Research Assistants, Sorting, and Career Outcomes: Evidence from the NBER Working Paper Series

Florian Caro*

September 2, 2024

Abstract

Research assistant (RA) positions play an increasingly important role in the economics profession, both for generating research and for nascent researchers to acquire skills, gather experience, and build professional networks. Despite this, we know little about the demographics of RAs, access to RA positions, and the impact of RA experience on downstream career outcomes. Using an original dataset on RAs collected from the acknowledgments of working papers published by the National Bureau of Economic Research, I (i) present novel, large-scale descriptive evidence on RAs, (ii) show that there exist strong sorting patterns between RAs and supervisors (PIs) along gender, race, and ethnicity, and (iii) provide evidence that gender alignment between RAs and PIs has a meaningful impact on the career outcomes of RAs.

*florian.caro@yale.edu. For helpful comments and suggestions I thank John Eric Humphries, Maria Clara Rodrigues, and Seth Zimmerman as well as seminar participants at Yale.

1 Introduction

Research assistant (RA) positions play an increasingly important role in the economics profession. As the discipline has become more empirical, the need to process large amounts of data and manage experiments has grown, and thus, so has the demand for RAs. At the same time, RA positions have a long tradition as an opportunity for college students, recent graduates, and PhD candidates to accumulate human and social capital in the form of research experience and professional networks. Beyond that, RA positions also serve as a source of additional funding. Despite this, RA positions remain an understudied part of the academic career path. A small body of research suggests that the financial benefits of RA positions increase PhD completion rates (Colander and Klamer, 1987; Stock, Siegfried and Finegan, 2011; Stock and Siegfried, 2014, 2015) and that the human and social capital component improves the career prospects of graduates (Bryan, 2019; Hansen, 1991; Krueger, 1991). These findings are, however, based on small samples and largely descriptive.

The need for more research on RA positions is amplified by evidence of persistent bias against women as well as racial and ethnic minorities in academia and economics in particular (see e.g. Dupas et al., 2021; Marschke et al., 2018; Sarsons et al., 2021). Existing research has shown that such bias exists at the undergraduate, graduate, and post-graduate levels, but has not investigated the RA stage. RA positions demand separate analysis as the net effects of bias are *a priori* unclear in this setting: Assuming homophilic preferences, e.g. men prefer to work with men, access to RA positions might be unequal because of the present underrepresentation of minorities among faculty. At the same time, minority faculty members might invest more resources in minority RAs, improving the outcomes of those who manage to find an RA position.

In this paper, I show that there exist clear sorting patterns between RAs and their supervisors (PIs) along both gender and race. Additionally, I provide the first descriptive evidence on the prevalence and composition of RAs over time and present evidence on the importance of gender alignment between RAs and PIs for downstream career outcomes of RAs. My analysis is based on a novel data set on over 16,000 RAs and their respective PIs. I collect this data from the acknowledgments section of publications in the Working Paper Series of the National Bureau of Economic Research (NBER), which has been running since 1973. I supplement this data with information on over 50,000 doctorates in economics to analyze the career outcomes of economists who worked as RAs during their graduate studies.

I first show that both the share of papers acknowledging RAs and the average number of RAs per paper have steadily increased between 1973 and 2023. Whereas only around 15% of papers published in the 1970s and 1980s mentioned any RAs in their acknowledgments, almost 40% did so in 2023. Over the same period, the average number of RAs mentioned on papers with any RAs rose from about 1.5 to over 3, suggesting that the increased share of papers with RAs does not only reflect trends in reporting. I also find changes in the composition of RAs with the share of women, racial, and ethnic minorities growing considerably over time. Yet, with the exception of Asians, none of the minority groups reaches a level of equal representation.

Further analysis of the links between RAs and their respective PIs reveals strong and statistically significant sorting along both gender and race. The observed sorting is not explained by other observable attributes, such as the affiliation, seniority, or field of research of the PI, in line with existing evidence on sorting among coauthors (Davies, 2022; Sarsons et al., 2021). I find female PIs to be almost 33% more likely to work with a female RA than male PIs. Sorting patterns by race and ethnicity are even stronger with for example Asian PIs being almost 72% more likely to work with an Asian RA than White PIs.

