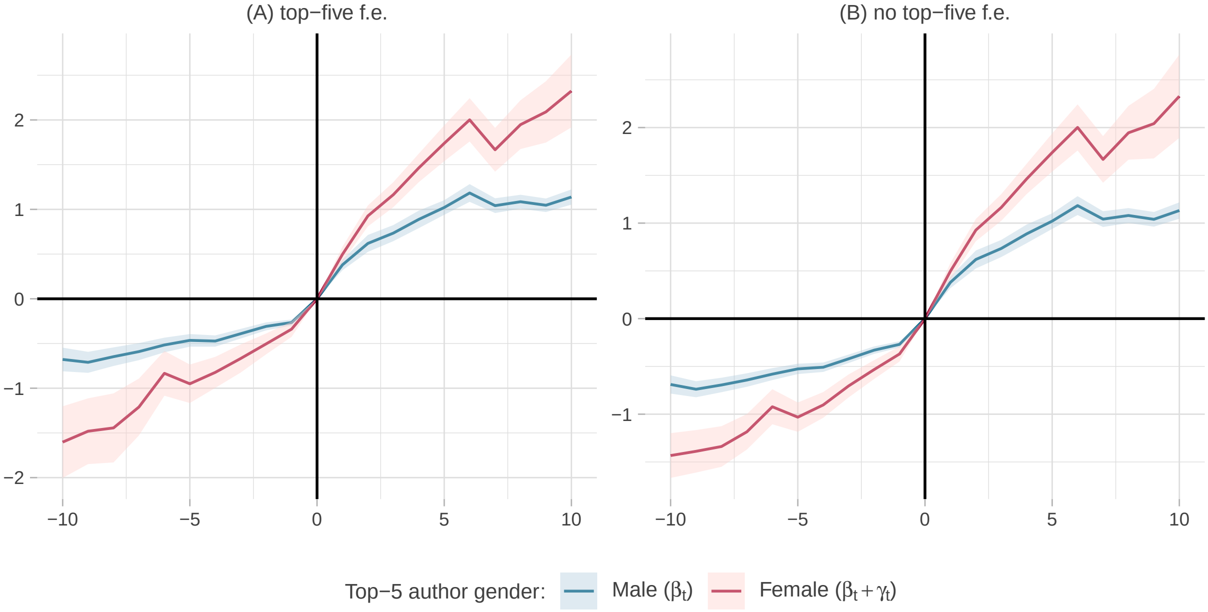

Figure 3: The visibility Matthew effects for men and women

Note. Figure 3 displays β_t and γ_t from estimating Equation (3) with OLS using a window spanning the 10 years before and after k's publication date. In graph (A), we account for shock$_k$ using fixed effects for each top-five article; in graph (B), we instead control for k's year and journal of publication (as fixed effects), co-author count, citation count and the number of previous top-five publications by k's most prolific co-author. Shaded areas represent 90 percent confidence intervals from standard errors clustered by j_g.

In both graphs (A) and (B), β_t is negative for $t < 0$ but then jumps at $t = 0$ and is positive for all $t > 0$. Assuming citations to previously published papers just after the event window would have evolved similarly to citations accrued just before it, this pattern suggests that publishing in a top-five journal brings attention—and consequently citations—to men's previously published work. According to our estimates, this "visibility Matthew effect" is worth, on average, about 1–2 extra citations a year per paper relative to the period before k was published.

$\beta_t + \gamma_t$ in Figure 3 plots the average visibility effect for women, where γ_t reflects deviations from β_t at each t. γ_t is negative and significant for all $t < 0$, suggesting that women's papers are cited relatively less than men's before they publish for the first time in a top-five journal. At around $t = 0$, however, γ_t turns positive, indicating a much larger visibility effect for women than for men; as a result, women's papers are cited relatively more than men's for all $t > 0$.

β_t and γ_t represent deviations from baseline gender differences in citations captured by δ (*i.e.*, the female fixed effect in Equation (3)). In Figure 3(A), δ is 0.44 (standard error 0.24); in Figure 3(B), it is 0.07 (standard error 0.11). Neither estimate is significant at traditional thresholds. Similarly, the coefficients on k's co-author count and the prominence of its most prominent author are very close to zero and not significantly different from it.

For robustness, we replicate Figure 3 using a Poisson likelihood function (Appendix C.1), in the samples of cited and citing papers published in economics and non-economics journals (Appendices C.3, C.4 and C.5), among authors who solo-authored k (Appendix C.6) and controlling for authors' institutional rank when k is published (Appendix C.8), k's primary and secondary *JEL* codes (Appendix C.9), authors' total number of pre-k publications (Appendix C.10), and j_g's journal impact factor and number of co-authors (Appendices C.11 and C.12). In all instances, results are similar to those reported in Figure 3.

We conclude by emphasising again that our estimates capture the visibility Matthew effect after an author publishes a high impact paper; they do not reflect *selection* into publishing a top-five paper as we do not adjust for a counterfactual in which a top-five worthy paper is published elsewhere. Indeed, the visibility effect and its gender difference are present—albeit smaller in magnitude—even when authors publish in non-top journals. In Figure F.1 (Appendix F.1), we reproduce Figure 3 but set $t = 0$ to five years before k's publication date and restrict the sample to authors who published in that year and their pre-k papers published before it. We find a visibility Matthew effect and gender difference that are about half the size of the effects shown in Figure 3.

4 The mechanisms behind the Matthew effect

4.1 Theoretical framework

In this section, we construct a simple theoretical framework to guide our interpretation of the Matthew effects observed in Figure 3. Let $j_g \in \mathcal{J}_g$ denote a paper by author i_g published before his first high-impact paper k, and $l \in \mathcal{L}$ a different paper (or idea for a paper) that potentially cites j_g, where $\mathcal{L}$ is the Dedekind complete body of research ideas.

Assume that l cites j_g only if it is sufficiently relevant to l, l's authors are sufficiently aware of it and they believe that the publication value of l is sufficiently high, where the latter two factors may depend on one another.[15] Under these conditions, l's decision to cite j_g is determined by Equation (4):

$$\mathbf{1}\Big[\theta_{j_g l} \geq \tilde{\theta}_{j_g l}\Big] \times \mathbf{1}\Big[\phi_l(\lambda_{j_g l}) \geq \tilde{\phi}_l\Big] \times \lambda_{j_g l}, \tag{4}$$

where $\mathbf{1}[\cdot]$ is the indicator function, $\theta_{j_g l}$ captures the relevance of j_g to l, $\lambda_{j_g l}$ is a binary variable equal to one if l's authors' are aware of j_g, $\phi_l(\lambda_{j_g l})$ absorbs l's authors' beliefs about the value of publishing l given their awareness of j_g and $\tilde{\theta}_{j_g l}$ and $\tilde{\phi}_l$ are the thresholds of relevance and value required for l to cite j_g and its authors to produce and publish l, respectively.

We additionally assume that: (i) $\phi_l(\lambda_{j_g l})$ is increasing in λ_{j_g}; (ii) $\lambda_{j_g l}$ and ϕ_l do not decline after k is published; and (iii) $\tilde{\theta}_{j_g l}$, $\lambda_{j_g l}$ and ϕ_l are otherwise constant in the absence of k. (i) and (ii) presume that the publicity generated from k does not reduce l's propensity to cite j_g or its authors' beliefs about the value of publishing; (iii) assumes that l's threshold of relevance to cite j_g, awareness of j_g and value of publishing given knowledge of j_g would not have changed had k not existed.

Given these assumptions, the Matthew effect is captured by Equation (5) in Theorem 1. Theorem 1 is proved in Appendix A.

Theorem 1.

Consider i_g, who authors two papers: the high impact paper k and a different (lower impact) paper j_g that was published before k. Let $\underline{\mathcal{L}}, \overline{\mathcal{L}} \subset \mathcal{L}$ denote the Dedekind complete subsets of research ideas that existed before and after k, respectively, and assume:

Assumption 1. *Equation (4) governs l's decision to cite j_g for all $l \in \mathcal{L}$.*

Assumption 2. $\phi'_l(\lambda_{j_g l}) > 0$, $\lambda_{j_g \underline{l}} \leq \lambda_{j_g \bar{l}}$ *and* $\phi_{\underline{l}} \leq \phi_{\bar{l}}$ *for all* $\hat{l} = (\underline{l}, \bar{l})$ *where* $\underline{l} \in \underline{\mathcal{L}}$ *and* $\bar{l} \in \overline{\mathcal{L}}$ *satisfy* $\theta_{j_g \underline{l}} = \theta_{j_g \bar{l}}$ *and* $\tilde{\phi}_{\underline{l}} = \tilde{\phi}_{\bar{l}}$.

Assumption 3. *In the absence of k,* $\tilde{\theta}_{j_g \underline{l}} = \tilde{\theta}_{j_g \bar{l}}$, $\phi_{\underline{l}} = \phi_{\bar{l}}$ *and* $\lambda_{j_g \underline{l}} = \lambda_{j_g \bar{l}}$ *for all* $\hat{l} = (\underline{l}, \bar{l})$.

[15] "Sufficient awareness" of j_g implies that l's authors are aware enough of j_g to: (i) make the logical connection of its relevance to l if there is one (see Kuhn, Younge, and Marco (2023) for a discussion in the context of patent citations); and (ii) have accurately determined that j_g's quality is high enough to warrant citation.

Then $\hat{l} \in \hat{\mathcal{L}}_g$ exists for all $\bar{l} \in \overline{\mathcal{L}}$ and Equation (5) is the Matthew effect for author i_g:

$$\Delta_{jk}^{Matthew} = \sum_{\hat{l} \in \hat{\mathcal{L}}'_g} \mathrm{M}_{j_g l}, \tag{5}$$

where

$$\mathrm{M}_{j_g \hat{l}} = \mathbf{1}\left[\theta_{j_g \bar{l}} \geq \tilde{\theta}_{j_g \bar{l}}\right] - \mathbf{1}\left[\theta_{j_g \underline{l}} \geq \tilde{\theta}_{j_g \underline{l}}\right] \times \mathbf{1}\left[\phi_{\underline{l}}(\lambda_{j_g \underline{l}}) \geq \tilde{\phi}_{\underline{l}}\right] \times \lambda_{j_g \underline{l}}, \tag{6}$$

and $\hat{\mathcal{L}}'_g$ is the subset of $\hat{\mathcal{L}}_g$ where $\lambda_{j_g \bar{l}} = 1$ and $\phi_{\bar{l}}(1) \geq \tilde{\phi}_{\bar{l}}$.

According to Equation (5), a positive Matthew effect is caused by three factors: (i) a decrease in $\tilde{\theta}_{j_g l}$ (*i.e.*, the threshold of relevance); (ii) an increase in $\lambda_{j_g l}$ (*i.e.*, awareness of j_g); or (iii) an increase in ϕ_l, independent of $\lambda_{j_g l}$ (*i.e.*, the publication value of l). (i) captures a rise in strategic citations to j_g after k is published; (ii) and (iii) reflect boosts in citations brought about by better and more widely disseminated information about i's papers—that is, either k made l's authors more aware of j_g ($\lambda_{j_g l}$ increased) or it caused them to positively update their beliefs about the publication value of l (ϕ_l increased, conditional on $\lambda_{j_g l}$), a phenomenon we refer to as "strategic publication".[16]

To compare the Matthew effects of men and women, suppose i_F and i_M co-author k and additionally assume that publishing k renders gender differences in $\lambda_{j_g l}$—and consequently ϕ_l—negligible. Or in other words, the information revealed to l after k is published is sufficient to equalise l's awareness of j_F and j_M and consequently also eliminate all remaining differences in the publication value of l that correlate with the genders of j_F and j_M's authors. When this assumption holds, then the gender difference in the Matthew effect is captured by Equation (7) in Lemma 2.

Lemma 2.

Suppose Theorem 1's assumptions are satisfied for k's co-authors i_F and i_M. Additionally assume:

Assumption 4. $\lambda_{j_F \bar{l}} = \lambda_{j_M \bar{l}}$ *for all* $\bar{l} \in \overline{\mathcal{L}}$.

Then Equation (7) is the gender difference in the Matthew effect:

$$\Delta\Delta_{jk}^{Matthew} = \sum_{\hat{l} \in \hat{\mathcal{L}}'} \left(\mathrm{M}_{j_F \hat{l}} - \mathrm{M}_{j_M \hat{l}}\right), \tag{7}$$

where $\mathrm{M}_{j_g \hat{l}}$ is defined in Equation (6) and $\hat{\mathcal{L}}'$ is the set $\hat{l} = (\underline{l}_M, \bar{l}) \cup (\underline{l}_F, \bar{l})$ for a given $\bar{l}$.

A positive Equation (7) is potentially caused by three factors: (i) a greater propensity to strategically cite j_F after k is published; (ii) a bigger post-k boost in strategic publications among the $l \in \mathcal{L}$ that are more relevant to j_F; or (iii) less familiarity with j_F before k is published, conditional on relevance.

Neither Equations (5) nor (7) specifically identify the individual impacts of strategic citations, strategic publications and better awareness of j_g; however, there are circumstances in which the former two factors are separable from the latter. In particular, suppose we observe citations in an environment where l's authors are aware of j_F and j_M before k is published.[17] A positive Equation (5) is therefore caused by strategic citations or strategic publications by l; a positive Equation (7) suggests a larger rise in strategic citations to j_F or a greater increase in strategic publications among papers that are more relevant to it.

Lemma 3.

Suppose the assumptions in Theorem 1 and Lemma 2 hold, and additionally assume:

[16] More specifically, "strategic publication" is captured by a change in $\phi_l(\lambda_{j_g l})$ after k is published, holding $\lambda_{j_g l}$ fixed. Intuitively, it reflects the narrow desire to produce papers that will publish well.

[17] Lemma 3 actually makes the slightly weaker assumption that l's authors' awareness of j_F and j_M is unaffected by publication of k—that is, either they were aware of both papers before *and* after k was published or they weren't aware of them in either period.

Assumption 5. $\lambda_{j_g\underline{l}} = \lambda_{j_g\bar{l}}$ *for all* $\hat{l} \in \hat{\mathcal{L}}$*.*

Then for all $\hat{l} \in \hat{\mathcal{L}}'$*, Equation* (6) *simplifies to*

$$\mathrm{M}_{j_g\hat{l}} = \mathbf{1}\left[\theta_{j_g\bar{l}} \geq \tilde{\theta}_{j_g\bar{l}}\right] - \mathbf{1}\left[\theta_{j_g\underline{l}} \geq \tilde{\theta}_{j_g\underline{l}}\right] \times \mathbf{1}\left[\phi_{\underline{l}}(1) \geq \tilde{\phi}_{\underline{l}}\right]. \tag{8}$$

When Lemma 3 applies, then potential citing papers are fully aware of j_g before and after k is published. As a result, an increase in $\lambda_{j_g l}$ does not cause the Matthew effect. Instead, it is thanks to a rise in strategic citations or an increase in strategic publications. Similarly, because k's publication does not change l's familiarity with either j_F or j_M, a higher Matthew effect among female authors exclusively reflects a larger rise in strategic behaviour among l more relevant to j_F.

In contrast, suppose Assumption 5 is satisfied and $\Delta^{\text{Matthew}}_{jk} = 0$ and/or $\Delta\Delta^{\text{Matthew}}_{jk} = 0$. In this setting, strategic behaviour by l either plays no role in forming the Matthew effect and/or its gender difference or its component parts perfectly offset one another.

4.2 Estimation strategy

Lemma 3 establishes circumstances in which the Matthew effect exclusively reflects the strategic behaviour of citing papers. To apply it, we estimate Equation (3) in treatment environments where potential citing papers ($l \in \mathcal{L}$) are aware of authors' earlier work ($j_g \in \mathcal{J}_g$) both before and after the latter publish their first high impact paper (k). We further compare them to counterfactual environments where this condition is not met.

The two treatment environments we study are (i) the subset of $\mathcal{J}_g$ cited by k and (ii) the subset of authors who first release k as an NBER working paper; their counterfactual environments are, respectively, (i) the subset of $\mathcal{J}_g$ not cited by k and (ii) the subset of authors who do not release k as an NBER working paper. To maximise the probability that Assumption 5 is satisfied in treatment environments—and ensure comparability with counterfactual environments—we always restrict l and j_g to papers published in economics journals within a six-year window before and after k's publication date.

Our justification for studying the first treatment environment relies on specific details of the publication, working paper and seminar/conference cultures in economics. Conditional on acceptance, the peer review process at top economics journals averages two years or more (Ellison, 2002; Hadavand, Hamermesh, and Wilson, 2024). Partially in response to this, economists almost always publicly release working paper versions of these manuscripts several years before publication; they also widely present pre-published work at seminars and conferences. As a consequence, most other economists with a paper related to k will be aware of it—as well as the papers that it cites—several years before k is actually published. In contrast, if j_g and k are unrelated, then the latter will not cite the former so releasing k as a working paper and/or presenting it at conferences should not, on their own, drive citations to j_g.[18]

The second treatment environment consists of papers by authors who specifically released k as an NBER working paper before publishing it.[19] Not only is the NBER series the most well-known and widest read working paper series in economics, but the papers released in it receive more downloads and abstract views than their published versions (Lusher, Yang, and Carrell, 2023).[20] As a result, k's authors—and consequently j_g—are probably exposed to more attention when k is released as an NBER working paper

[18] This implicitly assumes that publishing in a top-five journal is an important enough event to bring significant attention to authors' previously published work, regardless of its relevance to k. To justify this assumption, we note that economists are willing to give up half a thumb to publish in the *AER* (Attema, Brouwer, and van Exel, 2014)!

[19] See Hengel (2022, Online Appendix G.5) for evidence that almost all NBER working papers eventually accepted by *Econometrica* were submitted to peer review at the same time or before they were released as NBER working papers.

[20] NBER emails over 50,000 subscribers to its newsletter information on new working papers posted to its website; on average, NBER working paper abstract pages are viewed about 650,000 times every month (NBER, 2025).

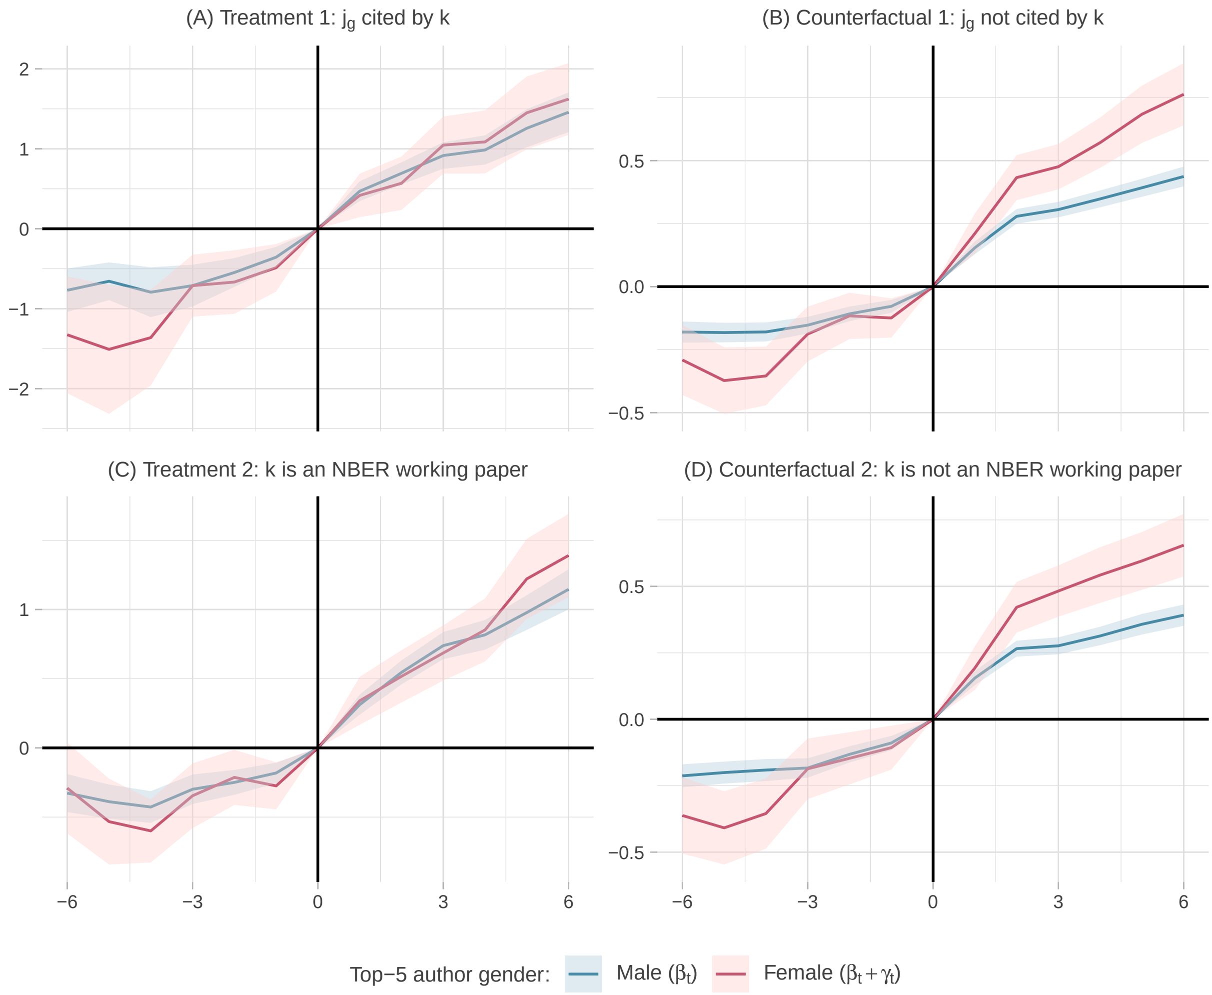

Figure 4: The Matthew effects in treatment and counterfactual environments

Note. Figure 4 displays β_t and γ_t from estimating Equation (3) with OLS and controlling for k fixed effects. Graphs in the first and second columns refer to treatment and counterfactual environments, respectively. Row 1 compares citations to j_g that were (graph (A)) and were not (graph (B)) cited by k. Row 2 plots citations to authors who did (graph (C)) and did not (graph (D)) release k as an NBER working paper. In all graphs, l and j_g are restricted to papers published in economics journals within the 6 years before and after k's publication date. Shaded areas represent 90 percent confidence intervals from standard errors clustered by j_g.

than when it is actually published. In contrast, the j_g associated to ks that were *not* NBER working papers likely receive their biggest publicity shock upon publication of k.

