The golden leap: gender differences in the Matthew dffect*

Jennifer Doleac† Erin Hengel‡ Myra Mohnen§

March 2026
(First version: March 2023)

Abstract

Scientific recognition is not determined by merit alone. Eminent researchers benefit disproportionately from the Matthew effect, while women's contributions are often overlooked, a phenomenon known as the Matilda effect. This paper is the first to examine how these forces interact. We exploit the timing of economists' first top-five journal publications—a salient status shock—to estimate its impact on citations to their earlier work. Male and female co-authors on the same top-five paper experience a significant citation boost following the top-five publication, consistent with a Matthew effect. Strikingly, the increase is larger for women only if their earlier work was less visible; when prior awareness is high, gender differences in the Matthew effect disappear. Guided by a model of the decision to cite, we show that this pattern is consistent with gendered differences in early-career visibility rather than differential strategic considerations—*e.g.*, a desire to cite well-known authors. These findings suggest that status shocks can help correct the Matilda effect, and that gender citation gaps evolve as information about researchers and their work accumulates.

*We are especially grateful to Hayoung Jo and Shyrley P. (Upwork) for their outstanding research assistance. We thank the Career Structures Initiative of the European Economic Association and the Hub for Equal Representation in the Economy for early funding support. We also appreciate valuable comments and suggestions from Giovanna Vallanti, Thomas Breda, Anja Prummer, Daniel Hamermesh, Imran Rasul and seminar/conference participants at the Société canadienne de science économique Canada, DRUID Université Côte d'Azur, Yosemite Retreat, SInnoPSis/TWIN4MERIT University of Cyprus, Career Structures in Economics Workshop at the LSE, Royal Holloway, Freie Universität Berlin, Brunel University of London, the Gender Gaps in Academia Workshop Centre d'Economie de la Sorbonne and the GRAPE Gender Gaps Conference 2025.

†Arnold Ventures, jdoleac@arnoldventures.org.

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

§University of Ottawa, mmohnen@uottawa.ca.

1 Introduction

Two well-documented phenomena—the Matthew and Matilda effects—describe how credit and merit can deviate from one another in science. The Matthew effect refers to the self-reinforcing cycle of success afforded to better known scientists.[1] The Matilda effect describes how women's scientific contributions are often undervalued and overlooked.[2] While the Matthew and Matilda effects have been extensively studied separately, their interaction remains unexplored. In this paper, we fill the gap by investigating whether the Matthew effect mitigates the Matilda effect once women achieve a certain level of scientific success.

A consequence of the Matthew effect is 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. 62). Although this outcome is usually criticised for penalising researchers who have "not yet made their marks", it may also serve a positive function by reducing gender disparities in the recognition awarded to those who *have* made their marks.

The story of Katalin Karikó is a case in point. Karikó, a biochemist, faced years of marginalisation and funding rejections. Despite producing ground-breaking work on in vitro-transcribed messenger RNA (mRNA), Karikó struggled to gain recognition, secure stable funding and advance in academia (Bansal, 2023), hallmarks of the Matilda effect. That all changed during COVID-19. The global urgency of the pandemic realigned scientific priorities and catapulted Karikó's long-overlooked research into the centre of vaccine development efforts. This triggered a Matthew effect that finally bestowed on Karikó the recognition she deserved—in 2023 she won a Nobel Prize in Physiology and Medicine.

In this paper, we empirically study whether Karikó's experience is more widespread. Our analysis focuses on economics, a discipline marked by persistent gender disparities (Costa-Dias *et al.*, 2024; Ginther and Kahn, 2021; Hamermesh, 2013; Lundberg and Stearns, 2019) and consistent evidence that women's research is under-valued relative to men's (see, *e.g.*, Card *et al.*, 2020; Hengel, Moon, and Tol, 2024; Krawczyk and Smyk, 2016; Sarsons *et al.*, 2021). Economics is also unique in the out-sized status and visibility it awards to manuscripts published in its so-called "top-five" journals—*i.e.*, the *American Economic Review* (*AER*), *Econometrica* (*ECA*), *Journal of Political Economy* (*JPE*), *Quarterly Journal of Economics* (*QJE*) and *Review of Economic Studies* (*REStud*).[3]

To study the interaction between the Matthew and Matilda effects, we apply two estimation strategies to a comprehensive dataset covering all economists who published their first top-five article after 1985. Our dataset includes almost 4,000 individuals, 14% of whom are women. By combining bibliometric data from Web of Science with hand-collected CVs, the dataset provides rich information on publication histories, citation patterns, institutional affiliations, grant funding, career milestones, seminar invitations and demographic characteristics. The data suggest that women are now better represented at top-ranked institutions and in prestigious editorial roles; nevertheless, they are promoted more slowly and are awarded smaller grants, particularly at less prestigious institutions.

Our two-pronged estimation strategy first assumes that publishing in a top-five economics journal creates a shock to perceived status and visibility that is sufficient to trigger the Matthew effect. To identify the impact of this shock, we implement an event study that examines citations to an author's body of research

[1] The term "Matthew effect" comes from a biblical passage in Matthew: "For every one who has will more be given, and he will have abundance; but from him who has not, even what he has will be taken away" (Matthew 25:29, Revised Standard Version).

[2] Margaret Rossiter coined the "Matilda effect"; she named it after the feminist and suffragist Matilda Joslyn Gage (Rossiter, 1993).

[3] Publications in these journals are widely recognised as a marker of research quality, and "[f]aculty 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" (Heckman and Moktan, 2020, p. 419). Indeed, the average economist would sacrifice half a thumb to publish in the *AER* (Attema, Brouwer, and van Exel, 2014)!

published *before* his first top-five paper in the ten-year window surrounding its publication date. Since the quality of the earlier work is set before this publication date—and other unobserved determinants of citations are arguably orthogonal to it—any sharp increase in citations around it plausibly signals an exogenous shift in how the author's work is perceived.[4]

To study the Matilda effect, we compare the citation trajectories of male and female co-authors on the *same* top-five paper. This eliminates heterogeneity in the visibility men and women receive from publishing for the first time in a top-five economics journal. Assuming that the per-manuscript quality of the female co-author's pre-top-five research is at least as high as comparable work by her male co-authors, fewer citations to her papers suggests a Matilda effect.

We find compelling evidence of both the Matthew and Matilda effects as well as their interaction. Before jointly co-authoring a top-five paper, women and men publish in journals of comparable quality, but women's research is cited 0.5–1 fewer times each year, consistent with a Matilda effect. Afterwards, both co-authors experience an increase in citations to this work—worth about 1–3 citations per year for each paper—evidence of a Matthew effect. Strikingly, the Matthew effect is larger for women than it is for men. Before their first top-five publications, female co-authors' pre-top-five research is cited less than their male co-authors' pre-top-five research. Afterwards, this gender gap *reverses*, suggesting that a Matthew effect at least partially offsets earlier under-recognition caused by a Matilda effect.

Our findings are robust across a variety of specifications, including among papers published in economics and non-economics journals, among authors who solo-authored their first top-five paper, and after controlling for *JEL* codes, number of co-authors, cited journals' impact factors, authors' institutional ranks, and their pre-top-five publication counts. The Matthew effect persists in a range of citation settings and is consistently stronger for women; however, it is specific to citations. We find no evidence that women experience a larger post-top-five Matthew effect on promotions to professor, editorial appointments, the numbers and dollar amounts of grants awarded, authors' employers' reputations and the numbers of seminar invitations received.

To better understand the mechanisms behind our results, we develop a model of citation behaviour that focuses on the decision of whether and when to cite previously published work by authors who later publish a top-five paper. The model predicts the emergence of a Matthew effect—and its gendered dimension—through three complementary channels: (i) increased awareness of an author's earlier work following his top-five publication (visibility/awareness); (ii) a lower threshold for relevance applied when deciding whether to cite that work, a behaviour known as "strategic citation" (see, *e.g.*, Rubin and Rubin, 2021); and (iii) a greater incentive to build upon that earlier work, a phenomenon we label "strategic publication".

As shown in Lemma 3, we can isolate the combined effects of strategic citation and strategic publication by restricting our sample of citing papers to those whose authors were demonstrably aware of the top-five author's earlier work prior to his first top-five publication. The intuition is straightforward: if a citing author was aware of the top-five author's earlier research before her first top-five publication, but only cited it afterwards, her citation cannot be attributed to increased visibility/awareness. Instead, it likely reflects a shift in strategic behaviour—*i.e.*, a relaxation of the relevance threshold (strategic citation) or a re-evaluation of the value of building on the earlier work (strategic publication). In both cases, the decision to cite is a direct response to the author's change in perceived status, revealing it as the primary mechanism behind the Matthew effect.

We implement this strategy by examining six settings that distinguish between comparable groups of

[4]Note that we only 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 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 journals outside the top five.)

citing authors who are more likely (high awareness treatment) *vs.* less likely (low awareness counterfactual) to have been familiar with an economist's earlier work prior to his first top-five publication. In the first setting, we focus on whether top-five authors cite one of their own earlier papers in their top-five articles (high awareness treatment), as these *self-cited* papers are probably more visible to the relevant citing literature before the top-five paper is published compared to other papers by the author that are not self-cited (low awareness counterfactual).[5] In the second setting, we compare citations from citing papers that were published in the same journal as the papers they cite (high awareness treatment) to citing and cited papers that were published in different journals (low awareness counterfactual), on the assumption that editors and referees draw submitting authors' attention to related work the journal has already published.[6]

In our third and fourth settings, we exploit the fact that economists who release their top-five articles as NBER working papers or enter the academic job market as highly visible "star" candidates (high awareness treatments) typically receive substantial attention well before their work appears in print.[7] These authors' earlier contributions are therefore likely to have been visible to peers prior to journal acceptance, so a Matthew effect following top-five publication most likely reflects a change in the incentives to cite. In contrast, when the top-five article is *not* disseminated through NBER or the author is a job market candidate but *not* a star (low awareness counterfactual), the journal publication itself constitutes the first major visibility shock for its author.

In our fifth and sixth settings, we exploit professional links between citing and top-five authors. The colleagues and co-authors of a top-five author (high awareness treatment) are almost certainly aware of his pre-existing research before his first top-five article is published. Thus, a spike in citations from these people following the top-five publication date cannot be explained by new information, and instead reflects a strategic response to his improved status within the profession. In contrast, researchers with no close connection to the top-five author (low awareness counterfactual) are far more likely to learn about his pre-existing research because the visibility shock of publishing in a top-five journal brought attention to it.

In all counterfactual and almost all treatment 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.[8]

Nevertheless, 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 all treatment environments, where citing papers are always familiar with authors' pre-top-five work, so the effect is caused by strategic behaviour. In contrast, we document an obvious gender difference in every counterfactual environment, 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 that each treatment and counterfactual pair only differ in ways that are independent of top-five author gender, our combined results suggest that the visibility shock of a top-five publication helps to offset an earlier Matilda effect caused by limited awareness. However, when awareness of women's research is already high, the Matthew effect confers them no

[5] This is especially likely to be true for economists, as they typically present and circulate pre-print versions of their future top-five papers years before formal publication, making these papers—and any earlier work they cite—visible to others well in advance.

[6] Anticipating this, authors may also pre-emptively identify and cite related papers published in journals they intend to submit to.

[7] We encourage particular caution when interpreting the results from environment four, as our samples of women in both the treatment and counterfactual settings are small.

[8] The only two environments without a Matthew effect at all are: (i) citing and cited papers published in the same journal; and (ii) citing and cited papers by co-authors. Its absence may suggest that citing authors in these settings are less able and/or have fewer incentives to engage in strategic behaviour.

advantage relative to comparable men.

We emphasise once again that differences between high visibility treatment environments and their low visibility counterfactual environments must be unrelated to top-five author gender in order to come to this conclusion. Nevertheless, available evidence supports this assumption. First, gender differences in the characteristics of citing, cited and top-five papers are roughly similar in treatment and counterfactual environments, particularly in the first, third, fourth, fifth and sixth settings. 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.

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; and (ii) re-emerge when the event date is the date a future top-five article was released in the NBER working paper series. Our empirical evidence supports both predictions: there are no gender differences in the Matthew effect around the publication dates of authors' second top-five papers; in contrast, they re-emerge around the dates authors release their future top-five papers in the NBER working paper series.

Our findings have important implications for how visibility shapes recognition and gender equity in research and academia. Because top-five publications disproportionately boost recognition of women's prior work *only* when that work was under-recognised, we believe shining a light on talented women can play a powerful equalising role in science while also lifting the overall quality of its research. More broadly, our results highlight that mechanisms like the Matthew effect are not inherently gender-neutral; they can reinforce or offset existing disparities depending on the visibility landscape in which they operate.

This paper makes several contributions to the literature. In particular, we are the first to theoretically and empirically examine the interaction between the Matthew and Matilda effects in shaping scientific recognition. Prior work has documented the Matthew effect in 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).[9] Meanwhile, a growing body of research finds gender disparities in recognition consistent with the Matilda effect: women receive fewer prizes and awards relative to their representation in nomination pools (Lincoln *et al.*, 2012); grant reviewers score women's applications lower than they score men's (Royal Statistical Society, 2025), which other evidence suggests may structurally disadvantage them in later funding competitions (Fiorentin *et al.*, 2023);[10] women are also subject to tougher evaluation standards in peer review (Alexander *et al.*, 2023; Card *et al.*, 2020; Hengel, 2022; Hengel, Moon, and Tol, 2024; Krawczyk and Smyk, 2016), and tenure committees discount their contributions to collaborative work (Sarsons *et al.*, 2021). Our paper brings these two strands together by offering novel evidence that processes generating cumulative advantage can mitigate existing gender disparities in academic recognition.

Our second contribution is specific to the literature on gender disparities in citation practices—a central mechanism through which the Matilda effect operates. Prior work consistently shows 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).[11] We add to this research by

[9] Most relevant to our paper, Azoulay, Stuart, and Wang (2014) estimate the impact of the Howard Hughes Medical Institute Investigators award on citations to 424 winning scientists' previously published papers. Our own study complements their research by (i) incorporating a sample of researchers that is almost 10-times larger; and (ii) theoretically and empirically analysing the gendered impact of the shock.

[10] Royal Statistical Society (2025) does not find a gender awarding gap for regular funding applications despite the tendency for grant reviewers to rate women's applications lower than men's.

[11] Among the population of research papers, men's under-citation of female-authored papers probably (although not necessarily) leads to lower per paper citation rates for women compared to men (see, *e.g.*, Larivière *et al.*, 2013). Nevertheless, some studies have observed per paper citation gaps favouring women in the (selected) population of articles published in

showing that poorer quality information about female-authored papers may partially or wholly drive men's under-citation of women.

Third, our study complements research exploring how uncertainty creates and exacerbates gender outcome gaps. Prior work suggests that decision-makers discriminate less against women as they gain information about their true performance. For instance, Bohren, Imas, and Rosenberg (2019) show that biased beliefs about women's ability drive initial gaps in evaluations, but they narrow as evaluators learn more.[12] 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 particular journals. 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 an emerging literature on strategic behaviour in academic knowledge production, particularly in decisions to cite and publish (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). Much of this work studies how incentives, norms, and reputation concerns shape citation decisions. Consistent with this research, we find that citing authors respond to status shocks not only by becoming more aware of prior work, but also by engaging more with the work they were already familiar with—a form of strategic behaviour. We extend the literature by showing that citing authors with strategic motivations do not discriminate by author gender; only when they previously lacked awareness of women's research do their citation patterns favour female-authored papers. These results underline the importance of considering gender not only in outcomes but also in the behavioural and informational mechanisms that shape them.

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 new and 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.[13] From the datasets analysed in Hengel (2022) and Hengel, Moon, and Tol (2024), 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,

certain journals (see, *e.g.*, Card *et al.*, 2020; Hengel, Moon, and Tol, 2024). These latter gaps were likely formed because female-authored papers were held to higher standards in the decision to *publish*—that is, female-authored papers were only accepted if their quality (and hence anticipated citations) was especially high—and this more than offset the bias in favour of men in the decision to *cite* (see Hengel, Moon, and Tol, 2024, for further discussion). Interestingly, our findings provide further nuance by suggesting that increases in information over the life-cycle of a paper can dramatically reduce the impact of bias in the decision to cite on women's per paper citation rates.

[12] 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 authors publish other papers in top-five journals).

[13] We define full-length, original research articles as any non-errata/corrigenda/editorial article published with an abstract (or in some instances a submit-accept date), excluding *Papers & Proceedings* issues of the *American Economic Review*. See Appendix B for further details.

awards, employment and seminar invites—were retrieved by hand-collecting and digitizing CVs.

Our final database covers the entire pre-top-five publication histories of 3,929 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.[14] First, women remain starkly under-represented among first-time top-five authors. Despite a modest increase in female representation over time, the gender gap has remained large and persistent. As shown in Graph (A), only 14 percent of authors are women; their shares range from 9 percent at *ECA*, 12 percent at the *JPE* to 15–16 percent at the *AER*, *QJE* and *REStud*. Graph (B) displays women's representation by primary *JEL* code of their first top-five paper. Even in areas where women are relatively more active—such as health, education, development, and environmental economics—their share among top-five authors rarely exceeds 25 percent. In traditionally male-dominated fields like macroeconomics, finance, and microeconomics, the proportion is significantly lower.

Second, gender disparities emerge not just in representation but also in career dynamics. Graphs (C) and (D) plot counts of years and papers between authors' first publications and their first top-five papers, respectively. On average, women produce 2.2 fewer pre-top-five papers and publish their first top-five paper roughly 1.2 years earlier than men.[15] 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.

Despite publishing fewer papers before their top-five breakthrough, women's earlier work is cited about the same—and even slightly more in later years—and published in higher-impact venues. In Graph (E) we plot citations per paper: in almost every year after 1995, women's papers are cited more than men's.[16] 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. Combined, these results suggest that top-five women's early research output is at least as strong as men's.

Finally, co-authorship and institutional differences highlight additional dimensions of stratification. Over time, co-author counts have increased for everyone, with no consistent gender difference (Graph (G)). In the 1990s, papers were co-authored by 1.8 people, on average; by 2015 that figure had risen to 2.5. Meanwhile, authors based at lower-ranked institutions tend to publish more before their top-five début (Graph (H)).[17] However, the pattern does not systematically vary by author gender.

Using the information we collected from authors' CVs, Figure 2 presents career trajectories and professional visibility of top-five authors over the past three decades. First, women's representation in senior academic roles has steadily improved (see also Costa-Dias *et al.*, 2024; Lundberg and Stearns, 2019). Graph (A) shows that the share of female professors among top-five authors has risen across all employer ranks, increasing from roughly 4–5 percent in 1990 to 12 percent by 2015. Meanwhile, both men and women have gravitated toward more prestigious institutions, but the trend is especially pronounced for women. In the early 1990s, men's employers ranked 16 positions ahead of women's; by 2001, women

[14] Although we focus on a highly selected group of economists, many of these facts are consistent with patterns found elsewhere (see, *e.g.*, Ayarza and Iriberri, 2025; Costa-Dias *et al.*, 2024; Doleac, Hengel, and Pancotti, 2021; Ductor, Goyal, and Prummer, 2023; Funk, Iriberri, and Venus, 2025; Ginther and Kahn, 2004; Hengel, 2022; Hengel, Moon, and Tol, 2024; Lundberg and Stearns, 2019).

[15] On average, men and women publish their first top-five papers 6 and 7 years after completing their PhD, respectively, although this gap has completely closed in later years (see Figure B.3 in Appendix B.2).

[16] Across all years, women's and men's pre-top-five journal articles are cited 69 and 72 times, respectively. Among authors who published their first top-five paper after 1990, men's pre-top-five articles are cited, on average, 66 times; the corresponding figure for women is 72.

[17] Institution in Graph (H), refers to the highest ranked institution among all institutions associated to an author's on her 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).

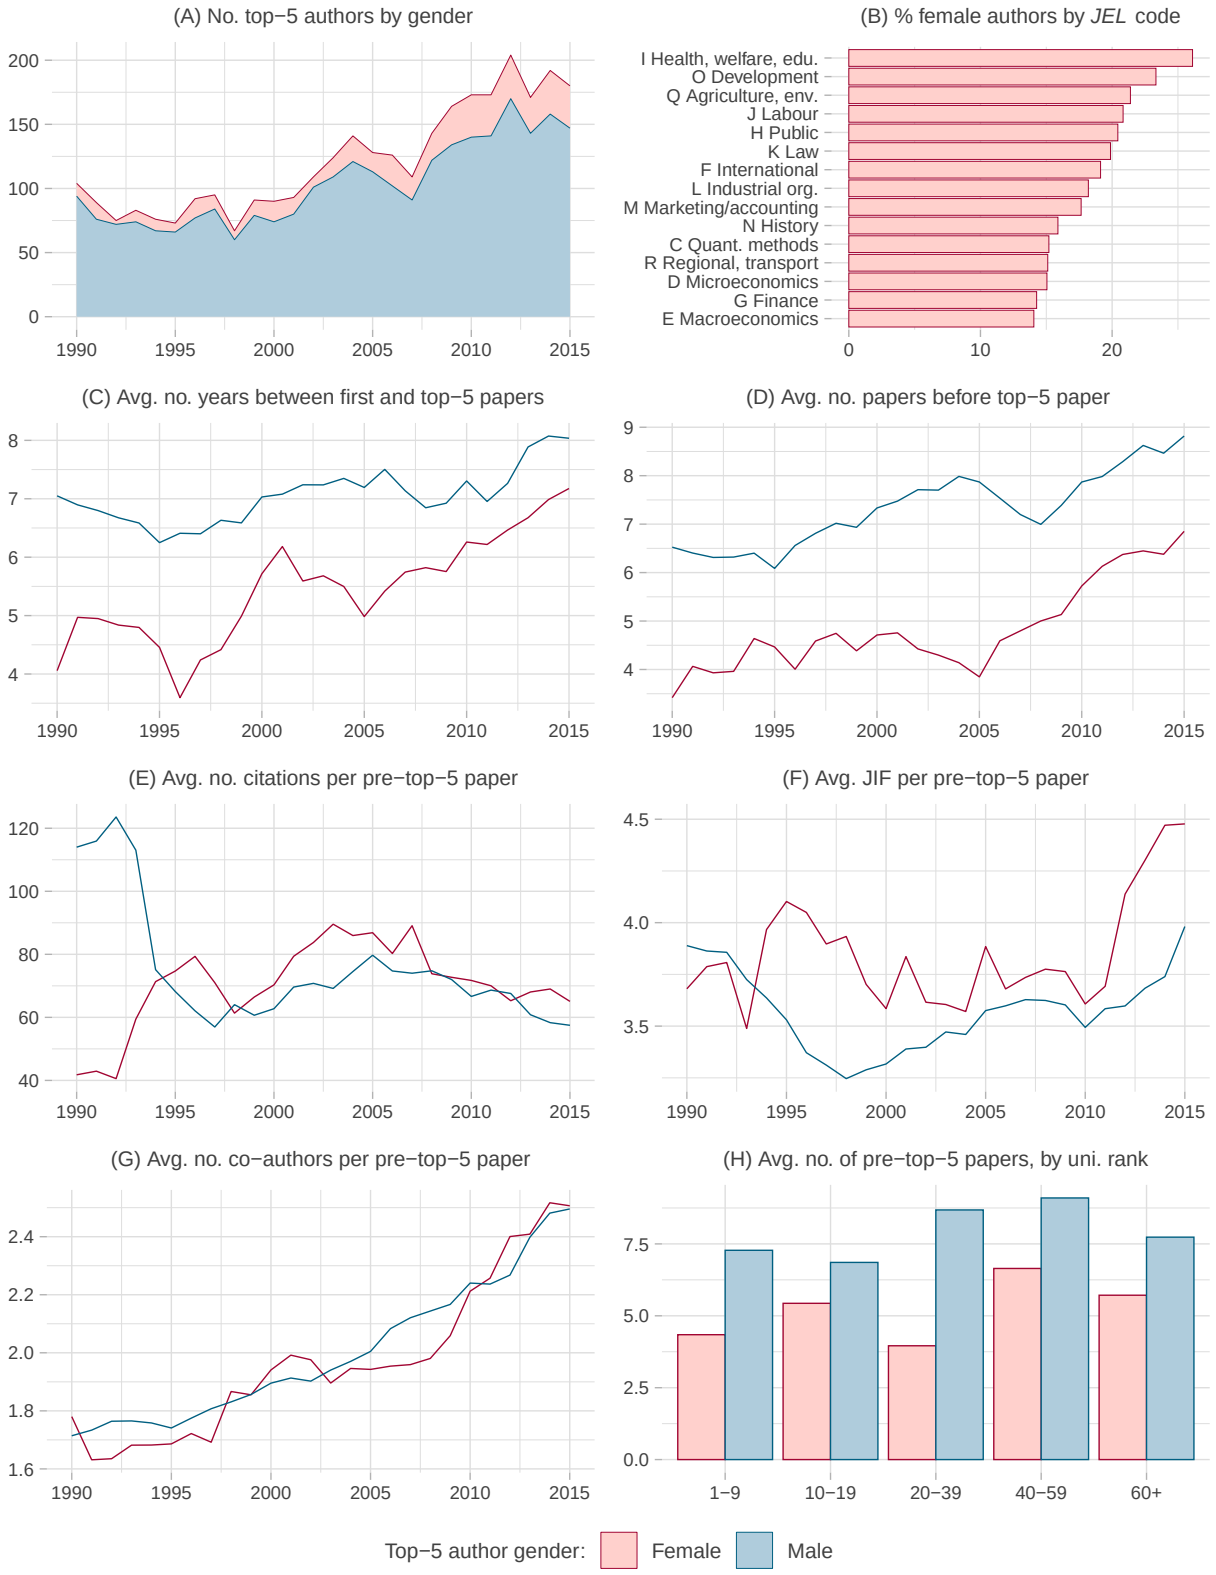


Figure 1: Characteristics of pre-top-five publications

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 share of women by the primary *JEL* code of their first top-five paper. Graphs (C) and (D) are the average numbers of years and published papers between authors' first publication and their first publication in a top-five journal. Graphs (E) and (F) are the average number of citations and journal impact factors per pre-top-five paper. Graph (G) shows the average number of co-authors per pre-top-five paper. Graph (H) is the average number of pre-top-five papers by authors' institutional rank. Time series data are five-year moving averages.