Finally, I present results that suggest that an RA working with a PI of the same gender achieves better short- and long-term career outcomes than their peers working with PIs of the opposite sex. An analysis of the underlying mechanism suggests that these results are driven in part by higher rates of coauthoring between RAs and PIs of the same gender. Specifically, my results suggest that a female RA with a female PI is over 24.4% more likely to publish a paper in the NBER Working Paper Series within 5 years after completing her PhD, compared to her female peer with a male PI. This finding provides additional evidence on the importance of female role models in promoting new generations of women in academia (Boustan and Langan, 2019; Carrell, Page and West, 2010; Hilmer and Hilmer, 2007; Neumark and Gardecki, 1998).

This paper is primarily related to two strands of literature. First, a large and growing body of work documents the obstacles faced by minority groups and in particular women in academia. Past research has shown that women get less credit for coauthored papers than their male peers (Sarsons et al., 2021), face harsher treatment in seminars (Dupas et al., 2021), and are more likely to leave academia (Spoon et al., 2023). Research focused on the economics profession also reports a chronic underrepresentation of minorities within the field (Bayer and Rouse, 2016; Fortin, Lemieux and Rehavi, 2021; Hale and Regev, 2014; Lundberg and Stearns, 2019) despite gradual improvements and effective small-scale interventions (Becker, Rouse and Chen, 2016). While inspired and informed by findings on discrimination by gender, race, and ethnicity in the general labor market (e.g. Bertrand and Mullainathan, 2004), this literature accounts for the peculiarities of the academic labor market, where different metrics of success (Galiani and Panizza, 2020) imply different ways that discrimination can manifest.

The second related strand of literature studies the idiosyncrasies of the economics profession, including the training, career paths, and work of economists. Within this literature, a small body of descriptive and qualitative research suggests that graduate students who work as an RA during their PhD exhibit higher degree completion rates and better career outcomes (Bryan, 2019; Hansen, 1991; Krueger, 1991; Stock and Siegfried, 2014). The notion that RA experience conveys valuable human capital is also reflected in the growing demand from top PhD programs for applicants with that background (Economist, 2020). A final factor highlighting the relevance of RA positions is the recent rise of coauthoring and research teams in economics (Jones, 2021) which suggests that the demand for RAs might have experienced a similar rise.

I contribute to the existing literature on four fronts: (i) First, I introduce and demonstrate the effectiveness of a novel, text-based method that facilitates the large-scale collection of data on RAs. (ii) Second, I use this method to provide new descriptive evidence on RAs in economics and document long-run changes in their prevalence and composition. (iii) Third, I show that there exist strong sorting patterns among RAs and PIs, similar to those found among coauthors (Davies, 2022; Sarsons et al., 2021). These sorting patterns suggest that minority groups have less access to RA positions because of the existing underrepresentation among faculty. (iv) Finally, I confirm existing evidence that RA positions are an important opportunity for PhD candidates to accumulate human and social capital, and present first evidence on the importance of RA-PI gender alignment for downstream career outcomes of RAs.

The remainder of this paper is structured as follows: Section 2 introduces the data sources I am using and describes the sample construction. Section 3 highlights important descriptive facts in the data. Section 4 and Section 5 present my analysis and results on sorting among RAs and PIs and academic career outcomes of RAs, respectively. Section 6 discusses the implications of the results presented in the previous sections and Section 7 concludes.

2 Data

2.1 Sources

My main data source on RAs and their respective PIs is the NBER Working Paper Series.[1] Since 1973 the NBER has been publishing working papers of its affiliates, which are circulated for discussion and comment. While these papers do not list RAs in a standardized format, authors often thank RAs in the acknowledgments section of the paper. I extract the names of RAs mentioned in the acknowledgments of each paper and assign at least one of the authors as their PI.

The assignment algorithm matches RAs with authors by selecting the pair(s) with the highest number of co-occurrences across all observed working papers. This algorithm does not always yield a unique RA-PI match, and it is also possible that the same RA has more than one PI, e.g., dual appointments of RAs or appointments in different educational stages, such as undergraduate and doctoral studies. To limit noise in the RA-PI pairs without discarding real matches, I restrict the sample to RAs that my algorithm assigns to no more

[1] `https://www.nber.org/papers`, retrieved November 18th.

than four different PIs. Appendix Sections A.1 and A.2 describe the data extraction and RA-PI pairing in detail. In addition to the names of RAs and PIs, I also collect information on the date of publication,[2] the current affiliation of all authors on the paper, and the fields of research that the paper is associated with.