4.3 Results

Figure 4 displays β_t and $\beta_t + \gamma_t$ from estimating Equation (3) on the treatment and counterfactual environments described in Section 4.2. Row one compares citations to j_g that were (graph (A)) and were not (graph (B)) cited by k; row two plots citations to authors who did (graph (C)) and did not (graph (D)) release k as an NBER working paper. Every environment controls for k fixed effects and restricts l and j_g to papers published in economics journals within the six years before at after k's publication date.

In all treatment and counterfactual environments, publishing for the first time in a top-five journal boosts citations to j_g. β_t and $\beta_t + \gamma_t$ are consistently negative for $t < 0$, jump around $t = 0$ and are then positive for all $t > 0$. In treatment environments (first column), citing papers were aware of j_g before k was published, so the bump is evidence that they respond strategically to k's acceptance in a top-five journal (Lemma 3).[21]

However, we find no evidence that this strategic behaviour affects male and female authors differently. When citing papers are already familiar with j_g, then gender differences in the Matthew effect are essentially zero (Figure 4, graphs (A) and (C)). This suggests that the gap's formation when citing papers are *not* familiar with j_g (graphs (B) and (D)) isn't from their strategic behaviour, but instead due to greater awareness of j_F after k is published.[22]

For robustness, we also re-define treatment and counterfactual environments to consist of authors who were and were not recent academic job market stars, on the assumption that potential citing papers should be more aware of stars' earlier work relative to comparable non-stars (Appendix D.5). Again, this is exactly what we observe. In the treatment environment, publishing k boosts citations to j_g, but the effect is similar for both men and women. In contrast, the Matthew effect in the counterfactual environment is noticeably larger for women than it is for men.[23]

To conclude that greater awareness of women's research drives their larger Matthew effect crucially assumes that gender differences in citing papers' strategic behaviour are unrelated to selection into treatment. Available evidence supports this assumption. Tables D.1 and D.2 (Appendix D.1) report gender differences in the characteristics of citing, cited and top-five papers across treatment and counterfactual environments. While there are meaningful gender differences on several dimensions—*e.g.*, journal impact factors and numbers of co-authors—their signs and magnitudes across environments usually match.

Furthermore, we find no evidence that citing papers face additional incentives to behave strategically in response to a woman (vs. a man) publishing in a top-five journal. In Figures 5(A) and E.1 (Appendix E.1) we plot the probability that men and women become editors of any journal, a top-five journal and a top-20 journal in the 10 years before and after k's publication date. In all graphs, there is a noticeable bump in this probability around $t = 0$; however, it does not differ by author gender for any $t > 0$. We also explore the impact publishing k has on promotion to professor (Appendix E.2), the number and dollar amount of grants awarded to authors (Appendix E.3), their employers' reputations (Appendix E.4) and the number of seminar invitations they receive (Appendix E.5). In all instances, we either do not observe a Matthew effect or it does not differ by author gender.

Finally, we additionally: (i) restrict the sample to authors who initially released k as an NBER working paper and set $t = 0$ to the release date (Figure 5(B)); and (ii) drop authors with only one top-five publication and set $t = 0$ to the date the remaining authors' *second* top-five paper was published (Appendix F.1).[24] The first environment mirrors the one in Figure 4(C), except that potential citing papers are no longer necessarily aware of authors' earlier work at $t < 0$; as a result, gender differences in the Matthew effect should re-emerge. The second environment resembles Figure 3, except that potential citing papers are now aware of j_g at $t = 0$; thus, gender differences in the Matthew effect should disappear. Our empirical evidence supports both predictions: when $t = 0$ is k's NBER release date, $\beta_t + \gamma_t$ closely tracks β_t for all $t < 0$, but then surpasses it afterwards; when $t = 0$ refers to authors' second top-five publication date, $\beta_t + \gamma_t$ resembles β_t for all t.

[21] We also reproduce Figure 3 using the Poisson likelihood function (Appendix D.2) and controlling for cited papers' journal impact factors and their authors total pre-k publication counts (Appendix D.4). Results are very similar to those reported in Figure 4.

[22] In Appendix D.2 we reproduce Figure 4, graphs (A) and (B) restricting the treatment and counterfactual samples to articles by authors who did and did not release k as an NBER working paper. We only observe a gender difference in the Matthew effect in the counterfactual environment under conditions where citing papers are not already familiar with j_g before k (namely, when k was not released as an working paper and did not cite j_g).

[23] Relatedly, we also estimate Equation (3) on the subsets of authors employed at institutions ranked 1–5, 6–9, *etc.* (see Appendix D.6). The gender difference in the viability Matthew effect is smallest among authors at the highest ranked institutions, consistent with the hypothesis that their faculty is exposed to more attention compared to faculty at lower ranked institutions.

[24] In (i), j_g are restricted to papers published before k's NBER working paper release date; in (ii), we only include j_g published before an author's first top-five publication date.

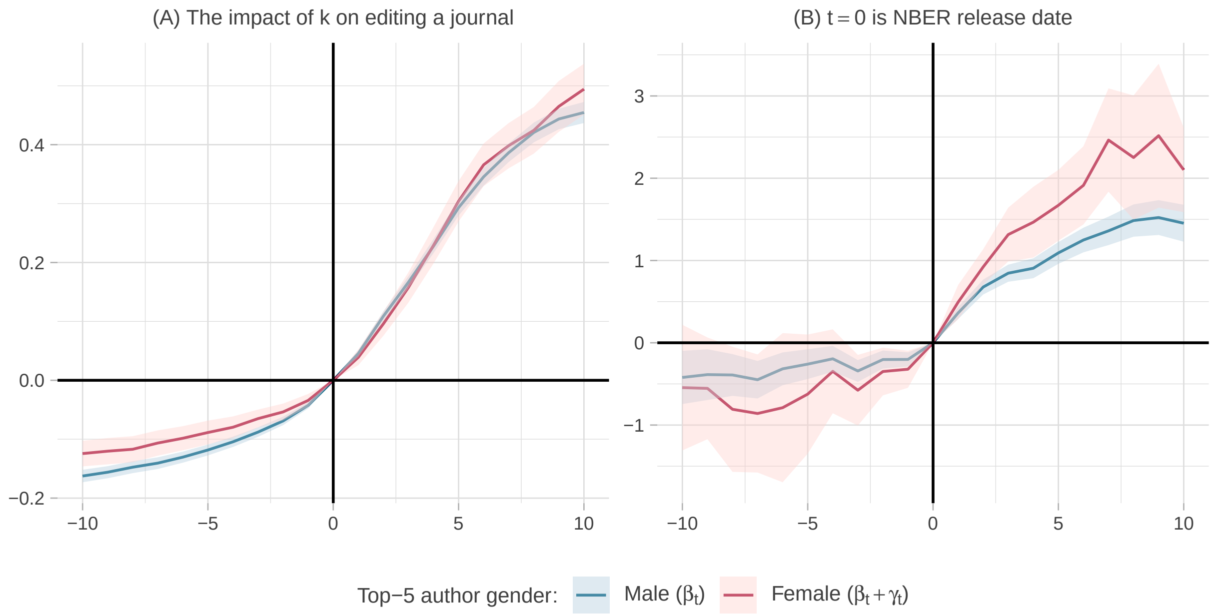

Figure 5: Alternative outcomes and treatments

Note. In Figure 5(A) we plot the probability that men and women become journal editors in the 10 years before and after k's publication date. In Figure 5(B) we restrict the sample to authors who initially released k as an NBER working paper and set $t = 0$ to the date of its release. In both graphs, β_t and γ_t are from estimating Equation (3) with OLS controlling for k fixed effects on the relevant outcome variables. Shaded areas represent 90 percent confidence intervals from standard errors clustered by j_g.

5 Conclusion

In this paper, we investigate whether the Matthew effect mitigates under-recognition of research by women. To identify the effect, we adopt an event study approach that estimates citations to authors' previously published papers around the date they first publish in a top-five journal, controlling for top-five paper fixed effects. Our results suggest that male and female co-authors on the same top-five paper experience a large bump in citations to their earlier work around the event date. However, the effect is much larger for women than it is for men: before their first top-five publications, women's research is cited less than research by their future male co-authors; afterwards, it is cited more.

To better understand the mechanisms driving our results, we construct a straightforward model of the decision to cite early research by authors who later publish a high impact paper. According to our model, the Matthew effect and its gender difference are caused by increases in: (i) the information that potential citing papers have about the existence and relevance of an author's pre-top-five research; (ii) the propensity to strategically cite this research; or (iii) the strategic value of publishing papers related to it.

Although we cannot separately identify (i), (ii) and (iii), we *can* isolate the combined impact of strategic citations and strategic publications by restricting the sample of citing papers to those that were always aware of authors' pre-top-five research. Intuitively, suppose citing papers are familiar with an author's early work before he publishes for the first time in a top-five journal, but they only cite it afterwards. Because they were aware of this research in both periods, their second period citations are necessarily motivated by strategic considerations related to the top-five publication (*e.g.*, a desire to cite well-known authors).

We apply this strategy to several plausible "treatment" environments—*i.e.*, environments in which citing papers were already aware of authors' previous work before they published in top-five journals—and compare them to their relevant counterfactual environments—*i.e.*, otherwise identical environments in which this condition is not met. In all treatment environments, we find clear evidence of a Matthew effect—suggesting that strategic considerations play a role in its general formation—although it does not differ by author gender. In contrast, women's Matthew effect is always larger than men's in counterfactual environments. Combined, these results imply that a lack of information among potential citing papers is a necessary condition for women to enjoy a larger effect; when awareness of their research is already high, the Matthew effect confers women no advantage relative to comparable men.

Unfortunately, our data cannot precisely identify *why* citing papers are initially under-informed about women's research. Gender differences persist despite accounting for field, number of co-authors, institutional rank, journal impact factors and authors' research productivity. Yet there are many remaining channels. Citing authors may under-estimate women's research until faced with significant evidence of its quality.[25] Alternatively, institutions may be more likely to promote the work of men compared to work by similarly credentialed women. Men probably also enjoy stronger ties to authors of related work (see, *e.g.*, Ductor, Goyal, and Prummer, 2023).

More generally, our results suggest that bringing attention to women's research may help overcome earlier tendencies to overlook it. We caution, however, that promotion activities alone may be insufficient to achieve a similar effect. In all of the environments we study, attention shocks are accompanied by credible signals that the publicised work is high quality—*e.g.*, because it was cited in another high quality paper or authored by a job market star, researcher closely connected to the NBER, or economist published in a top-five journal. Indeed, evidence from the NBER suggests that promoting more papers *without* making an additional effort to credibly signal their quality may lead to overcrowding and reduce the attention specific studies receive (Lusher, Yang, and Carrell, 2023).

In addition to spotlighting the role better information plays in gendered citation gaps, our results also highlight the dramatic changes they undergo as authors' circumstances evolve. For these reasons, we believe our evidence calls for extra caution when interpreting gender differences in citation counts among less prominent authors.

By failing to appropriately recognise research by women, we not only hinder women's careers, we also rob ourselves of the future innovations their ideas could have sparked.[26] We hope our study brings us closer to a future where women's ideas have the same chances to succeed on their merits as men's.

References

Alexander, Diane, Olga Gorelkina, Erin Hengel, and Richard Tol (2023). "Gender and the time cost of peer review". Mimeo.

Attema, Arthur E., Werner B. F. Brouwer, and Job van Exel (2014). "Your right arm for a publication in AER?" *Economic Inquiry* 52 (1), pp. 495–502.

Azoulay, Pierre, Toby Stuart, and Yanbo Wang (2014). "Matthew: effect or fable?" *Management Science* 60 (1), pp. 92–109.

Bikard, Michaël, Isabel Fernandez-Mateo, and Ronak Mogra (2025). "Standing on the shoulders of (male) giants: gender inequality and the technological impact of scientific ideas". *Administrative Science quarterly* (forthcoming).

Bohren, J. Aislinn, Alex Imas, and Michael Rosenberg (2019). "The dynamics of discrimination: theory and evidence". *American Economic Review* 109 (10), pp. 3395–3436.

[25] Relatedly, citing papers may rely on a paper's correspondence with some unobserved category to proxy for relevance and/or quality. If female economists are more likely to produce manuscripts that fall outside these category boundaries, then potential citing papers may penalise them until their quality is well established.

[26] Although it is impossible to know how science would have progressed had women's ideas always been fairly recognised, one obvious repercussion has been the evolution of a body of medical research that understudies health issues specific to women (for a discussion, see Holdcroft, 2007).

Bol, Thijs, Mathijs de Vaan, and Arnout van de Rijt (2018). “The Matthew effect in science funding”. *PNAS* 115 (9), pp. 4887–4890.

Bornmann, Lutz and Hans-Dieter Daniel (2007). “Multiple publication on a single research study: does it pay? The influence of number of research articles on total citation counts in biomedicine”. *Journal of the American Society for Information Science and Technology* 58 (8), pp. 1100–1107.

Card, David, Stefano DellaVigna, Patricia Funk, and Nagore Iriberri (2020). “Are referees and editors in economics gender neutral?” *Quarterly Journal of Economics* 135 (1), pp. 269–327.

Clarivate (2022). *Web of Science [database]*. Data accessed and downloaded July 2022.

Combes, Pierre-Philippe and Laurent Linnemer (2010). “Inferring missing citations: a quantitative multi-criteria ranking of all journals in economics”. GREQAM Working Paper no. 2010–28.

Dion, Michelle L., Jane Lawrence Sumner, and Sara McLaughlin Mitchell (2018). “Gendered citation patterns across political science and social science methodology fields”. *Political Analysis* 26 (3), pp. 312–327.

Ductor, Lorenzo, Sanjeev Goyal, and Anja Prummer (2023). “Gender and collaboration”. *Review of Economics and Statistics* 105 (6), pp. 1366–1378.

Dworkin, Jordan D., Kristin A. Linn, Erin G. Teich, Perry Zurn, Russell T. Shinohara, and Danielle S. Bassett (2020). “The extent and drivers of gender imbalance in neuroscience reference lists”. *Nature Neuroscience* 23 (8), pp. 918–926.

Ellison, Glenn (2002). “The slowdown of the economics publishing process”. *Journal of Political Economy* 110 (5), pp. 947–993.

Ferber, Marianne A. (1986). “Citations: are they an objective measure of scholarly merit?” *Signs* 11 (2), pp. 381–389.

– (1988). “Citations and networking”. *Gender and Society* 2 (1), pp. 82–89.

Franzoni, Chiara, Giuseppe Scellato, and Paula Stephan (2011). “Changing incentives to publish”. *Science* 333 (6043), pp. 702–703.

Hadavand, Aboozar, Daniel S. Hamermesh, and Wesley Wilson (2024). “Publishing economics: how slow? Why slow? Is slow productive? How to fix slow?” *Journal of Economic Literature* 62 (1), pp. 269–293.

Heckman, James J. and Sidharth Moktan (2020). “Publishing and promotion in economics: the tyranny of the Top Five”. 58 (2), pp. 419–470.

Hengel, Erin (2022). “Are women held to higher standards? Evidence from peer review”. *The Economic Journal* 132 (648), pp. 365–381.

Hengel, Erin and Euyoung Moon (2023). “Gender and equality at top economics journals”. Mimeo.

Holdcroft, Anita (2007). “Gender bias in research: how does it affect evidence based medicine?” *Journal of the Royal Society of Medicine* 100, pp. 2–3.

Koffi, Marlène (2025). “Innovative ideas and gender (in)equality”. *American Economic Review* (forthcoming).

Krawczyk, Michał and Magdalena Smyk (2016). “Author’s gender affects rating of academic articles: evidence from an incentivized, deception-free laboratory experiment”. *European Economic Review* 90, pp. 326–335.

Kuhn, Jeffrey, Kenneth Younge, and Alan Marco (2023). “Strategic citation: a reassessment”. *Review of Economics and Statistics* 105 (2), pp. 458–466.

Lampe, Ryan (2012). “Strategic citation”. *Review of Economics and Statistics* 94 (1), pp. 320–333.

Larivière, Vincent, Chaoqun Ni, Yves Gingras, Blaise Cronin, and Cassidy R. Sugimoto (2013). “Global gender disparities in science”. *Nature* 504 (7479), pp. 211–213.

Lundberg, Shelly and Jenna Stearns (2019). *Journal of Economic Perspectives* 33 (1), pp. 3–22.

Lusher, Lester R., Winnie Yang, and Scott E. Carrell (2023). “Congestion on the information superhighway: does economics have a working papers problem?” *Journal of Public Economics* 225, p. 104978.

Merton, Robert K. (1968). “The Matthew Effect in science: the reward and communication systems of science are considered”. *Science* 159 (3810), pp. 56–63.

NBER (2025). *About the NBER*. URL: %5Curl%7Bhttps://www.nber.org/about-nber%7D.

Perc, Matjaž (2014). “The Matthew effect in empirical data”. *Interface* 11 (98), p. 20140378.

Royal Statistical Society (2025). *EPSRC grant funding: statistical analysis of diversity in the portfolio and peer review*. Tech. rep.

Rubin, Amir and Eran Rubin (2021). “Systematic bias in the progress of research”. *Journal of Political Economy* 129 (9), pp. 2666–2719.

Sarsons, Heather, Klarita Gërxhani, Ernesto Reuben, and Arthur Schram (2021). “Gender differences in recognition for group work”. *Journal of Political Economy* 129 (1), pp. 101–147.

Seeber, Marco, Mattia Cattaneo, Michele Meoli, and Paolo Malighetti (2019). “Self-citations as strategic response to the use of metrics for career decisions”. *Research Policy* 48 (2), pp. 478–491.

Siler, Kyle and Vincent Larivière (2022). “Who games metrics and rankings? Institutional niches and journal impact factor inflation”. *Research Policy* 51 (10), p. 104608.

Simcoe, Timothy S. and Dave M. Waguespack (2011). “Status, quality, and attention: what’s in a missing name?” *Management Science* 57 (2), pp. 274–290.

Teich, Erin G. *et al.* (2022). “Citation equity and gendered citation practices in contemporary physics”. *Nature* 18 (10), pp. 1161–1170.

Appendices

A Proofs

Proof of Theorem 1. Due to the Dedekind completeness of $\underline{\mathcal{L}}$ and $\overline{\mathcal{L}}$, there exists an isomorphism that matches each $\bar{l} \in \overline{\mathcal{L}}$ to an $\underline{l} \in \underline{\mathcal{L}}$ that shares the same $\theta_{j_g l}$ and $\tilde{\phi}_l$. Let $\hat{l} \in \hat{\mathcal{L}}_g$ denote the resulting ordered pair $\hat{l} = (\underline{l}, \bar{l})$.

The difference in citations between $\underline{l}$ and $\bar{l}$ in the ordered pair $\hat{l}$ is:

$$\sum_{\hat{l} \in \hat{\mathcal{L}}_g} \mathrm{M}_{j_g \hat{l}}, \tag{A.1}$$

where

$$\mathrm{M}_{j_g \hat{l}} = \mathbf{1}\left[\theta_{j_g \bar{l}} \geq \tilde{\theta}_{j_g \bar{l}}\right] \times \mathbf{1}\left[\phi_{\bar{l}}(\lambda_{j_g \bar{l}}) \geq \tilde{\phi}_{\bar{l}}\right] \times \lambda_{j_g \bar{l}} - \mathbf{1}\left[\theta_{j_g \underline{l}} \geq \tilde{\theta}_{j_g \underline{l}}\right] \times \mathbf{1}\left[\phi_{\underline{l}}(\lambda_{j_g \underline{l}}) \geq \tilde{\phi}_{\underline{l}}\right] \times \lambda_{j_g \underline{l}}. \tag{A.2}$$

Suppose k had never existed. By Assumption 2, $\lambda_{j_g \bar{l}} = \lambda_{j_g \underline{l}}$, $\phi_{\underline{l}} = \phi_{\bar{l}}$ and $\tilde{\theta}_{j_g \underline{l}} = \tilde{\theta}_{j_g \bar{l}}$, so Equation (A.1) is zero.[1] Because Equation (A.1) is only non-zero if k is published, it represents the visibility Matthew effect for author i_g from publishing k.