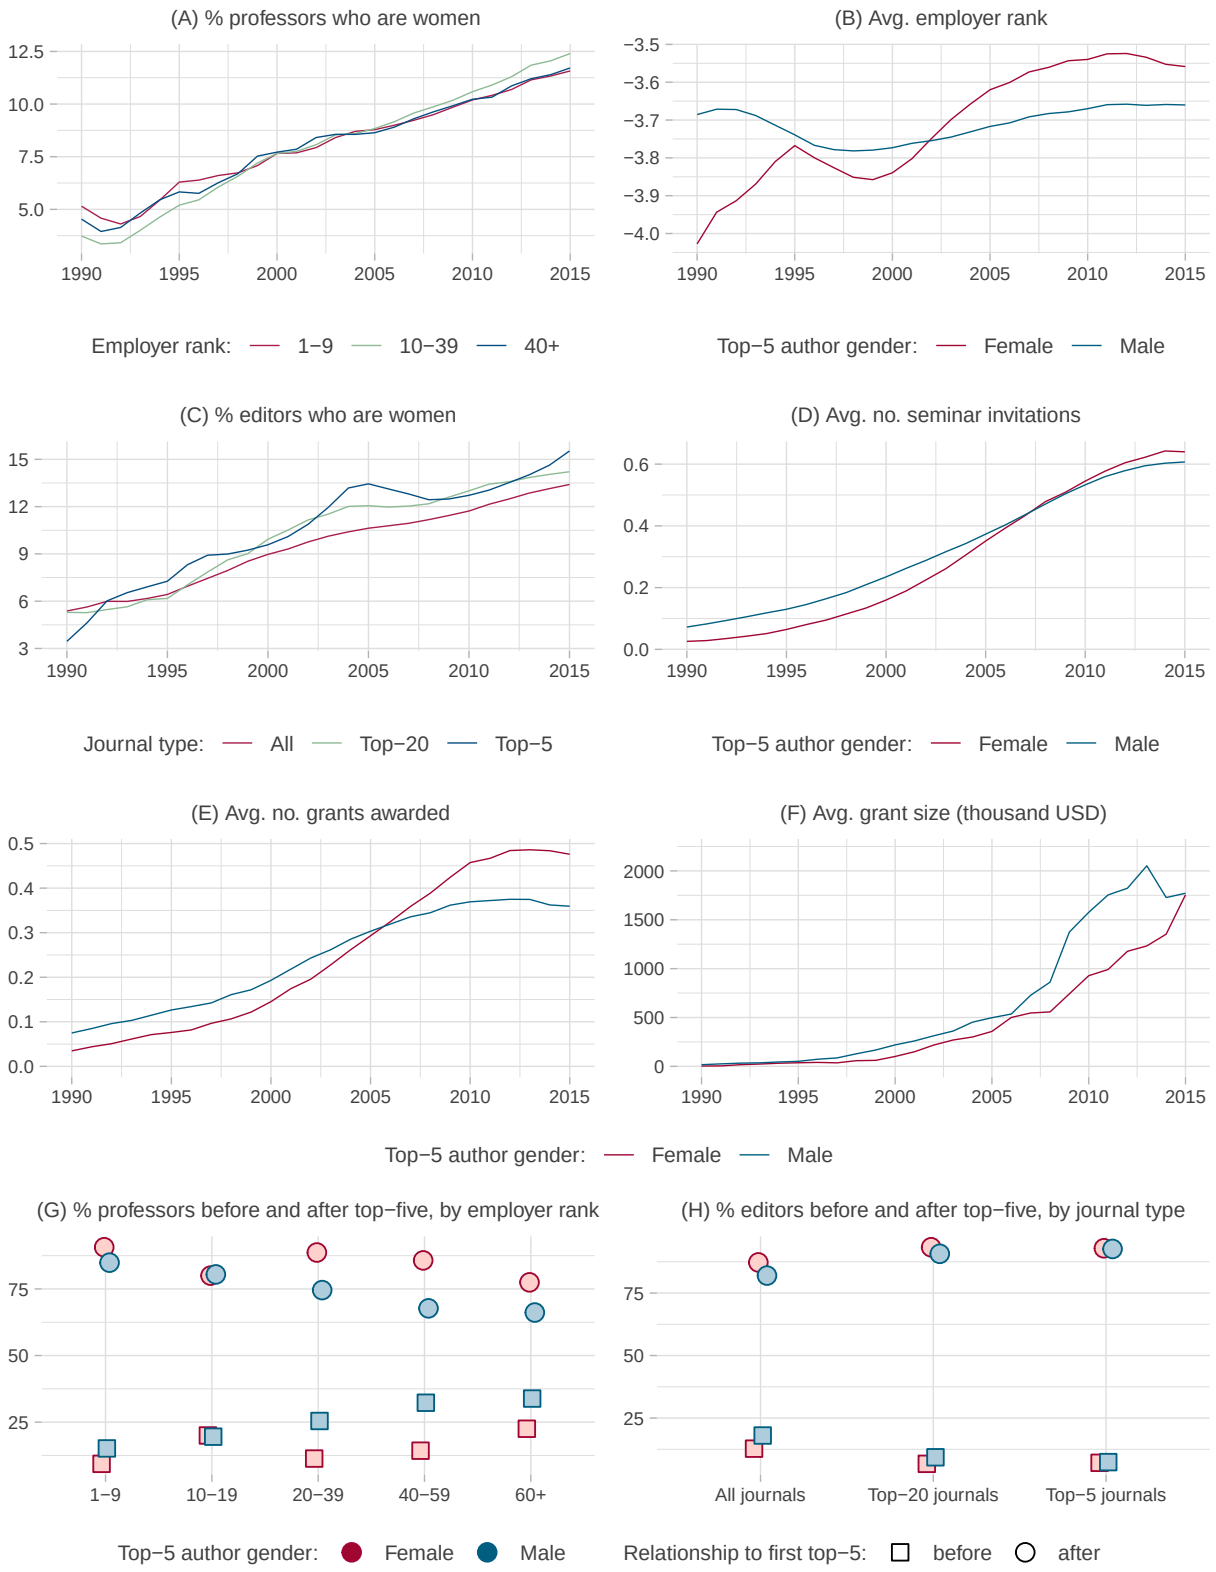


Figure 2: Characteristics of top-five authors

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. Time series data are five-year moving averages.

were at higher ranked institutions than men (Graph (B)).[18]

Second, the gender gap in visibility and recognition beyond publications has also improved over time. Graph (C) shows that women's representation among journal editors has risen across all tiers—from 5 percent of editors of any journal 1990 (and just 3 percent of top-five editors) to 13 percent by 2015.[19] Similarly, men gave more seminars in earlier years, but women have since caught up, and they now give slightly more (Graph (D)).

Third, 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.[20] 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.

Finally, career milestones tend to occur later for women. Graphs (G) and (H) plot the shares of authors who were promoted to professor and named journal editor, respectively, before *vs.* after their first top-five publication. Men are more likely to hold these roles beforehand, while women are more likely to attain them afterwards. Notably, the professorship gap widens as institutional rank declines institutions (Graph (G)), suggesting that women at less prestigious universities face greater barriers to early promotion—a dynamic consistent with the Matilda effect.[21]

3 The gendered Matthew effect

3.1 Conceptual framework

Top-five publications are widely regarded as signals of research quality and prestige (Heckman and Moktan, 2020); as a consequence, the increase in visibility and status that they confer likely triggers a Matthew effect. To study the effect, we examine citations to an author's pre-existing work around the date he publishes for the first time in a top-five journal.

While selection into top-five journals is clearly non-random—reflecting talent, effort, and opportunity—the exact timing of the breakthrough is arguably orthogonal to unobserved determinants of citations to earlier work, including quality and content. As a result, a discontinuous rise in citations to an authors' pre-existing research following the publication of his first top-five paper likely reflects a increase in recognition brought about by the event.

To quantify this effect, Equation (1) measures how citations to paper j, published before the top-five paper k, change after the author i publishes 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 citations_{jt} is the number of citations to paper j at time t, and $\text{citations}'_{jt}$ denotes the counterfactual citation count in the absence of k.

Equation (1) relies on the assumption that citations to j would not have changed between times $t = -1$ and $t = 0$ had k not been published. That is, any "bump" in citations at $t = 0$ is attributed to the publication of k. Under this assumption, $\Delta_{ijk}^{\text{Matthew}} > 0$ indicates a boost in recognition attributable to

[18] 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).

[19] 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 Journal*s.

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

[21] 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.

the author's top-five breakthrough. We interpret this as evidence of the 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).

The Matthew effect may arise through several channels. One possibility is that a top-five publication improves awareness of the author's earlier work. Before the high-profile success, some researchers may have overlooked or undervalued an author's previous papers simply because they were unfamiliar with her research or unsure of its relevance; afterwards, knowledge of the author's skill and talent increases, drawing retrospective attention to her earlier work and increasing its citations. In a second channel, the new publication shapes the direction of future research. For example, it may influence incoming PhD students to pursue related topics, thereby increasing demand for citations to the author's prior contributions. Finally, a top-five publication often enhances an author's professional standing and influence within his field. As a result, other researchers may have stronger incentives to build on and/or cite the author's earlier work—especially if the author is now more likely to review submissions, serve on hiring committees, or write tenure letters.

Equation (2) compares the Matthew effect of men and women who experience an equivalent shock from a top-five paper k:

$$\Delta\Delta_{ijk}^{\text{Matthew}} = \mathbb{E}\left[\Delta_{jFk}^{\text{Matthew}} - \Delta_{jMk}^{\text{Matthew}} \mid \text{shock}_k\right], \tag{2}$$

where shock_k captures variation in the degree of visibility or professional standing conferred by k.

Assuming shock_k is appropriately controlled for, Equation (2) is only non-zero if there are gender differences in how the publicity of a top-five paper k relates to characteristics of author i or her subsequently cited paper j. Such differences may arise if female-authored work is systematically under-recognised due to the Matilda effect. Alternatively, publishing a high impact paper could boost women's citations more than men's if there are greater career benefits from strategically citing them—*e.g.*, because women are more likely than men to become journal editors after publishing a top paper.

3.2 Estimation strategy

To estimate the effects in Equations (1) and (2), we adopt an event study approach that exploits the timing of an author's first publication in a top-five journal as a shock to perceived status and visibility. Specifically, we regress the annual number of citations received by author i's previously published work j—whose quality and contribution remains constant over the estimation period—on dummy variables for each year surrounding the publication of his first top-five paper k:

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

citations_{ijt} is the number of citations received in year t by paper j (excluding citations from k), published by author i prior to k.[22] female_i is an indicator equal to one if i is a woman, and captures baseline gender differences in citations. The event-time dummies $\mathbf{1}[t = t']$ are centred on the year of the top-five publication $t = 0$.[23] shock_k captures heterogeneity in the shock of visibility from publishing k. We estimate Equation (3) using OLS over a window spanning the 10 years before and after k's publication

[22] To ensure citations do not rise mechanically around $t = 0$ because the top-five paper k itself cites j_g, citations from k are excluded from citations_{ijt}.

[23] We normalize $t = 0$ to the publication date of the author's first top-five article, which we interpret as the timing of the shock. While some researchers may become aware of the article before formal publication—through working paper circulation, seminars, or conferences—others may only discover it afterwards. Thus, the actual timing of exposure is heterogeneous. As a robustness check (see Section F.1), we re-estimate our event study using $t = -1$ as the normalization point and obtain nearly identical results, suggesting that our findings are not sensitive to the precise timing of the event.

date, with standard errors clustered by j.

In our preferred specification, we account for shock_k using fixed effects for each top-five publication; thus, γ_0 compares male and female co-authors on the *same* top-five paper. To cover a broader sample of authors—including those who solo-authored k—we also estimate specifications that control instead for observable attributes of the top-five k, including its publication year, journal, number of co-authors, citation count (asinh-transformed), and a proxy for author prominence based on the number of prior top-five publications by the most prolific co-author.

The coefficients β_0 and $\beta_0 + \gamma_0$ capture the average male and female Matthew effects defined in Equation (1). As discussed in Section 3.1, these effects are identified under the assumption that citations to earlier work would have remained stable at time $t = 0$ had the top-five paper not existed. 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 is arguably orthogonal to unobserved determinants of citations to earlier work. As a result, sharp changes in citations to previously published papers that coincide with $t = 0$ are driven by the shock of attention from publishing k rather than, *e.g.*, changes in the content and merit of the cited material.[24]

Crucially, γ_t captures gender disparities in recognition patterns. Assuming that the per-unit quality of a female co-author's pre-top-five research is at least as high as comparable work by her male co-authors, γ_t captures baseline gender disparities in recognition.[25] A negative γ_t suggests that women's research is cited less than comparable research by their male counterparts, pointing to a Matilda effect—*i.e.*, the under-recognition of equally meritorious work by women.

3.3 Results

Figure 3 illustrates how citations to authors' earlier work evolve in the years before and after their first top-five publication, separately for men and women. Two specifications are shown: Graph (A) includes fixed effects for the top-five paper, enabling comparisons between male and female co-authors who experience an identical visibility shock; Graph (B) instead controls for characteristics of the shocks (*e.g.*, co-author count, citation count, author prominence, and journal-year fixed effects). Three key patterns emerge.

First, we find clear evidence of a Matthew effect. For both men and women, citations to earlier research increase sharply around the year of the first top-five publication ($t = 0$). Prior to this event, β_t and $\beta_t + \gamma_t$ are negative. Immediately afterwards, both coefficients turn positive and remain so for the rest of the estimation period. Assuming citations to previously published papers just after the event window would have evolved similarly to citations accrued just before it, this discontinuity suggests that publishing in a top-five journal brings attention—and consequently citations—to men's previously published work. For male authors, this effect is worth approximately 1–2 additional citations per paper per year.

Second, we document a substantial gender gap in citations prior to the top-five publication date, consistent with a Matilda effect. In the years preceding $t = 0$, female authors receive significantly fewer citations to their earlier work relative to comparable men. In Graph (A)—where we compare male and female co-authors on the same top-five publication—our results suggest that women's prior contributions are systematically less recognised relative to work by their male co-authors. Assuming both co-authors

[24] Our estimates capture the 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 Matthew effect and its gender difference are present—albeit much smaller in magnitude—even when authors publish in non-top journals (see Figure F.1 in Appendix F.1).

[25] We assume that shock_k appropriately absorbs heterogeneity in k (*e.g.*, manuscript quality). If this assumption fails, then gender differences in the characteristics of high impact papers create gender differences in the attention their authors subsequently receive.

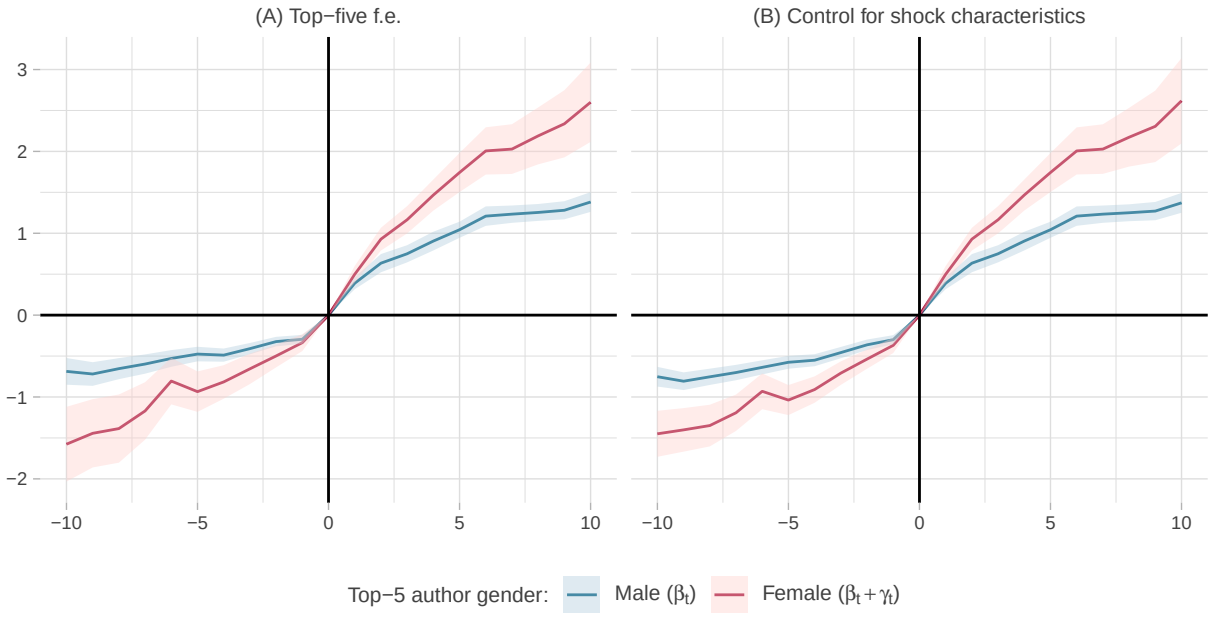


Figure 3: The Matilda and Matthew effects: baseline evidence and interaction

Note. 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. In Graph (A), we account for shock_k using a top-five fixed effects; in Graph (B), we instead control for the first top-five 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 95 percent confidence intervals from standard errors clustered by j_g.

produce comparable quality pre-top-five research, this evidence suggests a Matilda effect for all $t < 0$.[26]

Third, the Matthew effect is markedly larger for women (*i.e.*, γ_t is positive and statistically significant for all $t > 0$).[27] Around the year of the top-five publication, citations to women's earlier work increase more sharply than citations to comparable work by men. After publication, women's pre-top-five papers are cited at higher rates than men's. This reversal in the gender citation gap indicates that publishing in a top-five journal benefits women more than men. It is consistent with a correction in informational frictions: when women's high-quality work is initially overlooked, a status shock like a top-five publication can substantially increase its salience. We explore this idea further in Section 4.

We observe non-parallel trends in the citation trajectories for men's and women's earlier work leading up to $t = 0$: women's citation trajectories rise more steeply than men's in the years preceding their first top-five publication.[28] As we show in Section 4, the diffusion process of information is likely delayed but more concentrated for women. The steep rise in citations before $t = 0$ therefore probably reflects a catch-up process, whereby women's prior research—initially overlooked—gains recognition more rapidly as their visibility increases. Indeed, these pre-trends become notably flatter once we condition on environments where knowledge diffusion is plausibly faster and more complete—such as among co-authors or when an author's top-five paper cites his earlier research.

Our estimates are robust to a wide range of alternative specifications and sample restrictions. Results

[26] In Appendix C.1, we show gender differences in citations and journal impact factors (JIFs) to pre-top-five work with the same sample used to estimate Figure 3. Our results suggest that women's pre-top-five papers are cited more on a per-paper basis compared to their male top-five co-authors' work; we find no consistent or significant gender differences in co-authors' pre-top-five JIFs. We also replicate Figure 3 controlling for the quantity and quality of papers j published prior to the first top-five (Appendices C.12 and C.13); results are almost identical to those shown in Figure 3.

[27] β_t and γ_t represent deviations from baseline gender differences in citations captured by δ. In Figure 3(A), δ is 0.59 (standard error 0.28); in Figure 3(B), it is -0.02 (standard error 0.12). Neither estimate is significant at traditional thresholds.

[28] While we observe some pre-trends in Figure 3, this pattern is likely due to the slow diffusion of exposure—*e.g.*, via working paper circulation, seminar presentations and word-of-mouth—in the run-up to publishing in a top-five journal.

are similar when we estimate Poisson models (Appendix C.3) and persist across various sub-samples and control sets, as shown in Appendix C. Specifically, the pattern holds among authors who solo-authored their first top-five (Appendix C.7), and in samples of cited and citing papers published in both economics and non-economics journals (Appendices C.5, C.6, and C.7). We also obtain similar estimates when adjusting for the first top-five's primary and secondary *JEL* codes (Appendix C.11), the author's institutional rank at the time the first top-five is published (Appendix C.10), the author's total number of pre-top-five publications (Appendix C.12), and characteristics of the earlier paper, including its journal impact factor and number of co-authors (Appendices C.13 and C.14). In all cases, results remain closely aligned with those reported in Figure 3.

4 Decomposing the Matthew effect

4.1 Theoretical framework

In this section, we construct a simple theoretical framework to guide our interpretation of the Matthew effect across genders 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 (*i.e.*, first top-five), where $g = F, M$ denotes the gender of the author. Additionally, let $l \in \mathcal{L}$ denote a different paper (or idea for a paper) that potentially cites j_g, where $\mathcal{L}$ is the Dedekind complete body of research ideas.

An article l cites another article j_g only if three conditions are met: (i) j_g is sufficiently relevant to l, (ii) the authors of l are sufficiently aware of j_g, and (iii) they believe that citing j_g strengthens the publication value of their own work.[29] The second and third conditions—awareness and perceived value—can themselves depend on the visibility and perceived quality of j_g, which may be shaped by earlier citations or publications. 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, $\theta_{j_g l}$ is the relevance of j_g to l, $\lambda_{j_g l}$ is a binary variable equal to one if the authors of l are aware of j_g, and $\phi_l(\lambda_{j_g l})$ captures the authors of l's beliefs about the value of publishing l given their awareness of j_g. The thresholds of relevance and value required for l to cite j_g and its authors to produce and publish l are $\tilde{\theta}_{j_g l}$ and $\tilde{\phi}_l$ respectively.

We additionally assume that: (i) the perceived value of citing j_g ($\phi_l(\lambda_{j_g l})$) increases with the visibility of j_g, λ_{j_g} (*e.g.*, via prior citations or publications); (ii) awareness of j_g ($\lambda_{j_g l}$) and perceived value of citing it (ϕ_l) do not decline after the publication of the first top-five paper k; and (iii) in the absence of k, all other citation-relevant parameters—relevance ($\tilde{\theta}_{j_g l}$), awareness ($\lambda_{j_g l}$), and perceived value (ϕ_l)—remain unchanged. (i) and (ii) presume that the publicity generated from the first top-five k does not reduce l's propensity to cite j_g or its authors' beliefs about the value of publishing. (iii) assumes that the authors' thresholds and beliefs are stable, and any increase in the likelihood of citing j_g after k's publication is attributable to increased awareness or perceived value driven by k.

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 $l \in \mathcal{L}$.

[29] "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.

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* $\hat{l} = (\underline{l}, \bar{l})$.

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 perceived 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 the top-five k made the authors of l 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".[30]

To compare the Matthew effects of men and women, suppose i_M and i_F co-author k and additionally assume that publishing a first top-five k eliminates gender differences in awareness $\lambda_{j_g l}$—and, by extension, in the perceived value of citing ϕ_l. In other words, once k is published, citing authors l become equally aware of j_F and j_M, and any residual gender-based differences in the perceived value of citing either author vanish. 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 first top-five* 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-shock 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 the citing authors l are aware of previous papers j_g before the first top-five k is published.[31] A positive Equation (5) is therefore caused by strategic citations or strategic publications by l; and a positive

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

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

Equation (7) suggests a larger rise in strategic citations to women's previous papers 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:*

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*

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

When Lemma 3 applies, potential citing papers are already aware of j_g before the top-five author publishes her first first top-five paper k. As a result, increases in awareness $\lambda_{j_g l}$ cannot explain the Matthew effect. Instead, the effect must arise from a rise in strategic citations or strategic publications. Similarly, because the publication of the first top-five k does not alter citing paper l's familiarity with either j_F or j_M, a stronger Matthew effect for female authors must reflect a larger increase in strategic behaviour by citing authors l directed toward j_F.

In contrast, suppose Assumption 5 holds and $\Delta^{\text{Matthew}}_{jk} = 0$ and/or $\Delta\Delta^{\text{Matthew}}_{jk} = 0$. In this case, strategic behaviour by citing paper l either does not contribute to the Matthew effect or its gender difference, or its component parts cancel each other out.

4.2 Estimation strategy

Lemma 3 characterises environments in which the Matthew effect is attributable entirely to strategic behaviour by citing authors. To empirically assess the contribution of this mechanism—and its gender-specific component—we estimate Equation (3) across environments that differ in the likelihood that a citing paper l was already aware of an author's early work j_g before she published her first top-five paper k. That is, we compare citation responses in settings where awareness of j_g is plausibly held constant around k's publication date (the high-visibility "treatment" environment) to those where the publication event instead acts as a discovery shock (the low visibility "counterfactual"). This design isolates strategic citation responses by shutting down the awareness channel in the high-visibility treatment setting and comparing it to a low-visibility counterfactual that is observationally similar but retains the awareness channel.

We start by identifying characteristics of the earlier paper j_g that make it more likely that a potential citing paper l was already aware of it before its author published his first top-five paper k. In our first high-visibility treatment environment, we study the sub-sample of $j_g \in \mathcal{J}_g$ cited by k. Given common seminar and working paper dissemination practices in economics, the j_g cited in k are more likely to be known by a related citing paper l in the field of economics compared to the j_g that are *not* cited in k. 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.[32] Conversely, earlier j_g not cited in k serve as a lower visibility counterfactual in which publication of k likely operates as a discovery shock.

In our second high-visibility treatment, we restrict our sample to citing papers l that are published in the same journal as the papers they cite, j_g. The editors and referees at a journal where l is under submission

[32] 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)!

are probably familiar with related work that the journal has already published; consequently, they are more likely to flag these papers to l's authors during the review process. Anticipating this, authors may also pre-emptively identify and cite related papers published in journals they intend to submit to (for a discussion, see, *e.g*, Chorus and Waltman, 2016; Gazni and Didegah, 2021). Thus, every l published in the same journal as a relevant j_g should have been made aware of j_g—and given the opportunity to cite it—before being published. In contrast, l's awareness of a relevant j_g is not necessarily guaranteed when l and j_g are published in different journals.

We next identify characteristics of a top-five author that make it more likely that her past publications were well-known to a relevant citing paper l before her first top-five paper k was published. In our third treatment-counterfactual pair, we compare citations to j_g by authors who disseminated k as an NBER working paper prior to publication (the high-visibility treatment) with citations to j_g by authors who did not (the low-visibility counterfactual). 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).[33] As a result, k's authors—and consequently j_g—probably received a bigger visibility boost among economists when k was released as an NBER working paper than they did when it was actually published.

In a fourth high-visibility treatment, we study academic job market stars in economics. Each year, graduating PhD students in economics compete for academic positions through a highly visible process involving department "flyouts", faculty meetings and job-market paper presentations. While all job market candidates receive some exposure, a small number of "stars" capture a disproportionate share of attention and are more likely to be recognised by their peers early on in their careers.[34] As a result, potential citing authors should be more aware of job market stars' earlier work (high visibility treatment) compared to non-stars who are also on the market (low visibility counterfactual). To identify comparable job market candidates, we restrict the sample 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 define candidates as "stars" if they are affiliated to institutions ranked 1–9 at the time their top-five paper was published; authors employed at lower-ranked institutions are assumed to be "non-stars".[35]

Finally, we exploit professional links between cited and citing authors to identify situations in which authors' pre-existing work is already well known to relevant citing papers when their first top-five paper k is published. In the fifth and sixth treatments, we study citing authors who collaborate with or are colleagues of the top-five author, respectively. We assume citing and top-five authors are colleagues if an author of citing paper l was affiliated with any of the institutions listed on k on or before l's publication date;[36] they collaborate if they co-authored at least once before l was published. Professional links to a top-five author make it highly likely that a citing author was familiar with her earlier work j_g before k was published (high visibility treatment). In contrast, citing scholars who have never worked with her are more likely to first become aware of j_g after k is published (low visibility counterfactual).

If gender differences in the Matthew effect operate mainly through increased awareness, then high visibility treatment environments should consistently reveal little or no gender difference in the Matthew

[33] NBER regularly sends emails summarizing its new working papers to over 50,000 subscribers to its newsletter; on average, NBER working paper abstract pages are viewed 650,000 times every month (NBER, 2025).

[34] 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).

[35] We describe our ranking methodology in Appendix C.10. We also urge additional caution when interpreting results from the fourth environment, as the number of women who are job market candidates (either stars or non-stars) according to our definition is very small.

[36] Because Web of Science does not always map papers' co-authors to their specific affiliations, we technically define a citing and cited author colleague relationship to be one where the citing author is affiliated with any of the institutions listed on the cited author's first top-five paper.

effect, while low visibility counterfactual environments should instead reveal large differences. To ensure comparability between treatment and counterfactual environments, we estimate Equation (3) with fixed effects for the top-five article k.[37] Because identification in the first, third and fourth environments rely on research dissemination norms that are specific to economics, we restrict their treatment and counterfactual samples to the cited j_g and citing l published in economics journals; to focus the domain colleagues to people with the most relevant work, we similar restrict fifth environment samples.