I obtain gender and race information for the individuals in my NBER sample from a mix of data sources. First, Davies (2022) provides information on the gender of NBER authors up to 2020.[3] Second, Meade, Starr and Bansak (2021) provide a dataset with information on gender among presenters at the annual meetings of the American Economic Association. Finally, for any names in my sample that could not be classified with the sources above, I use data from the SSA to infer the gender of a given RA or PI based on their first name. Race and ethnicity are similarly inferred using US Census data on last names. While name-based inference is not perfect, especially in the case of race and ethnicity, past research has shown that it can recover gender and racial disparities (Rieke et al., 2022) as intended in the present context.

The working papers do not provide any information on RAs beyond their name and the researchers they worked with, which prevents an analysis of their academic career outcomes. To overcome this shortcoming, I collect records on PhD graduates from several sources, including EconLit, hand-collected data on job market placements, the RePEc Genealogy project, and the Mathematics Genealogy Project. Combining all of these records yields a data set with over 50,000 unique PhD graduates which I then link to RAs based on name and year of graduation.

The data on PhD completions allows me to identify RAs who were likely graduate students during their time as RA. In particular, if the paper an RA is mentioned on was published up to 5 years before I observe the same RA complete their PhD, I classify them as a "graduate RA." Manual inspection of CVs for a random subsample of RAs suggests that this cutoff is appropriate and robustness checks show it to be not sensitive to small variations.[4] I subsequently select all identified graduate RA spells, refine the RA-PI match by requiring that the university affiliation of the PI and the PhD-granting university of the RA are the same, and keep only those RAs whose PIs are all of the same gender, i.e. if an RA has more than one PI assigned to them, I require that all PIs are either male or female. The benefit of this final sample restriction step is that correct identification of the gender of the PI does not require precise identification of the RA-PI link, a non-trivial problem given that I observe many RAs only once.

[2]In case of papers that were eventually updated at a later date, I still take the date of initial publication. Using the most recent date that a paper was updated instead does not change my results in a meaningful way. This is in line with the observation that most updates to papers take place within 1-2 years after initial publication.

[3]Davies (2022) uses baby name data from the United States (US) Social Security Administration (SSA). The author augments this data with information from Facebook and through manual verification.

[4]Note that the papers I am looking at here are working papers and will thus take less time on average to be published as such. The same cutoff would most likely be less appropriate in the context of peer-reviewed publications.

The descriptive analysis in Section 3 and the sorting analysis in Section 4 are based on my full sample of RAs and their respective PIs, while the analysis of career outcomes among former RAs in Section 5 is based on the restricted sample of graduate RAs. The reasons for focusing on graduate RAs in the outcome analysis are threefold. First, the composition of the full sample of RAs is heterogeneous and contains observations on undergraduates, pre-doctoral fellows ("predocs"), and doctoral students. Comparing outcomes between these groups is difficult in the best circumstances and becomes virtually impossible without information on who belongs to which group. Focusing on graduate RAs resolves this issue. Second, working with graduate RAs gets as close as possible to capturing the direct effect of RA experience on academic career outcomes like the number of publications. The third reason for working with graduate RAs is that I can control for the institution that a given individual received their PhD from. In the absence of information on undergraduate education, test scores, and other proxies for ability, this enables me to control for unobserved heterogeneity among RAs that might affect sorting and downstream career outcomes.