It remains to show that Equation (A.1) is equivalent to Equation (5). Define $\hat{\mathcal{L}}'_g$ as the subset of $\hat{\mathcal{L}}_g$ where $\lambda_{j_g \bar{l}} = 1$ and $\phi_{\bar{l}}(1) \geq \tilde{\phi}_{\bar{l}}$. Then, for all $\hat{l} \in \hat{\mathcal{L}}'_g$, Equation (A.2) simplifies to Equation (6).

Now consider any $\hat{l} \in \hat{\mathcal{L}}_g \setminus \hat{\mathcal{L}}'_g$. Because $\hat{l} \notin \hat{\mathcal{L}}'_g$, either $\lambda_{j_g \bar{l}} = 0$ or $\phi_{\bar{l}}(1) < \tilde{\phi}_{\bar{l}}$. Suppose first $\lambda_{j_g \bar{l}} = 0$. By assumption, $\lambda_{j_g \underline{l}} \leq \lambda_{j_g \bar{l}} = 0$, thus $\lambda_{j_g \underline{l}} = 0$, as well. As a result, Equation (A.2) is zero.

Suppose now $\lambda_{j_g \bar{l}} = 1$ but $\phi_{\bar{l}}(1) < \tilde{\phi}_{\bar{l}}$. Thus, $\mathbf{1}[\phi_{\bar{l}}(1) \geq \tilde{\phi}_{\bar{l}}] = 0$. Suppose $\lambda_{j_g \underline{l}} = 1$. Because $\phi_{\underline{l}}(1) \leq \phi_{\bar{l}}(1) < \tilde{\phi}_{\bar{l}} = \tilde{\phi}_{\underline{l}}$ (Assumption 2 and the fact that $\tilde{\phi}_{\underline{l}} = \tilde{\phi}_{\bar{l}}$ for all $\hat{l} \in \hat{\mathcal{L}}_g$), $\mathbf{1}[\phi_{\underline{l}}(1) \geq \tilde{\phi}_{\underline{l}}] = 0$ so Equation (A.2) is zero. Suppose $\lambda_{j_g \underline{l}} = 0$. Because $\phi'_{\underline{l}}(\lambda) > 0$ and $\phi_{\underline{l}}(1) < \phi_{\bar{l}}(1)$ (Assumption 2), $\phi_{\underline{l}}(0) \leq \phi_{\underline{l}}(1) \leq \phi_{\bar{l}}(1)$ so $\phi_{\underline{l}}(0) < \tilde{\phi}_{\underline{l}}$ and $\mathbf{1}[\phi_{\underline{l}}(0) \geq \tilde{\phi}_{\underline{l}}] = 0$; hence, Equation (A.2) is zero. Thus, Equation (A.2) is zero for all $\hat{l} \in \hat{\mathcal{L}}_g \setminus \hat{\mathcal{L}}'_g$.

Because (A.2) is equal to Equation (6) for all $\hat{l} \in \hat{\mathcal{L}}'_g$ and zero for all $\hat{l} \in \hat{\mathcal{L}}_g \setminus \hat{\mathcal{L}}'_g$, Equation (A.1) is equivalent to Equation (5). Thus, all is proved. □

Proof of Lemma 2. From Theorem 1, $\hat{\mathcal{L}}_F$ and $\hat{\mathcal{L}}_M$ exist for all $\bar{l} \in \overline{\mathcal{L}}$. Thus, for any $\hat{l}_F = (\underline{l}_F, \bar{l}_F)$ in $\hat{\mathcal{L}}_F$ there exists a corresponding $\hat{l}_M = (\underline{l}_M, \bar{l}_M)$ in $\hat{\mathcal{L}}_M$ where $\bar{l}_F = \bar{l}_M$ (and vice versa).

It remains to show that $\bar{l}_F$ exists in some $\hat{l}_F \in \hat{\mathcal{L}}'_F$ if and only if it also exists in some $\hat{l}_M \in \hat{\mathcal{L}}'_M$. Without loss of generality, consider any $\hat{l}_F \in \hat{\mathcal{L}}'_F$. By the definition of $\hat{\mathcal{L}}'_F$, $\lambda_{j_F \bar{l}_F} = 1$ and $\phi_{\bar{l}_F}(1) \geq \phi_{\bar{l}_F}$. From Assumption 4, $\lambda_{j_F \bar{l}_F} = \lambda_{j_M \bar{l}_F} = 1$. Thus, $\bar{l}_F$ must be in some $\hat{l}_M \in \hat{\mathcal{L}}'_M$, as well. An analogous proof similarly establishes that for any $\hat{l}_M \in \hat{\mathcal{L}}'_M$, $\bar{l}_M$ must also exist in some $\hat{l}_F \in \hat{\mathcal{L}}'_F$. This concludes the proof.[2] □

Proof of Lemma 3. From Assumption 5, $\lambda_{j_g \underline{l}} = \lambda_{j_g \bar{l}} = 1$ for all $\hat{l} \in \hat{\mathcal{L}}_g$. Thus, Equation (6) simplifies to Equation (8), as required. □

[1] Recall that $\theta_{j_g \bar{l}} = \theta_{j_g \underline{l}}$ by definition and $\phi_{\bar{l}}(\lambda) = \phi_{\underline{l}}(\lambda)$ by assumption.

[2] Note that it is not necessarily true that $\underline{l}_F = \underline{l}_M$ for the $\hat{l}_F$ and $\hat{l}_M$ that share the same $\bar{l}$.

B Data

B.1 Additional data descriptions

Pre-top-five publications. Our first database contains the entire early publication histories of every economist who, between 1986–2015, published (for the first time) a full-length, original research article in one of the following "top-five" economics journals: the *American Economic Review*, *Econometrica*, *Journal of Political Economy*, *Quarterly Journal of Economics* and *Review of Economic Studies*. We follow Hengel and Moon (2023) and define "full-length, original research" as regular issue articles published with abstracts and/or submit-accept dates. This definition excludes most book reviews, editorials, reports and non-peer-reviewed essays; it also excludes conference proceedings such as the May issue of the *American Economic Review*.

Our final database contains 3,897 authors. We manually assigned each a gender based on (i) obviously gendered given names (*e.g.*, "James" or "Brenda"); (ii) photographs on personal or faculty websites; (iii) personal pronouns used in text written about the individual; and (iv) by contacting the author himself or people and institutions connected to him.

To obtain pre-top-five publication histories, we first matched authors to their Web of Science records and downloaded the lists of publications attributed to them; with the help of a research assistant, we subsequently manually verified and corrected these records. Our final sample of pre-top-five publications contains 29,406 articles (roughly 8 articles per person). Data on these articles and the 1,806,675 articles that cite them—including journal, publication date, number of authors, category, language, manuscript type (*e.g.*, article or book review) and journal impact factor—are from Web of Science (Clarivate, 2022).

Authors' career trajectories. Our second dataset contains digitised data from authors' CVs for 4,335 authors with a top-five paper included in the database analysed in Hengel (2022). CVs were collected by googling all 8,187 authors analysed in Hengel (2022) and downloading (or creating) a PDF of their CV if one could be found. With the help of a research assistant, we then digitised the following information from each CV: education, employment, editorial positions, awards, seminars and personal information (citizenship, date of birth, children and languages spoken). Among authors with a CV, 3,037 satisfy the condition for inclusion in our first database—namely, that they published their first full-length original research article in a top-five journal between 1986–2015. Unless otherwise mentioned, this is the sample of authors analysed in Sections 2 and 4 and Appendix E.

B.2 Additional summary statistics

B.2.1 Pre-k paper counts by *JEL* code

In Figure B.1 we present average paper counts by *JEL* code.[3] While there are noticeable differences in publication counts across fields, we observe no meaningful gender differences.

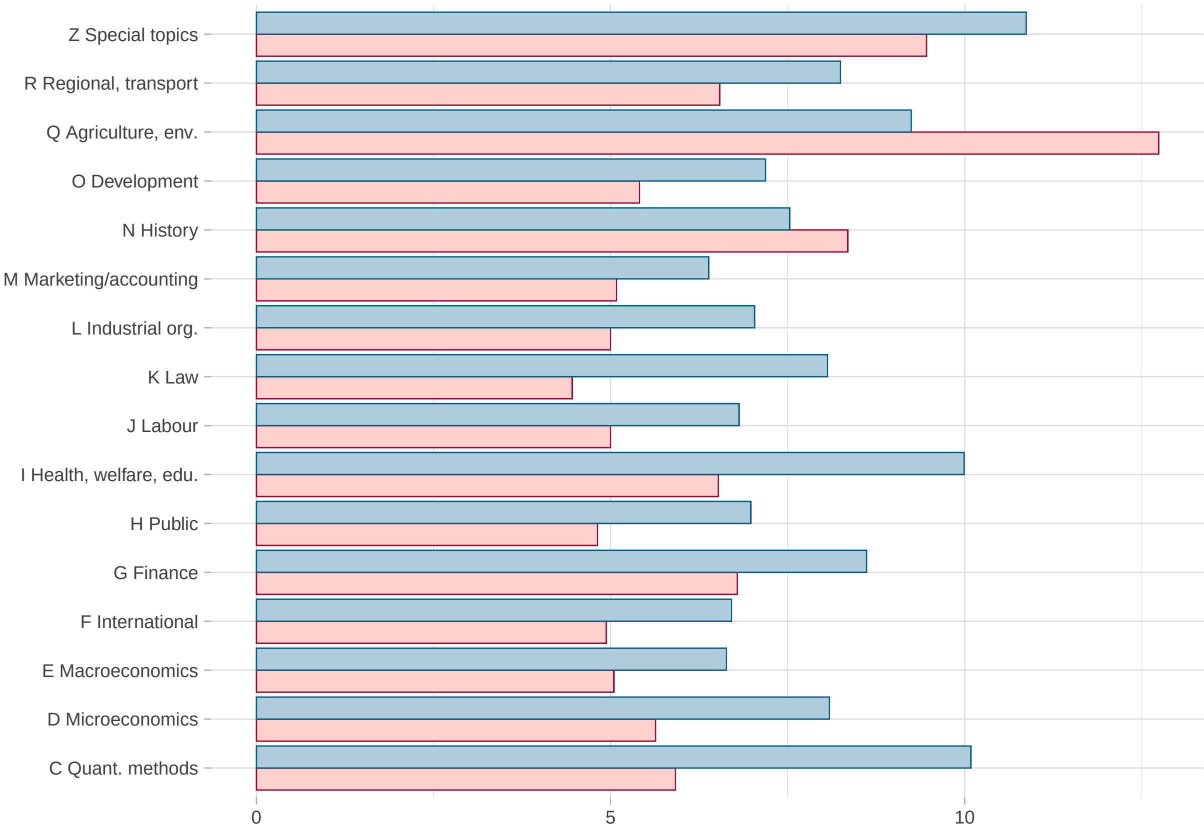

Figure B.1: Avg. no. pre-top-5 papers, by *JEL* code

Note. Figure displays the average number of pre-top-five paper counts for authors by top-five paper *JEL* code.

[3]Figure B.1 excludes *JEL* codes A, B and P due to small samples, particularly of female authors.

B.2.2 Grants before and after first top-five

The left-hand-side graph in Figure B.2 plots the percentage of authors with a grant before and after they publish their first top-five paper, across employer rank. The right-hand-side graph similarly plots the amount awarded (in thousand USD). Institution refers to the author's highest ranked institutions among all institutions listed on his first top-five paper. Institutions are annually ranked (in descending order) by the number of top-five articles affiliated to them, smoothed over a five-year period.

Unlike graphs (G) and (H) in Figure 2, neither graph in Figure B.2 suggests an obvious pattern in the gender composition of grants awarded before and after an author's first top-five paper.

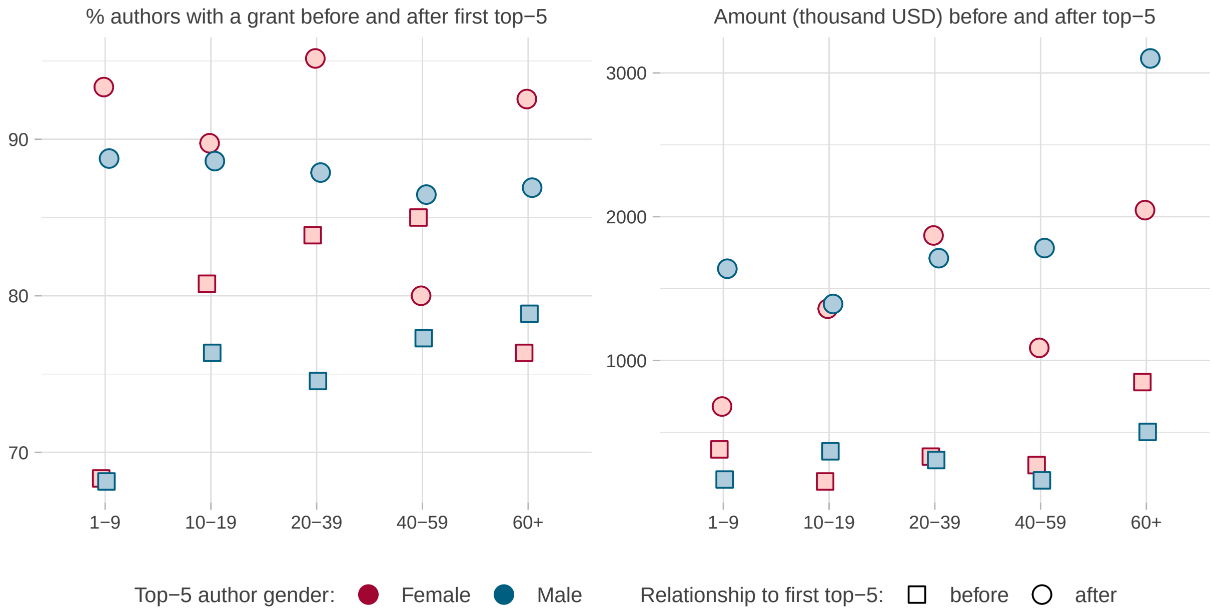

Figure B.2: Grants to authors, before and after their first top-five paper

Note. Left-hand-side graph plots the percentage of each gender's authors who were awarded a grant before and after they published their first top-five paper. Right-hand-side graph plots the amount awarded (in thousand USD).

B.2.3 Years since PhD

In Figure B.3 we show the distribution of the distance (in years) between an author's first top-five publication date and the date he obtained a PhD (graph (A)) and by year of top-five publication (graph (B)). PhD year refers to the year an author obtained his first PhD.[4] Estimates in graph (B) are shown as 5-year moving averages.

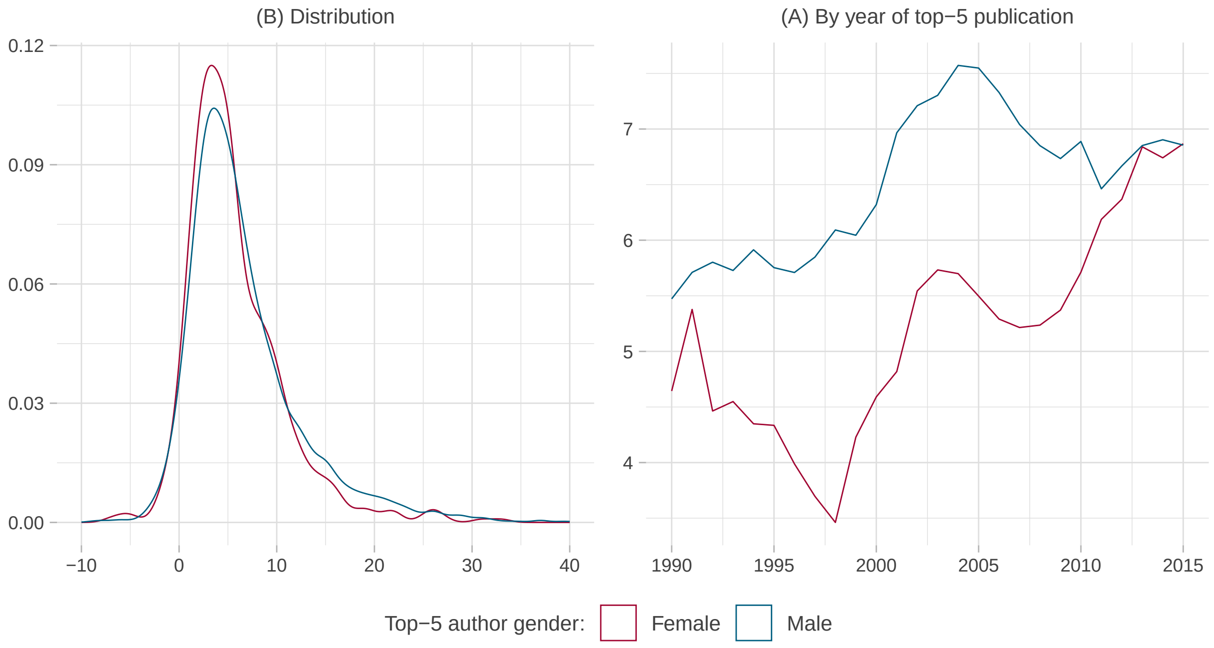

Figure B.3: Years between PhD and first top-five

Note. Figure shows the distance (in years) between the date an author published his first top-five paper and the date an author obtained his PhD. Graph (A) is the distribution of this distance for men and women; graph (B) plots their average distances over time.

On average, women publish their first top-five paper 5 years after obtaining their PhD, while men publish 6 years after obtaining their PhD (graph (A)). However, this gap has completely closed in recent years (graph (B)), suggesting that women's publication records—or at least the publication records of very prominent women—are starting to look more like the publication records of comparable men.

[4]Some authors in our data have obtained multiple PhDs.

B.2.4 Distribution of citations to j_g

In Figure B.4 we plot the distribution of total citation counts to cited papers (graph (A)) and the average number of citations papers receive in the years following their publication (graph (B)). Graph (A) suggests that male-authored j_g are disproportionately likely to receive 0 or only a few citations; female-authored j_g are more commonly found in the left-tail of the distribution. In the first few years after publication, graph (B) suggests that women's papers are cited more than men's papers. However, that pattern reverses around the 25 year mark, likely because men have disproportionately authored the small number of highly cited papers that continue to be cited year-after-year.

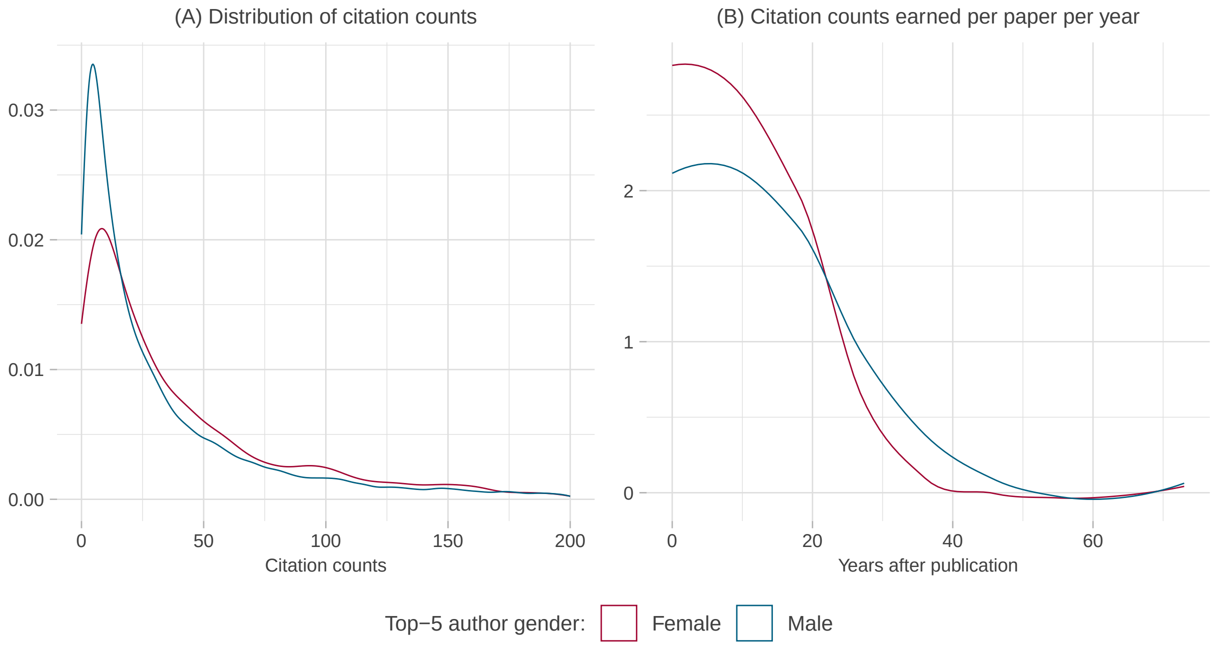

Figure B.4: Distribution of citations to j_g

Note. Graph (A) is the distribution of total citation counts to cited papers. Graph (B) is the average number of per paper citations earned in each year following its publication.