4.3 Results

Figures 4, 5 and 6 estimate Equation (3) in "treatment" and "counterfactual" environments that differ in the degree to which citing authors were aware of economists' earlier work before they published their first top-five papers. Graphs in the left-hand panels are high visibility "treatment" environments; graphs in the right-hand panels are low visibility "counterfactual" environments. As argued in Section 4.2, Matthew effects in both environments could be caused by the strategic behaviour of citing papers, but only in low visibility counterfactual environments could they also be caused by a lack of information. Assuming that treatment and counterfactual environments' only meaningful gender difference is in the information citing papers have on women's earlier work, comparing the two establishes whether awareness or strategic behaviour drive their larger Matthew effects.[38]

In Figure 4 we identify high-visibility "treatment" environments using characteristics of earlier cited paper j_g that make it more likely that a potential citing paper l was already aware of it before its author published his first top-five paper k. In our first treatment-counterfactual pair, we compare the sub-sample of j_g cited by k (Graph (A)) to the sub-sample of j_g not cited by k (Graph (B)). In the second, we compare the sub-sample of citing papers l that are published in the same journal as the papers they cite j_g (Graph (C)) to the sub-sample of l and j_g published in different journals (Graph (D)). In both high-visibility treatment environments (Graphs (A) and (C)), men and women experience similarly sized Matthew effects after their first top-five papers are published. In contrast, both low-visibility counterfactual environments (Graphs (B) and (D)) suggest the Matthew effect is much larger for women than it is for men. Or in other words, the visibility shock of publishing in a top-five journal disproportionately benefits women *only* when awareness of their earlier work was limited; when awareness of it is high, gender differences disappear.

Figure 5 identifies high- and low-visibility environments using characteristics of top-five authors; it similarly suggests that women experience a greater Matthew effect because citing papers initially ignored their research. In high-visibility treatment environments, we study authors who disseminated their first top-five article as an NBER working paper (Graph (A)) or were job market stars (Graph (C)); our low-visibility counterfactual environments consist of authors who did not release their first top-five paper as an NBER working paper (Graph (B)) and job market candidates who were not stars (Graph (D)). In treatment environments, awareness of authors' work is already high at $t < 0$, meaning strategic behaviour drives the Matthew effect observed at $t = 0$; however, the effect does not differ by gender, suggesting that strategic behaviour does *not* affect women more than men. In contrast, gender differences return in low visibility counterfactual environments (Graphs (B) and (D)), where publishing in a top-five journal is the biggest visibility shock authors experience. As larger Matthew effects for women are only observed post top-five publication, they are likely caused by poor awareness of women's research before that.

Figure 6 investigates how professional proximity shapes awareness of authors' earlier work; it again emphasises that women disproportionately benefit from a Matthew effect, but *only* when awareness of their

[37] By definition, all authors in the fourth treatment-counterfactual pair solo-authored k, preventing us from accounting for shock_k in Equation (3) using fixed effects for k; instead, we control for k's year and journal of publication, co-author count and citation count (asinh).

[38] More specifically, we assume that any other dimension on which treatment and counterfactual environments differ is unrelated to author gender.

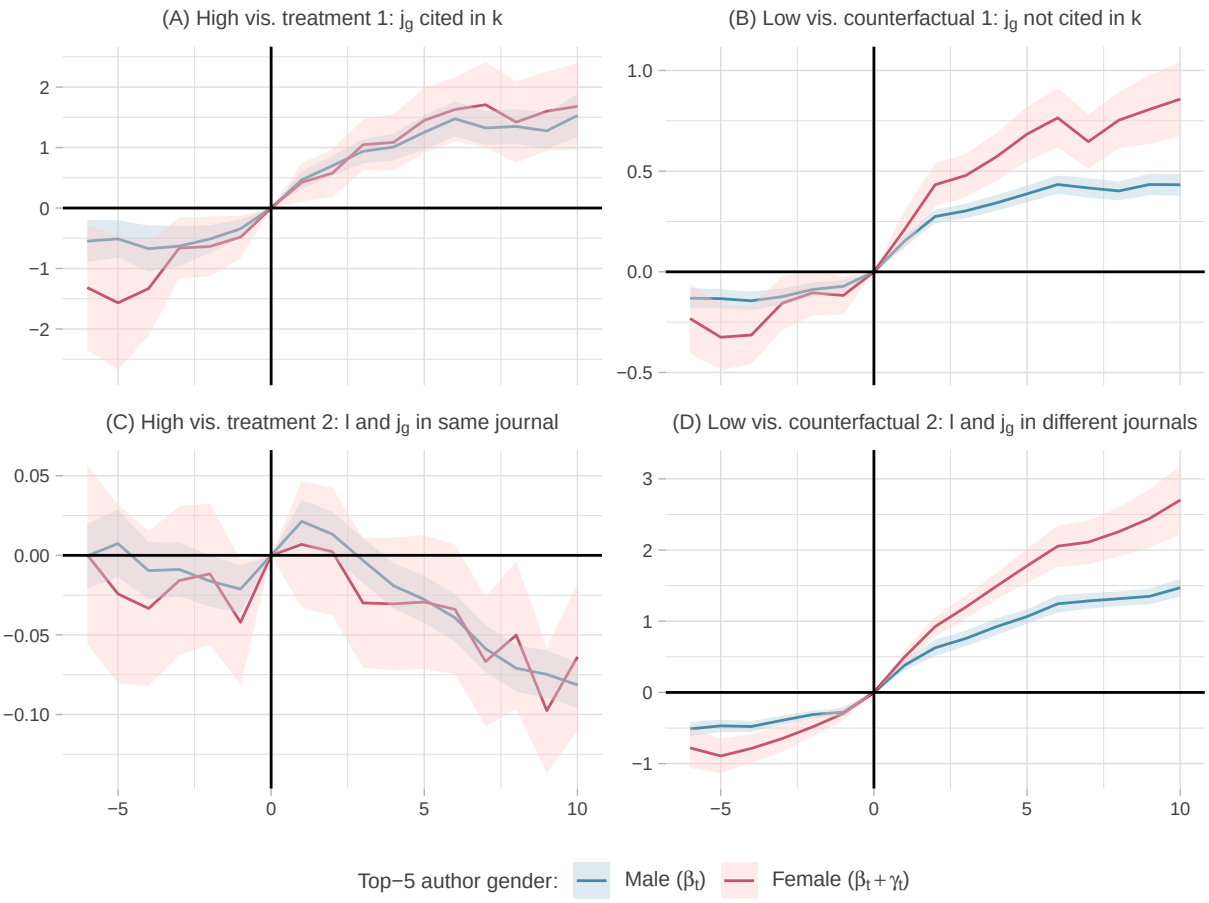


Figure 4: High vs. low visibility environments based on characteristics of j_g

Note. Figure 4 displays β_t and $\beta_t + \gamma_t$ from estimating Equation (3) with OLS and controlling for first top-five k fixed effects. The top row of graphs compare the sub-sample of cited papers j_g cited in k (Graph (A)) to the sub-sample of j_g not cited in k (Graph (B)). The bottom row of graphs compare the sub-sample of citing papers l that are published in the same journals as the papers they cite j_g (Graph (C)) to the sub-sample of l and j_g published in different journals (Graph (D)). In Graphs (A) and (B), l and j_g are restricted to papers published in economics journals; in all graphs, the estimation window is 6 years before and 10 years after k's publication date. Shaded areas represent 95 percent confidence intervals from standard errors clustered by j_g.

research is low. Our high-visibility treatment environments consist of citations from colleagues and co-authors, who should be more aware of authors' research before they publish in a top-five journal (Graphs (A) and (C)); their counterfactual environments analyse citations from less-connected researchers, for whom prior awareness was plausibly weaker (Graphs (B) and (D)). In high-visibility treatments, there are no gender differences in the Matthew effect—and indeed, no Matthew effect at all among citations from co-authors.[39] In contrast, citations in the low-visibility counterfactuals increase sharply around $t = 0$, and the spike for women is much larger than it is for men.

Taken together, Figures 4, 5 and 6 paint a consistent picture. In settings where prior awareness is high—whether because earlier work was directly linked to the top-five article, because it was widely disseminated through NBER or the job market, or because potential citing papers were personally connected to the cited author—the Matthew effect operates purely through strategic channels, and men and women benefit equally. In contrast, when prior awareness is low, women experience significantly larger citation increases. These results suggest that gender differences in the Matthew effect are driven not by strategic responses to status—which appear gender-neutral—but by differential visibility in the

[39] The lack of a Matthew effect stands in contrast to most other settings with the exception of citations from papers published in the same journal (Figure 4, Graph (C)). It may suggest that these settings reduce citing papers' incentives and/or scope to engage in strategic behaviour.

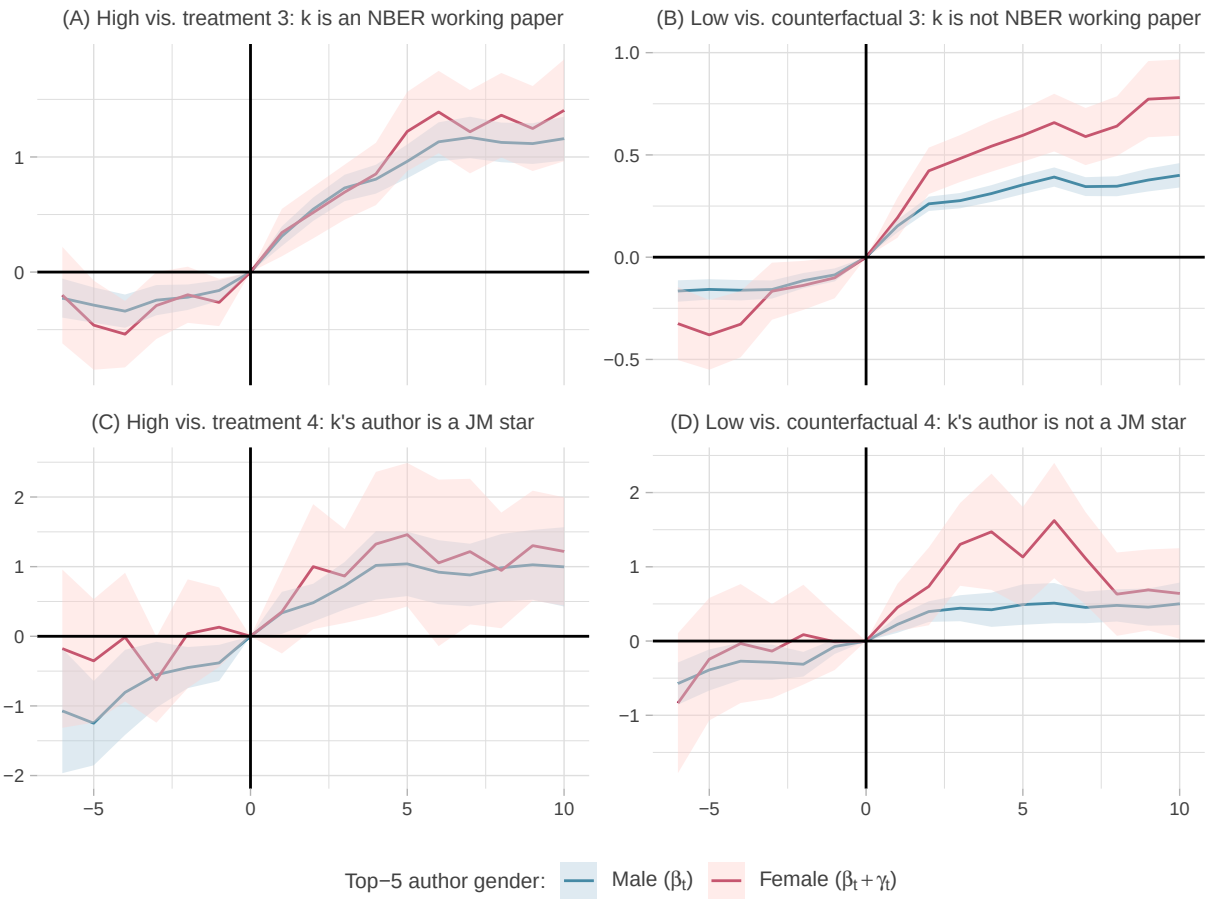


Figure 5: High vs. low visibility environments based on characteristics of k/k's author

Note. Figure 5 displays β_t and $\beta_t + \gamma_t$ from estimating Equation (3) with OLS. In the top row of graphs, we compare the citation dynamics of authors who released k as an NBER working paper (Graph (A)) to authors who did not (Graph (B)). In the bottom row, we compare job market (JM) stars (Graph (C)) to job market candidates who were not stars (Graph (D)). (See Section 4.2 for details on how we identify job market candidates and stars.) In the top row of graphs, we account for shock_k using fixed effects for k; in the bottom row, we instead control for k's year and journal of publication, co-author count and citation count (asinh) (see Footnote 37). In all graphs, l and j_g are restricted to papers published in economics journals, and the estimation window is 6 years before and 10 years after k's publication date. Shaded areas represent 95 percent confidence intervals from standard errors clustered by j_g.

early stages of authors' careers. In this way, the Matthew effect interacts with the Matilda effect: it helps correct prior under-recognition of women by advertising their overlooked work.

To conclude that greater awareness of women's research drives their larger Matthew effects crucially assumes no gender differences between treatment and counterfactual environments except in the information citing papers have on women's earlier work at $t < 0$. Available evidence supports this assumption. Tables D.1–D.6 (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, particularly in the first, third, fourth, fifth and sixth environments.

Gender differences in citing papers' strategic behaviour are also unrelated to selection into treatment—that is, 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 Figure 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

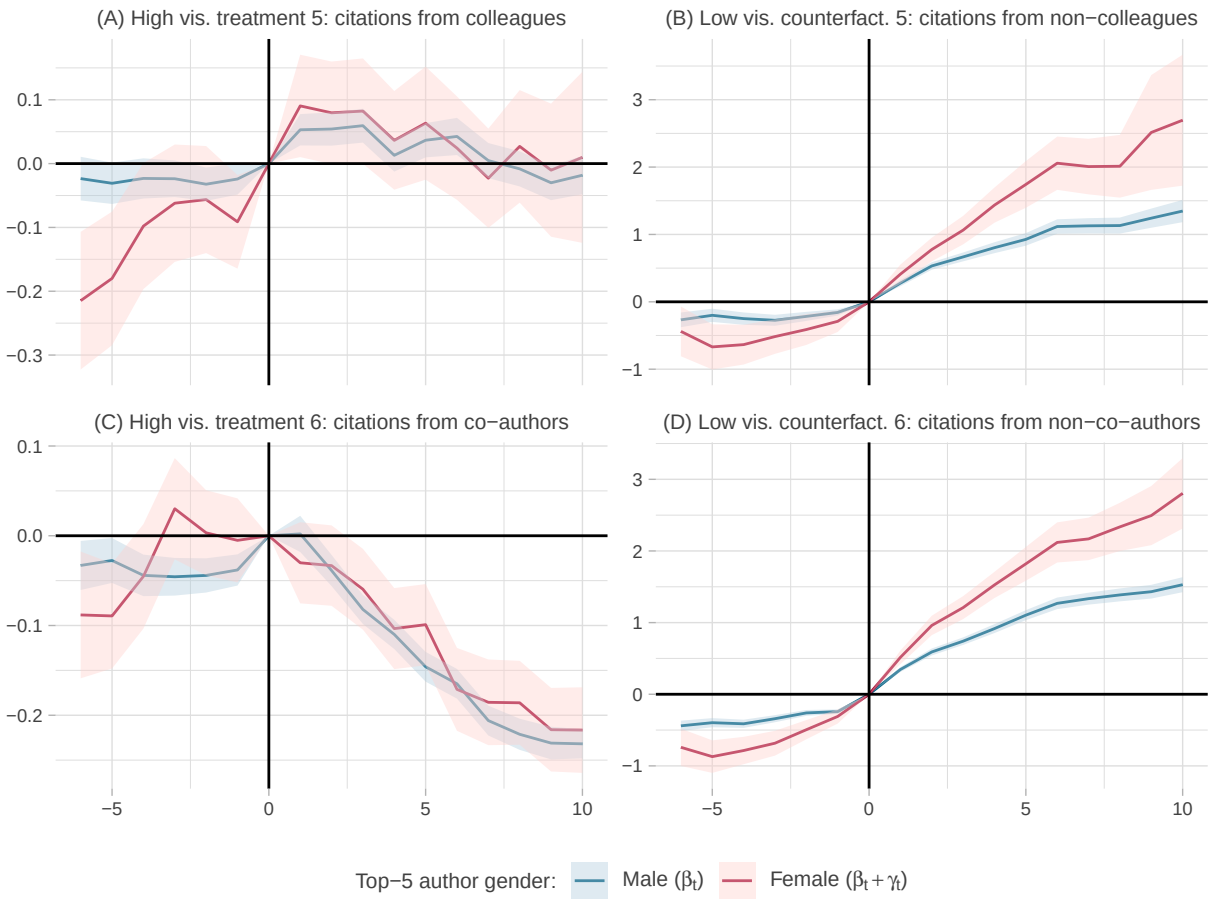


Figure 6: High vs. low visibility environments based on professional links between l's and j_g's authors

Note. Figure 6 displays β_t and $\beta_t + \gamma_t$ from estimating Equation (3) with OLS and controlling for the first top-five k fixed effects. Graphs distinguish citations received from current or former colleagues (Graph (A)) *vs.* non-colleagues (Graph (B)) and from co-authors (Graph (C)) *vs.* non-co-authors (Graph (D)). (See Section 4.2 for further details on how we identify co-authors and colleagues.) In Graphs (A) and (B), l and j_g are restricted to papers published in economics journals; in all graphs, the estimation window is 6 years before and 10 years after k's publication date. Shaded areas represent 95 percent confidence intervals from standard errors clustered by j_g.

in this probability around $t = 0$; however, it does not differ by author gender for any $t > 0$. We further 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.

Three key placebo checks also reinforce our conclusion that a lack of information about women's research causes their greater Matthew effects. In Appendix F.2 we bring the $t = 0$ event date forward in time to the date an author released her future first top-five article as an NBER working paper. This environment mirrors the one in Figure 5(A), except that potential citing papers are no longer necessarily aware of the author's earlier work at $t < 0$; as a result, gender differences in the Matthew effect should re-emerge. Indeed, this is precisely what we observe: Figure F.2 suggests a citation response that is much larger for women than it is for men.[40]

In Appendix F.3, we push the $t = 0$ event date back in time to the date authors publish their *second* top-five publications. Because potential citing papers were made aware of an author's earlier work when

[40] In Appendix D.2 we also 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).

she published for the first time in a top-five journal, visibility is high when she publishes there for a second time, and the subsequent Matthew effect will be driven by strategic channels, where men and women are treated equally. Again, this is exactly what we find: Figure F.3 suggests no gender difference in the Matthew effect at $t = 0$.

In Appendix D.5, we further estimate Equation (3) on the subsets of authors employed at institutions ranked 1–5, 6–9, *etc.* Gender differences in the Matthew effect are 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.

Finally, conclusions are similar when we reproduce Figures 4, 5 and 6 using Poisson models (Appendix D.3). They also continue to hold after controlling for journal impact factors and pre-k publication counts (Appendix D.4).

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; we then compare citation trajectories for male and female co-authors on the same top-five paper. We find that: (i) before the shock of publishing in a top-five journal, women's research is systematically under-cited relative to comparable research by their future male co-authors, suggesting an initial Matilda effect; and (ii) both co-authors experience a large bump in citations around the top-five publication date—evidence of a Matthew effect—but it is especially large for women.

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 six 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 most treatment environments we find clear evidence of a Matthew effect—suggesting that strategic considerations play a role in its general formation—although it never differs by author gender. In contrast, women's Matthew effects are always larger than men's in counterfactual environments. This pattern suggests that gender differences in the Matthew effect arise *not* because citing authors respond more strategically when women gain status, but instead because they were simply less likely to know about women's work to begin with.[41]

[41] Additional results in contexts where prior awareness is particularly limited sharpen this conclusion: the Matthew effect is especially pronounced for solo-authored women's papers (Appendix C.8) and when the top-five article represents the first such publication for all co-authors (Appendix C.9).

More broadly, our findings also indicate that the Matthew effect is not inherently gender-neutral. Prestige-driven mechanisms can either reinforce or mitigate existing disparities depending on the informational environment in which they operate. In settings where awareness is already high, status gains primarily activate strategic responses and do not differentially benefit women. Where awareness is limited, however, visibility shocks disproportionately raise awareness of women's prior work, helping to offset earlier under-recognition. This has implications for how citations accumulate over academic careers. It also suggests that policies aimed at increasing early-stage visibility—such as wider dissemination of working papers, inclusive seminar invitations, and institutional support for pre-publication exposure—may be especially effective at reducing the under-recognition of women's work.[42]

But although a future Matthew effect can shine a light on women's *existing* research, it cannot retroactively award them the advantages and opportunities they should have been given had their talents been appropriately recognised from the very start. Or in other words, a Matthew effect cannot revive whatever research is lost when a Matilda effect stymies women from pursuing their best ideas. How many more innovations would Katalin Karikó have generated had we fully appreciated and funded her word 30 years ago? And how much faster would the field of mRNA research have developed had others been able to build on these innovations in the intervening years?

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.

Ayarza, Alaitz and Nagore Iriberri (2025). "Gender gaps among scholars in economics: analysis across cohorts". Mimeo.

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

Bansal, Amit (2023). "From rejection to the Nobel Prize: Karikó and Weissman's pioneering work on mRNA vaccines, and the need for diversity and inclusion in translational immunology". *Frontiers in Immunology* 14 (1306025).

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

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.

Chorus, Caspar and Ludo Waltman (2016). "A large-scale analysis of impact factor biased journal self-citations". *PLOS ONE* 11 (8), e0161021.

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.

[42] We caution, however, that visibility alone may not be enough. The first top-five publication functions not just as a visibility shock, but also as a credible signal of quality. 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).

Costa-Dias, Monica, Erin Hengel, Melanie Jones, and Henrique Neves (2024). *A UK-based report on the status of women in academic economics*. Tech. rep.

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.

Doleac, Jennifer L., Erin Hengel, and Elizabeth Pancotti (2021). "Diversity in economics seminars: who gives invited talks?" *AEA Papers and Proceedings* 111, pp. 55–59.

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.

Fiorentin, Florencia, Mariano Pereira, Diana Suarez, and Alexis Tcach (2023). "When Matthew met Matilda: the Argentinean gender gap in the allocation of science and technology public grants". *Economics of Innovation and New Technology* 32 (5), pp. 700–712.

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

Funk, Patricia, Nagore Iriberri, and Nicole Venus (2025). "Women in editorial boards: an investigation of female representation in top economic journals". Mimeo.

Gazni, Ali and Fereshteh Didegah (2021). "Journal self-citation trends in 1975–2017 and the effect on journal impact and article citations". *Learned Publishing* 34 (2), pp. 233–240.

Ginther, Donna K. and Shulamit Kahn (2004). "Women in economics: moving up or falling off the academic career ladder?" *Journal of Economic perspectives* 18 (3), pp. 193–214.

– (2021). "Women in academic economics: Have we made progress?" *AEA Papers and Proceedings* 111, pp. 138–142.

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.

Hamermesh, Daniel S. (2013). "Six decades of top economics publishing: Who and how?" *Journal of Economic literature* 51 (1), pp. 162–172.

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, Euyoung Moon, and Richard Tol (2024). "Gender and equality at economics and finance journals". Mimeo.

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.

Lincoln, Anne E., Stephanie Pincus, Janet Bandows Koster, and Phoebe S Leboy (2012). "The Matilda Effect in science: awards and prizes in the US, 1990s and 2000s". *Social Studies of Science* 42 (2), pp. 307–320.

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: https://www.nber.org/about-nber.

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

Rossiter, Margaret W. (1993). "The Matilda effect in science". *Social Studies of Science* 23 (2), pp. 325–341.

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.

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.

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

Appendices

A Proofs **3**

B Data **4**

B.1 Additional data descriptions 4

B.2 Additional summary statistics 5

C Section 3.3, robustness **11**

C.1 Gender differences in the quality of pre-top-five work 11

C.2 Using $t = -1$ as the normalisation point 13

C.3 Poisson likelihood function 14

C.4 Extensive margin 15

C.5 Cited papers published in economics vs. non-economics journals 16

C.6 Citing papers published in economics vs. non-economics journals 17

C.7 Cited and citing papers published in econ. vs. non-econ. journals 18

C.8 Solo- vs. co-authored k 19

C.9 Comparing estimates from different collaboration patterns on k 20

C.10 Controlling for authors' institutional rank when k is published 22

C.11 Controlling for the first top-five k's primary and secondary *JEL* codes 23

C.12 Controlling for authors' pre-first-top-five k publication counts 24

C.13 Controlling for previous paper j_g's journal impact factor 25

C.14 Controlling for previous paper j_g's number of co-authors 26

C.15 Early vs. late top-five publication 27

C.16 Including only k that are the first top-five publications for all co-authors 28

C.17 Excluding authors with pre-k top-five publications 29

D Section 4.3, robustness **30**

D.1 Balance tables for treatment and counterfactual environments 30

D.2 Treatment 1, with and without NBER k 37

D.3 Poisson likelihood function 39

D.4 Controlling for journal impact factors and pre-k publication counts 42

D.5 The Matthew effect by employer rank 45

E Alternative outcomes **47**
E.1 Editorships 47
E.2 Promotion to professor 49
E.3 Grants 50
E.4 Employer institutional rank 51
E.5 Seminar invitations 52

F Alternative treatments **53**
F.1 Publication in a non-top-five journal 53
F.2 NBER working paper release date 54
F.3 Publication of the second top-five paper 55

References **56**

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 3, $\lambda_{j_g \bar{l}} = \lambda_{j_g \underline{l}}$, $\phi_{\underline{l}} = \phi_{\bar{l}}$ and $\tilde{\theta}_{j_g \underline{l}} = \tilde{\theta}_{j_g \bar{l}}$, so Equation (A.1) is zero.[a] 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.[b] □

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. □

[a] 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.

[b] 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, Moon, and Tol (2024) 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 3,929 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. Unless otherwise mentioned, our dataset of cited papers includes all English language pre-top-five: (i) manually verified (*e.g.*, by cross-checking an author's CV) articles, book chapters, *etc.* attributed by Web of Science to an author; and (ii) peer reviewed journal articles that Web of Science *did not* attribute to the author but were indexed by it and found on the author's CV.[c] Our final sample of pre-top-five publications contains 30,435 articles (roughly 8 articles per person). Data on these articles and the 2,133,539 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 4,334 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, 3,073 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.

[c] In other words, our dataset does not include non-journal-article outputs *unless* they were attributed to an author in the list of publications we originally downloaded from Web of Science *and* the output was clearly listed on the relevant author's CV. (That is, we did not add a non-journal-article output to an author's publication list if it was not already attributed to the author by Web of Science, but we also did not remove it unless he did not actually author it.) We took this decision because Web of Science indexes so few non-journal-article outputs to begin with—they constitute only 8.4 percent of cited papers in our data—and including or excluding these outputs has no meaningful impact on any of the analyses presented in this paper.

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.[d] In most *JEL* categories, men publish more pre-top-five papers than women. The exceptions are Q (agricultural and natural resource economics) and N (economic history).

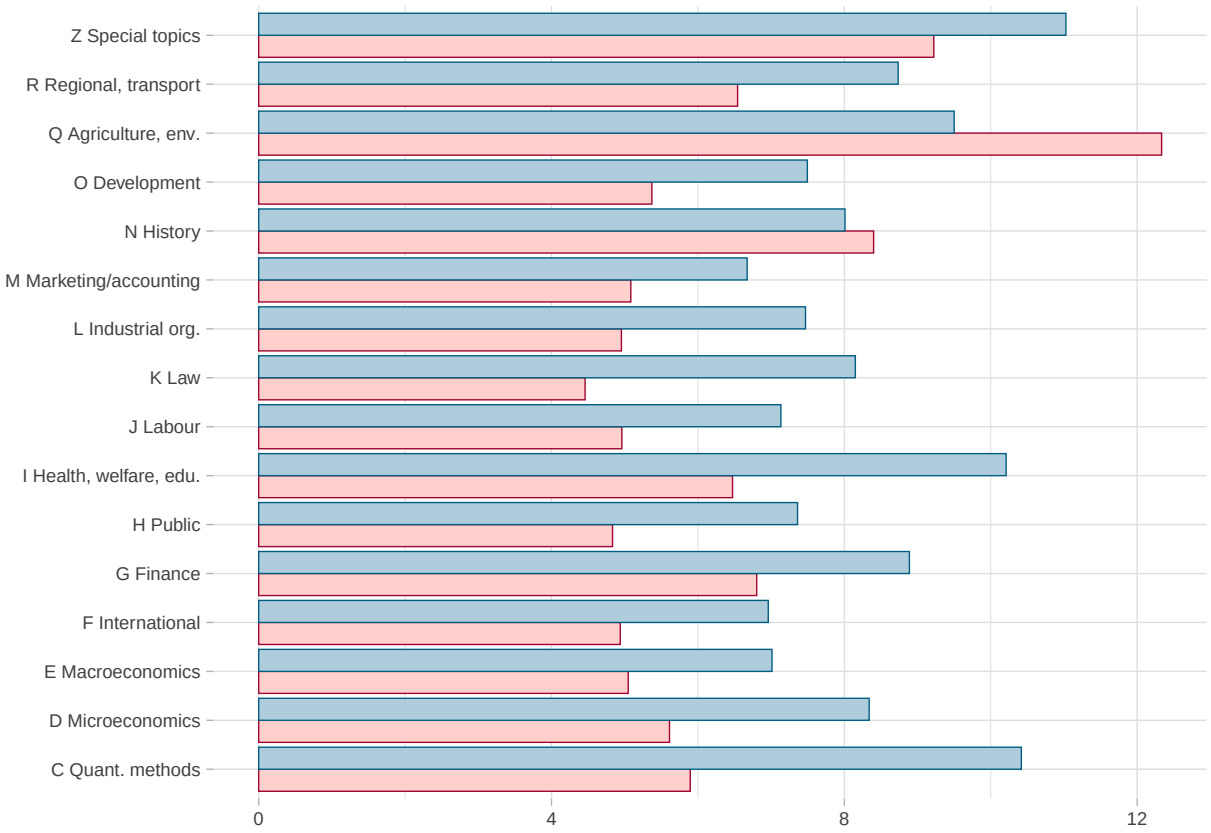


Figure B.1: Avg. no. pre-top-5 papers, by *JEL* code

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

[d] 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 institution among all institutions attributed to him 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.

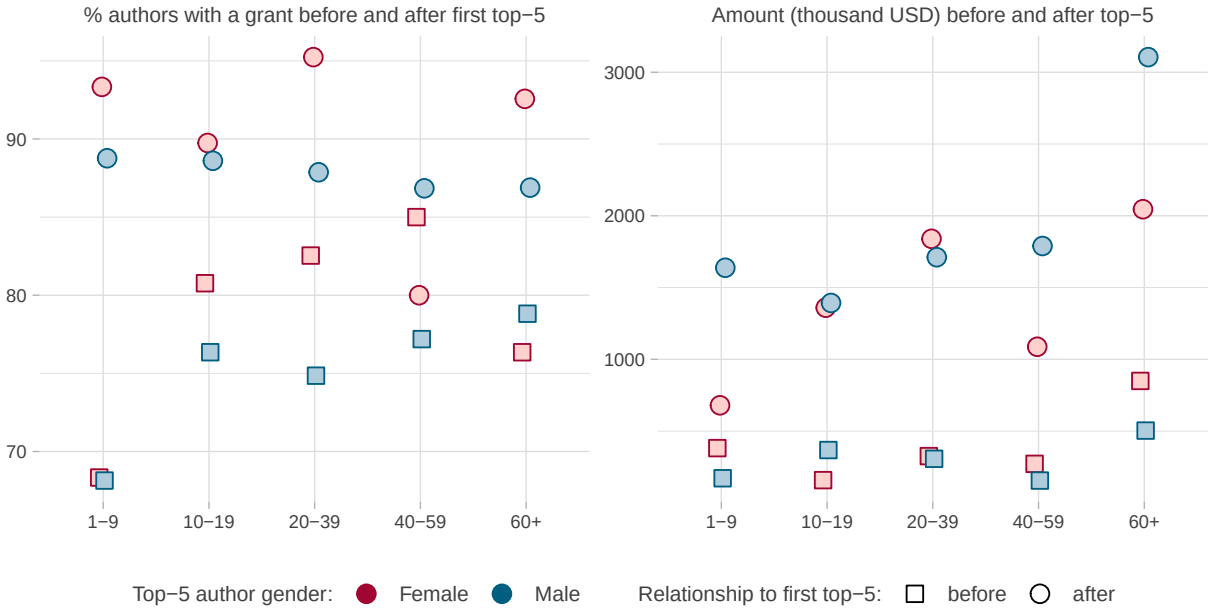


Figure B.2: Grants to authors, before and after their first top-five paper

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.[e] 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.

[e] 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 disproportionately authored the small number of highly cited papers that were published *before* 1990 and continue to be cited year-after-year.[f]

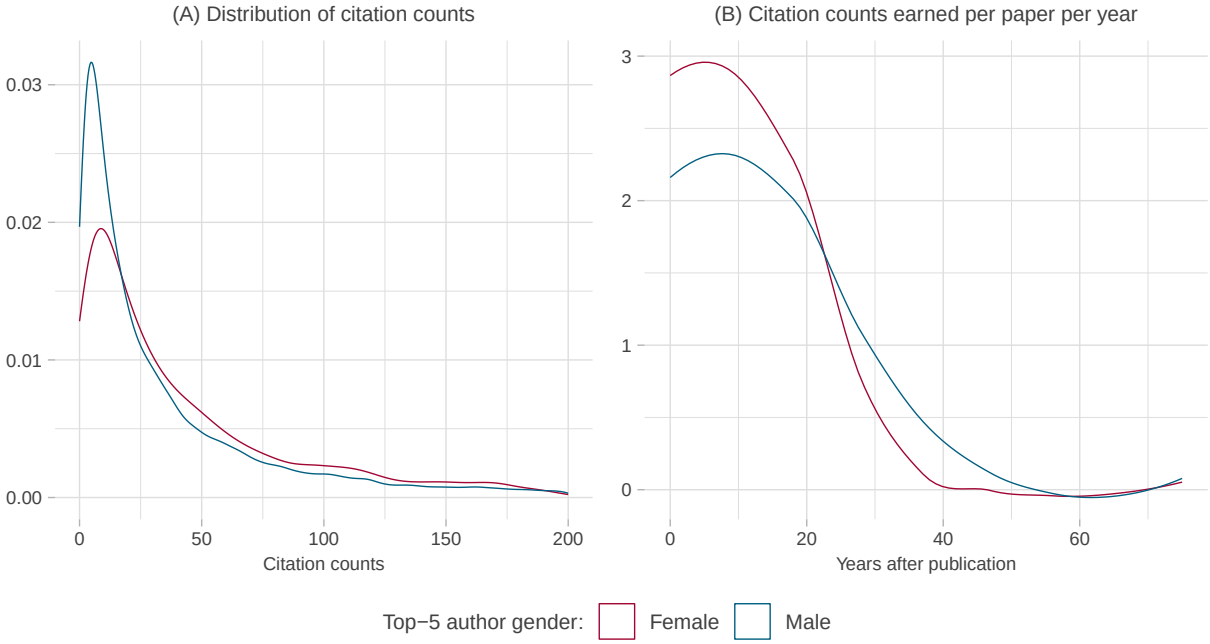


Figure B.4: Distribution of citations to j_g

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.

[f]Recall that every cited paper in our data was published on or before 2015.

B.2.5 Distribution of pre-k publication counts

In Figures B.5 and B.6, we plot the distributions of authors with 0, 1, 2, *etc.* pre-k publications. Figure B.5 plots the distribution for all authors in the data; Figure B.6 shows the distribution among authors who published k before 1990 and from 1991–2000, 2001–2010 and 2011–2015.

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 right-tail of the distribution.[g]

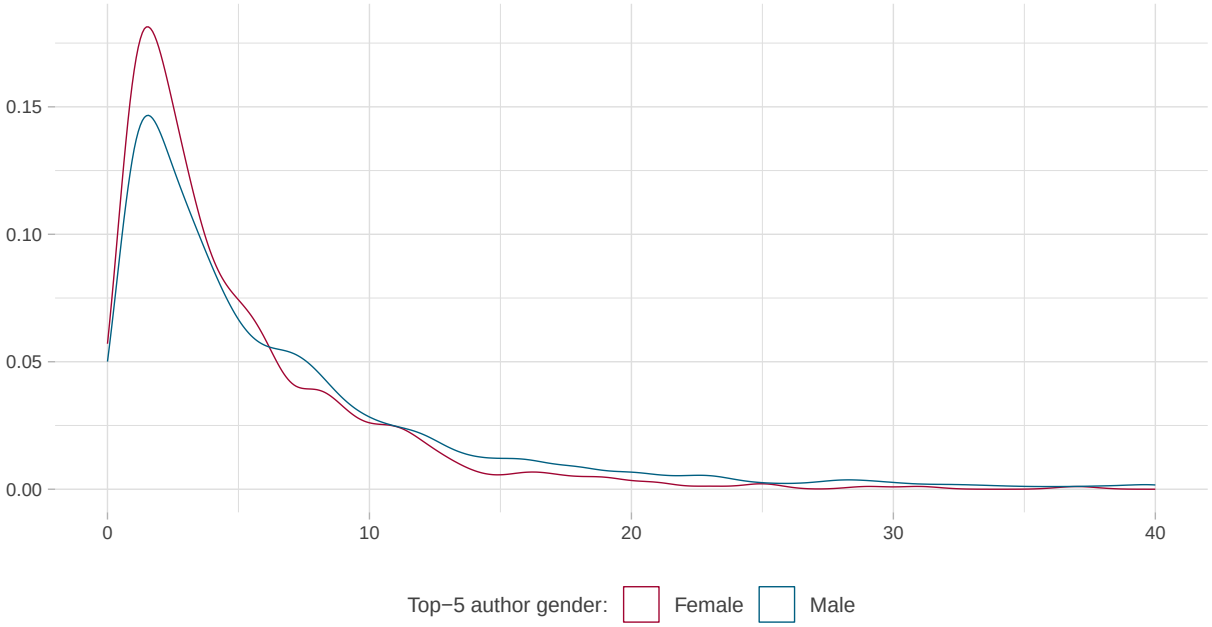


Figure B.5: Distribution of authors by pre-k publication count

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

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

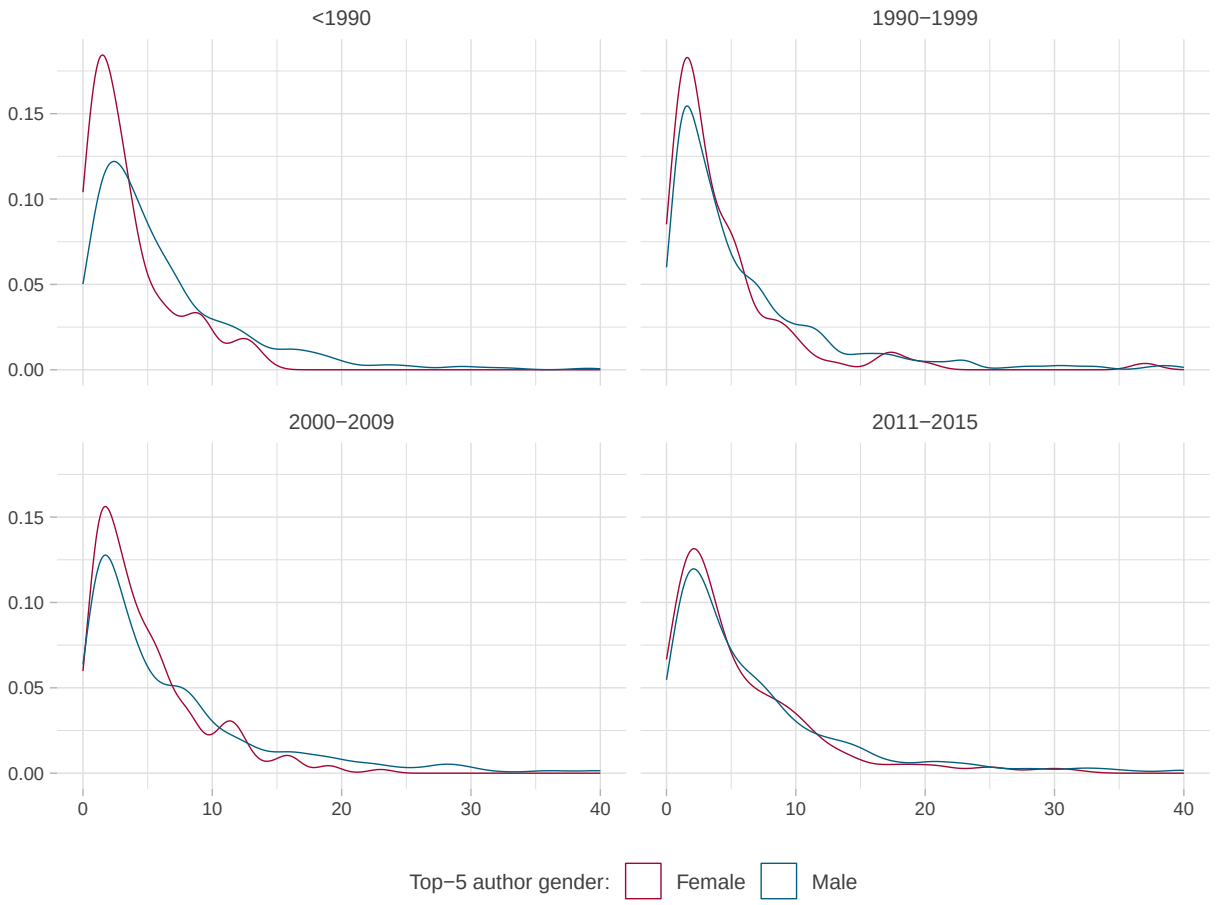


Figure B.6: Figure B.5 over time

Note. Figure replicates B.5 among the set of authors whose first top-five publication was published before 1990, during 1990–1999, 2000–2009 and 2011-2015.

C Section 3.3, robustness

C.1 Gender differences in the quality of pre-top-five work

In Tables C.1 and C.2 we regress top-five author gender on each cited paper's cumulative citations (asinh), raw citations and the journal impact factor (JIF) of the journal it was published in. Table C.1 is estimated exclusively on the sample used to estimate Figure 3; Table C.2 includes all pre-top-five papers published by the author and accumulates all citations to a paper.[h] In columns (1), (3) and (5) we control for fixed effects for the top-five paper—thus, results compare citations to pre-top-five papers of male and female co-authors on the same top-five paper—as is done in Figure 3(A). Columns (2), (4) and (6) instead control for number of authors on the top-five paper, the prominence of the most prominent author on the top-five paper (as proxied for by total number of top-five papers) and citations (asinh) to the top-five paper, as is done in Figure 3(B).

Both Tables C.1 and C.2 suggest that women's papers are cited more on a per-paper basis compared to their male top-five co-authors, although the association is weaker (and insignificant in Table C.2) when raw citations are used as the dependent variable. Neither Table C.1 nor C.2 suggest a consistent gender difference in the JIF of pre-top-five cited papers.

Table C.1: Gender differences in citations and JIFs to pre-top-five work, sample from Figure 3

	Citations (asinh)		Citations		JIF	
	(1)	(2)	(3)	(4)	(5)	(6)
Top-five author female	0.228***	0.160***	11.616***	3.779*	−0.155	0.065
	(0.054)	(0.026)	(4.403)	(2.263)	(0.244)	(0.112)
No. top-five authors		0.031***		1.103		0.369***
		(0.01)		(0.832)		(0.068)
Top-five author prominence		−0.007***		−0.213		0.013*
		(0.002)		(0.135)		(0.008)
Top-five paper citations (asinh)		0.213***		8.948***		0.311***
		(0.008)		(0.659)		(0.029)
No. obs.	29,357	29,357	29,357	29,357	28,166	28,166
R^2	0.325	0.111	0.152	0.027	0.185	0.027
Top-5 f.e.	✓		✓		✓	
Top-5 journal		✓		✓		✓
Top-5 year		✓		✓		✓

Note. Coefficients from a regression of per-paper citations (asinh) (columns 1 and 2), raw citations (columns 3 and 4) and journal impact factors (JIFs) (columns 5 and 6) on top-five author gender. In columns (1), (3) and (5) we control for fixed effects for the top-five paper as is done in Figure 3(A); in columns (2), (4) and (6) we control for the numbers of co-authors on the top-five article, the prominence of the most prominent author on the top-five paper, and the number of citations (asinh) to the top-five paper, as is done in Figure 3(B). Sample of pre-top-five papers is restricted to the sample used to estimate Figure 3. Standard errors clustered at the cited paper level in parentheses. ***, ** and * statistically significant at 1%, 5% and 10%, respectively.

[h] Recall that Figure 3 only includes citations to papers occurring in the ten years before and after the author's first top-five paper. Thus, it excludes pre-top-five papers that accumulated no citations during this 20-year window and citations that occurred before or after this period.

Table C.2: Gender differences in citations and JIFs to pre-top-five work, all papers and citations

	Citations (asinh)		Citations		JIF	
	(1)	(2)	(3)	(4)	(5)	(6)
Top-five author female	0.244***	0.180***	9.206	1.031	−0.141	0.047
	(0.056)	(0.028)	(9.395)	(3.798)	(0.241)	(0.111)
No. top-five authors		0.021*		−2.133		0.367***
		(0.011)		(1.623)		(0.066)
Top-five author prominence		−0.007***		−0.360		0.013*
		(0.002)		(0.251)		(0.008)
Top-five paper citations (asinh)		0.246***		25.570***		0.302***
		(0.008)		(3.255)		(0.029)
No. obs.	30,435	30,435	30,435	30,435	29,106	29,106
R^2	0.302	0.064	0.200	0.012	0.185	0.027
Top-5 f.e.	✓		✓		✓	
Top-5 journal f.e.		✓		✓		✓
Top-5 year f.e.		✓		✓		✓

Note. Table is identical to Table C.1 except that the estimation sample includes all pre-top-five papers by an author and all citations to them (see Footnote h). ***, ** and * statistically significant at 1%, 5% and 10%, respectively.

C.2 Using $t = -1$ as the normalisation point

In Figure C.1, we replicate Figure 3 but use $t = -1$ as the normalisation point. Results are nearly identical to those shown in Figure 3, suggesting that our results are not sensitive to the precise timing of the event.

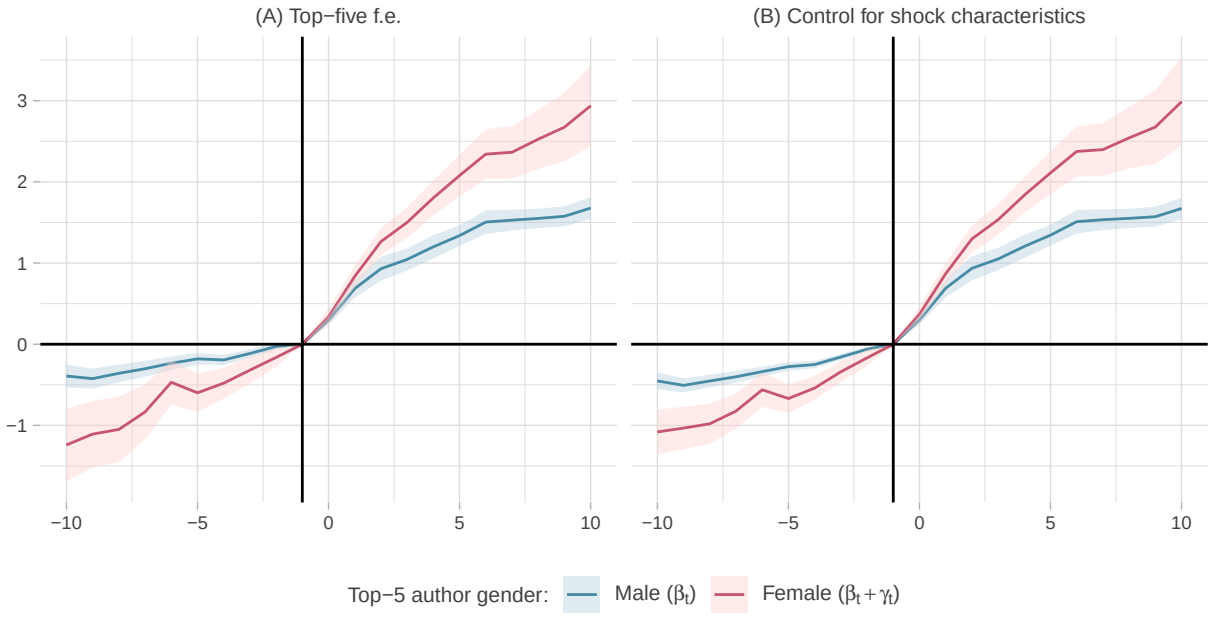


Figure C.1: Figure 3, using $t = -1$ as the normalisation point

Note. Figure C.1 replicates Figure 3 but it uses $t = -1$ as the normalisation point.

C.3 Poisson likelihood function

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

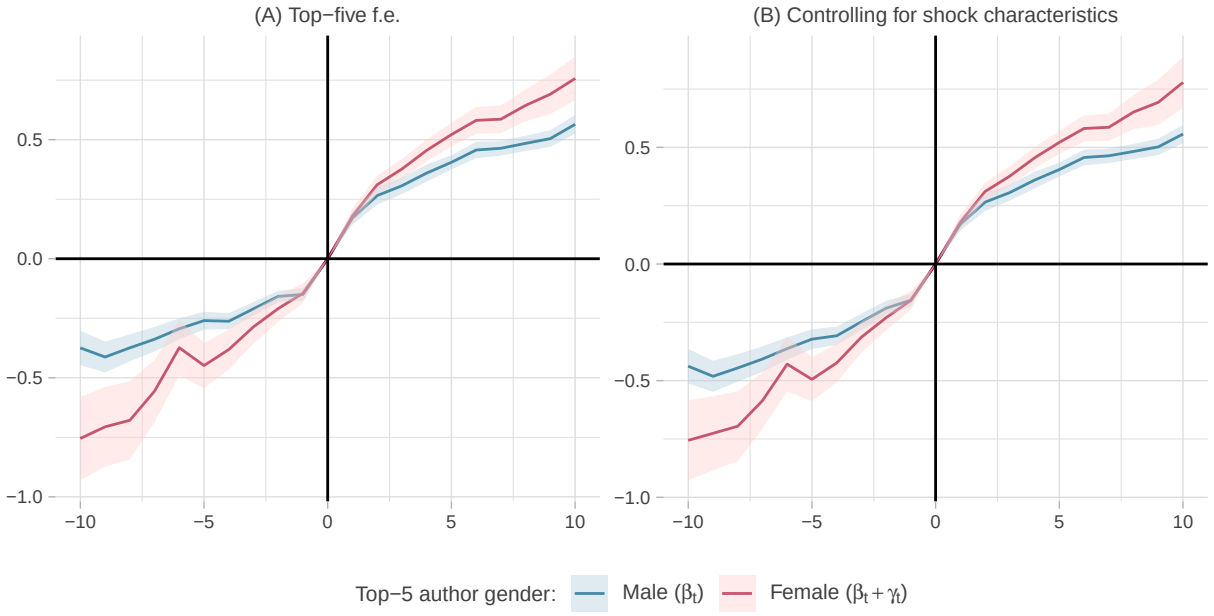


Figure C.2: Figure 3, Poisson likelihood function

Note. Figure C.2 replicates Figure 3 but it estimates Equation (3) using a Poisson likelihood function. Shaded areas reflect 95 percent confidence intervals from standard errors clustered by j_g.

C.4 Extensive margin

In Figure C.3, 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.

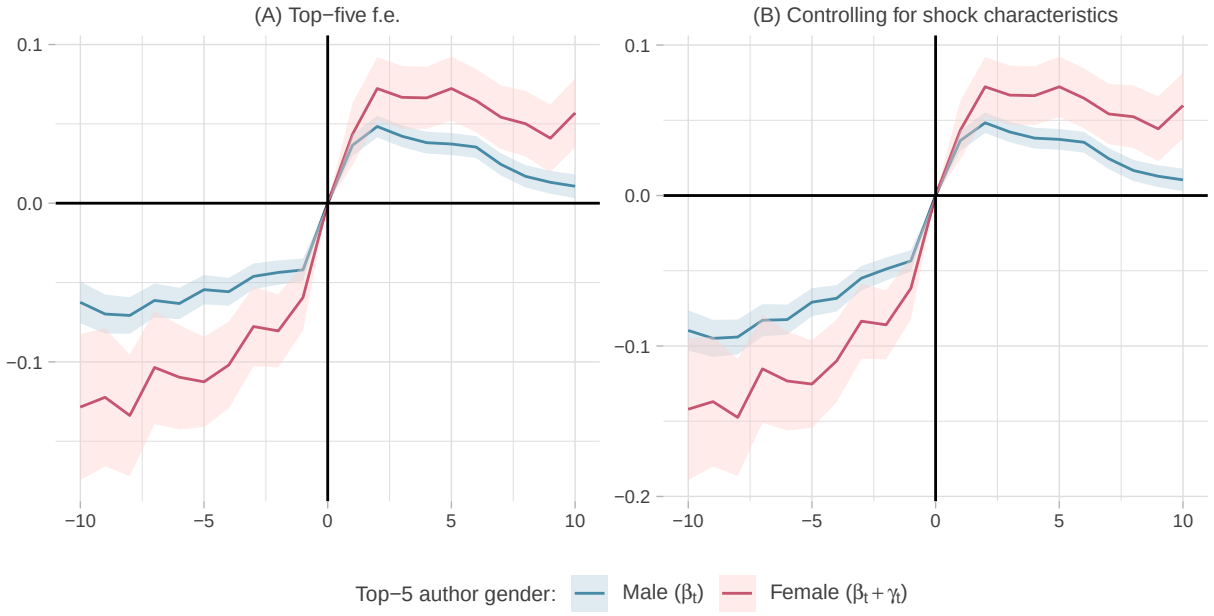


Figure C.3: Figure 3, extensive margin

Note. Figure C.2 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 95 percent confidence intervals from standard errors clustered by j_g.

C.5 Cited papers published in economics vs. non-economics journals

In Figure C.4 we reproduce Figure 3 but restrict the sample to cited papers (j_g) published in economics (top row) and non-economics journals (bottom row) previous to the first top-five. 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.4 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.