B.2.5 Distribution of authors by pre-k publication count

In Figure B.5 we plot the distribution of authors with 0, 1, 2, *etc.* pre-k publications. Women are disproportionately more likely to have had zero or only a few pre-top-five papers published before k was published. Men are more common in the left-tail of the distribution.[5]

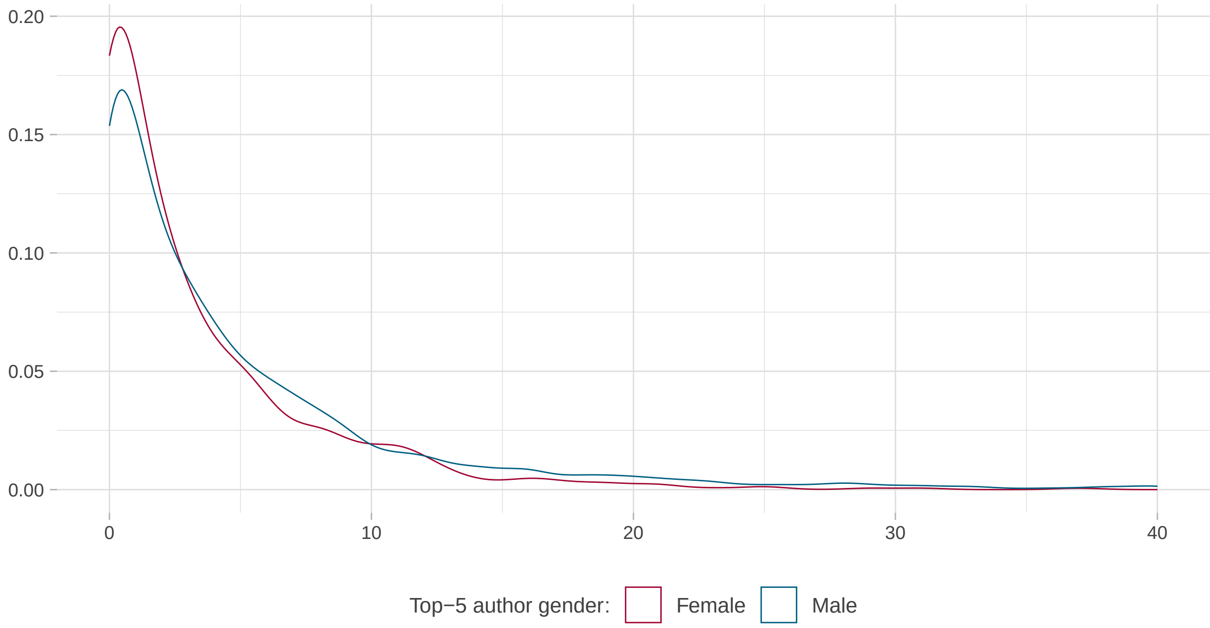

Figure B.5: Distribution of authors by pre-k publication count

Note. Figure shows the distribution of authors with 0, 1, 2, *etc.* pre-k publications.

[5] A small number of authors have a very high pre-top-five publication count (*e.g.*, on author published 535 pre-k papers). These are largely non-economists in fields like medicine and physics.

C Section 3.3, robustness

C.1 Poisson likelihood function

In Figure C.1, we reproduce Figure 3, but estimate Equation (3) using a Poisson likelihood function. Results are very similar to those shown in Figure 3.

Figure C.1: Figure 3, Poisson likelihood function

Note. Figure C.1 replicates Figure 3 but it estimates Equation (3) using a Poisson likelihood function. Shaded areas reflect 90 percent confidence intervals from standard errors clustered by j_g.

C.2 Extensive margin

In Figure C.2, we reproduce Figure 3 but replace the dependent variable in Equation (3) with a dummy variable equal to one if the paper was cited at all at time t (and zero otherwise). We find a clear Matthew effect around $t = 0$ as well as a larger effect for women than men. However, the propensity to cite papers on the extensive margin declines more rapidly, suggesting that the sustained higher citation counts for all $t > 0$ observed in Figure 3 are driven by a small number of papers.

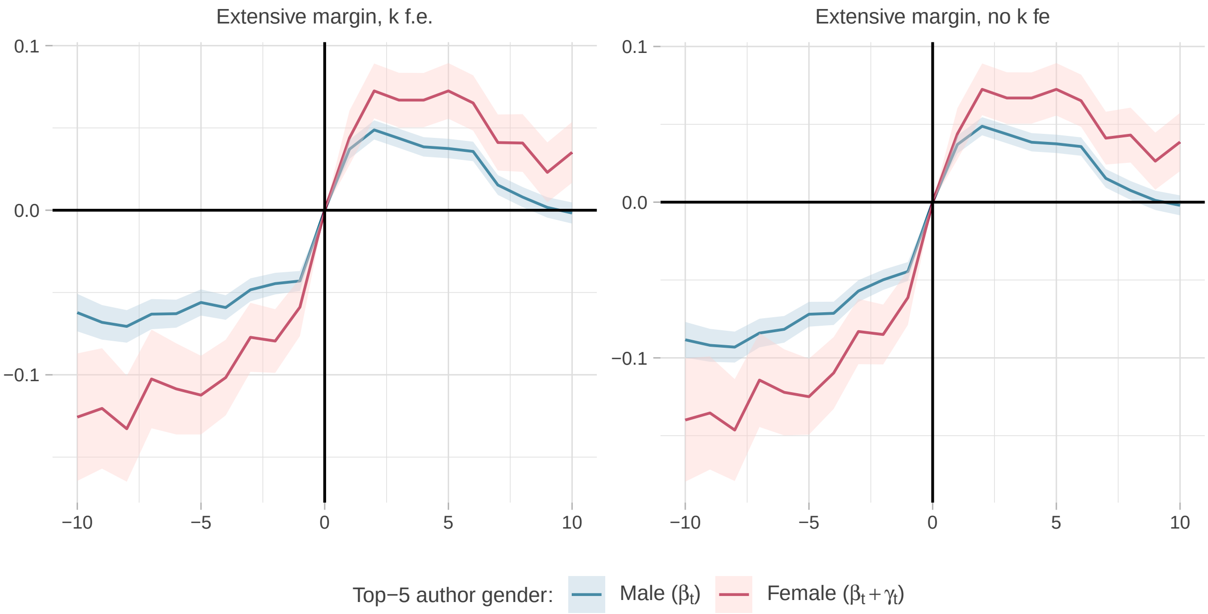

Figure C.2: Figure 3, extensive margin

Note. Figure C.1 replicates Figure 3 but it replaces the dependent variable in Equation (3) with a dummy variable equal to one if the paper was cited at time t. Shaded areas reflect 90 percent confidence intervals from standard errors clustered by j_g.

C.3 j_g published in economics vs. non-economics journals

In Figure C.3 we reproduce Figure 3 but restrict the sample to cited papers (j_g) published in economics (top row) and non-economics journals (bottom row). The first column shows graphs controlling for shock$_k$ using fixed effects for each top-five article k; the second column displays graphs that instead control for k's year and journal of publication, co-author count, citations (asinh) and author prominence.

Figure C.3 suggests that the Matthew effect and its gender difference originally observed in Figure 3 are present among citations to papers published in both economics and non-economics journals.

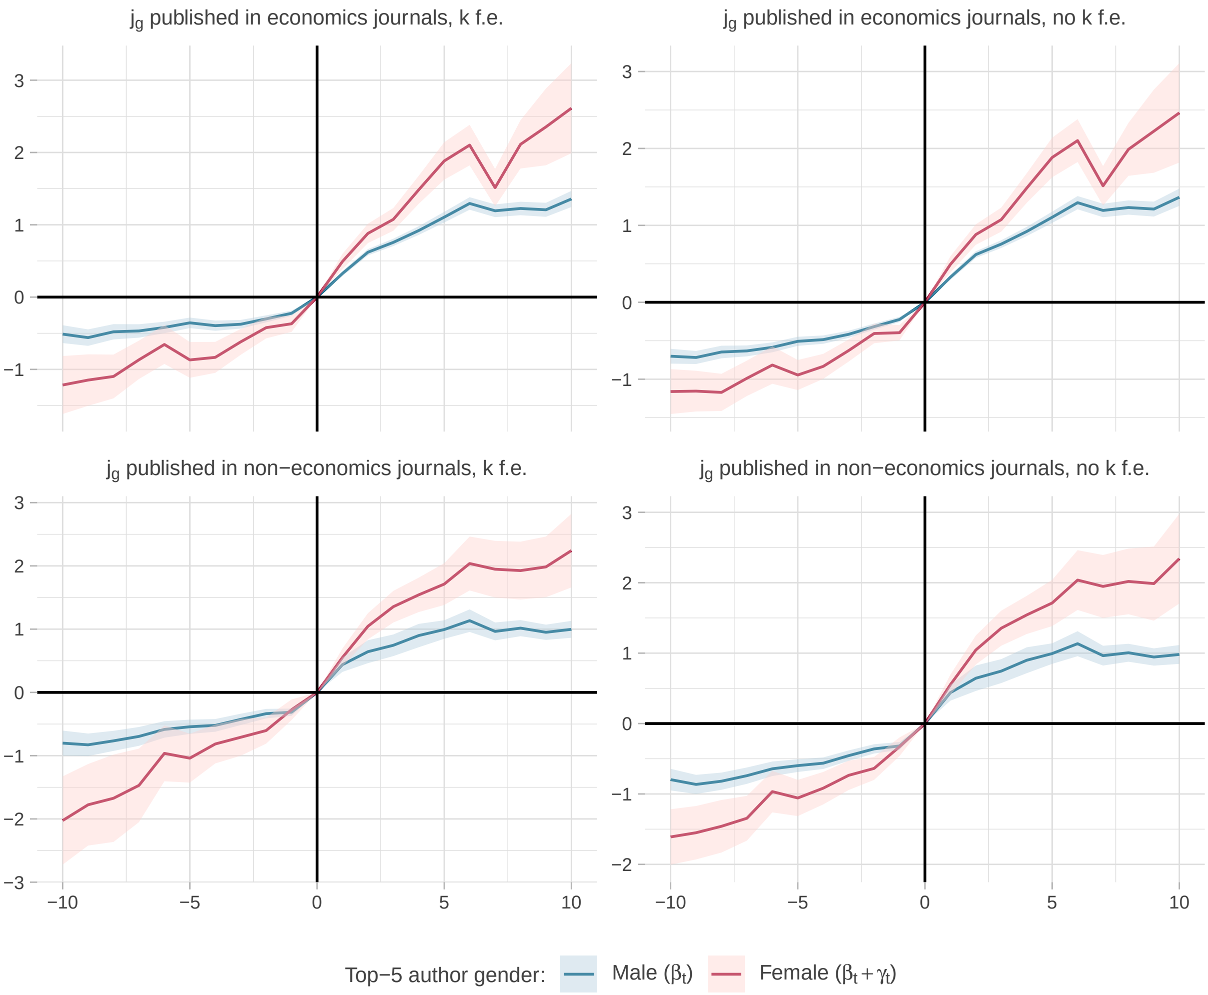

Figure C.3: Figure 3, j_g published in economics vs. non-economics journals

Note. Figure C.3 reproduces Figure 3 on the sample of cited papers published in economics journals (top row) and non-economics journals (bottom row).

C.4 l published in economics vs. non-economics journals

In Figure C.4 we reproduces Figure 3 but restrict the sample to citations (l) from papers published in economics journals (top row) and non-economics journals (bottom row). As in Figure 3, the first column shows results controlling for shock_k using fixed effects for each top-five article; the second column instead shows results controlling for shock_k using k's year and journal of publication, co-author count, citations (asinh) and author prominence.

Again, we see clear evidence of a Matthew effect and its gender difference in both subsets of citations.[6]

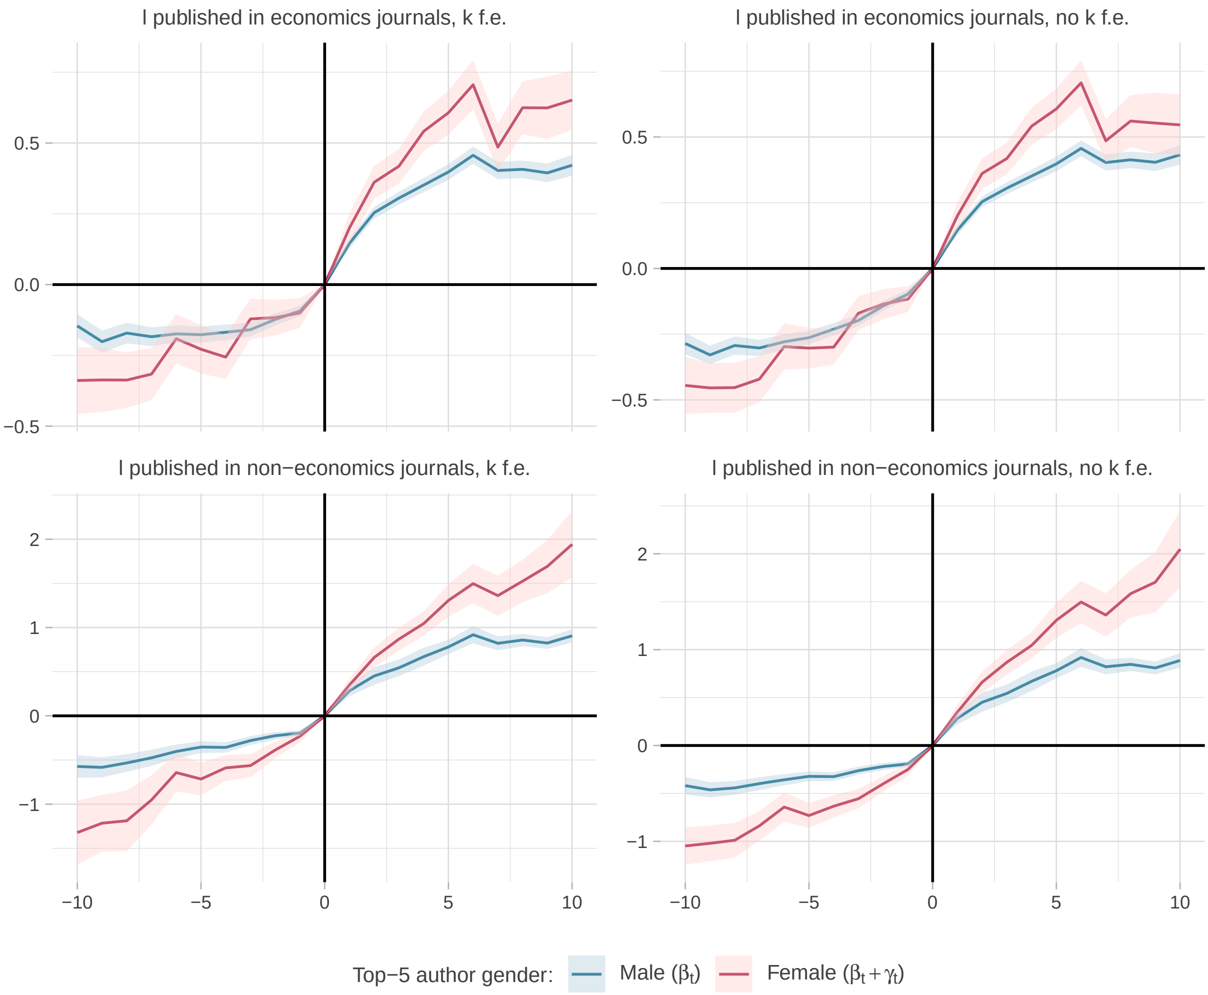

Figure C.4: Figure 3, l published in economics vs. non-economics journals

Note. Figure C.4 reproduces Figure 3 on the sample of citing papers published in economics journals (top row) and non-economics journals (bottom row).

[6]Both the Matthew effect and its gender difference are smaller among citations from papers published in economics journals. This is because there are fewer citations from these papers at any given t compared to citations from papers published in non-economics journals. (Citations from papers published in economics journals only represent about 29 percent of total citations to a paper.)

C.5 j_g and l published in economics vs. non-economics journals

In Figure C.5 we reproduce Figure 3 among the sample of cited *and* citing papers published in economics (top row) and non-economics (bottom row) journals, respectively. Results are very similar to those shown in Figures C.3 and C.4.

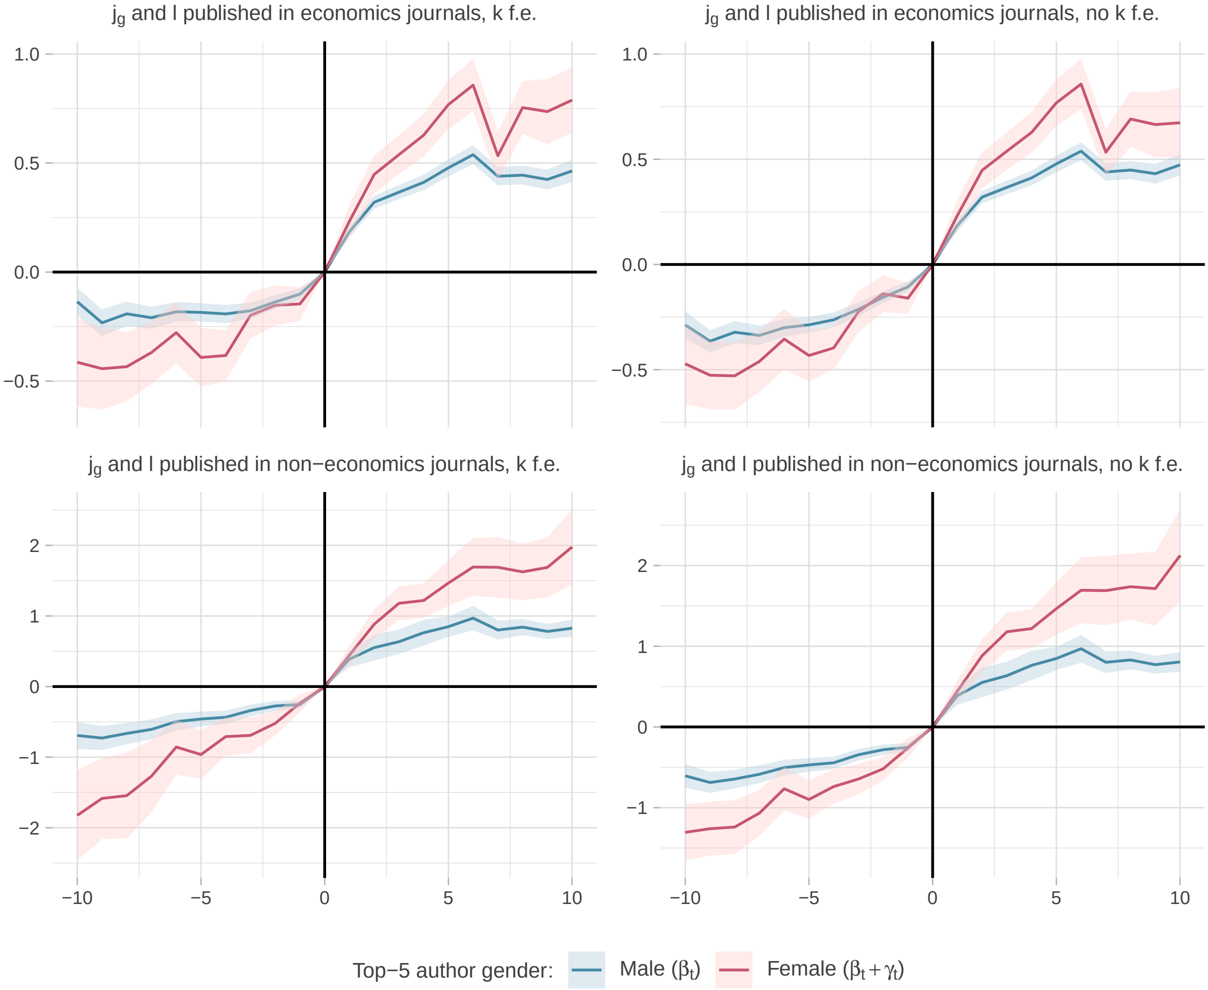

Figure C.5: Figure 3, j_g and l published in economics vs. non-economics journals

Note. Figure C.5 reproduces Figure 3 on the sample of cited and citing papers published in economics journals (top row) and non-economics journals (bottom row).

C.6 Solo- vs. co-authored k

In Figure C.6 we reproduce Figure 3(B) on the sample of top-five papers that were solo-authored (left-hand side) and co-authored (right-hand side). Both graphs display clear evidence of a Matthew effect and its gender difference.

The Matthew effect is similarly sized in both graphs. Its gender difference, on the other hand, is much larger among authors with solo-authored k than it is among authors with co-authored k. Previously published papers by women with solo-authored top-five articles receive, on average, three fewer citations per year before k is published; afterwards, they receive about three *more* citations a year compared to previously published papers by men with solo-authored k. In contrast, the gendered jump in citations is more muted among authors who co-author k—before k, women receive about one fewer citation per paper per year compared to men; after k, they receive about 1–1.5 more citations per paper per year.