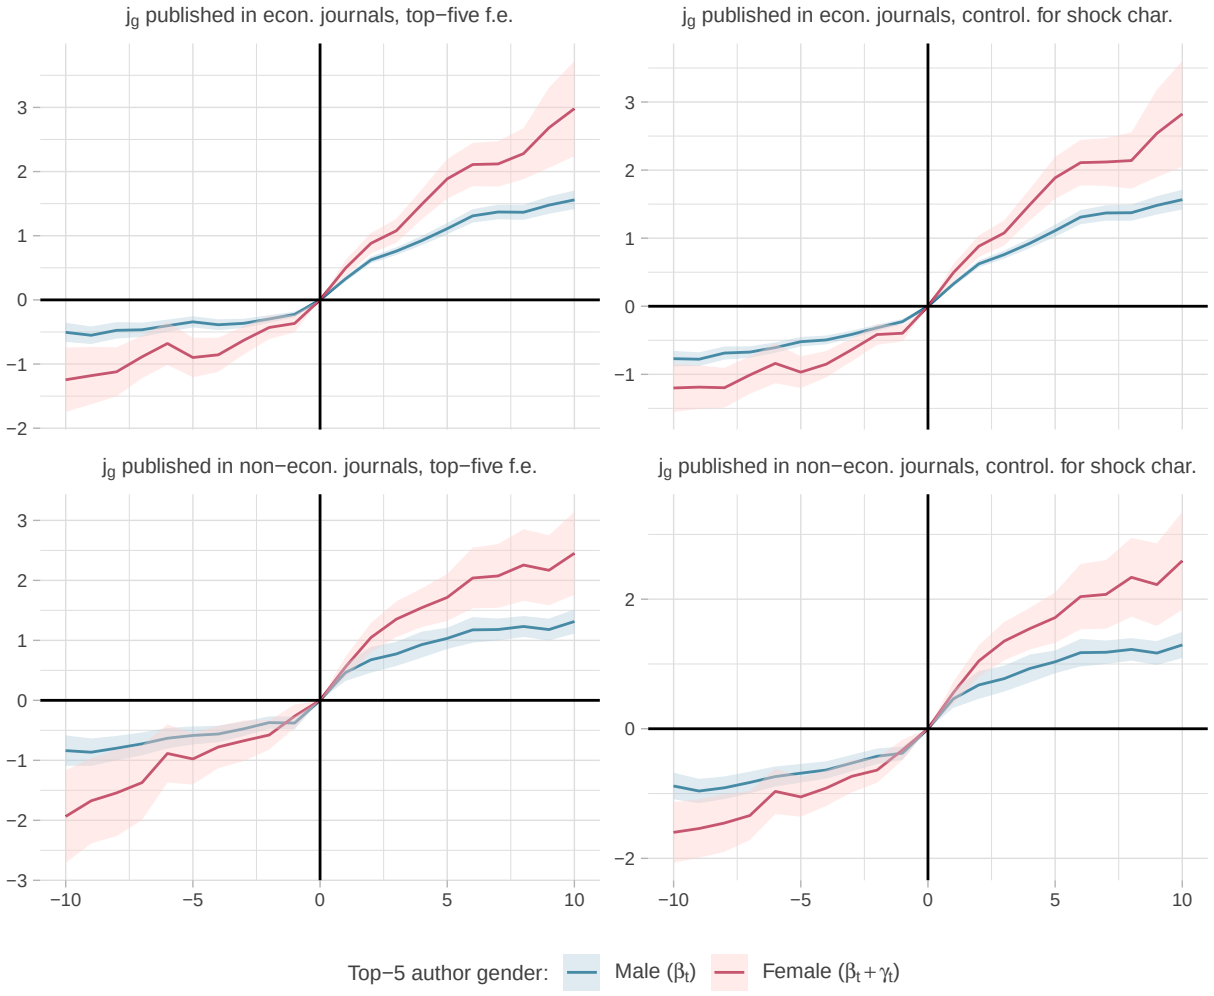


Figure C.4: Figure 3, j_g published in economics vs. non-economics journals

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

C.6 Citing papers published in economics vs. non-economics journals

In Figure C.5 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.[i]

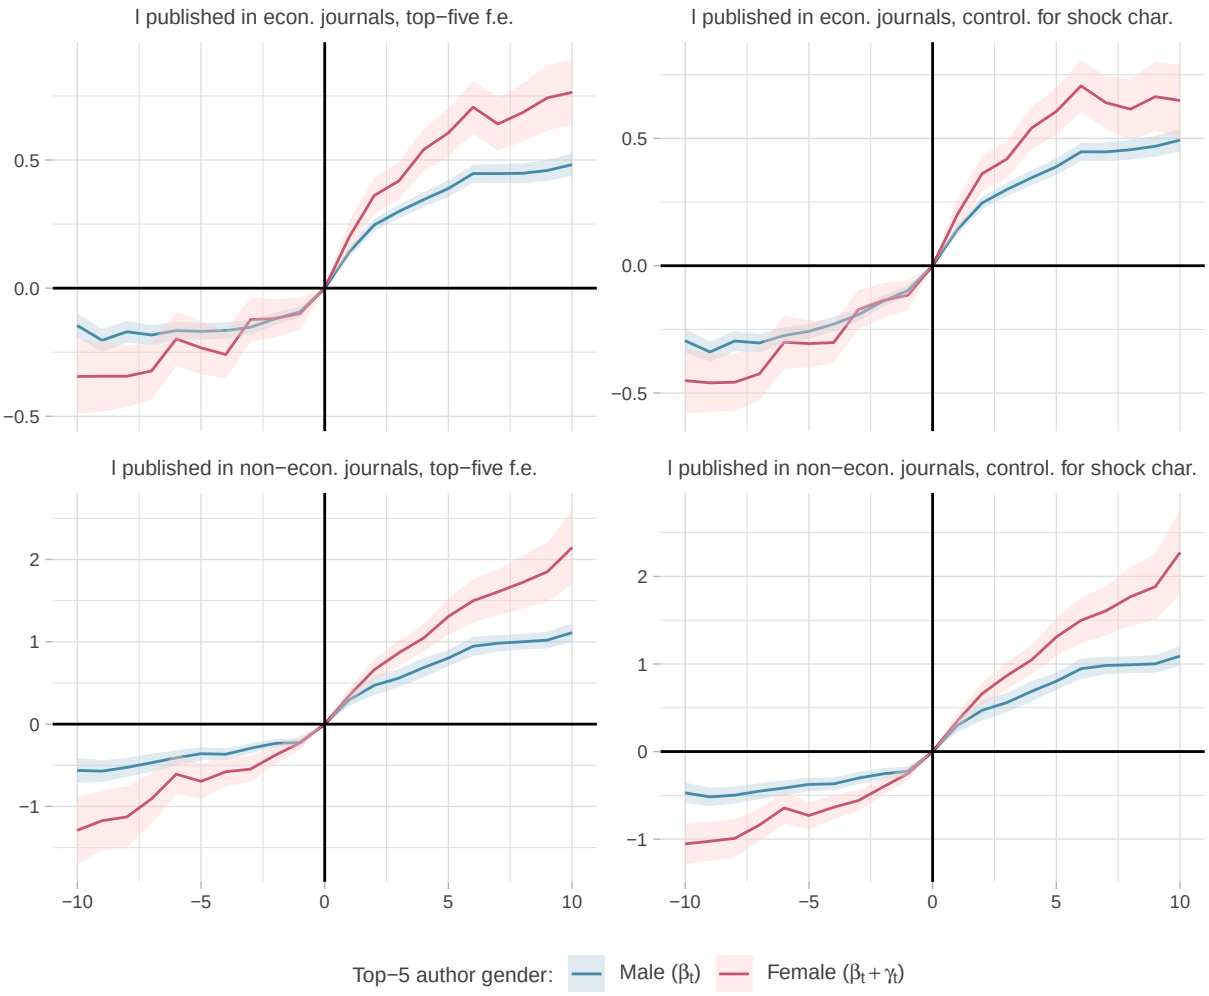


Figure C.5: Figure 3, l published in economics vs. non-economics journals

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

[i]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. (On average, citations from papers published in economics journals only represent about 28 percent of total citations to a paper.)

C.7 Cited and citing papers published in econ. vs. non-econ. journals

In Figure C.6 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.4 and C.5.

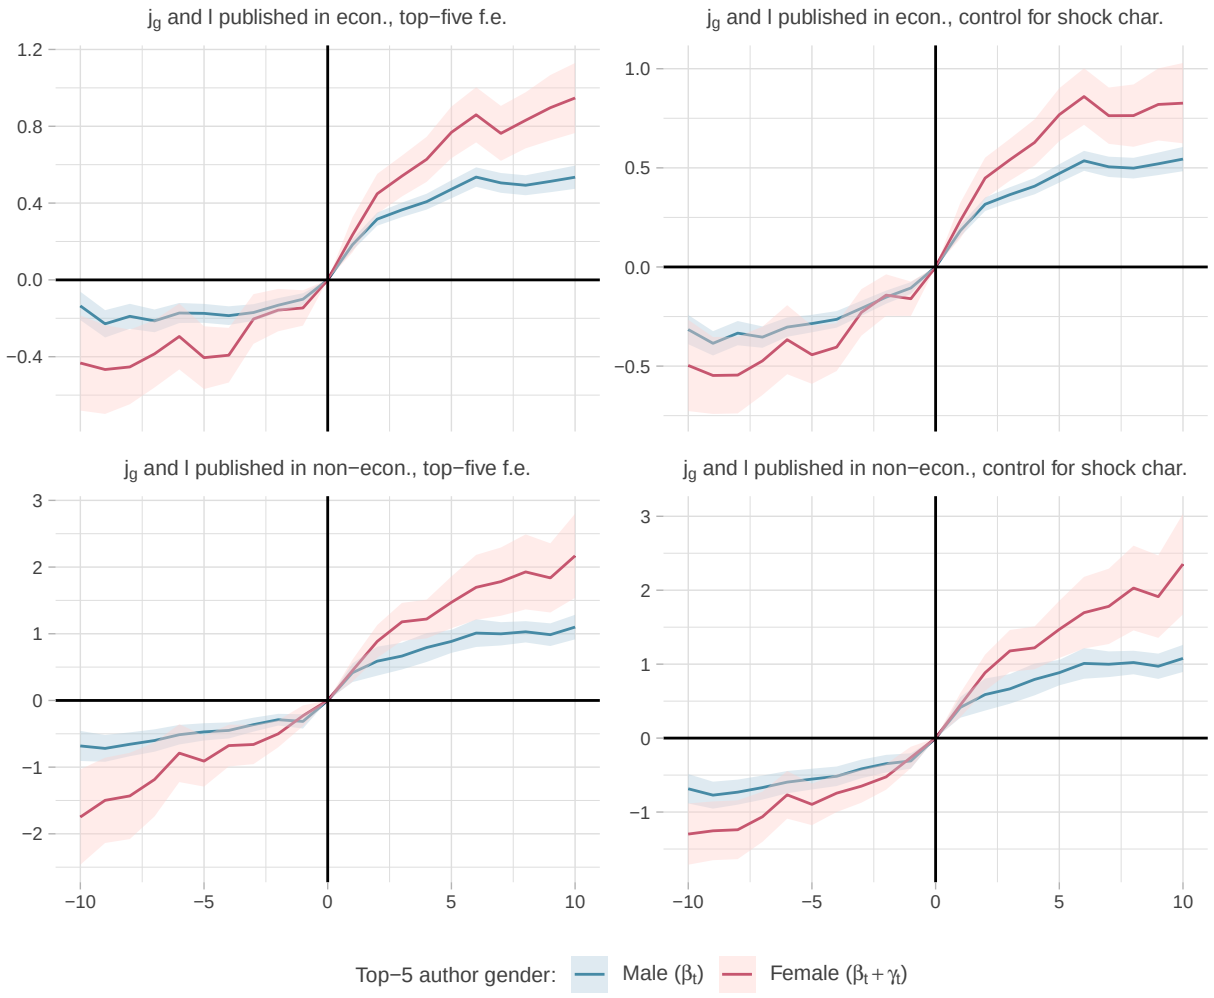


Figure C.6: Figure 3, j_g and l published in economics vs. non-economics journals

Note. Figure C.6 reproduces Figure 3 on the sample of cited and citing papers published in economics journals (top row) and non-economics journals (bottom row).

C.8 Solo- vs. co-authored k

In Figure C.7 we reproduce Figure 3(B) on the sample of authors with a top-five paper that was 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 among men is similarly sized in both graphs. The Matthew effect for women, 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.[j]

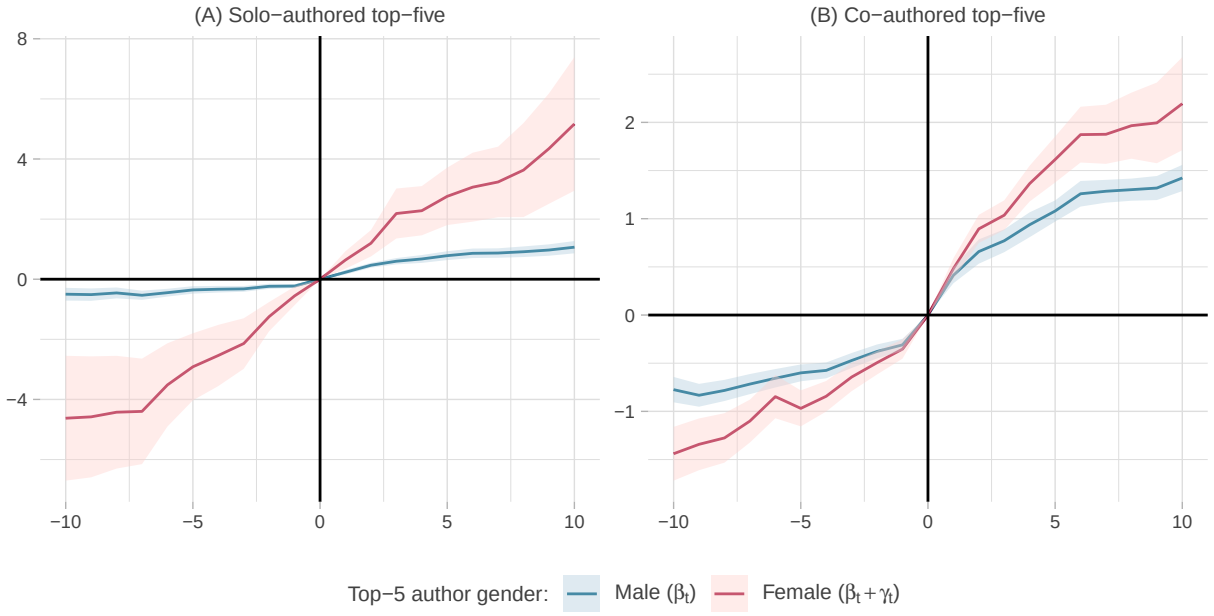


Figure C.7: Figure 3(B), solo- vs. co-authored k

Note. Figure C.7 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).

[j] We do not replicate Figure 3(B) on the sample of authors with a top-five paper that was exclusively co-authored by two or more authors of the same gender as only 36 women (co-authoring 28 top-five articles) satisfy this criterion.

C.9 Comparing estimates from different collaboration patterns on k

In Figure C.8, we reproduce Figure 3 with different top-five paper k 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.[k] Graph (C) defines the estimation sample to include only authors who co-authored k with exactly one more senior male co-author.[l] 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 for 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.8, 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)).

[k] We do not estimate results using the sample of junior authors who co-authored k with a member of the same sex as there are only 15 women (co-authoring 9 top-five articles) who satisfy this criterion.

[l] 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.

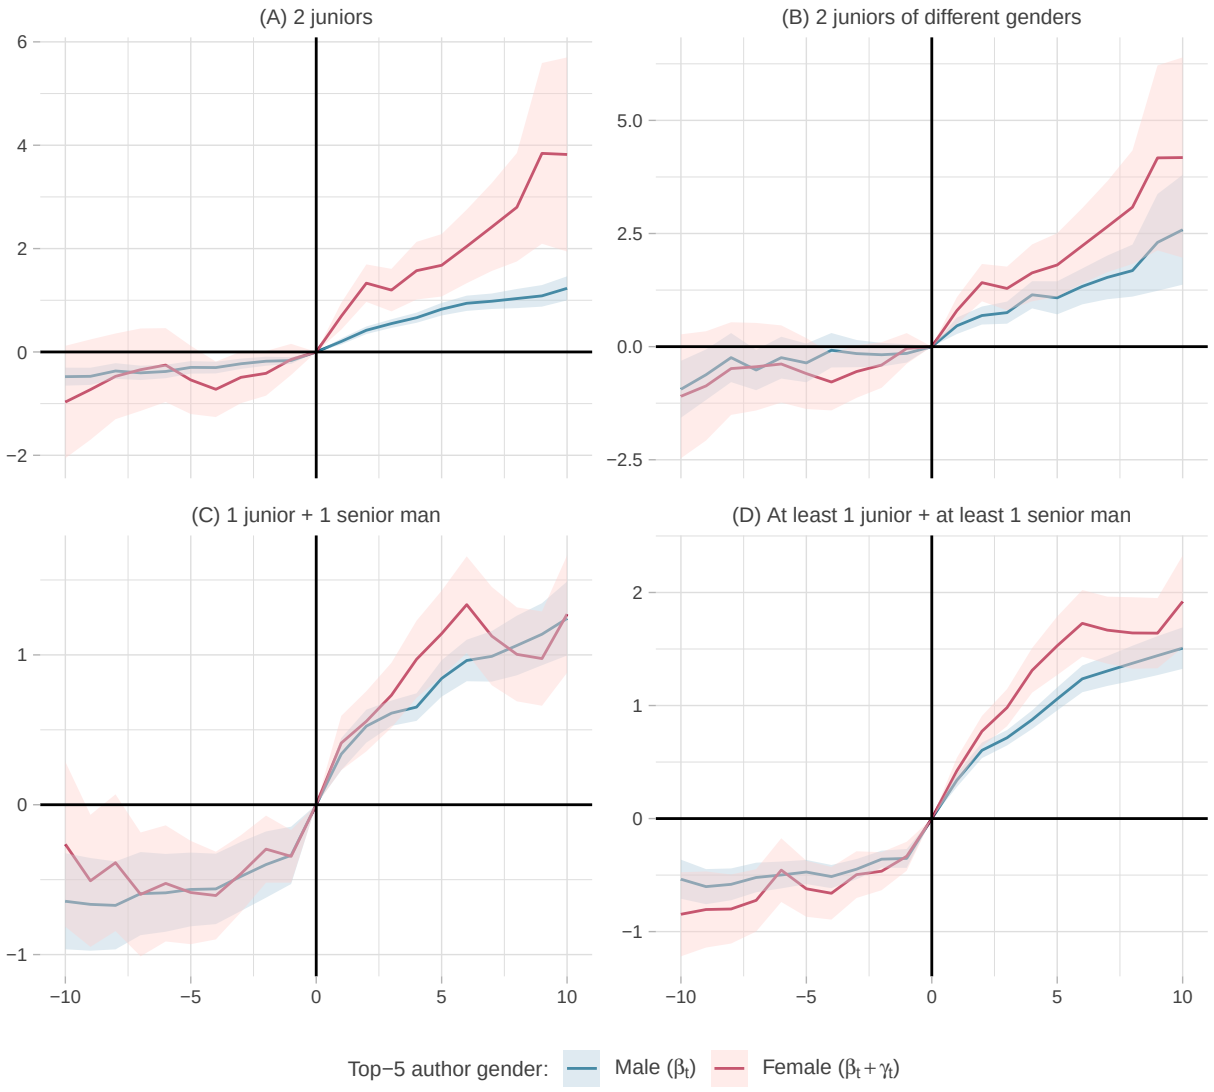


Figure C.8: Figure 3 across different collaboration patterns on k

Note. Figure C.8 replicates Figure 3 on the sample of authors whose first top-five paper was co-authored with exactly one other first-time top-five author (Graph (A)), exactly one other first-time top-five author of the opposite sex (Graph (B)), exactly one senior male co-author (Graph (C)) and at least one senior male co-author (Graph (D)).

C.10 Controlling for authors' institutional rank when k is published

Figure C.9 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 author on an article, only the dummy variable of the highest ranked institution is used.

The Matthew effect and its gender difference in Figure C.9 resemble the estimates shown in Figure 3.

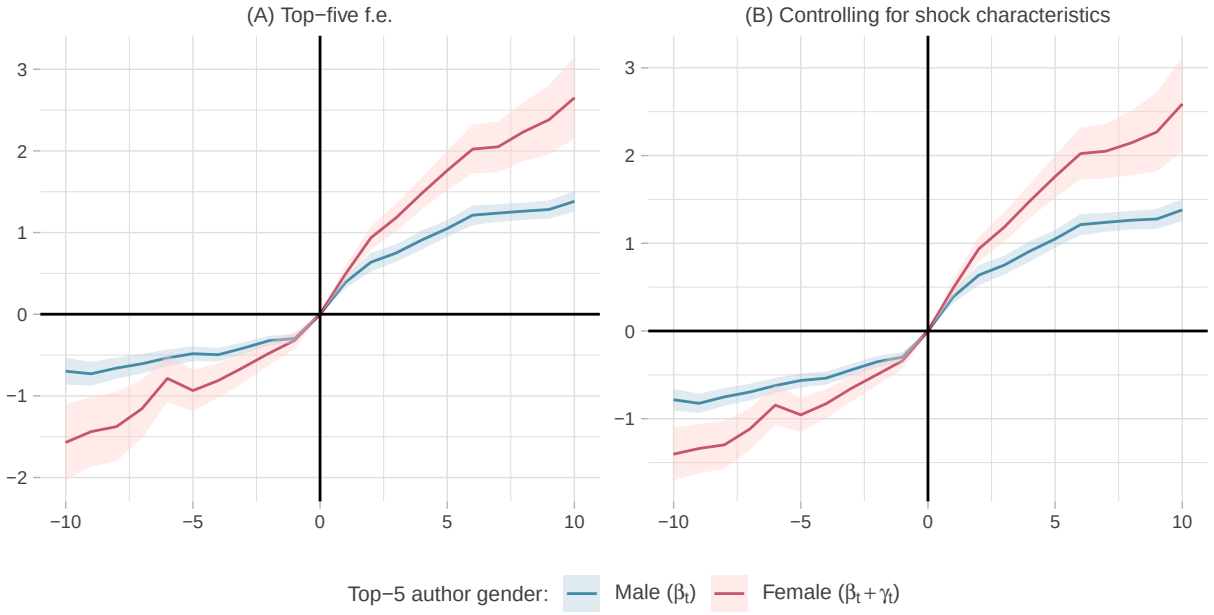


Figure C.9: Figure 3, controlling for authors' institutional rank when k is published

Note. Figure C.9 reproduces Figure 3 controlling for authors' institutional rank at the time k was published.

C.11 Controlling for the first top-five k's primary and secondary *JEL* codes

Figure C.10 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.10.

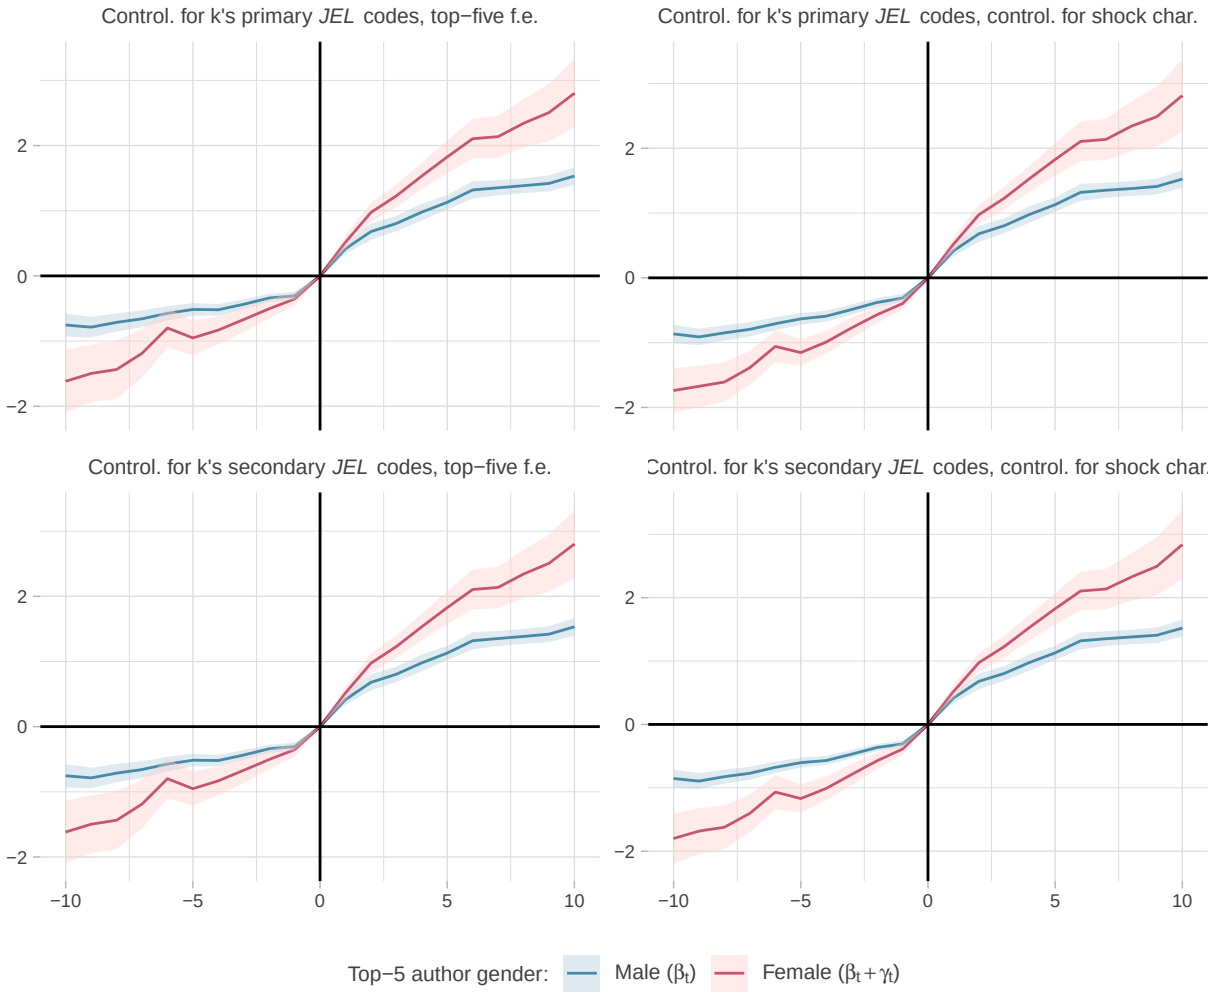


Figure C.10: Figure 3, controlling for k's primary and secondary *JEL* codes

Note. Figure C.10 reproduces Figure 3, controlling for k's primary (top row) and secondary (bottom row) *JEL* codes.

C.12 Controlling for authors' pre-first-top-five k publication counts

In Figure C.11, we reproduce Figure 3 controlling for the total number of publications j_g an author had before he published his first top-five article, k. Both graphs in Figure C.11 very closely resemble the graphs shown in Figure 3.

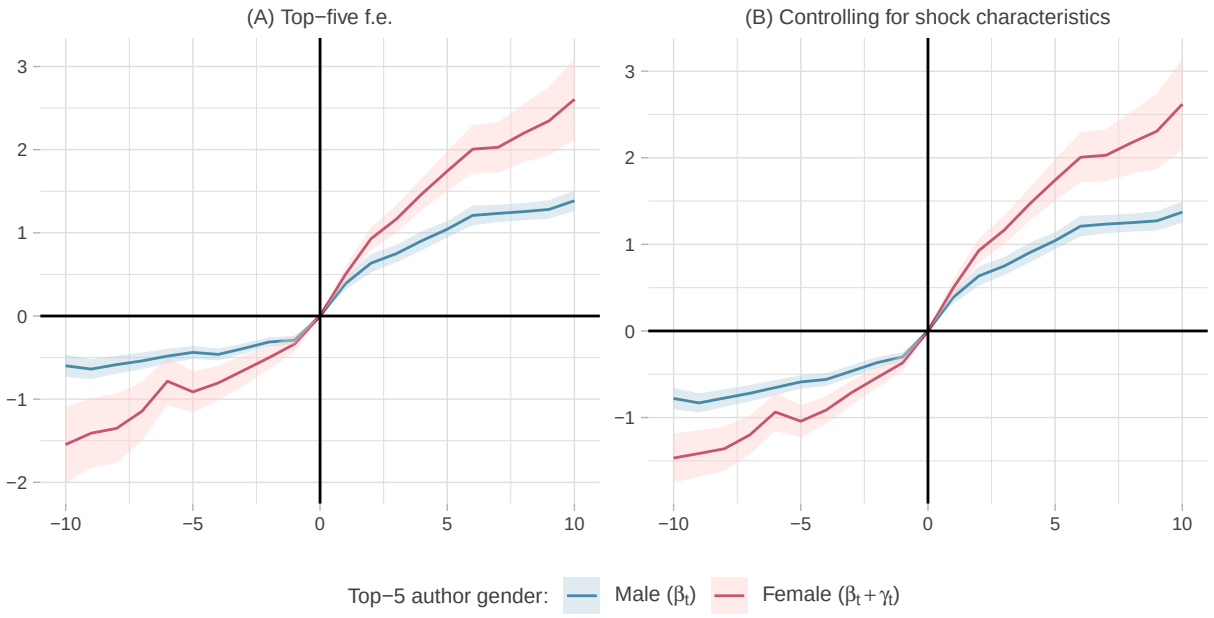


Figure C.11: Figure 3, controlling for authors' total j_g count when k is published

Note. Figure C.11 reproduces Figure 3, controlling for authors' total number of pre-k publications.