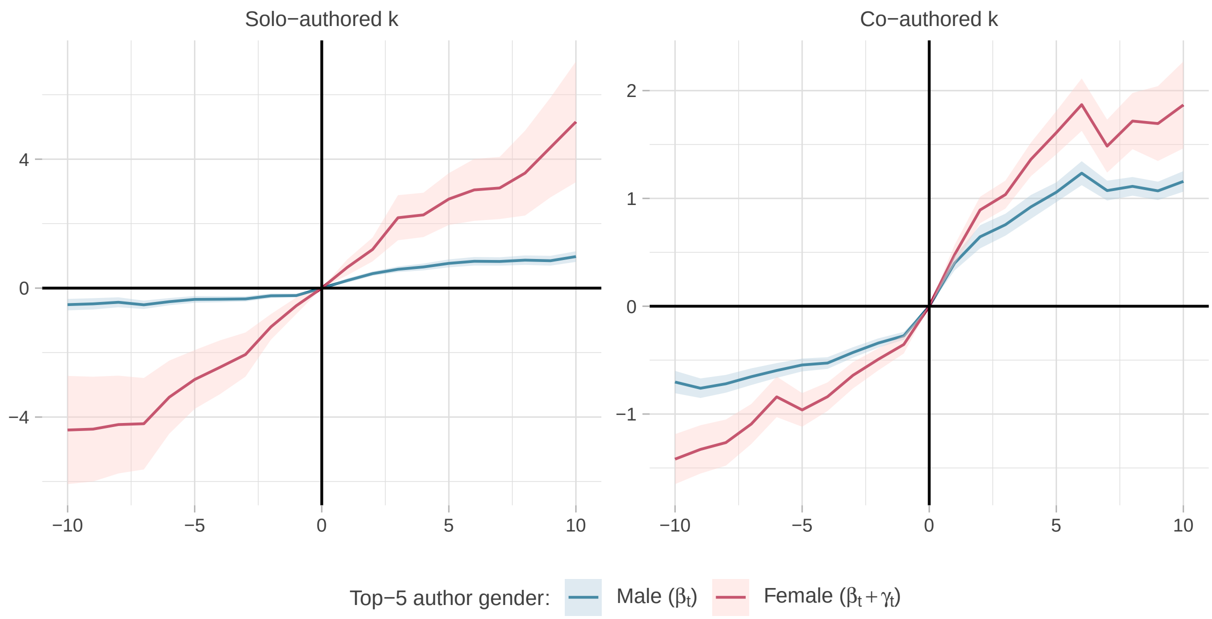

Figure C.6: Figure 3(B), solo- vs. co-authored k

Note. Figure C.6 replicates Figure 3(B) on the sample of authors whose first top-five article was solo-authored (left-hand graph) and co-authored (right-hand graph).

C.7 Comparing estimates from different collaboration patterns on k

In Figure C.7, we reproduce Figure 3 with different top-five paper collaboration patterns. In graph (A), we restrict the sample of top-five authors to those who co-authored k with exactly one other person for whom k was also a first top-five publication. In graph (B), we use the same sample used to estimate graph (A) but further limit it to authors who co-authored k with a member of the opposite sex. Graph (C) defines the estimation sample to include only authors who co-authored k with exactly one more senior male co-author.[7] Graph (D) includes all authors who co-authored k with at least one more senior male co-author. Estimates in graphs (A), (B) and (D) account shock$_k$ using fixed effects for each top-five article k' in graph (C), we control for k's co-author count, citation count (asinh), author prominence and year and journal fixed effects.

In all graphs of Figure C.7, there is a noticeable Matthew effect for both men and women around $t = 0$, and the effect is always larger for women than it is for men. However, the gender difference is smaller when junior co-authors are co-authoring with senior men (graphs (C) and (D)) than it is when k is co-authored by two junior authors (graphs (A) and (B)).

[7] We define seniority according to previous top-five publication counts. That is, a person with 2 previous top-five publications is assumed to be the more senior co-author on a paper with two other co-authors, neither of whom have previously published in a top-five journal.

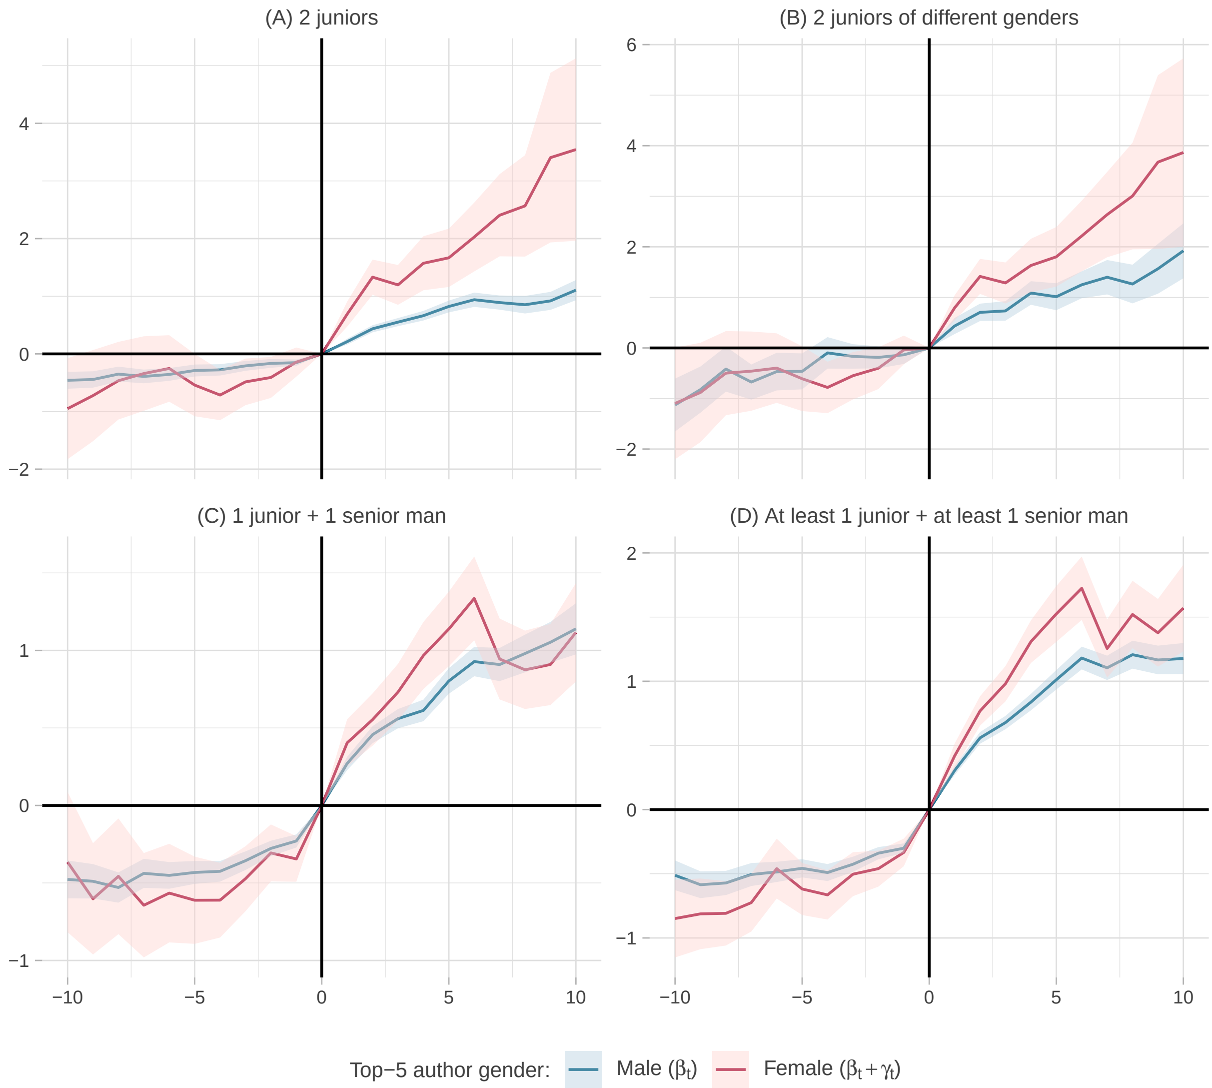

Figure C.7: Top-five papers collaboration patterns

Note. Figure C.7 replicates Figure 3 on the sample of authors whose first top-five paper was co-authored with exactly one other person, for whom it was also a first top-five publication. Graphs (A) and (B) include all such teams; graphs (C) and (D) exclude same-sex teams.

C.8 Controlling for authors' institutional rank when k is published

Figure C.8 reproduces Figure 3 controlling for authors' institutional rank at the time k was published. To determine institutional rank, we follow the same procedure as Hengel (2022). That is, for each institution, we count the number of top-five articles in which it was listed as an affiliation in a given year and smooth the average over a five-year period. Institutions were then ranked on an annual basis using this figure and grouped to create fifteen dynamic dummy variables. Institutions ranked in positions 1–9 are assigned individual dummy variables. Those in positions 10–59 are grouped in bins of 10 to form six dummy variables. Institutions ranked 60 or above were collectively grouped to form a final dummy variable. When multiple institutions are associated with an article, only the dummy variable of the highest ranked institution is used.

The Matthew effect and its gender difference in Figure C.8 resemble the estimates shown in Figure 3.

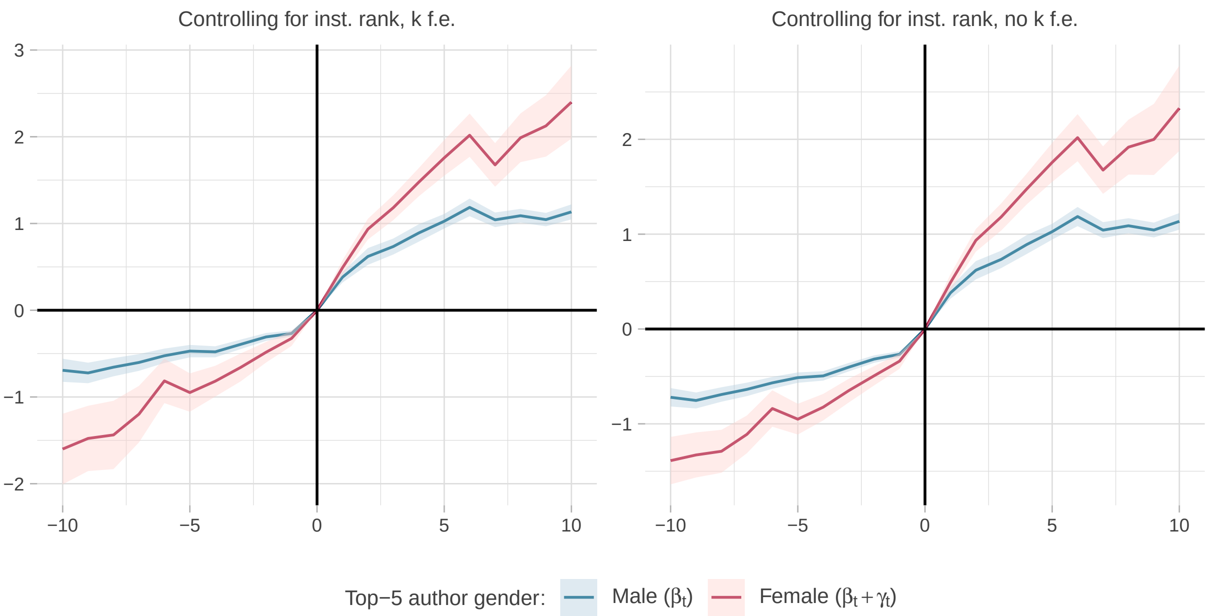

Figure C.8: Figure 3, controlling for authors' institutional rank when k is published

Note. Figure C.8 reproduces Figure 3 controlling for authors' institutional rank at the time k was published.

C.9 Controlling for k's primary and secondary *JEL* codes

Figure C.9 reproduces Figure 3 controlling for fixed effects for each of k's primary (top row) and secondary (bottom row) *JEL* codes. Estimates are almost identical to those shown in Figure C.9.

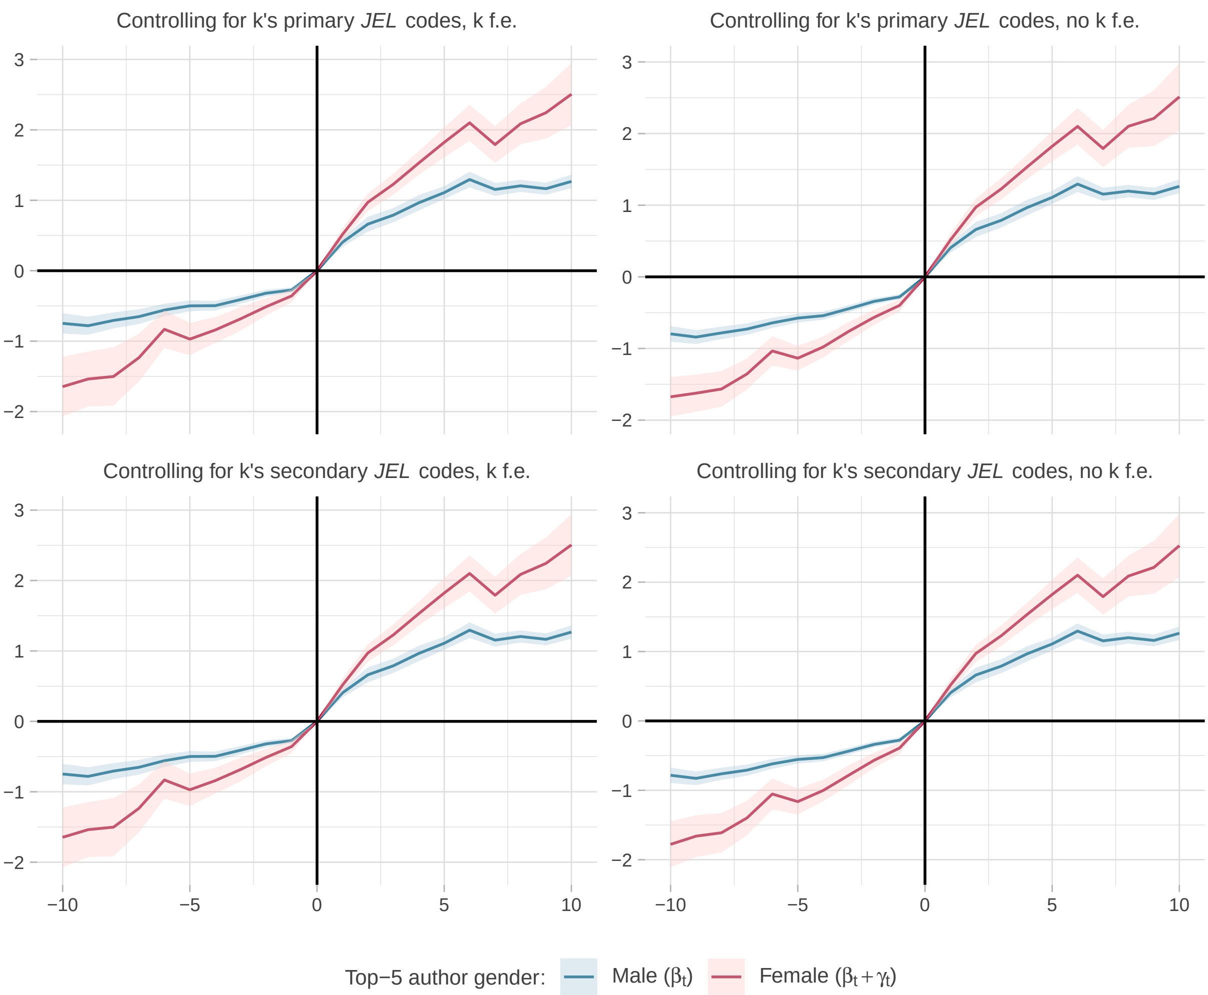

Figure C.9: Figure 3, controlling for k's primary and secondary *JEL* codes

Note. Figure C.9 reproduces Figure 3, controlling for k's primary (top row) and secondary (bottom row) *JEL* codes.

C.10 Controlling for authors' pre-k publication counts

In Figure C.10, we reproduce Figure 3 controlling for the total number of publications an author had before he published his first top-five article, k. Both graphs in Figure C.10 very closely resemble the graphs shown in Figure 3.

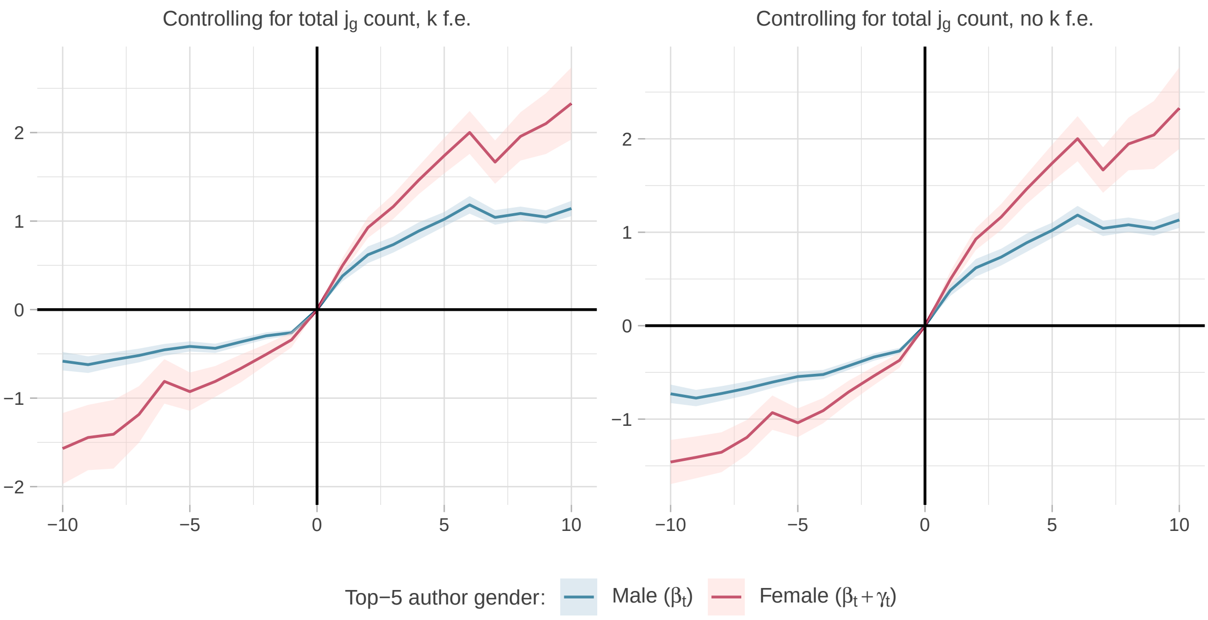

Figure C.10: Figure 3, controlling for authors' total j_g count when k is published

Note. Figure C.10 reproduces Figure 3, controlling for authors' total number of pre-k publications.

C.11 Controlling for j_g's journal impact factor

In Figure C.11, we reproduce Figure 3 controlling for cited journal impact factor (JIF).[8] The visibility Matthew effect and its gender difference are very similar to their corresponding estimates in Figure 3.

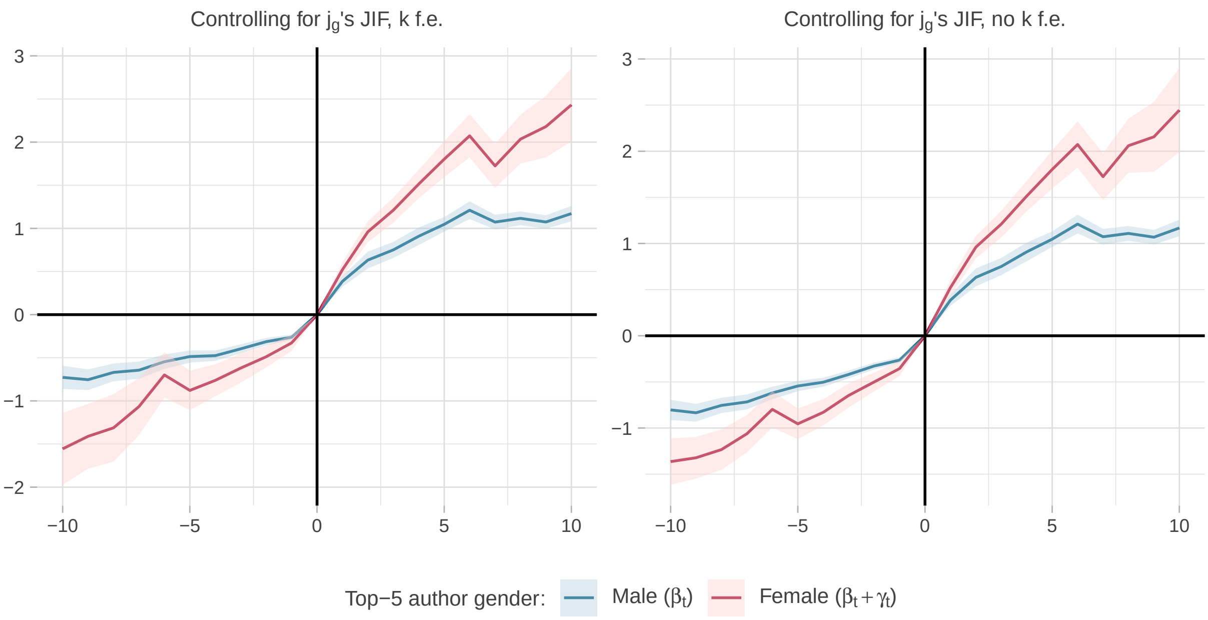

Figure C.11: Figure 3, controlling for j_g's journal impact factor

Note. Figure C.11 reproduces Figure 3, controlling for j_g's journal impact factor.

[8] Web of Science's journal impact factor measures the frequency with which articles in a journal are cited. It is intended to reflect the impact a journal has in its respective field.

C.12 Controlling for j_g's number of co-authors

In Figure C.12, we reproduce Figure 3 controlling for j_g's number of co-authors. Results are very similar to those shown in Figure 3.

Figure C.12: Figure 3, controlling for j_g's number of co-authors

Note. Figure C.12 reproduces Figure 3, controlling for the number of co-authors on each j_g.

C.13 Early vs. late top-five publication

In Figure C.13 we restrict the sample to authors with five or fewer publications before their first top-five publication (graph (A)) and 10 or more (graph (B)). We see clear evidence of a Matthew effect and a gender difference in the Matthew effect in both graphs. However, the gender difference is smaller among authors who published fewer pre-top-five papers than it is among authors with more pre-top-five papers.

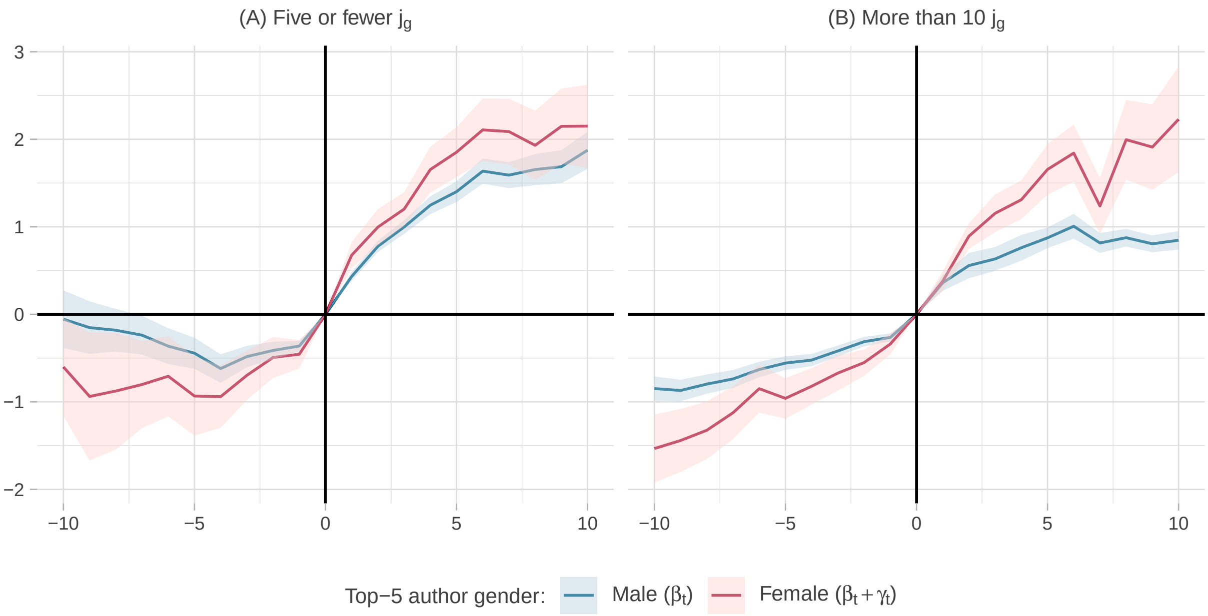

Figure C.13: Early vs. late top-five publication

Note. Figure C.13 reproduces Figure 3 except graph (A) restricts the sample to authors with five or fewer pre-k publications; graph (B) restricts the sample to authors with 10 or more pre-k publications.

D Section 4.3, robustness

D.1 Balance tables for treatment and counterfactual environments

Tables D.1 and D.2 show gender differences in author and paper characteristics in the treatment and counterfactual environments defined in Section 4.2. Estimates are from OLS regression of the relevant dependant variable on dummy variables equal to one if the author is a woman, the observation falls in the treatment category, their interaction and fixed effects for k (among the sample of cited and citing papers, only).

Table D.1: Treatment 1, j_g is cited by k

	Treatment	Counterfactual	Difference
Cited papers			
Publication year	0.7876*	1.2448***	−0.4573
	(0.4456)	(0.3529)	(0.3876)
No. co-authors	−0.0908	0.0525	−0.1433*
	(0.0855)	(0.0579)	(0.0820)
Journal impact factor	0.5130*	0.5016***	0.0114
	(0.2924)	(0.1424)	(0.2933)
No. citations	2.5306	1.5624	0.9682
	(2.6947)	(1.3371)	(2.5445)
Citing papers			
Publication year	0.1353	0.2611**	−0.1258
	(0.1394)	(0.1077)	(0.1381)
No. co-authors	0.0602***	0.0307*	0.0295
	(0.0224)	(0.0181)	(0.0214)
English language	0.0031	0.0014	0.0017
	(0.0022)	(0.0013)	(0.0024)
Article	0.0040	−0.0005	0.0045
	(0.0034)	(0.0025)	(0.0036)

Note. Figures in the treatment and counterfactual columns are gender differences in paper characteristics when j_g is and is not cited by k, respectively; figures in the last column are their differences. Estimates are from OLS regression of the variables listed in the first column on dummy variables equal to one if the author is a woman, the observation falls in the treatment category, their interaction and fixed effects for k. Standard errors (in parentheses) are clustered at the level of the cited paper. ***, ** and * statistically significant at 1%, 5% and 10%, respectively.

Table D.1 reports estimates for treatment and counterfactual environments when the treatment is defined as the $j_g \in \mathcal{J}_g$ cited by k. Female-authored cited papers are published 0.8–1.2 years after male-authored cited papers; papers citing the latter are published 0.1–0.3 years before papers citing the former. Among cited papers, women are published in journals with a somewhat higher journal impact factor and their papers earn slightly more citations. However, similarly sized gender differences are found in both treatment and counterfactual environments. We observe no gender differences at all in the language of the citing paper or its type (*e.g.*, article or book review).[9]

Indeed, the only significant difference between treatment and counterfactual environments in Table D.1 is the gender difference in cited papers' number of co-authors. Conditional on treatment, women's pre-top-five publications are co-authored with slightly fewer people than men's; in contrast, women co-author with slightly more people in the counterfactual environment. Although neither gender difference is statistically significant on its own, their difference (weakly) is.

[9]Almost all cited papers are in English and were published in a journal.

Table D.2 reports estimates for the treatment and counterfactual environments when the treatment restricts k to articles first released as NBER working papers. In the first panel, we show gender differences among authors themselves.[10] Differences are rarely significant, conditional on treatment; they are larger and more often significant in the counterfactual environment. Nevertheless, their signs and magnitudes generally match; the only significant differences between the two environments is in gender differences in co-author prominence and the probability of publishing in the *Review of Economic Studies* (*REStud*). Women in the counterfactual group co-author with slightly more prominent co-authors compared to men; their first top-five paper is also more often published in *REStud*. In contrast, there are no significant gender differences, conditional on treatment.

The second two panels of Table D.2 show gender differences among cited and citing papers. Estimates for treatment and counterfactual environments generally match one another. The exceptions are year of publication and citing paper type. In the treatment environment, women's papers are more often cited by journal articles compared to men's papers relative to the counterfactual environment. In the counterfactual environment, women's papers (and the papers that cite them) are published later than men's, relative to the treatment environment.

[10] We do not show similar estimates in Table D.1 because the treatment is defined at the cited paper (rather than the author) level; as a result, the treatment and counterfactual environment can include papers from the same author.

Table D.2: Treatment 2, k is an NBER working paper

	Treatment	Counterfactual	Difference
Top-five papers			
Publication year	1.5099	2.8325***	−1.3225
	(0.9268)	(0.5244)	(0.9152)
No. co-authors	0.1222	0.1302**	−0.0081
	(0.0920)	(0.0555)	(0.1121)
No. citations	−38.3507	−4.7791	−33.5716
	(38.9896)	(16.0692)	(56.5420)
Co-author prominence	−0.0690	1.5175***	−1.5865**
	(0.5046)	(0.3297)	(0.7038)
No. j_g	−1.6719	−2.2519***	0.5800
	(1.5397)	(0.5421)	(1.4213)
Share of *AER* papers	−0.0104	0.0702**	−0.0806
	(0.0499)	(0.0293)	(0.0594)
Share of *Econometrica* papers	−0.0128	−0.0651***	0.0524
	(0.0387)	(0.0202)	(0.0371)
Share of *JPE* papers	0.0094	−0.0489**	0.0583
	(0.0394)	(0.0201)	(0.0464)
Share of *QJE* papers	0.0137	−0.0028	0.0165
	(0.0374)	(0.0192)	(0.0528)
Share of *REStud* papers	0.0000	0.0465*	−0.0465*
	(0.0393)	(0.0266)	(0.0266)
Institutional rank	−0.8631*	−0.3739	−0.4892
	(0.4437)	(0.2484)	(0.6102)
Cited papers			
Publication year	−0.4491	1.5925***	−2.0417***
	(0.6738)	(0.3927)	(0.7799)
No. co-authors	0.0095	0.0432	−0.0337
	(0.1435)	(0.0603)	(0.1556)
Journal impact factor	0.0889	0.6059***	−0.5169
	(0.3551)	(0.1506)	(0.3862)
No. citations	−0.1186	2.2061*	−2.3247
	(4.6673)	(1.2505)	(4.8319)
Citing papers			
Publication year	−0.1590	0.3726***	−0.5316*
	(0.2530)	(0.1021)	(0.2728)
No. co-authors	0.0270	0.0397**	−0.0126
	(0.0381)	(0.0185)	(0.0424)
English language	0.0003	0.0022	−0.0020
	(0.0023)	(0.0014)	(0.0027)
Article	0.0080**	−0.0023	0.0103**
	(0.0039)	(0.0027)	(0.0048)

Note. Figures in the treatment and counterfactual columns are gender differences when k is restricted to papers released as NBER working papers and papers not released as NBER working papers, respectively; figures in the last column are their differences. Estimates are from OLS regression of the variables listed in the first column on dummy variables equal to one if the author is a woman, the observation falls in the treatment category, their interaction and fixed effects for k (sample of cited and citing papers only). Standard errors (in parentheses) are robust in panel 1 and clustered at the level of the cited paper in panels 2 and 3. ***, ** and * statistically significant at 1%, 5% and 10%, respectively.

D.2 Treatment 1, with and without NBER k

Figure D.1 reproduces Figure 4, graphs (A) and (B) but restricts the treatment and counterfactual samples to articles by authors who did (top row) and did not (bottom row) release k as an NBER working paper.[11]

Consider first the graphs in row one of Figure D.1. These graphs reflect citations to j_g cited by k (left-hand-side graph) and j_g not cited by k (right-hand-side graph) among the sample of authors who first released k as an NBER working paper. As argued in Section 4.2, the act of releasing k in the NBER series should alert potential citing papers to authors' previously published work before k is published, irrespective of whether that that work is cited in k. As a result, neither graph in the first row of Figure D.1 should display a gender difference in the Matthew effect.

This is indeed what we observe. Both graphs suggest a jump in citations around $t = 0$; however, there is no gendered difference in that jump.

In Figure D.1's second row, we restrict the estimation sample to authors who *did not* release k as an NBER working paper. Now, potential citing papers at $t < 0$ should only be aware of authors' previously published work when that work is cited in k. Thus, the second row of graphs in Figure D.1 should mimic graphs (A) and (B) of Figure 4—*i.e.*, we should only observe a gendered difference in the Matthew affect among j_g not cited by k.

Again, this is precisely what we observe. In the counterfactual environment of the second row of graphs, women's Matthew effect is noticeably larger than men's.

[11] Or equivalently, Figure D.1 reproduces Figure 4, graphs (C) and (D) (*i.e.*, treatment and counterfactual group 2) but restricts estimation samples to j_g cited (column one) and not cited (column two) by k.

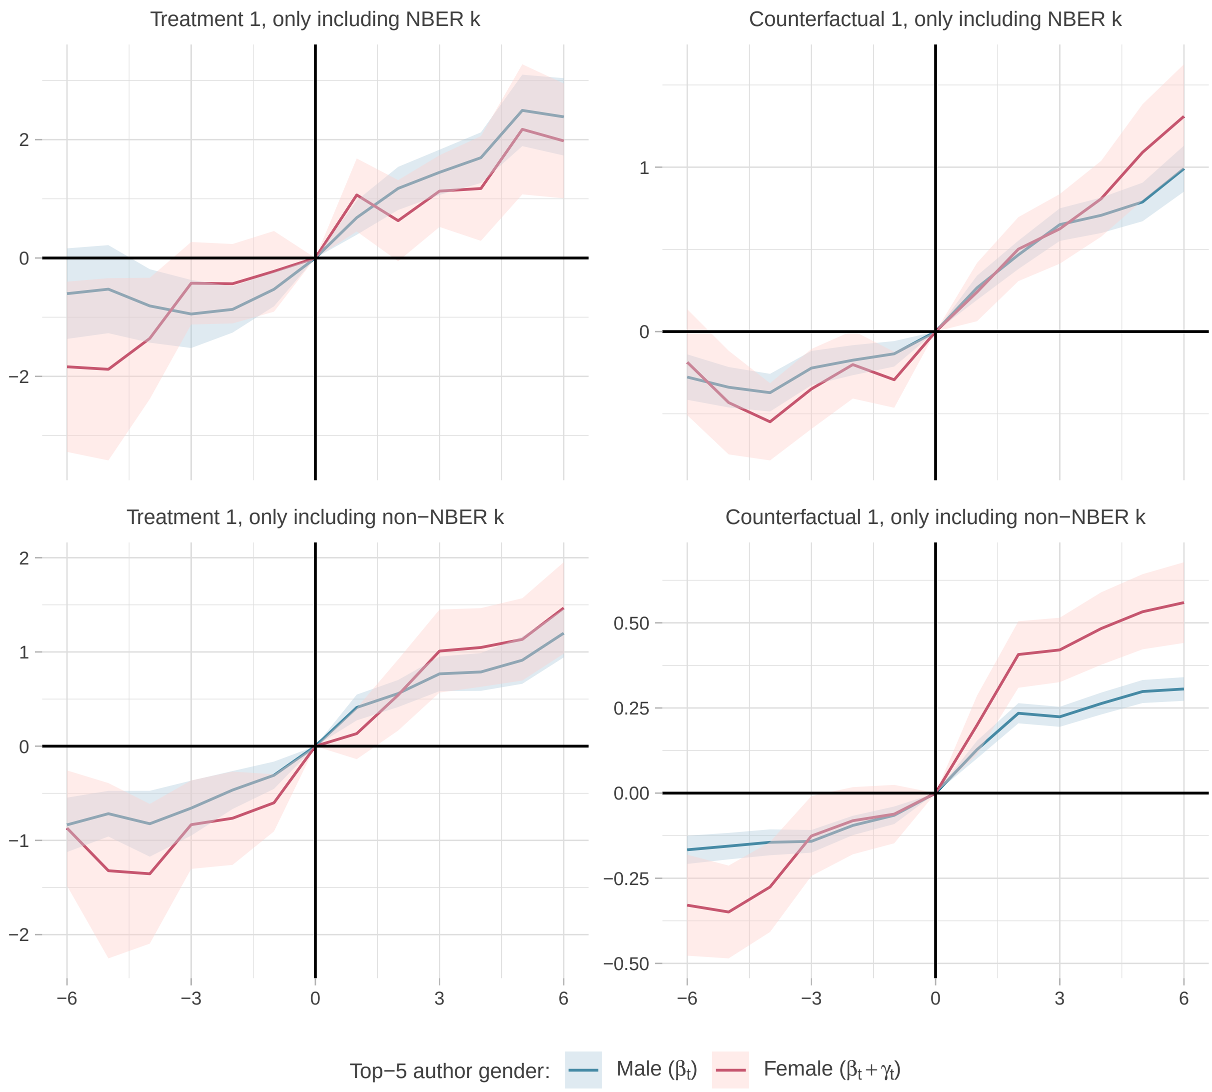

Figure D.1: Treatment and counterfactual group 1, with and without NBER k

Note. Graphs reproduce Figure D.1, graphs (A) and (B) but restrict the treatment and counterfactual samples to articles by authors who did (top row) and did not (bottom row) release k as an NBER working paper.

D.3 Poisson likelihood function

In Figure D.2, we reproduce Figure 4 but estimate Equation (3) using a Poisson likelihood function. Results are similar to those shown in Figure 4.

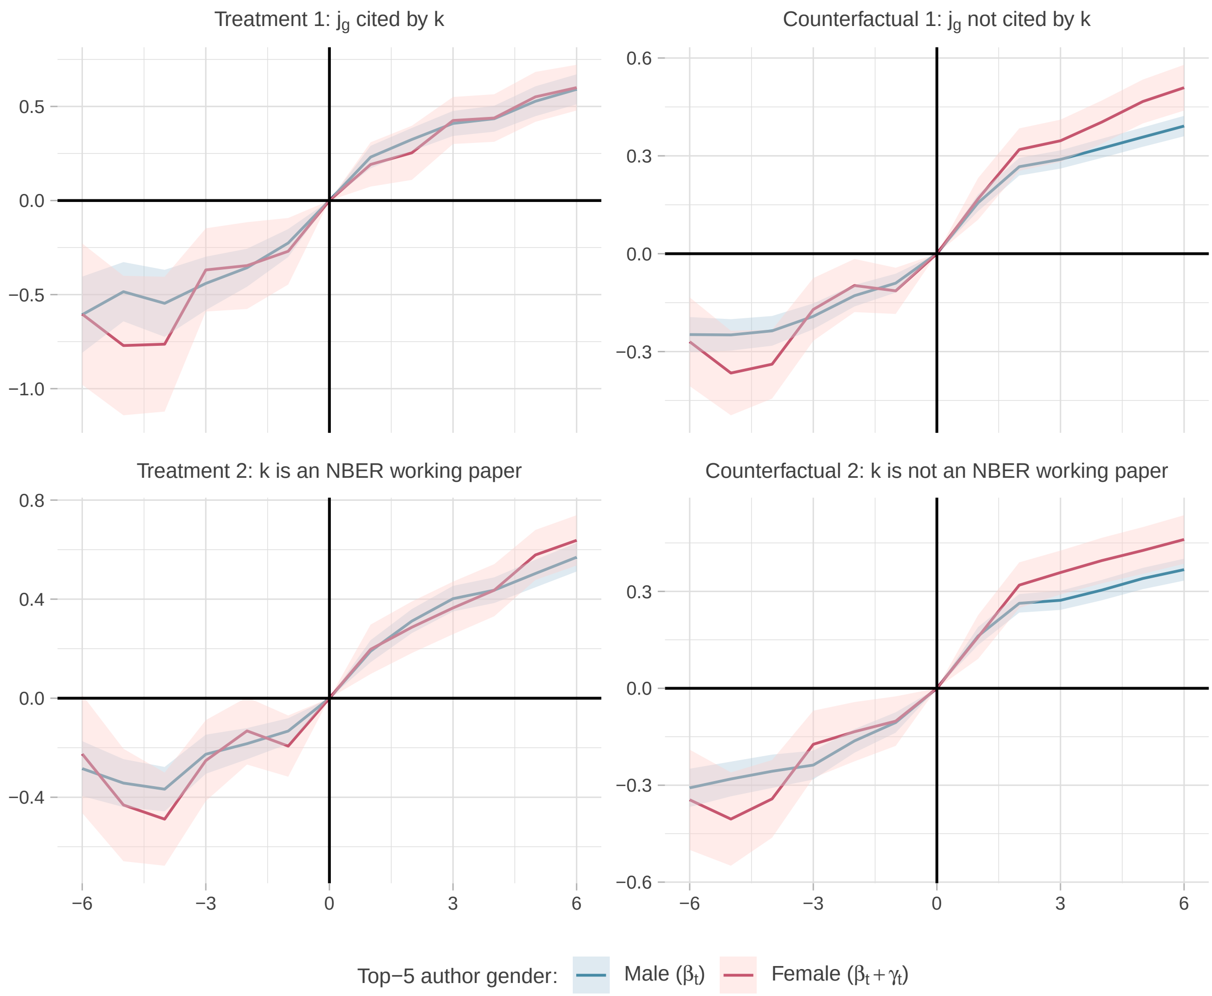

Figure D.2: Figure 4, Poisson likelihood function

Note. Graphs reproduce Figure 4, except they estimate Equation (3) using a Poisson likelihood function. Shaded areas reflect 90 percent confidence intervals from standard errors clustered by j_g.

D.4 Controlling for journal impact factors and pre-k publication counts

Figure D.3 reproduces Figure 4 controlling for cited papers' journal impact factors and authors' total pre-top-five publication counts. Results are very similar to those reported in Figure 4.