C.13 Controlling for previous paper j_g's journal impact factor

In Figure C.12, we reproduce Figure 3 controlling for cited journal impact factor (JIF).[m] The visibility Matthew effect and its gender difference are very similar to their corresponding estimates in Figure 3.

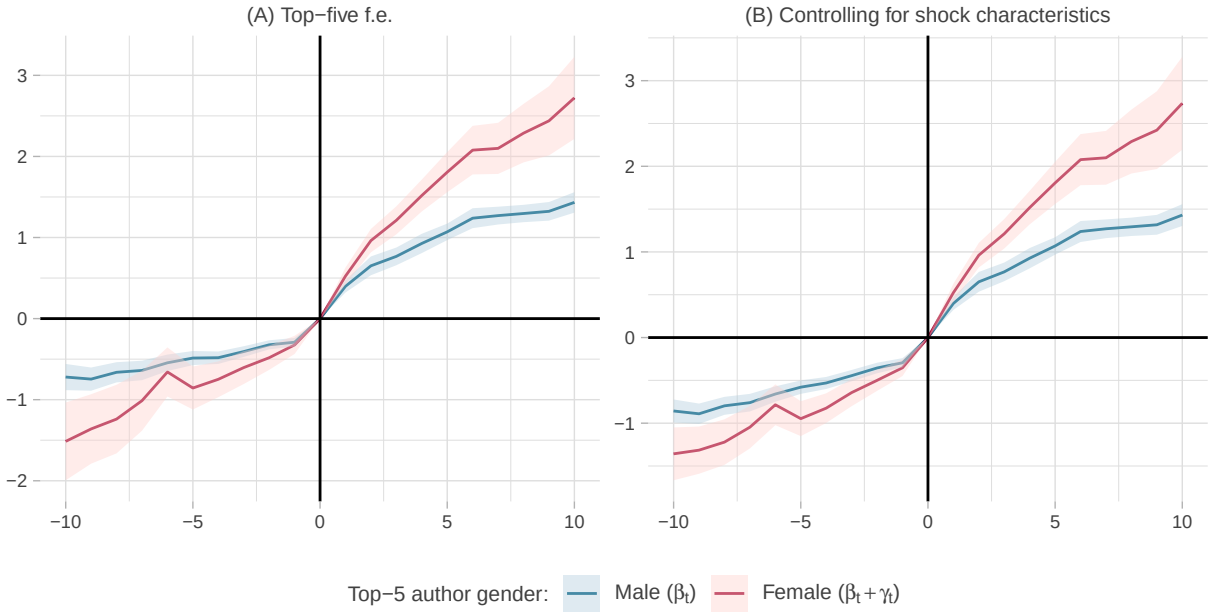


Figure C.12: Figure 3, controlling for j_g's journal impact factor

Note. Figure C.12 reproduces Figure 3, controlling for j_g's journal impact factor.

[m] 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.14 Controlling for previous paper j_g's number of co-authors

In Figure C.13, we reproduce Figure 3 controlling for j_g's number of co-authors. Results are very similar to those shown in Figure 3.

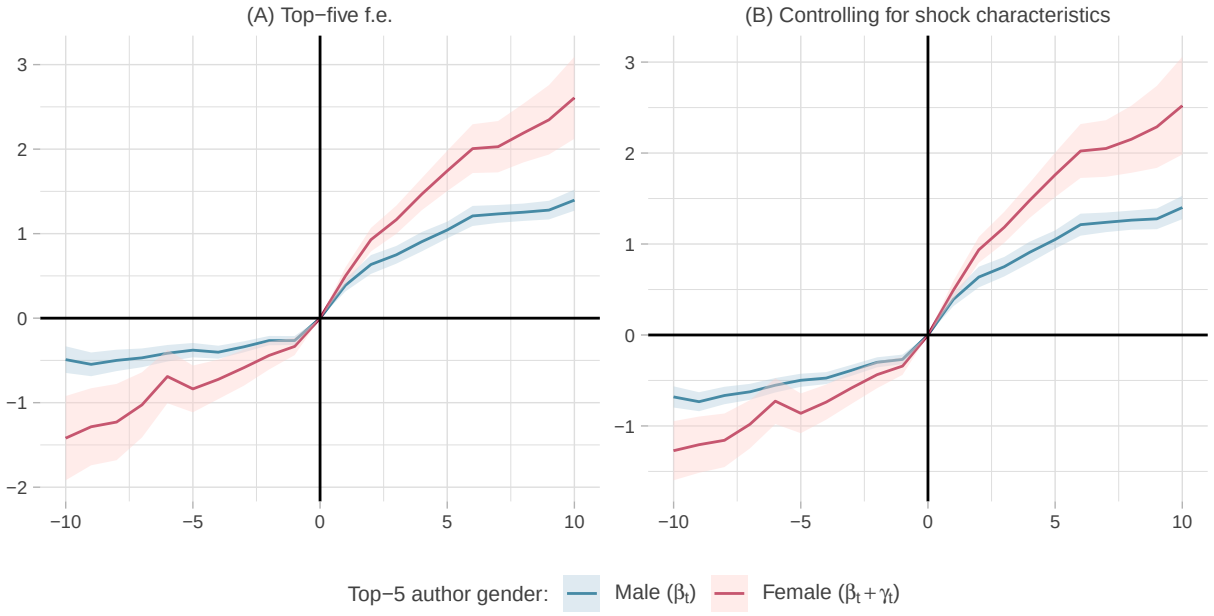


Figure C.13: Figure 3, controlling for j_g's number of co-authors

Note. Figure C.13 reproduces Figure 3, controlling for the number of co-authors on each j_g.

C.15 Early vs. late top-five publication

In Figure C.14 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 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.

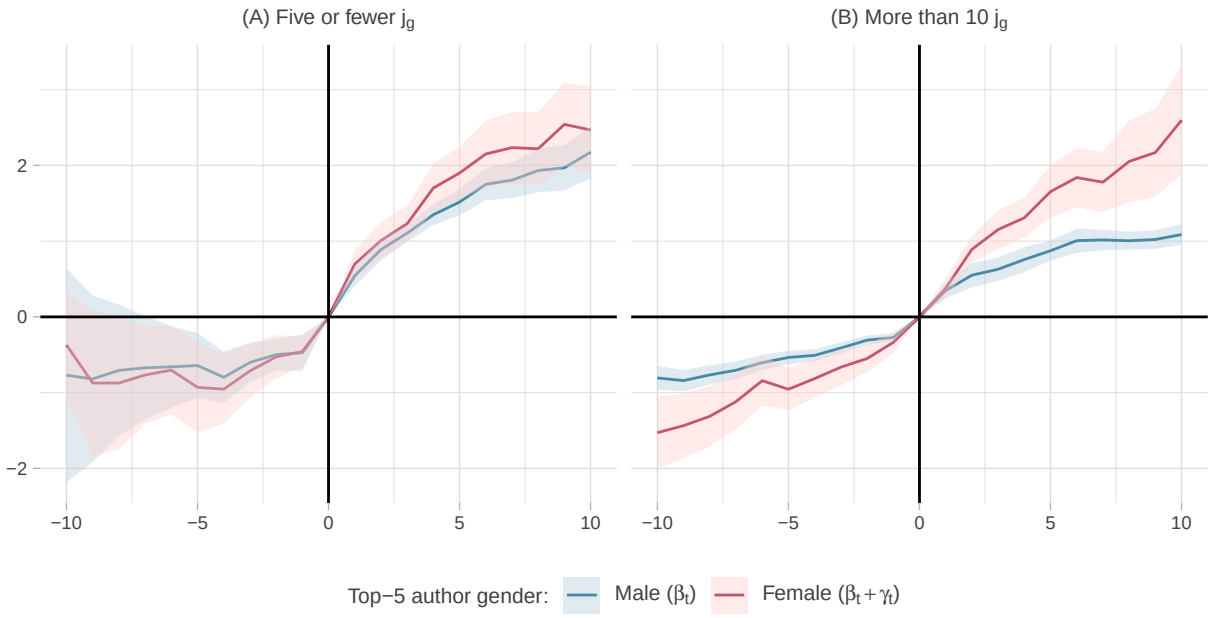


Figure C.14: Early vs. late top-five publication

Note. Figure C.14 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.

C.16 Including only k that are the first top-five publications for all co-authors

In Figure C.15, we replicate Figure 3 except that we include only k that are the first top-five publications for all co-authors. (That is, we exclude authors who co-authored k with someone who had previously published a top-five paper.) Figure C.15 suggests that, before k is published, women in this sample are cited 1–2 fewer times relative to their male peers; afterwards, they are cited 2–4 times more. Or in other words, the Matilda and Matthew effects are much more pronounced among authors for whom k was the first top-five paper for *all* of its co-authors.

The Matilda and Matthew effects in Figure C.15 are of similar magnitude to those in Figure C.7(A), which replicates Figure 3(B) on the sample of top-five papers that were solo-authored. Together, Figures C.15 and C.7(A) tentatively suggest that the Matilda and subsequent Matthew effects may be larger when the signal of publishing in a top-five journal is much clearer about the quality of women's research and their contributions toward it.

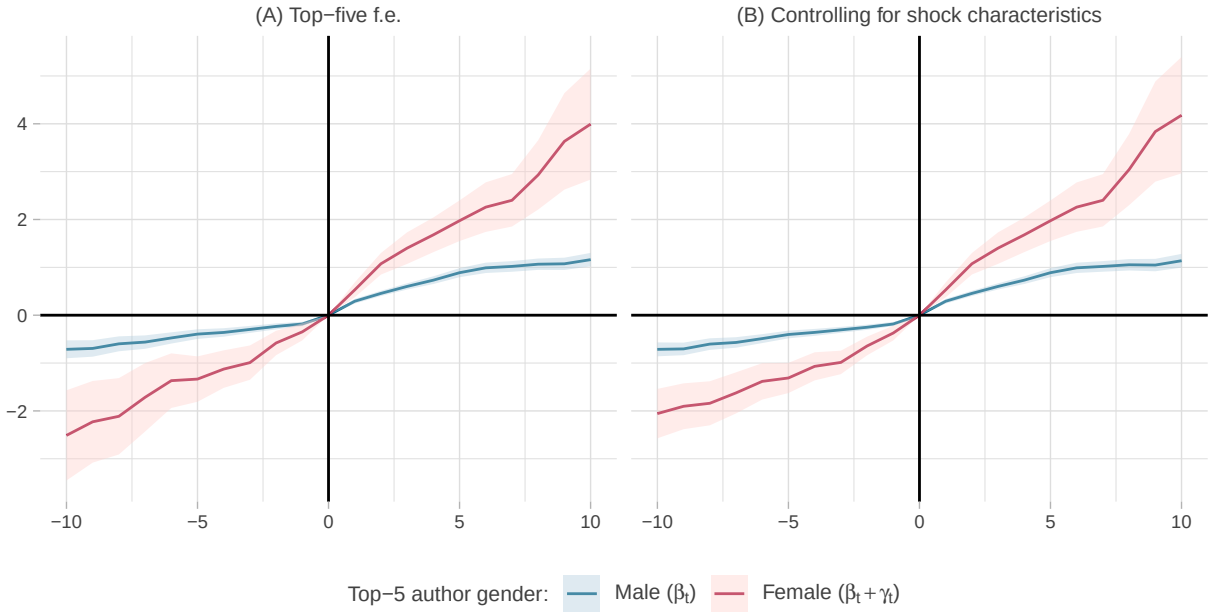


Figure C.15: Figure 3, including only k that are the first top-5 for all co-authors

Note. Figure C.15 replicates Figure 3 except that it includes only k that are the first top-five publications for all co-authors.

C.17 Excluding authors with pre-k top-five publications

A small number of authors included in the estimation sample used to generate Figure 3 actually published a paper in a top-five journal *before* their first top-five publication as recorded in Hengel (2022). In most instances, these "pre-k" top-five publications are excluded from Hengel (2022) because they are not full length articles (*e.g.*, short notes, corrections or book reviews).[n]

Below, we reproduce Figure 3 excluding authors with any pre-k top-five publication. Figure C.16 is very similar to Figure 3.

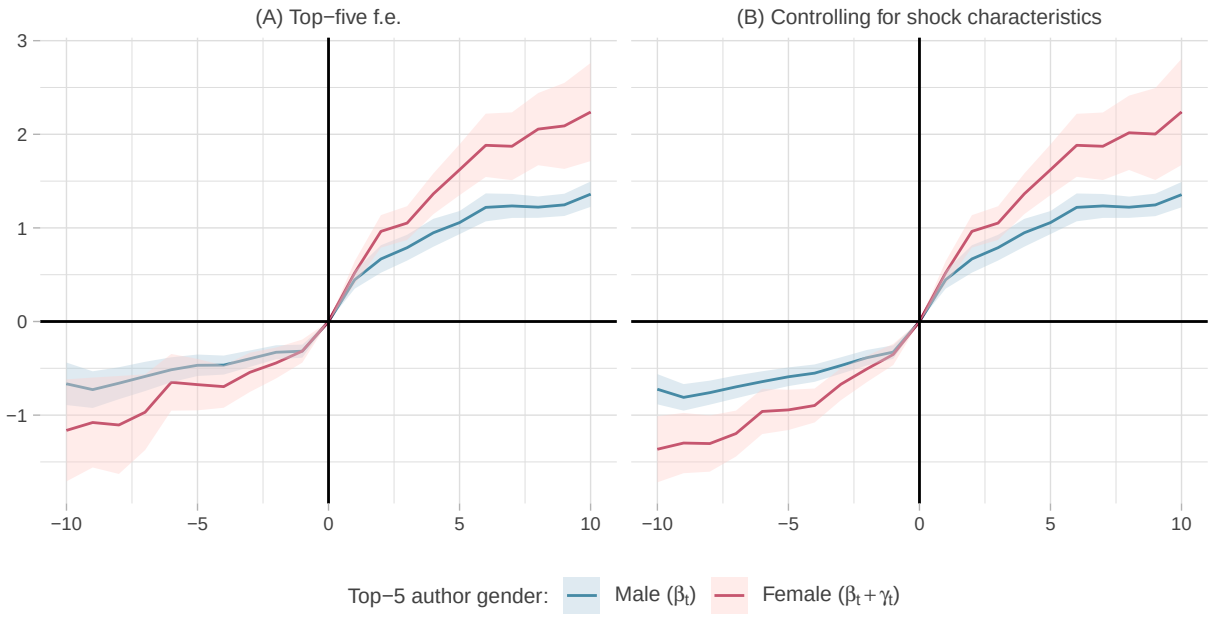


Figure C.16: Figure 3, excluding authors with pre-k top-five publications

Note. Figure C.16 replicates Figure 3 except that it excludes the small number of authors who published a top-five paper *before* their first official full length top-five paper as registered in Hengel (2022).

[n] In a small number of cases, the author's first top-five paper was not included in Hengel (2022) for other reasons, *e.g.*, the first article was published before the journal began publishing abstracts. Hengel (2022) only includes full length articles as proxied for by the existence of an abstract or submit-accept date; see Appendix 2 and Hengel (2022) for further details.

D Section 4.3, robustness

D.1 Balance tables for treatment and counterfactual environments

Tables D.1–D.6 show gender differences in paper (and, where relevant, author) characteristics in the treatment and counterfactual environments defined in Section 4.2. Estimates are from OLS regression of the dependant variables (listed in the first column of each table) on dummy variables equal to one if the author is a woman, the observation falls in the treatment category, their interaction and, unless otherwise mentioned, fixed effects for authors' first top-five papers.

Table D.1 reports estimates for treatment and counterfactual environments when the treatment is defined as the pre-top-five research j_g that authors' cite in their first top-five paper k. Although female-authored j_g are published somewhat later and in slightly higher impact journals compared to male-authored j_g, gender differences are similarly sized in both environments. No other estimates in Table D.1 significantly differ from zero.[o]

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

	Treatment	Counterfactual	Difference
Cited papers			
Publication year	1.2984***	0.8061*	0.4923
	(0.3461)	(0.4391)	(0.3850)
No. co-authors	0.0390	−0.0384	0.0774
	(0.0624)	(0.0913)	(0.0868)
Journal impact factor	0.4615***	0.4659	−0.0044
	(0.1416)	(0.2903)	(0.2920)
No. citations	1.4495	2.3164	−0.8669
	(1.3211)	(2.6572)	(2.5264)
Citing papers			
Publication year	0.2651**	0.1225	0.1426
	(0.1064)	(0.1377)	(0.1377)
No. co-authors	0.0224	0.0377	−0.0153
	(0.0185)	(0.0268)	(0.0261)
English language	0.0012	0.0031	−0.0019
	(0.0013)	(0.0022)	(0.0024)
Journal article	−0.0015	0.0035	−0.0050
	(0.0026)	(0.0035)	(0.0037)

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.

In Table D.2, we report estimates for our second treatment-counterfactual pair, which compares citations from citing papers l published in the same journal as j_g to citations from l published in different journals. In contrast to Table D.1, there are several meaningful differences between environments. In both environments, women's cited papers are published slightly later than men's, but the gender difference is significantly larger in the treatment environment. Female-authored j_g in the treatment environment are also co-authored with more co-authors and cited more whereas the opposite is true in the counterfactual environment. Relative to the counterfactual environment, the papers citing female-authored j_g in the treatment environment are published significantly later, with more co-authors and are more likely to be

[o]Our database of cited papers includes only English language papers published in a journal.

Table D.2: Treatment 2: j_g and l published in the same journal

	Treatment	Counterfactual	Difference
Cited papers			
Publication year	1.7359***	0.7033**	1.0327***
	(0.2135)	(0.3404)	(0.3116)
No. co-authors	0.2003**	−0.5821*	0.7824**
	(0.0830)	(0.3187)	(0.3420)
Journal impact factor	−0.1574	0.4234	−0.5807
	(0.2386)	(0.4866)	(0.4604)
No. citations	8.0631**	0.1475	7.9156*
	(3.3492)	(4.2044)	(4.4652)
Citing papers			
Publication year	0.3328***	0.1574	0.1754**
	(0.0841)	(0.1115)	(0.0699)
No. co-authors	0.1268***	−0.4782*	0.6049**
	(0.0387)	(0.2528)	(0.2681)
English language	−0.0004	−0.0019*	0.0015*
	(0.0008)	(0.0010)	(0.0008)
Journal article	0.0006	0.0151*	−0.0145
	(0.0037)	(0.0088)	(0.0092)

Note. Figures in the treatment and counterfactual columns are gender differences in paper characteristics when j_g and l are and are not published in the same journal, 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.

in English. We observe no significant differences between treatment and counterfactual environments in cited papers' journal impact factors and the probability that citing papers are journal articles (as opposed to, say, book chapters).

Table D.3 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.[p] Differences are rarely significant in the counterfactual environment; they are larger and more often significant, conditional on treatment. Nevertheless, their signs and magnitudes generally match; the only significant differences between the two environments is in gender differences in co-author prominence—relative to the counterfactual environment, women in the treatment group co-author with slightly more prominent co-authors compared to men.

The second two panels of Table D.3 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. Relative to the treatment environment, women's papers in the counterfactual environment are more often cited by journal articles compared to men's papers. Relative to the counterfactual environment, women's papers (and the papers that cite them) in the treatment environment are published later than men's papers.

Table D.4 reports estimates for the treatment and counterfactual environments when the treatment consists of cited papers by job market stars. We observe no meaningful differences between treatment and counterfactual environments in the characteristics of top-five papers. Relative to the counterfactual

[p] We do not show similar estimates in Tables D.1, D.2, D.5 and D.6 because the treatment is defined at the cited paper (rather than the author) level; as a result, treatment and counterfactual environments can include papers from the same author.

Table D.3: Treatment 3, k is an NBER working paper

	Treatment	Counterfactual	Difference
Top-five papers			
Publication year	3.0069***	1.3340	1.6729
	(0.5375)	(0.9308)	(1.0748)
No. co-authors	0.1300**	0.1158	0.0142
	(0.0535)	(0.0926)	(0.1070)
No. citations	−2.0763	−33.4924	31.4161
	(22.7596)	(39.4155)	(45.5146)
Co-author prominence	1.4523***	−0.0101	1.4625**
	(0.2950)	(0.5108)	(0.5899)
No. j_g	−2.5381***	−1.9314	−0.6068
	(0.9108)	(1.5773)	(1.8213)
Institutional rank	−0.3791	−0.8601*	0.4810
	(0.2566)	(0.4489)	(0.5171)
Share of *AER* papers	0.0771***	−0.0151	0.0923
	(0.0291)	(0.0504)	(0.0581)
Share of *Econometrica* papers	−0.0681***	−0.0094	−0.0586
	(0.0224)	(0.0388)	(0.0448)
Share of *JPE* papers	−0.0524**	0.0131	−0.0655
	(0.0232)	(0.0401)	(0.0463)
Share of *QJE* papers	−0.0057	0.0115	−0.0171
	(0.0218)	(0.0377)	(0.0436)
Share of *REStud* papers	0.0490**	0.0000	0.0490
	(0.0229)	(0.0397)	(0.0458)
Cited papers			
Publication year	1.6299***	−0.4174	2.0473***
	(0.3838)	(0.6715)	(0.7735)
No. co-authors	0.0089	0.1231	−0.1142
	(0.0668)	(0.1409)	(0.1560)
Journal impact factor	0.5514***	0.0803	0.4711
	(0.1495)	(0.3537)	(0.3845)
No. citations	1.9759	−0.2123	2.1882
	(1.2329)	(4.6591)	(4.8195)
Citing papers			
Publication year	0.3630***	−0.1544	0.5173*
	(0.1008)	(0.2521)	(0.2715)
No. co-authors	0.0113	0.0647	−0.0535
	(0.0184)	(0.0433)	(0.0470)
English language	0.0020	0.0003	0.0017
	(0.0013)	(0.0023)	(0.0026)
Journal article	−0.0029	0.0062	−0.0091*
	(0.0028)	(0.0040)	(0.0049)

Note. Figures in the treatment and counterfactual columns are gender differences when k is restricted to papers that are and are 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.

environment, female-authored cited and citing papers in the treatment environment are co-authored with more co-authors and more likely to be in English compared to male-authored cited and citing papers, respectively. However, both differences are small. We observe no other significant differences between environments in Table D.4.

In Tables D.5 and D.6 we report estimates for treatment and counterfactual environments when the treatment is defined as citing papers by authors who are colleagues (Table D.5) and co-authors (Table D.6) with the top-five author.[q] Among cited papers, gender differences in treatment environments resemble gender differences in counterfactual environments. Relative to the counterfactual, treatment environment papers citing female-authored papers are published slightly later than papers citing male-authored papers. In Table D.6, treatment environment papers citing female-authored papers are also co-authored with more co-authors compared to papers citing male-authored papers; the reverse is true in the counterfactual environment. Nevertheless, all differences are small.

[q]We technically define a citing and cited author colleague relationship to be one where the citing author is affiliated with any of the institutions listed on the cited author's first top-five paper. See Section 4.2, Footnote 36 for further details.

Table D.4: Treatment 4, k's author is a job market star

	Treatment	Counterfactual	Difference
Top-five papers			
Publication year	4.8300**	2.6501	2.1799
	(1.9168)	(2.0981)	(2.8418)
No. citations	2.6886	−108.2340	110.9226
	(95.2480)	(104.2577)	(141.2156)
No. j_g	−0.9858	−0.6398	−0.3460
	(0.6102)	(0.6679)	(0.9046)
Institutional rank	−0.4981	−0.1892	−0.3089
	(0.4621)	(0.5058)	(0.6851)
Share of *AER* papers	0.2931***	0.1139	0.1792
	(0.0918)	(0.1005)	(0.1362)
Share of *Econometrica* papers	−0.1041	−0.0832	−0.0208
	(0.0783)	(0.0858)	(0.1162)
Share of *JPE* papers	−0.0622	0.0025	−0.0647
	(0.0802)	(0.0878)	(0.1190)
Share of *QJE* papers	−0.0114	−0.0070	−0.0043
	(0.0746)	(0.0817)	(0.1106)
Share of *REStud* papers	−0.1154	−0.0261	−0.0893
	(0.0900)	(0.0986)	(0.1335)
Cited papers			
Publication year	0.1174	0.3695	−0.2521
	(0.2996)	(0.3303)	(0.4254)
No. co-authors	0.1591	−0.2789*	0.4381**
	(0.1458)	(0.1578)	(0.2178)
Journal impact factor	−0.1221	−0.0160	−0.1061
	(0.3716)	(0.6153)	(0.7114)
No. citations	2.8632	−1.4981	4.3613
	(2.9535)	(3.9503)	(5.0966)
Citing papers			
Publication year	0.2141	0.3815	−0.1674
	(0.1706)	(0.2588)	(0.3125)
No. co-authors	0.0177	−0.0077	0.0254
	(0.0462)	(0.0612)	(0.0772)
English language	−0.0042	−0.0167***	0.0125*
	(0.0039)	(0.0050)	(0.0067)
Journal article	0.0112	−0.0007	0.0119
	(0.0070)	(0.0117)	(0.0143)

Note. Figures in the treatment and counterfactual columns are gender differences among top-five authors who were and were not job market stars, 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.

Table D.5: Treatment 5: citations from colleagues

	Treatment	Counterfactual	Difference
Cited papers			
Publication year	1.1842***	1.1664***	0.0178
	(0.3493)	(0.4033)	(0.3234)
No. co-authors	0.0548	−0.0419	0.0967
	(0.0632)	(0.1013)	(0.1050)
Journal impact factor	0.4458***	0.4899**	−0.0440
	(0.1482)	(0.1943)	(0.1859)
No. citations	1.8704	0.6943	1.1761
	(1.4140)	(1.8854)	(1.6881)
Citing papers			
Publication year	0.2185**	0.1492	0.0693
	(0.1009)	(0.1165)	(0.0723)
No. co-authors	0.0253	0.0579**	−0.0326
	(0.0185)	(0.0258)	(0.0223)
English language	0.0019*	0.0007	0.0012
	(0.0011)	(0.0014)	(0.0013)
Journal article	−0.0014	−0.0020	0.0006
	(0.0024)	(0.0035)	(0.0033)

Note. Figures in the treatment and counterfactual columns are gender differences in paper characteristics when citing paper l's author is and is not a colleague of the top-five author, respectively; figures in the last column are their differences. Estimates 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.6: Treatment 6: citations from co-authors

	Treatment	Counterfactual	Difference
Cited papers			
Publication year	1.3880***	0.9856***	0.4024
	(0.2099)	(0.2859)	(0.2578)
No. co-authors	0.1648**	−0.2226	0.3875
	(0.0834)	(0.2793)	(0.3005)
Journal impact factor	0.1030	−0.0501	0.1531
	(0.2399)	(0.3307)	(0.3102)
No. citations	7.7286**	6.1783	1.5503
	(3.4098)	(4.2836)	(3.8469)
Citing papers			
Publication year	0.2675***	0.0627	0.2048***
	(0.0799)	(0.1019)	(0.0677)
No. co-authors	0.1352***	−0.6381***	0.7733***
	(0.0363)	(0.2114)	(0.2268)
English language	−0.0004	−0.0018*	0.0014
	(0.0008)	(0.0010)	(0.0009)
Journal article	0.0043	0.0011	0.0032
	(0.0036)	(0.0062)	(0.0052)

Note. Figures in the treatment and counterfactual columns are gender differences in paper characteristics when citing paper l's author has and has not co-authored with the top-five author, respectively; figures in the last column are their differences. Estimates 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.