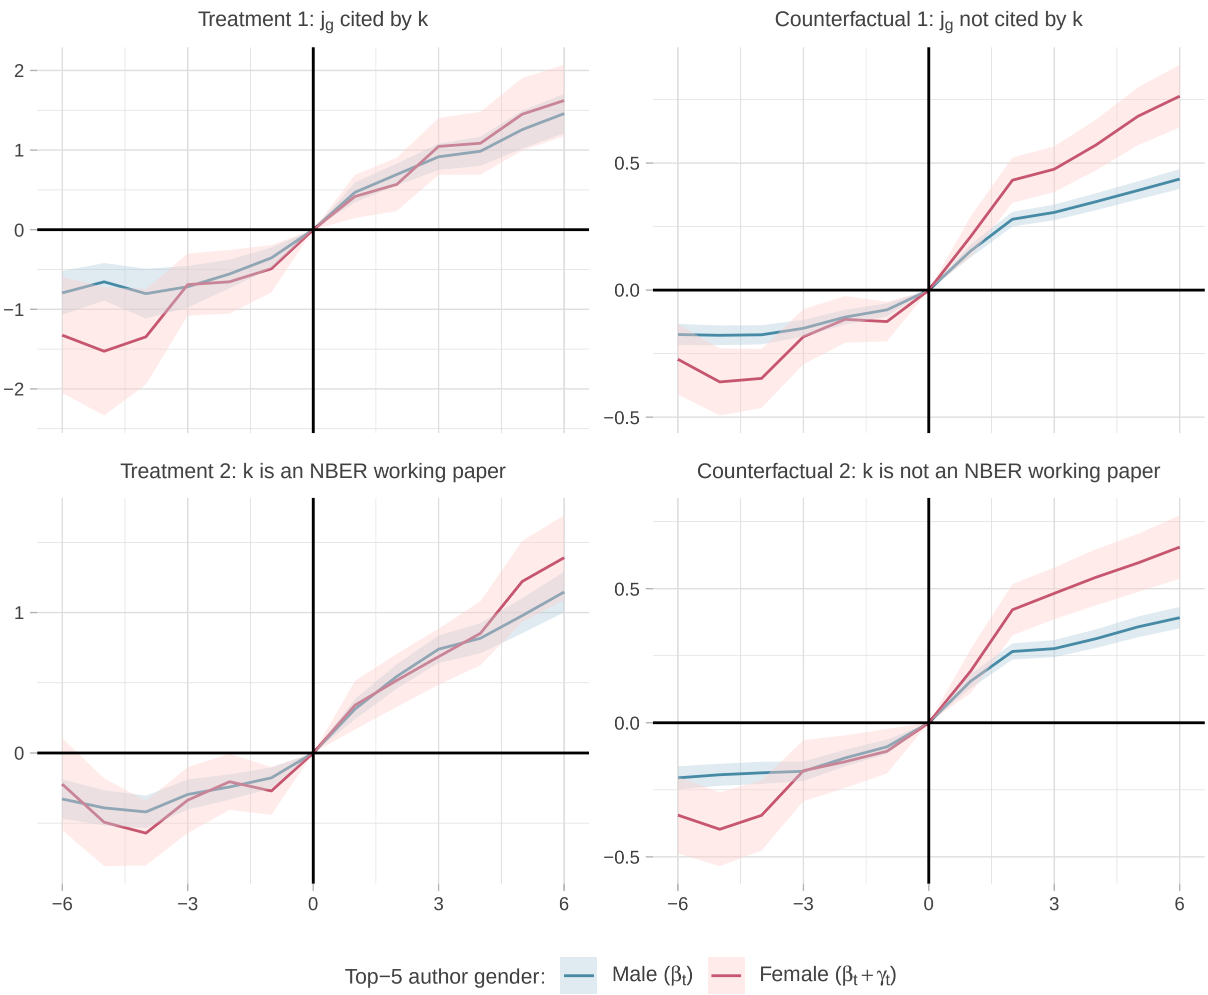

Figure D.3: Controlling for evaluation histories

Note. Graphs reproduce Figure 4 controlling for the journal impact factors of cited papers and authors' total pre-top-five publication counts.

D.5 Treatment 3: job market stars

We define additional treatment and counterfactual environments to consist of authors who were and were not recent academic job market stars, respectively. Every year, graduating economics PhD students compete for academic jobs at universities and other research institutions. During the process, departments invite candidates to campus "flyouts", where they are interviewed by committees, meet members of the faculty and present their job-market papers.[12]

The entire job market process exposes all candidates to a lot of attention, although a small number of "stars" capture the lion share of it.[13] Indeed, a common worry is that "The best candidates from [top-ranked] PhD programs may end up interviewing in so many places that they crowd out interviews and fly-outs for middle-of-the-class candidates" (Smeets, Warzynski, and Coupé, 2006, p. 170). As a result, potential citing papers should be more aware of job market stars' earlier work compared to non-stars who were also on the market.

To identify comparable job market candidates, we restrict both treatment and counterfactual environments to authors whose first top-five paper was probably their job-market paper—*i.e.*, solo-authored and published no more than six years after their first paper was published. We assume "stars" are affiliated to institutions ranked 1–9 at the time their top-five paper was published (see Appendix C.8); authors employed at other institutions are assumed to be "non-stars". To maximise the probability that Assumption 5 in Lemma 3 is satisfied, we again restrict citing and cited papers in both treatment and counterfactual environments to papers published in economics journals within a six-year window before and after the top-five paper's publication date.

In the treatment environment, publishing for the first time in a top-five journal boosts citations to j_g, but the effect is similar for both men and women.[14] Since potential citing papers should be aware of authors' previous work before the top-five publication, this is again evidence that strategic behaviour plays an important role in determining the citations that authors receive, but it is not impacted by the gender of the author being cited. In contrast, among authors who were at lower-ranked institutions at the time their first top-five paper was published, the visibility Matthew effect is noticeably larger for women than it is for men. Assuming that gender differences in citing papers' strategic behaviour is unrelated to selection into treatment, these results suggest that the gender gap in the visibility Matthew effect formed exclusively in the counterfactual environment is due to greater awareness of women's work after their first top-five paper is published.

[12] In economics, the job market paper is generally the strongest paper from a student's PhD thesis. Historically, this paper has almost always been solo-authored, although co-authored job market papers have become more common in the last decade.

[13] An even smaller number are also invited to participate in the *Review of Economic Studies* tour, where they meet and present to faculty at several European universities.

[14] In contrast to estimates shown in Figure 4, Figure D.4 displays β_t and $\beta_t + \gamma_t$ from estimating Equation (3) controlling for shock$_k$ using fixed effects for k's year and journal of publication, co-author count and citation count (asinh). (Given k is by definition solo-authored, we cannot account for shock$_k$ using fixed effects for k.)

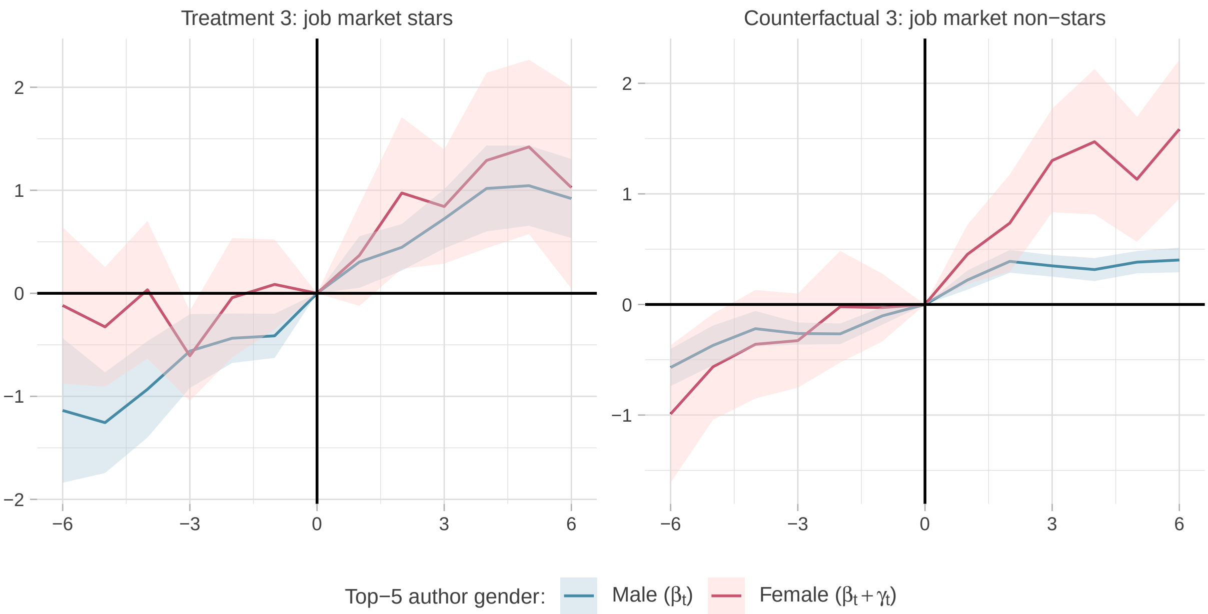

Figure D.4: Treatment 3: job market stars

Note. Graphs reproduce Figure D.1, graphs (A) and (B) but: (i) account for shock$_k$ using fixed effects for k's year and journal of publication, co-author count and citation count (asinh); (ii) restrict both environments to comparable recent job market candidates (*i.e.*, authors for whom k was likely their job market paper); and (iii) define the treatment environment as "stars" (*i.e.*, authors employed at institutions ranked 1–9 when k was published) and the counterfactual environment as "non-stars" (*i.e.*, authors employed at institutions ranked 10+ when k was published).

D.6 The Matthew effect by employer rank

In Figure D.5 we display β_t and $\beta_t + \gamma_t$ from estimating Equation (3) on the subsets of authors employed at institutions ranked 1–5 (graph (A)), 6–9 (graph (B)), *etc.* at the time their first top-five paper was published. Although employment at a top-ranked institution does not guarantee wide exposure to one's research, it probably generates more exposure compared to employment at lower ranked institutions. Thus, the gender gap in the visibility Matthew effect should probably be smaller at top-ranked institutions than it is at lower-ranked institutions.[15]

This is indeed what we observe. The graphs in Figure D.5 suggest that the smallest gender difference in the visibility Matthew effect is among authors employed at institutions ranked 1–9. At lower ranked institutions, there is a larger and significant gender difference.

[15]For example, many researchers affiliated to high-ranked institutions are students and/or research assistants whose research is likely to receive less attention relative to a tenure-track professor at a somewhat lower ranked institution.

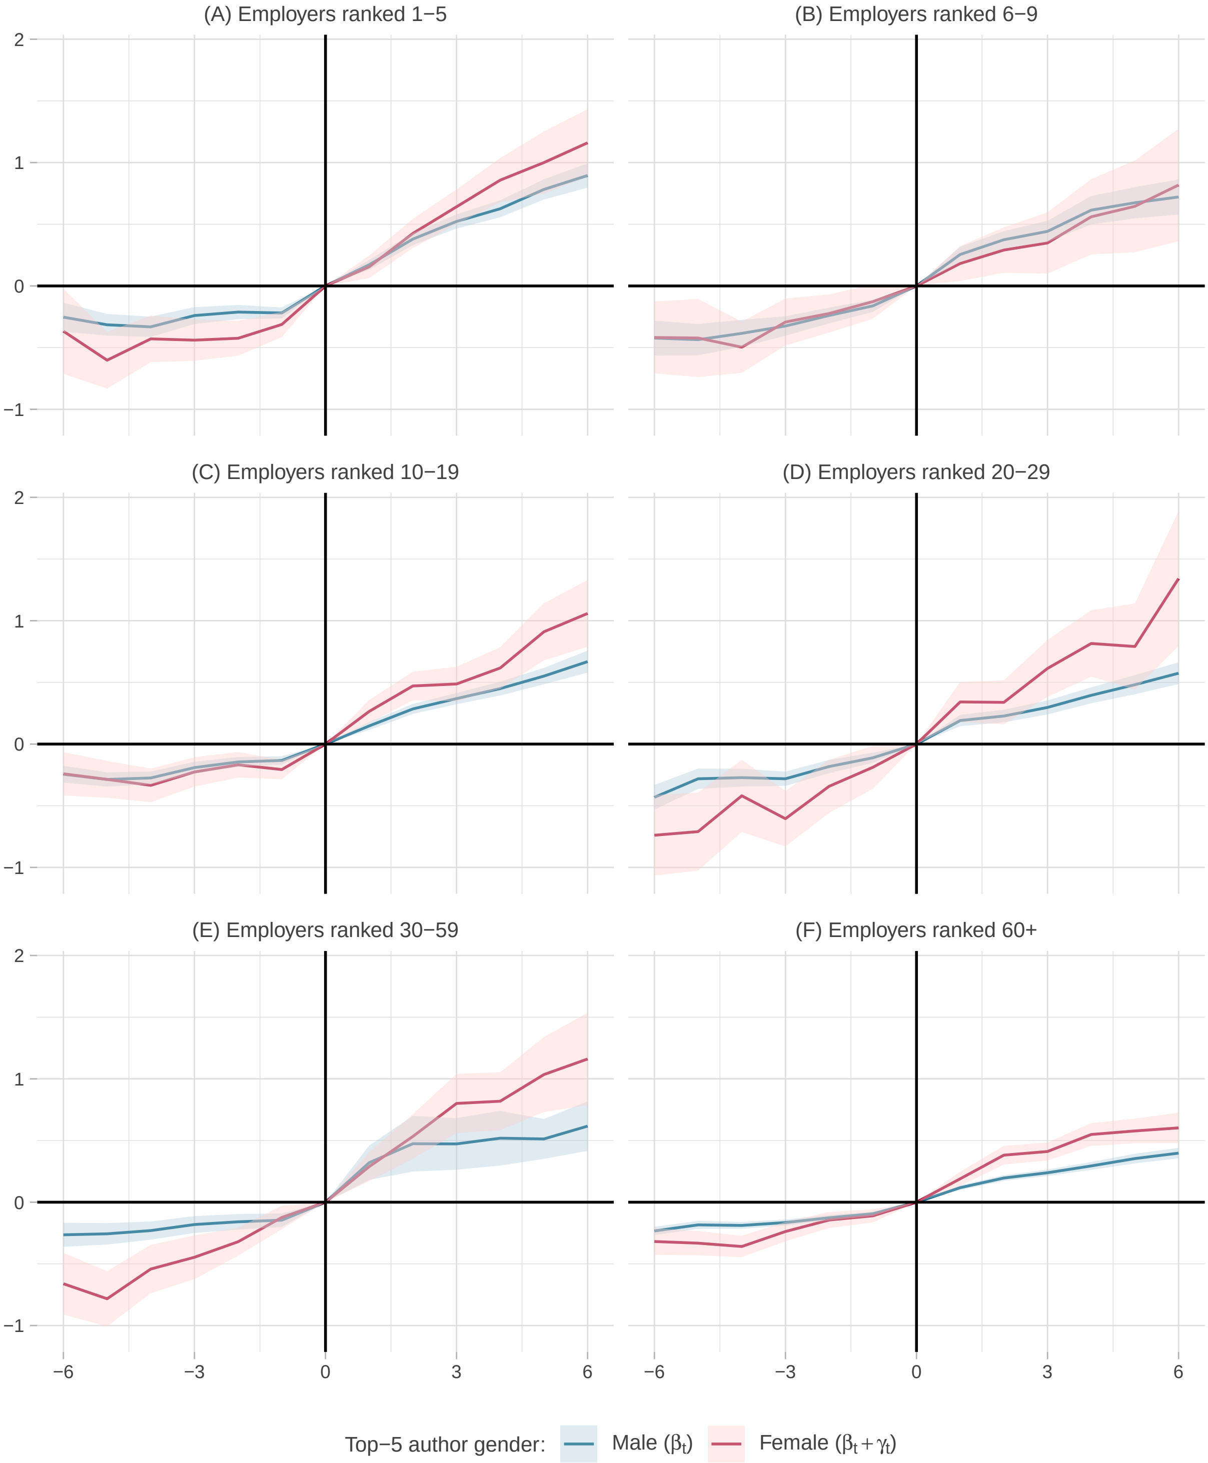

Figure D.5: The Matthew effect by employer rank

Note. Figure displays β_t and $\beta_t + \gamma_t$ from estimating Equation (3) controlling for k fixed effects on the subset of authors employed at institutions ranked 1–5 (graph (A)), 6–9 (graph (B)), 10–19 (graph (C)), 20–29 (graph (D)), 30–59 (graph (E)), and 60 or higher (graph (F)). Shaded areas represent 90 percent confidence intervals from standard errors clustered by j_g.

E Alternative outcomes

E.1 Editorships

In Figure E.1 we plot the probability that men and women become editors (as deviations from the mean) of the top 5 and top 21 economics/finance journals, where the latter include the 14 economics and finance journals ranked 1–15 in Combes and Linnemer (2010, p. 19), plus the *Journal of the European Economic Association* and the four *American Economic Journals*.[16] The first column shows graphs controlling for fixed effects for each top-five article k; the second column displays graphs that instead control for k's year and journal of publication, co-author count, citations (asinh) and author prominence.

As in Figure 5(A), there is a noticeable bump in the probability that men and women become top journal editors after k is published. However, we do not observe a difference in this jump by author gender.

[16]We define "editor" to include all regular issue editors, including associate editors. The top 15 journals listed in Combes and Linnemer (2010, Table 10, p. 19) is the *Quarterly Journal of Economics*, *American Economic Review*, *Journal of Political Economy*, *Econometrica*, *Review of Economic Studies*, *Journal of Financial Economics*, *Journal of Monetary Economics*, *Review of Economics and Statistics*, *Journal of Economic Theory*, *Journal of Finance*, *Journal of Econometrics*, *Economic Journal*, *RAND Journal of Economics*, *Journal of Public Economics* and *Journal of International Economics*. We add to that list the *American Economic Journal: Applied Economics*, *American Economic Journal: Economic Policy*, *American Economic Journal: Macroeconomics*, *American Economic Journal: Microeconomics*, and *Journal of the European Economics Association*.

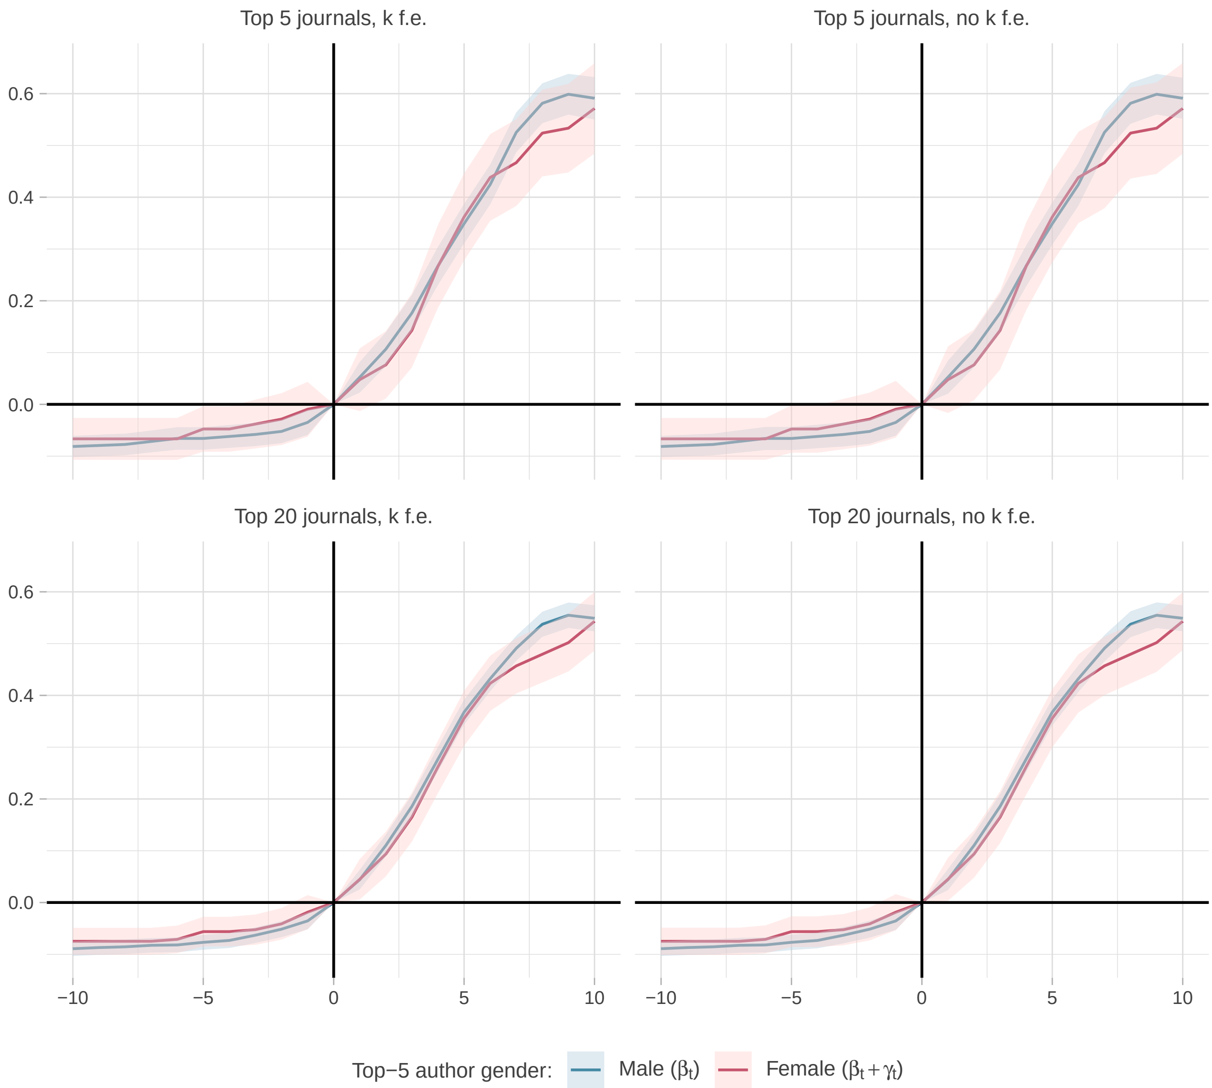

Figure E.1: The impact of k on editing a journal

Note. Graphs plot the probability that men and women become journal editors (as deviations from the mean) in the top 5 (top row) and top 21 (bottom row) economics/finance journals in the 10 years before and after k's publication date. In all graphs, β_t and γ_t are from estimating Equation (3) with OLS on a dummy variable equal to one if the author was an editor of a relevant journal in a given year. Shaded areas represent 90 percent confidence intervals from robust standard errors.

E.2 Promotion to professor

In Figure E.2 we plot the probability that men and women are promoted to professor (as deviations from the mean) in the ten years before and after their first top-five economics paper. The left-hand-side graph controls for fixed effects for each top-five article k; the right-hand-side graph instead controls for k's year and journal of publication, co-author count, citations (asinh) and author prominence.

We observe a noticeable bump in the probability of promotion around $t = 0$; however, this jump does not differ by author gender.

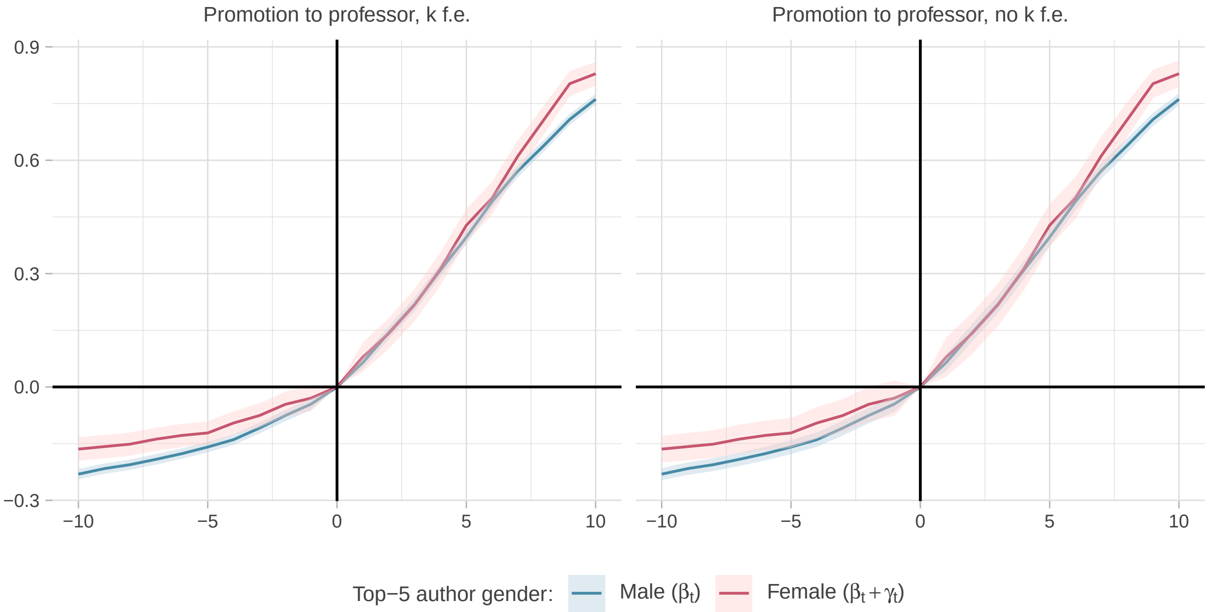

Figure E.2: The impact of k on being promoted to professor

Note. Graphs plot the probability that men and women are promoted to professor (as deviations from the mean) in the 10 years before and after k's publication date. In both graphs, β_t and γ_t are from estimating Equation (3) with OLS on a dummy variable equal to one if the author was a professor in a given year and zero otherwise. Shaded areas represent 90 percent confidence intervals from robust standard errors.

E.3 Grants

In Figure E.3 we plot grant counts (top row) and total dollar amounts (bottom row) by author gender in the ten years before and after publication of k. (All estimates represent deviations from the mean.) The first column shows graphs controlling for fixed effects for each top-five article k; the second column displays graphs that instead control for k's year and journal of publication, co-author count, citations (asinh) and author prominence.

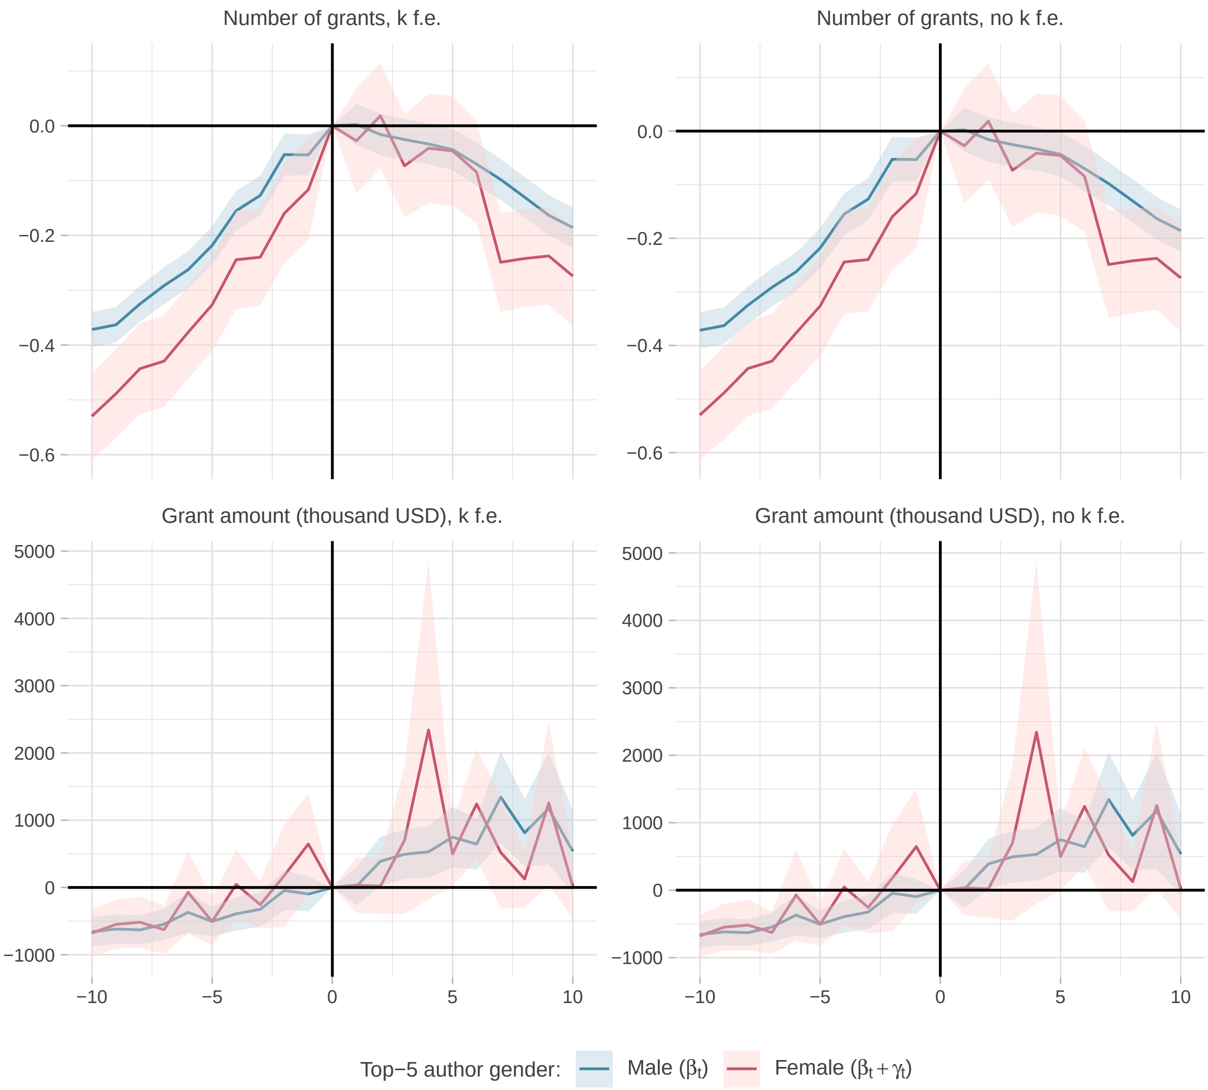

Figure E.3: Number and size of awarded grants (in thousand USD)

Note. Top row of graphs plot the number of grants received in the 10 years before and after k's publication date; bottom row of graphs plot the total grant amounts received in thousand USD (adjusted for inflation to reflect 2024 dollars). β_t and γ_t are from estimating Equation (3) with OLS and reflect deviations from the mean. Shaded areas represent 90 percent confidence intervals from robust standard errors.

Number of grants peaks around $t = 0$ and declines afterwards. Compared to men, women obtain fewer grants (relative to their mean) before k is published, but we do not observe a gender difference in the five to six years afterwards.[17] We observe no obvious boost in the dollar amounts awarded at $t = 0$ nor do we see any evidence of a gender difference on this dimension.

[17] On average women obtain slightly more grants, overall: our estimate of δ from Equation (3) is 0.17 (standard error 0.04).

E.4 Employer institutional rank

In Figure E.4 we plot the institutional rank of authors' employers (as deviations from the mean) in the ten years before and after publication of k, where institutions are annually ranked (in descending order) by the number of top-five articles affiliated to them, smoothed over a five-year period.[18] To limit the impact of outlier observations and impose ascending order, we additionally apply the log transformation and multiply the entire result by -1.

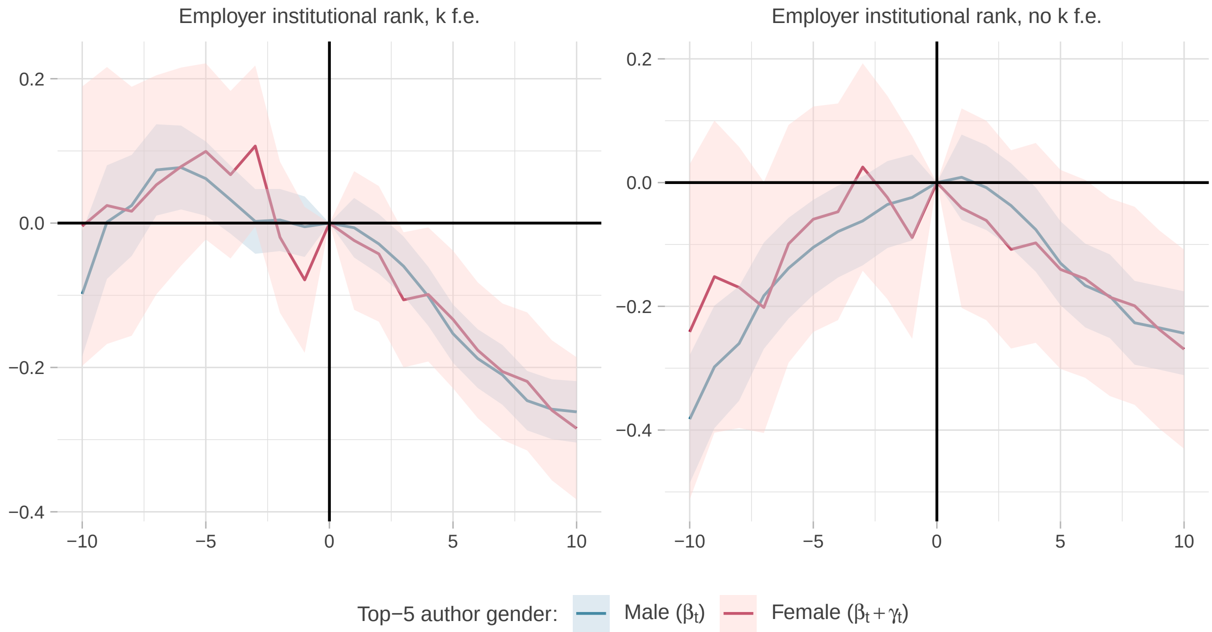

Figure E.4: Institutional rank of authors' employers

Note. Graphs plot the institutional rank of authors' employers (as deviations from the mean) in the 10 years before and after k's publication date, where higher numbers correspond to higher ranked institutions. In all graphs, β_t and γ_t are from estimating Equation (3) with OLS. Shaded areas represent 90 percent confidence intervals from robust standard errors.

Figure E.4 suggests weak evidence that institutional rank peaks at $t = 0$ but declines afterwards. We observe no gender differences—in either overall means or their deviations over time—at all.

[18]This ranking is identical to the one described in Appendix C.8, except that we do not aggregate institutions into 15 groups.

E.5 Seminar invitations

In Figure E.5 we plot the number of seminars given (as deviations from the mean) in the ten years before and after authors' first top-five publications. The first graph controls for fixed effects for each top-five article k; the second graph instead controls for k's year and journal of publication, co-author count, citations (asinh) and author prominence.

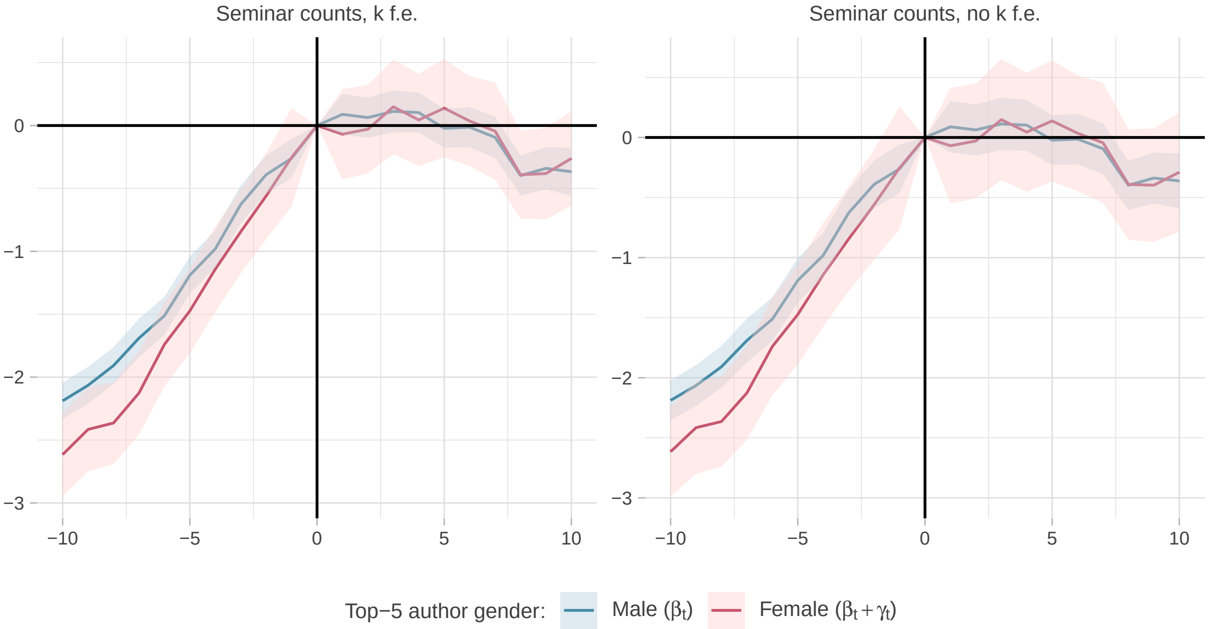

Figure E.5: Seminar invitations

Note. Graphs plot the number of seminars (as deviations from the mean) given by authors before and after k's publication date. In all graphs, β_t and γ_t are from estimating Equation (3) with OLS. Shaded areas represent 90 percent confidence intervals from robust standard errors.

Figure E.5 suggests a steep increase in the number of seminars authors give before k is published and a plateau afterwards. We observe no meaningful gender difference in seminar counts after k is published, although women may give fewer seminars (relative to their mean) beforehand.[19]

[19] Compared to their male co-authors on k, women give 0.48 (standard error 0.17) more seminars overall according to the regression underlying the left-hand-side graph in Figure E.5. We find no significant gender difference in average seminar counts in the regression underlying the right-hand-side graph.

F Alternative treatments

F.1 Publication in a non-top-five journal

In Figure F.1 we reproduce Figure 3 but set $t = 0$ as the date of publication of one random paper published by the author exactly five years before his first top-five paper was published, and restrict the sample of cited papers to those published before this paper. Thus, Figure F.1 identifies the Matthew effect and its gender difference for non-top papers.[20]

Figure F.1 suggests that authors experience a Matthew effect even when publishing non-top papers. However, the effect is about half the size as the one they enjoy with their first top-five paper. The gender difference is also much smaller.

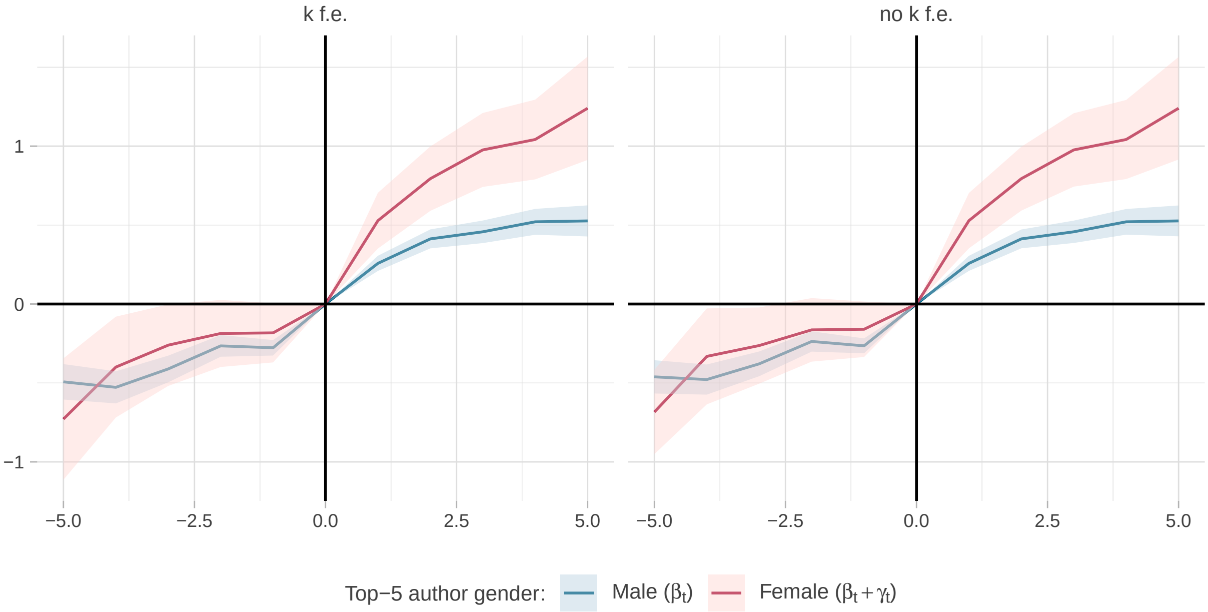

Figure F.1: The Matthew effect for non-top publications

Note. Graphs reproduce Figure 3 but set $t = 0$ as the date of publication of a random paper by the author published exactly 5 years before his first top-five paper.

[20]Figure F.1 only shows β_t and $\beta_t + \gamma_t$ in the five year window before and after the chosen j_g was published. This is because the sample of authors with j_g more than 10 years before their first top-five paper is small—recall $t = -5$ in the above graph corresponds to $t = -10$ in Figure 3—and $t = 5$ is the date the first top-five paper was published.

F.2 Publication of the second top-five paper

In Figure F.2 we replicated Figure 3 but restrict the sample to authors who also published a second top-five publication within the time period studied and set $t = 0$ to the second top-five paper's publication date. Under these conditions, potential citing papers should already be aware of authors' earlier work at $t < 0$ from the publicity generated from the authors' first top-five publication; as a result, Lemma 3 predicts that there should be no gender differences in the visibility Matthew effect.

This is indeed what we observe: in Figure F.2, $\beta_t + \gamma_t$ closely tracks β_t for all t.

Figure F.2: The visibility Matthew effects for authors' second top-five publication

Note. Graphs reproduce Figure 3 but set $t = 0$ as the date an author published his *second* top-five paper.

References

Clarivate (2022). *Web of Science [database]*. Data accessed and downloaded July 2022.

Combes, Pierre-Philippe and Laurent Linnemer (2010). “Inferring missing citations: a quantitative multi-criteria ranking of all journals in economics”. GREQAM Working Paper no. 2010–28.

Hengel, Erin (2022). “Are women held to higher standards? Evidence from peer review”. *The Economic Journal* 132 (648), pp. 365–381.

Hengel, Erin and Euyoung Moon (2023). “Gender and equality at top economics journals”. Mimeo.

Smeets, Valérie, Frédéric Warzynski, and Tom Coupé (2006). “Does the academic labor market initially allocate new graduates efficiently?” *Journal of Economic Perspectives* 20 (3), pp. 161–172.