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.[r]

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.

[r] Or equivalently, Figure D.1 reproduces Figure 5, Graphs (A) and (B) (*i.e.*, treatment and counterfactual group 3) but restricts estimation samples to j_g cited (column one) and not cited (column two) by k.

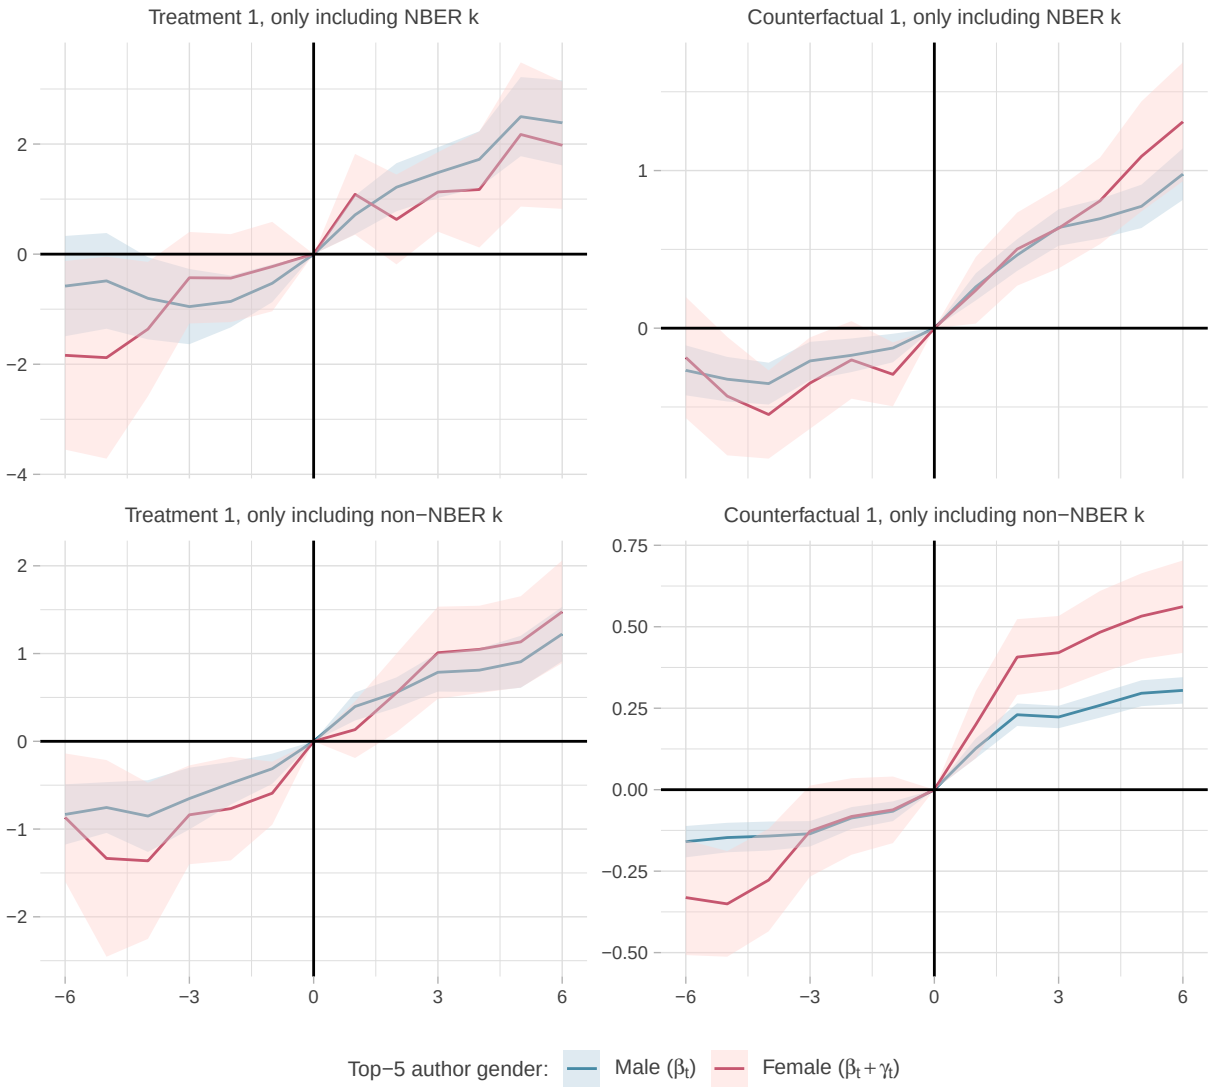


Figure D.1: Treatment and counterfactual group 1, with and without NBER k

Note. Graphs reproduce Figure 4, 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 Figures D.2, D.3 and D.4 we reproduce Figures 4, 5 and 6, respectively, but estimate Equation (3) using a Poisson likelihood function.[s] Results in most figures are similar to those shown in Figures 4, 5 and 6, although Figure D.3 may suggest an insignificant gender difference in the Matthew effect among the treatment category. However, we caution that the sample of female job market candidates (and especially female job market stars) is very small.

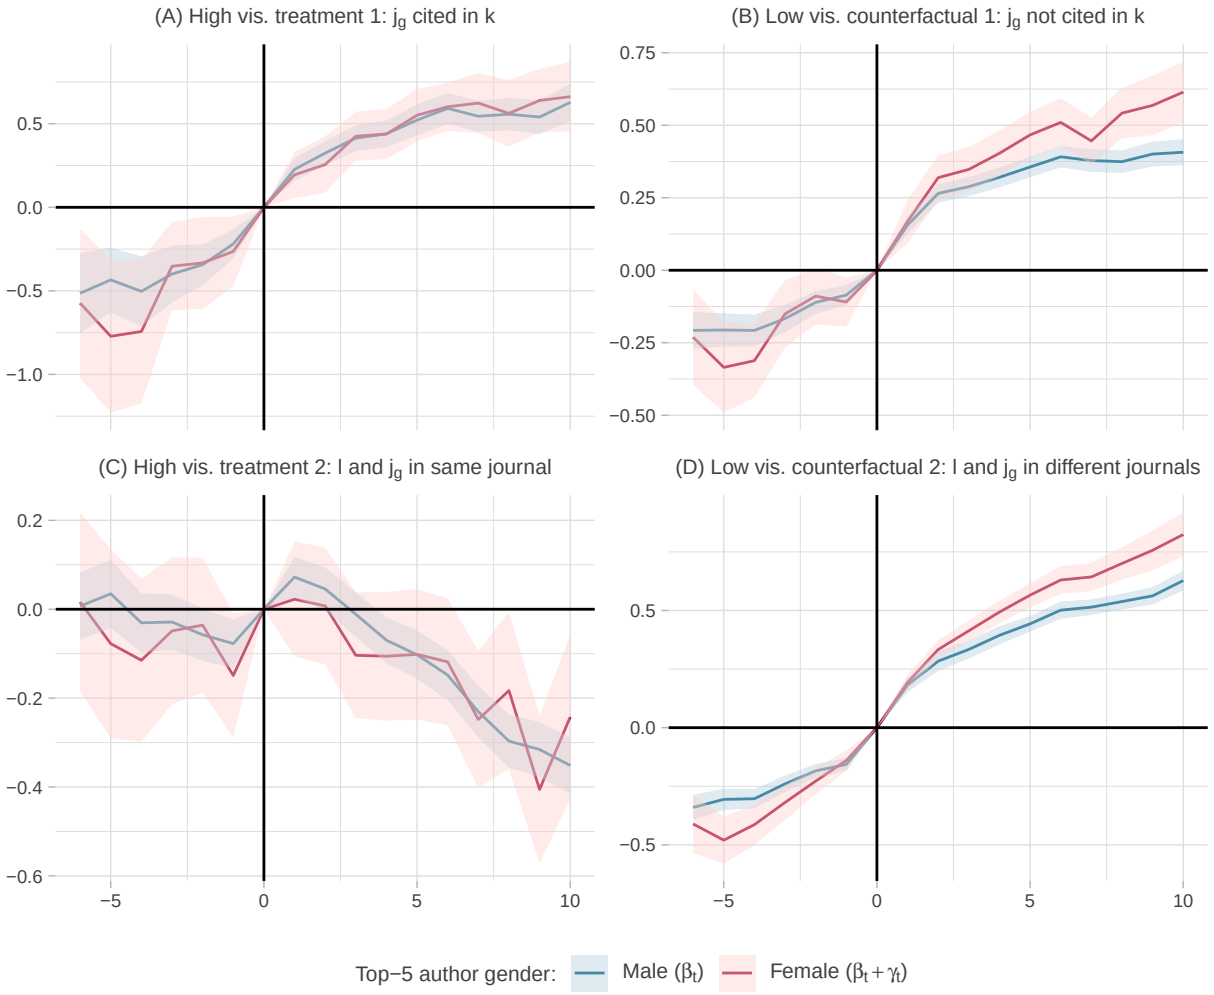


Figure D.2: Figure 4, Poisson likelihood function

Note. Graphs reproduce Figure 4, except they estimate Equation (3) using a Poisson likelihood function.

[s] Given small samples used to estimate Figure D.3, Graphs (C) and (D), we restrict the estimation window to 4 years before k's publication date.

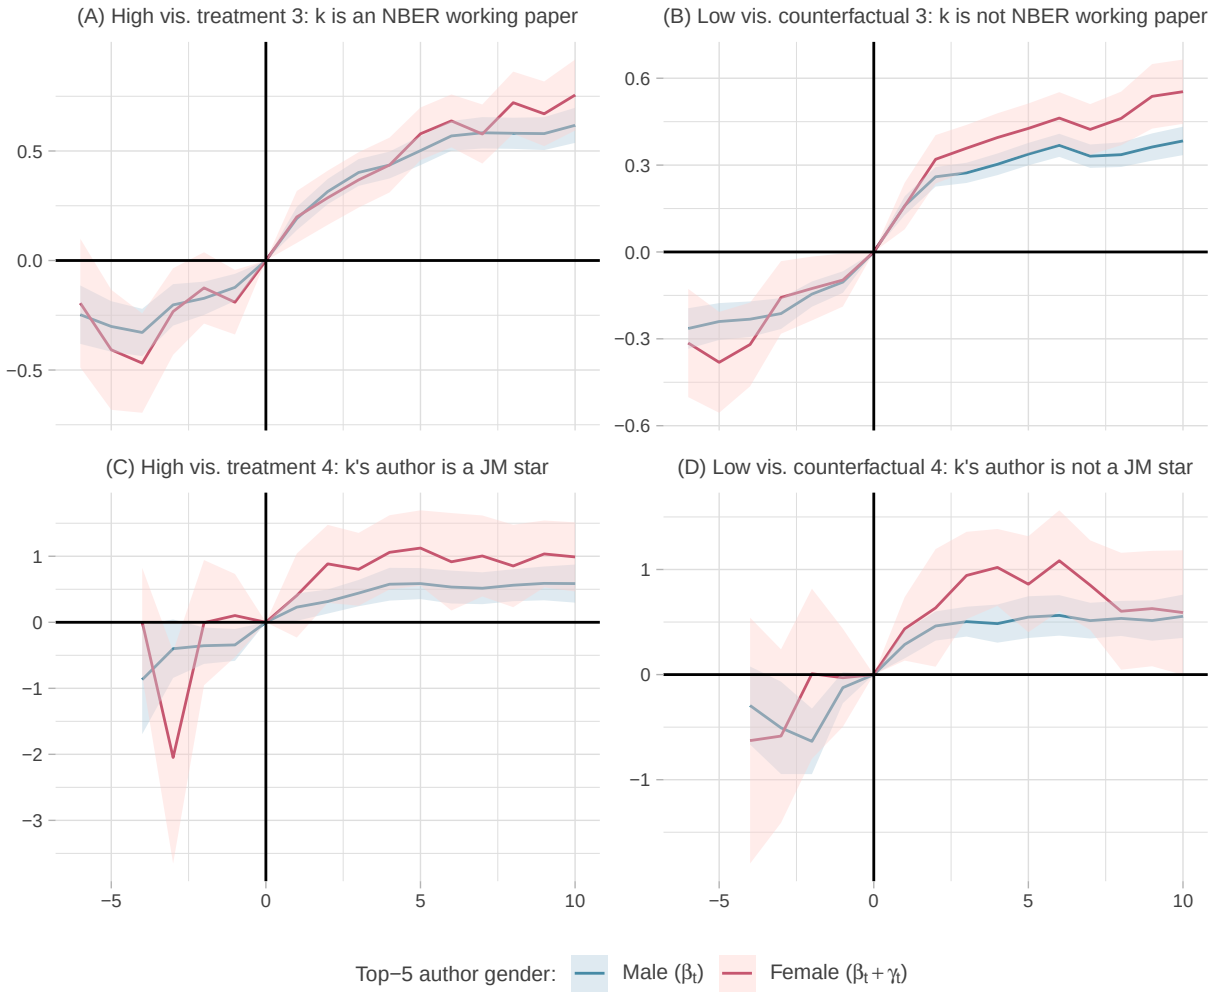


Figure D.3: Figure 5, Poisson likelihood function

Note. Graphs reproduce Figure 5, except they estimate Equation (3) using a Poisson likelihood function. Given small samples used to estimate Graphs (C) and (D), we restrict their estimation windows to 4 years before k's publication date.

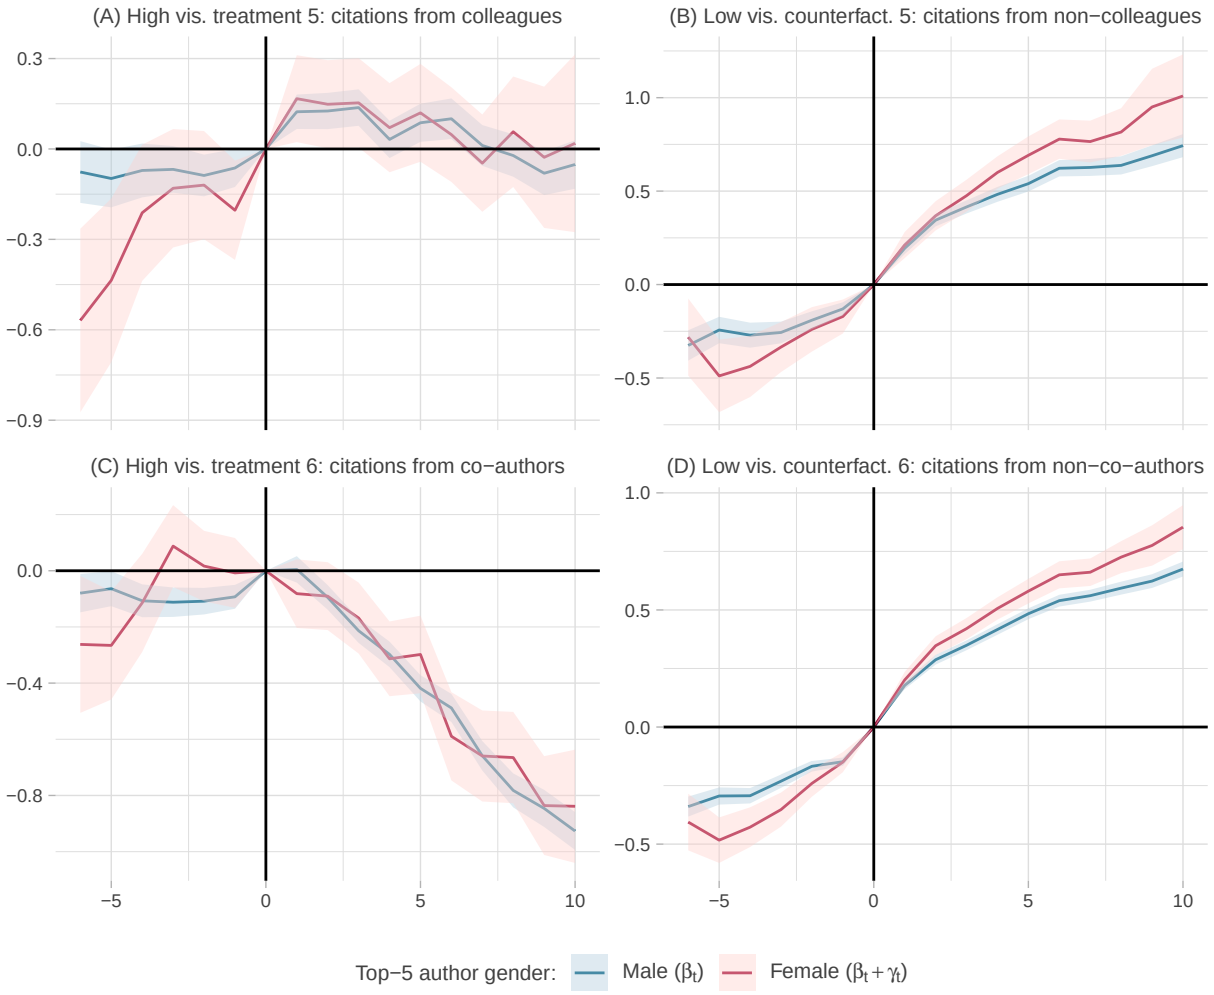


Figure D.4: Figure 6, Poisson likelihood function

Note. Graphs reproduce Figure 6, except they estimate Equation (3) using a Poisson likelihood function.

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

Figures D.5, D.6 and D.7 reproduce Figures 4, 5 and 6, respectively, except they also control for cited papers' journal impact factors and authors' total pre-top-five publication counts. Results in all figures are similar to the results reported in Figures 4, 5 and 6.

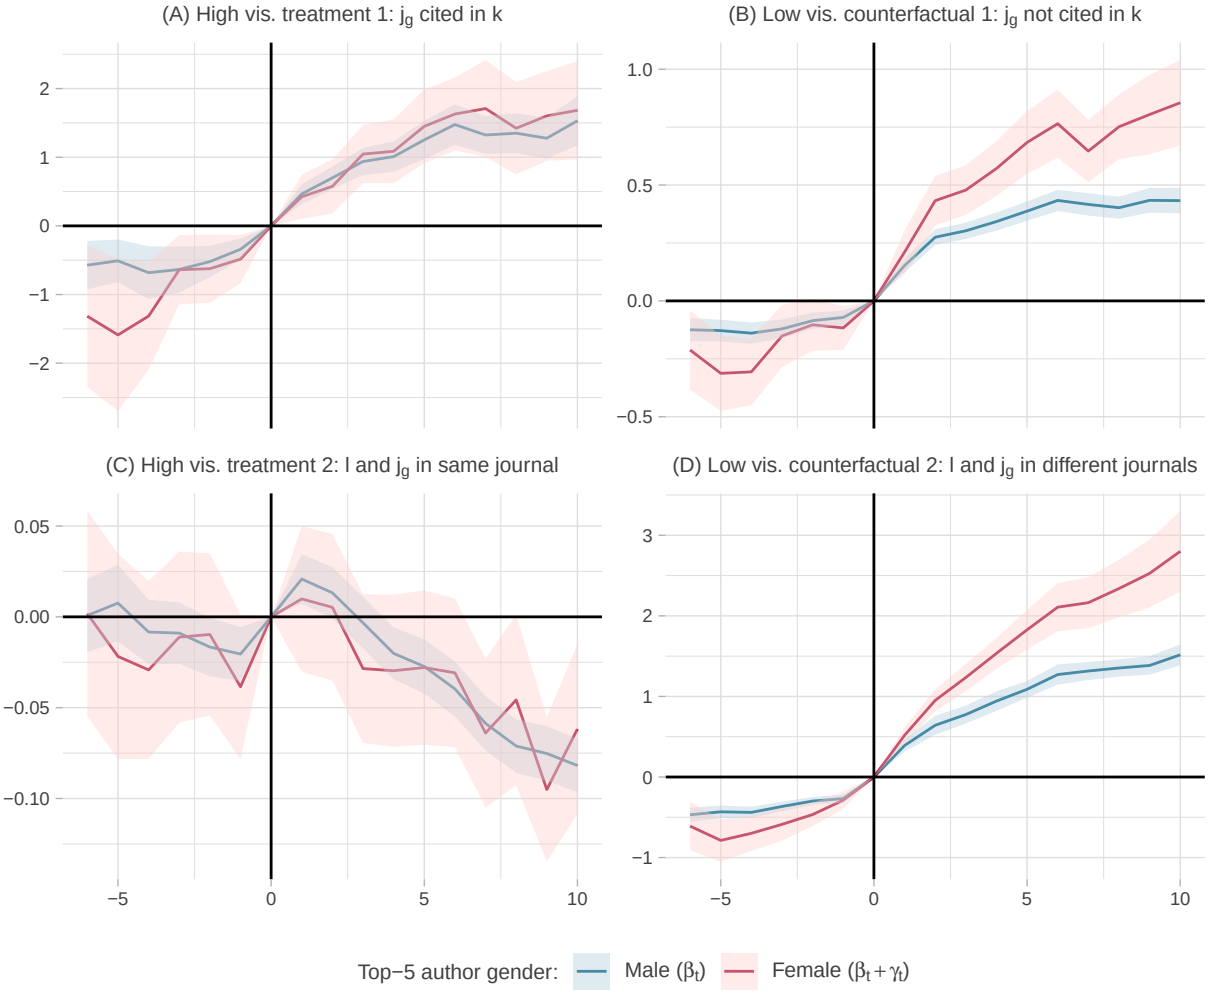


Figure D.5: Figure 4, controlling for evaluation histories

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

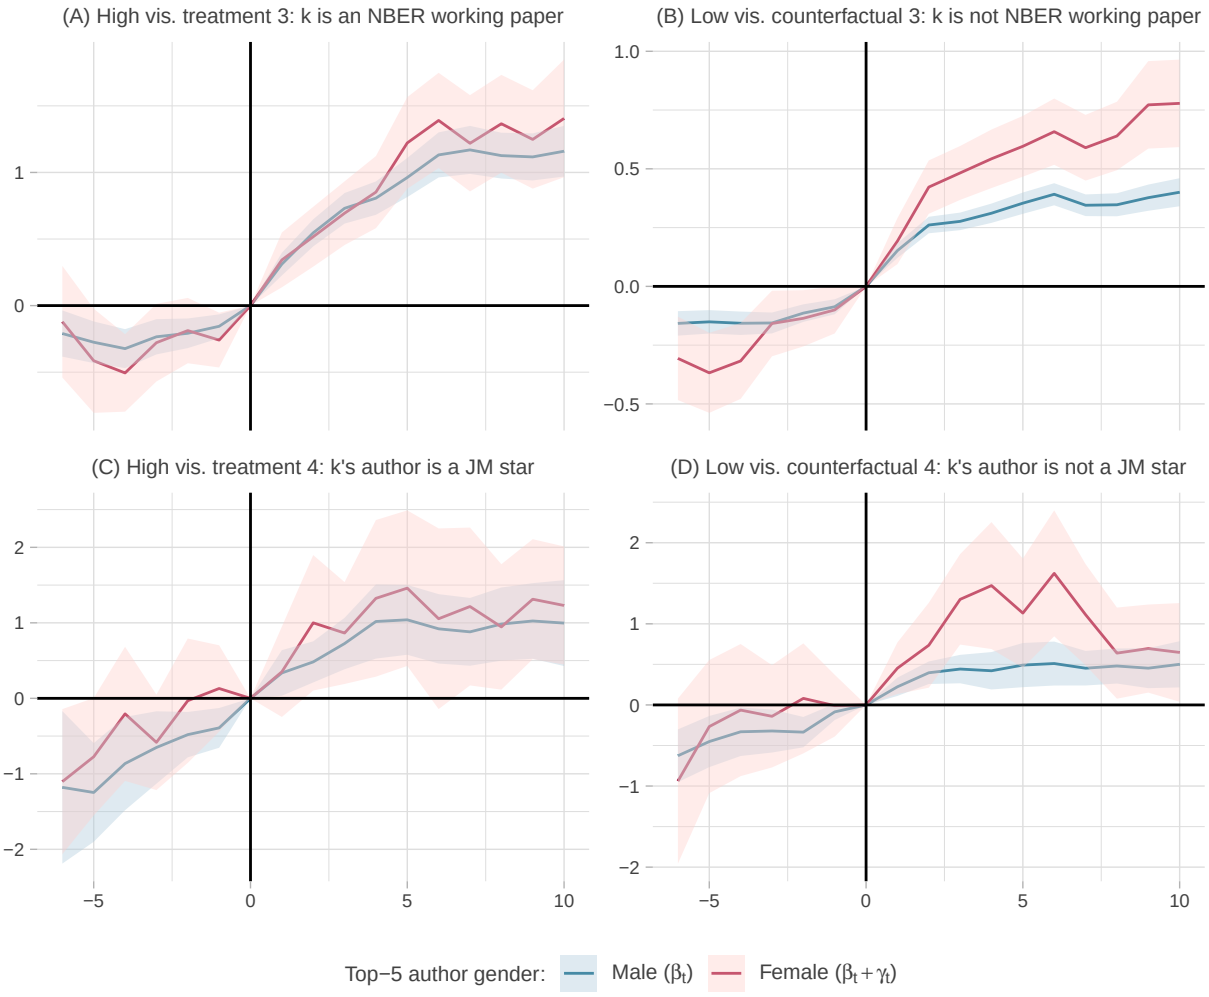


Figure D.6: Figure 5, controlling for evaluation histories

Note. Graphs reproduce Figure 5 controlling for the journal impact factors of cited papers and authors' total pre-top-five publication counts.

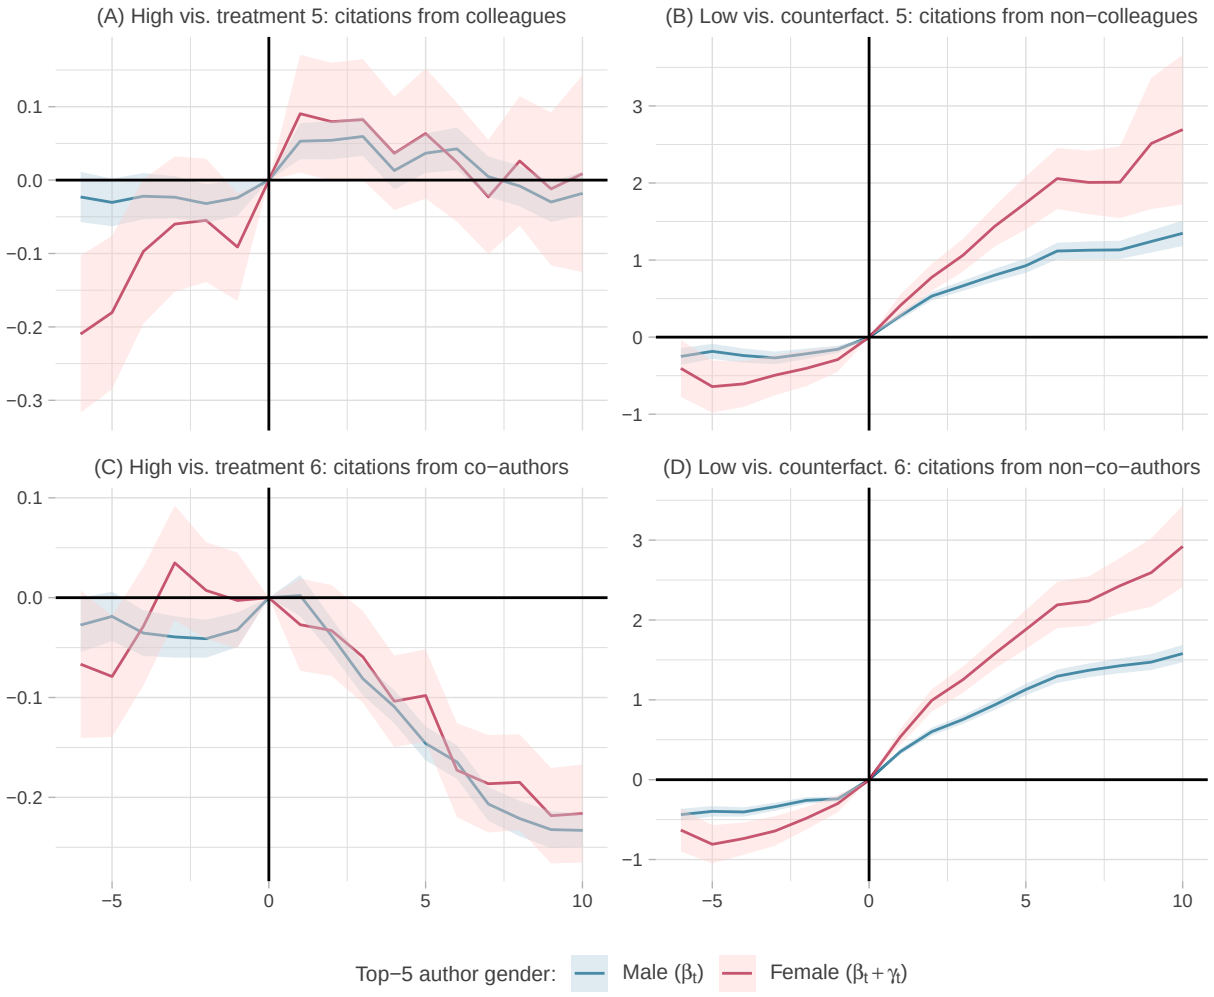


Figure D.7: Figure 6, controlling for evaluation histories

Note. Graphs reproduce Figure 6 controlling for the journal impact factors of cited papers and authors' total pre-top-five publication counts.

D.5 The Matthew effect by employer rank

In Figure D.8 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.[t]

This is indeed what we observe. The graphs in Figure D.8 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.

[t] 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.

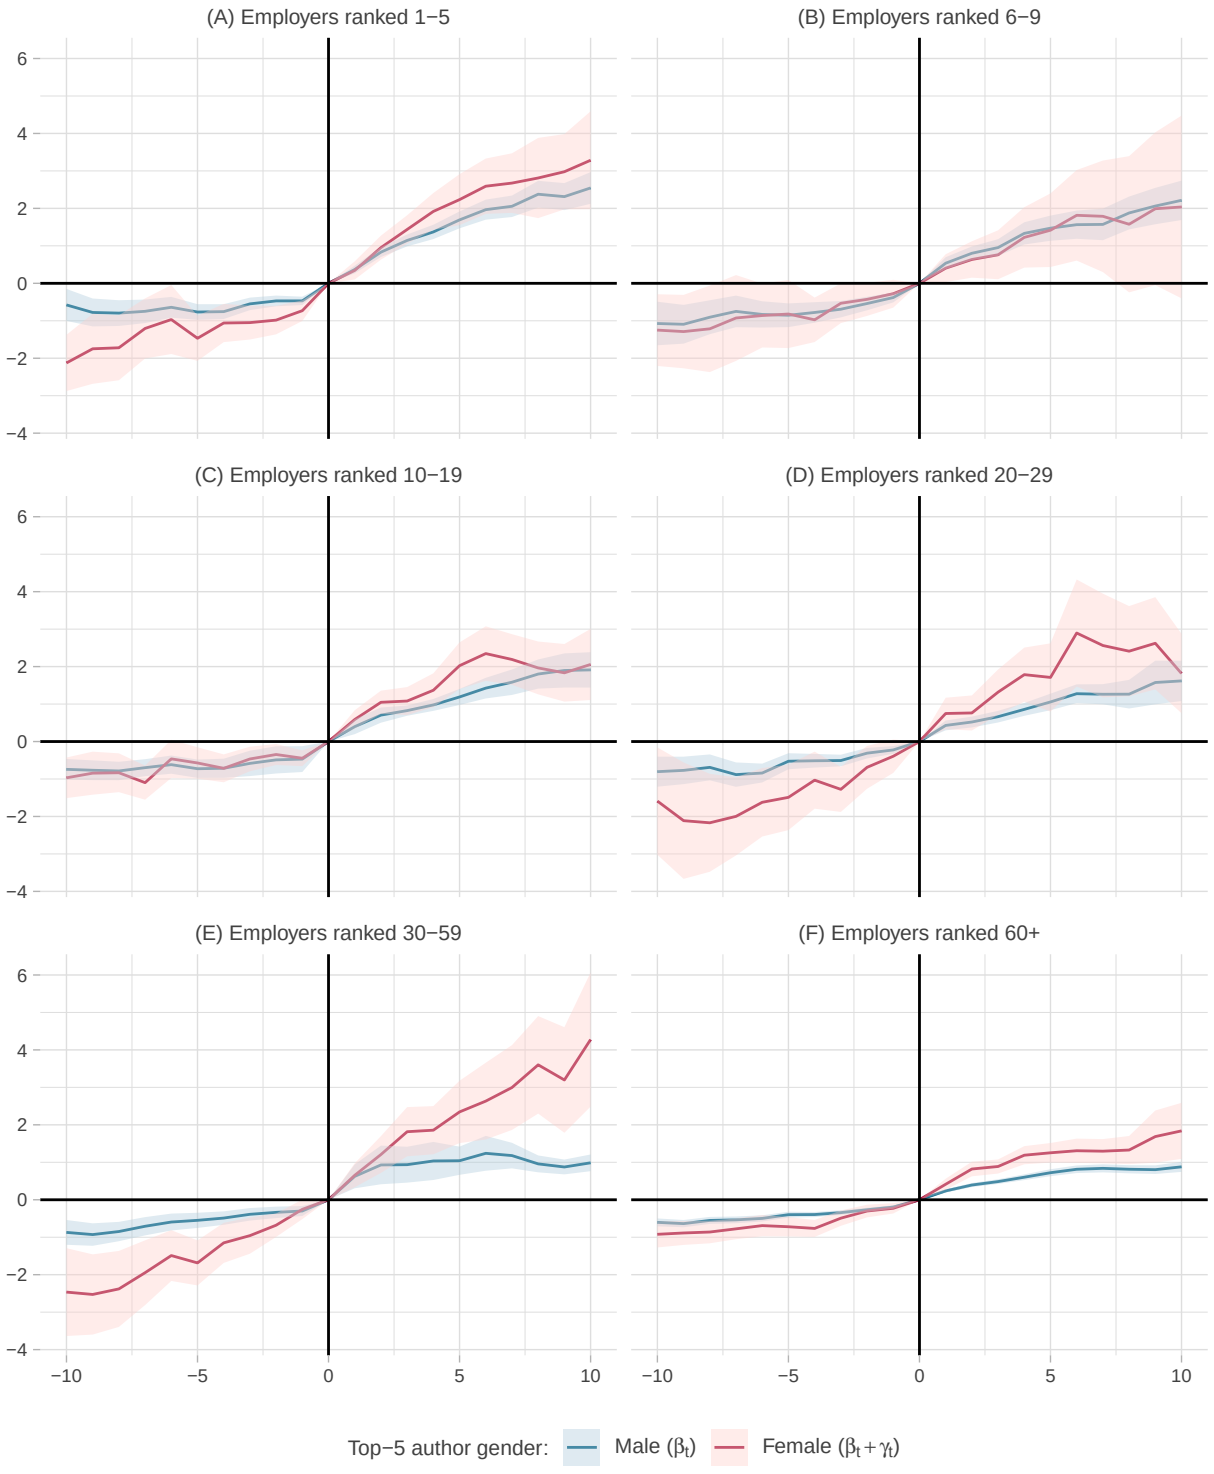


Figure D.8: The Matthew effect by employer rank

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 95 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 any journal, the top five economics journals and the top 20 economics/finance journals, where the latter include the 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*.[u] 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.

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.

[u] 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*.

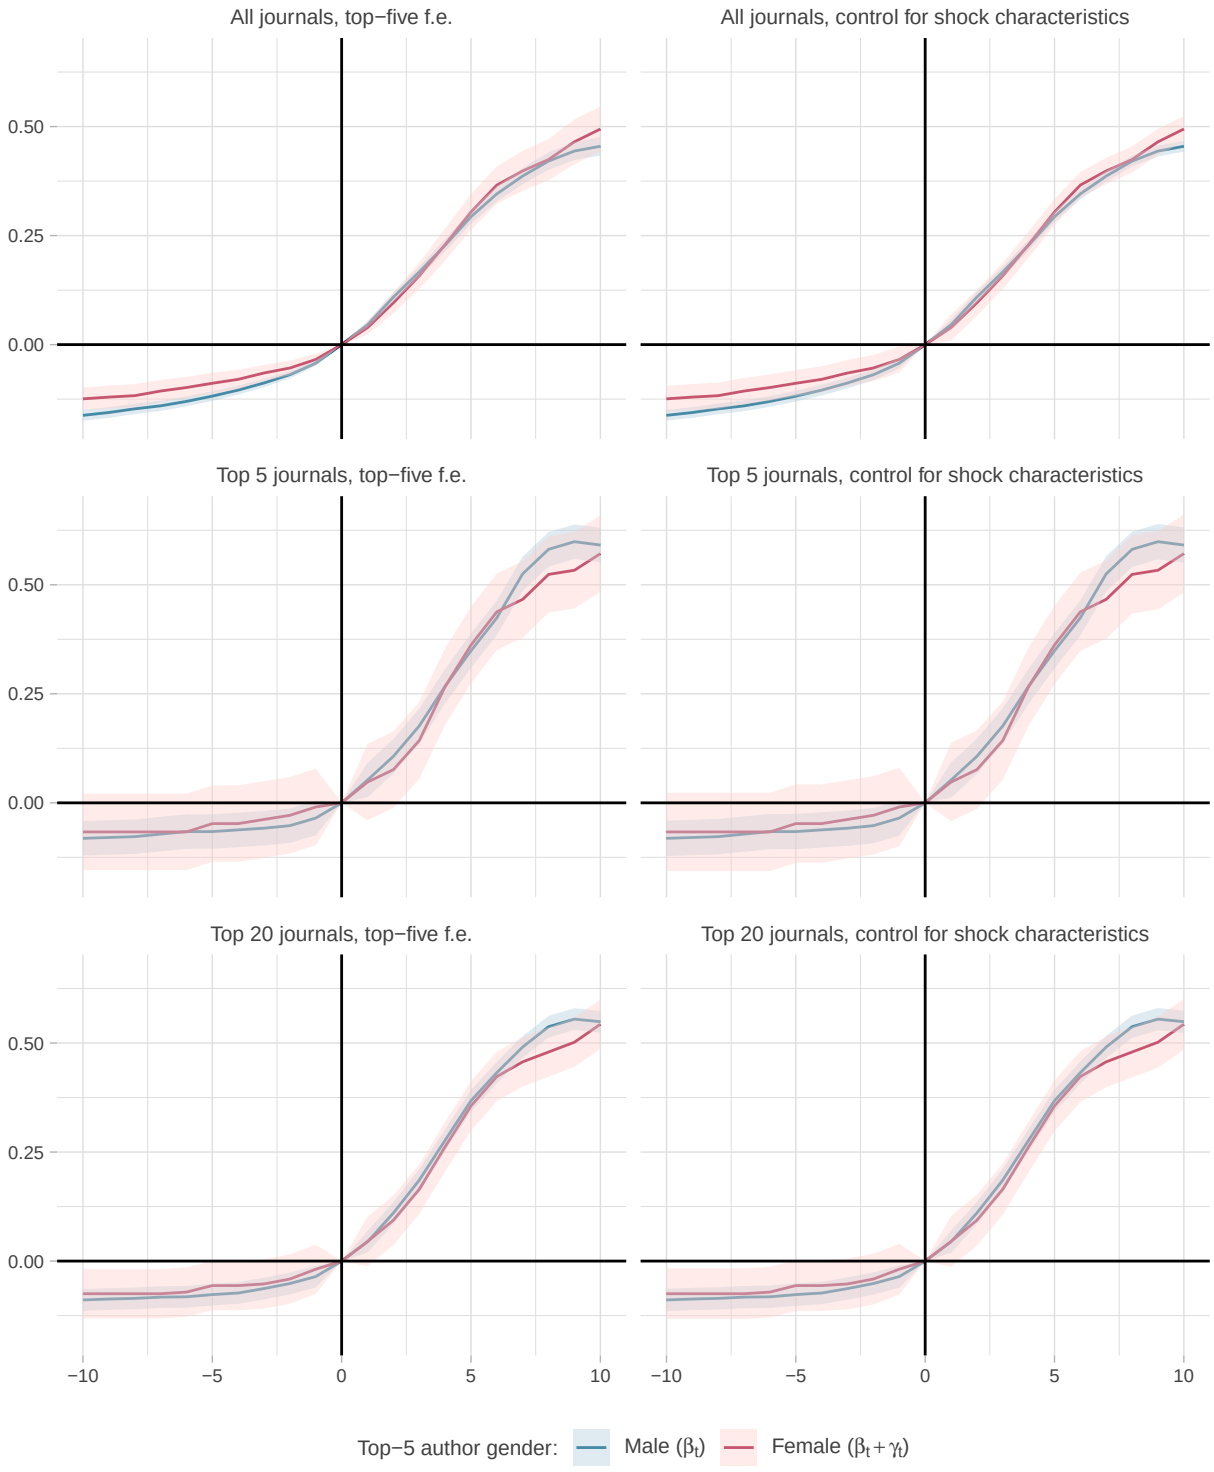


Figure E.1: The impact of k on editing a journal

Note. Graphs plot the probability that men and women become journal editors (as deviations from the mean) in any journal (top row), the top five economics journals (middle row) and the top 20 economics/finance journals (bottom row) 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 95 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.

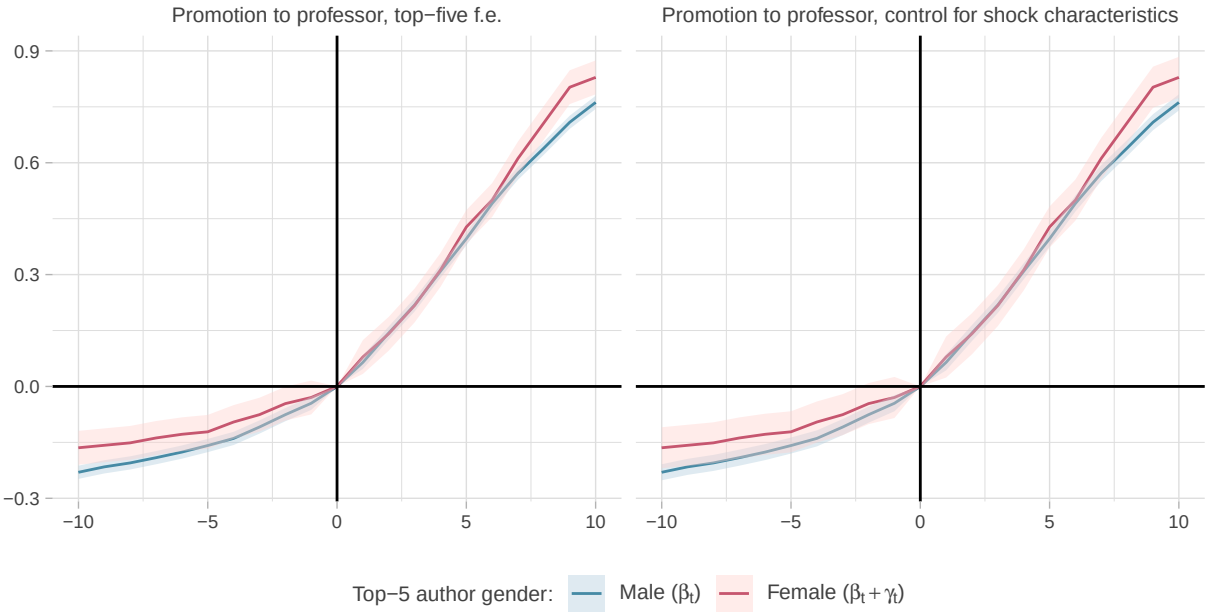


Figure E.2: The impact of k on being promoted to professor

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 95 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.

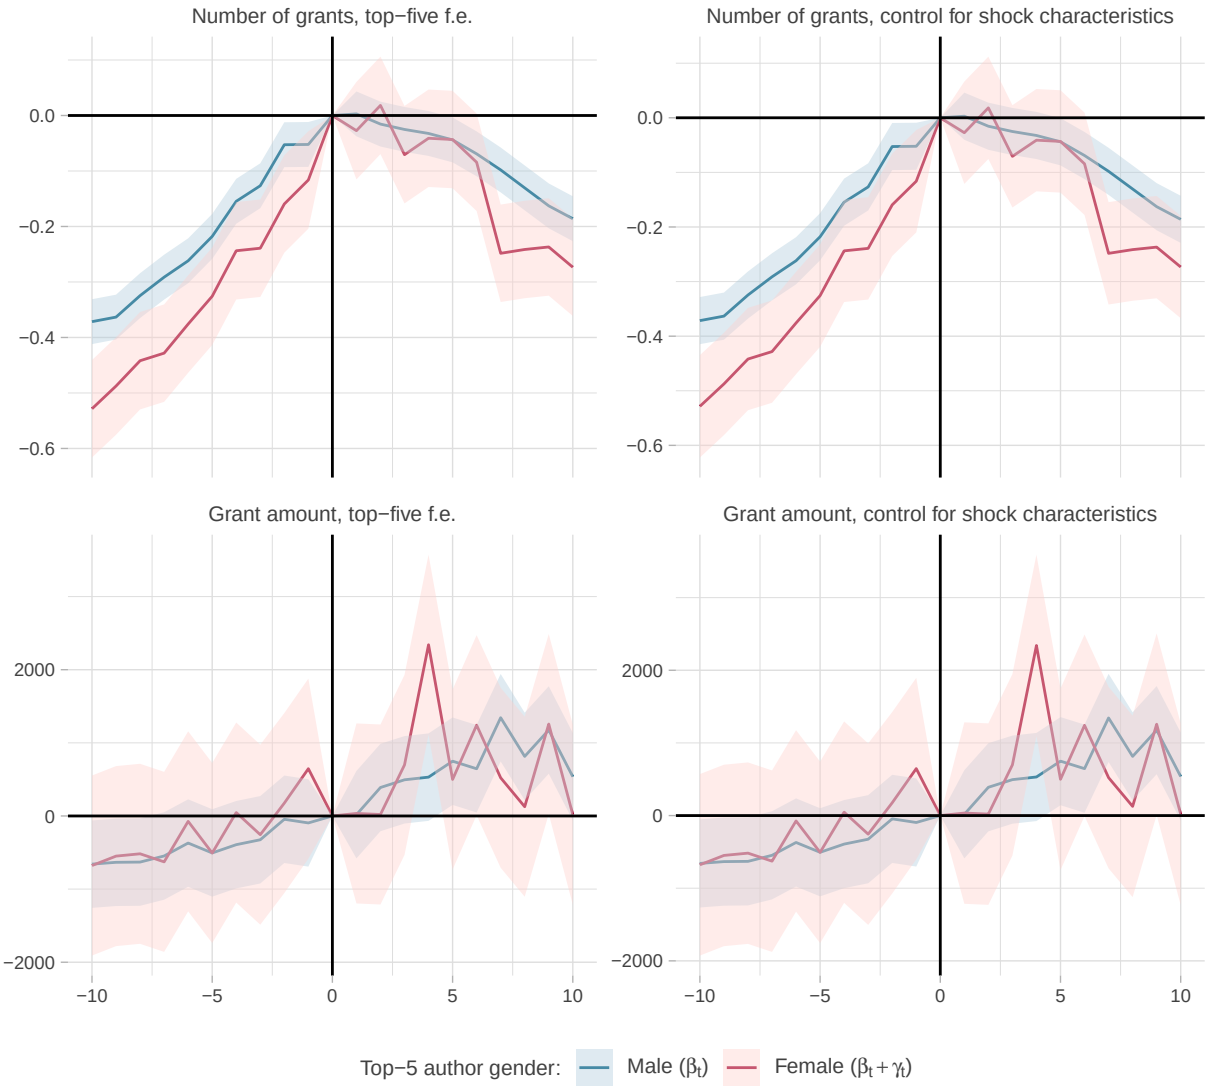


Figure E.3: Number and size of awarded grants (in thousand USD)

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 the relevant dependent variables using OLS and reflect deviations from the mean. Shaded areas represent 95 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.[v] 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.

[v] 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.[w] 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 employer institutional rank as the dependent variable using OLS. Shaded areas represent 95 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.

[w] This ranking is identical to the one described in Appendix C.10, 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 seminar counts as the dependent variable using OLS. Shaded areas represent 95 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.[x]

[x]Compared to their male co-authors on k, women give 0.5 (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.[y]

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.

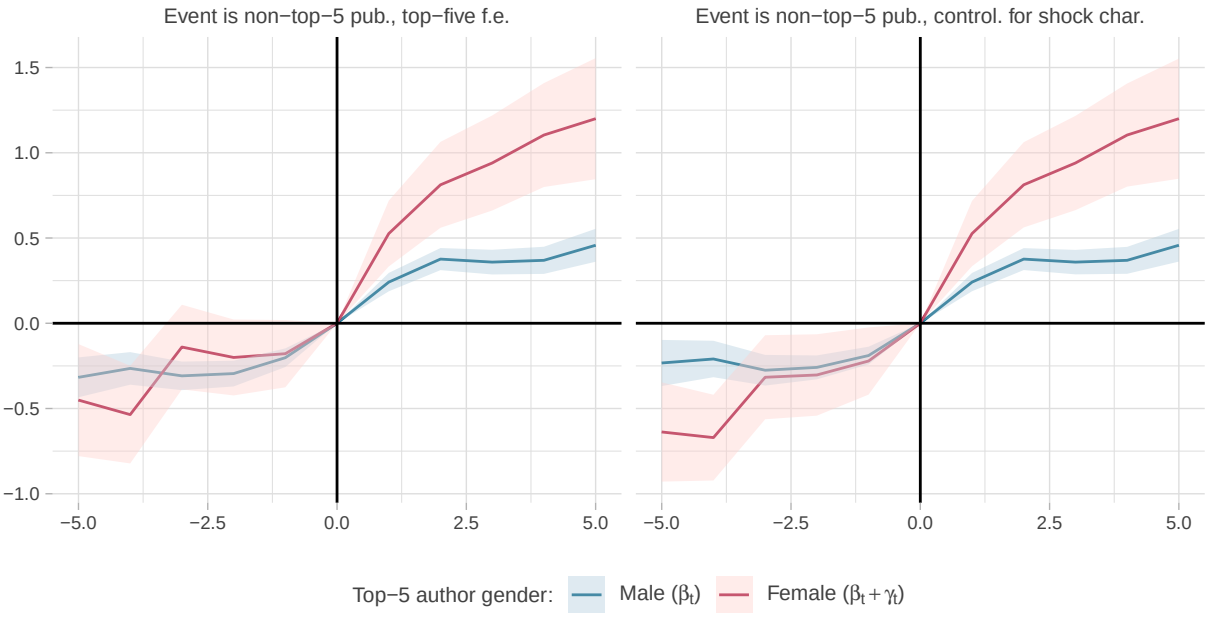


Figure F.1: The Matthew effect for non-top publications

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.

[y] 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 NBER working paper release date

In Figure F.2 we replicate Figure 3 but (a) restrict the sample to authors who initially released their first top-five paper as an NBER working paper; (b) bring the $t = 0$ event date forward in time to the NBER release date; and (c) restrict the sample of previous papers j_g to papers published before k's NBER working paper release date. This environment mirrors the one in Figure 4(A), 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. Indeed, this is precisely what we observe: Figure F.2 suggests a citation response that is much larger for women than it is for men.

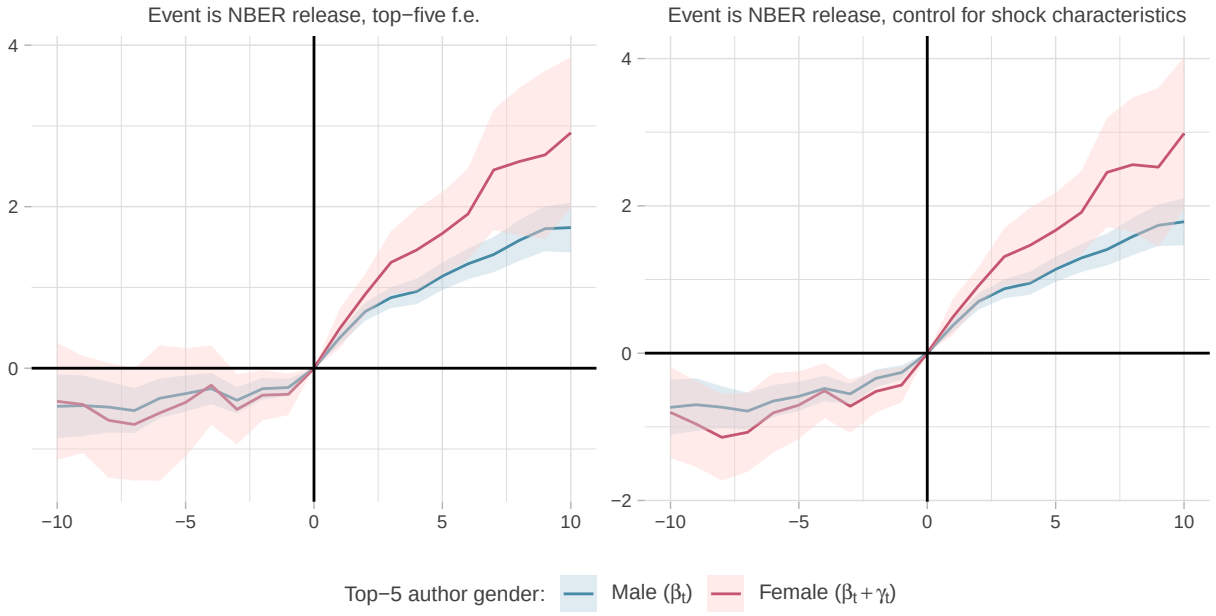


Figure F.2: The Matthew effect for NBER working papers

Note. Graphs reproduce Figure 3 but restrict the sample to authors who initially released k as an NBER working paper, set $t = 0$ to this release date and restrict j_g to papers published before $t = 0$.

F.3 Publication of the second top-five paper

In Figure F.3 we replicate 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.[z] Thanks to the publicity generated from an author's first top-five publication, potential citing papers should already be aware of his earlier work when he publishes his second top-five paper; as a result, Lemma 3 anticipates zero gender differences in the Matthew effect at $t = 0$. This is indeed what we observe: in Figure F.3, $\beta_t + \gamma_t$ closely tracks β_t for most of t.



Figure F.3: The Matthew effect when authors publish a second top-five paper

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

[z] Note that we continue to only include j_g published before an authors *first* top-five publication date.

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, Euyoung Moon, and Richard Tol (2024). "Gender and equality at economics and finance journals". Mimeo.