2.2 Summary statistics

Table 1: Summary Statistics

	RAs			PIs	
RA Variables	(1) All RAs	(2) Graduate RAs	PI Variables	(3) All PIs	(4) PIs of Graduate RAs
Female RA	0.38 (0.49)	0.27 (0.45)	Female PI	0.24 (0.43)	0.13 (0.34)
Asian	0.37 (0.48)	0.34 (0.47)	Asian	0.19 (0.39)	0.12 (0.32)
Black	0.01 (0.09)	0.00 (0.06)	Black	0.00 (0.06)	0.01 (0.09)
Hispanic	0.09 (0.28)	0.09 (0.29)	Hispanic	0.07 (0.25)	0.05 (0.21)
White	0.53 (0.50)	0.57 (0.50)	White	0.75 (0.44)	0.83 (0.38)
PhD Rank	28.32 (41.65)	22.76 (31.80)	Affiliation Rank	40.15 (58.18)	27.47 (38.73)
PhD Top 5	0.34 (0.47)	0.37 (0.48)	Affiliation Top 5	0.28 (0.45)	0.33 (0.47)
N PIs	2.53 (0.93)	1.30 (0.58)	N RAs	5.09 (7.62)	12.52 (13.45)
Any Female PI	0.40 (0.49)	0.32 (0.47)	Any Female RA	0.67 (0.47)	0.80 (0.40)
Any Male PI	0.94 (0.24)	0.95 (0.21)	Any Male RA	0.83 (0.37)	0.96 (0.19)
N	14,415	2,095	N	5,557	896

Notes: This table shows summary statistics for the different samples of RAs and PIs used in my analyses in Section 3 through Section 5. Each column reports sample means, standard deviations, and the total number of observations. Missing values are excluded where applicable. Columns (1) and (3) report values for the full sample of RAs and PIs and columns (2) and (4) report values for the graduate RAs sample and the corresponding PIs. The construction of both samples as well as the data sources used are described in Section 2.1. The period of observation is 1973-2023.

Table 1 shows summary statistics for unique RAs and PIs. Columns (1) and (3) report statistics for the full sample of RAs and their respective PIs, whereas columns (2) and (4) report statistics for the sample of graduate RAs and their respective PIs. Starting with the distribution of gender, race, and ethnicity, Table 1 shows that female RAs are outnumbered by male RAs in both samples, with 38% of the full sample being female and 27% of the graduate RAs sample. In terms of race and ethnicity, White RAs make up the majority with 53% in the full sample (57% in the graduate RAs sample), followed by Asians with 37% (34%), Hispanics with 9% (9%), and finally Black RAs with only 1% (< 1%). For PIs, I find the same general pattern although minority groups constitute an even smaller share in this population, with for example women making up only 24% of PIs in the full sample and 13% of PIs in the graduate RAs sample. The steady decline of individuals from minority groups at higher rungs of the academic career ladder reflected in these numbers is consistent with the literature on the "leaky pipeline" in academia (Ginther and Kahn, 2004; Spoon et al., 2023).

As pointed out earlier, the name-based indicators for race and ethnicity are imperfectly inferred and might underestimate the true number of racial and ethnic minorities (Rieke et al., 2022). This is especially true for African Americans whose racial identity is often hard to infer based on name alone. However, Table 1 paints a clear picture of the relative distribution of race and ethnicity among RAs and PIs even in the presence of moderate measurement error.

In terms of academic characteristics, among those RAs who go on to earn a PhD in economics, the average RA attends a doctoral program at a university ranked 28th internationally (23rd for the sample of graduate RAs) and 34% (37%) attend a university ranked among the international top 5. This high representation of top universities is not surprising considering that the NBER Working Paper Series represents are strongly selected sample. This point is also reflected in the high share of top 5 affiliations among PIs, 28% and 33%, respectively.

Finally, Table 1 provides information on the average number and gender distribution of PIs per RA in my data, and vice versa. The average RA is assigned to between two and three PIs in the full sample and around one PI in the graduate RAs sample, while the average PI is assigned to approximately five RAs and almost twelve RAs, respectively. Moreover, about 40% (32% in the graduate RAs sample) of RAs are assigned at least one female PI compared to 94% (95%) with at least one male PI. Among PIs, 67% (80% in the graduate RAs sample) have at least one female RA and 83% (96%) have at least one male RA.

3 Descriptive Facts

This section establishes a novel set of descriptive facts about research assistantships for the period 1973-2023. I document a considerable increase in the prevalence of RAs as well as big shifts in the composition of RAs and their PIs, although minority groups remain underrepresented among both groups. Figure 1 captures the rise of research assistants over the last five decades: whereas only about 10% of NBER working papers between 1973-1983

Figure 1: The Rise of Research Assistants

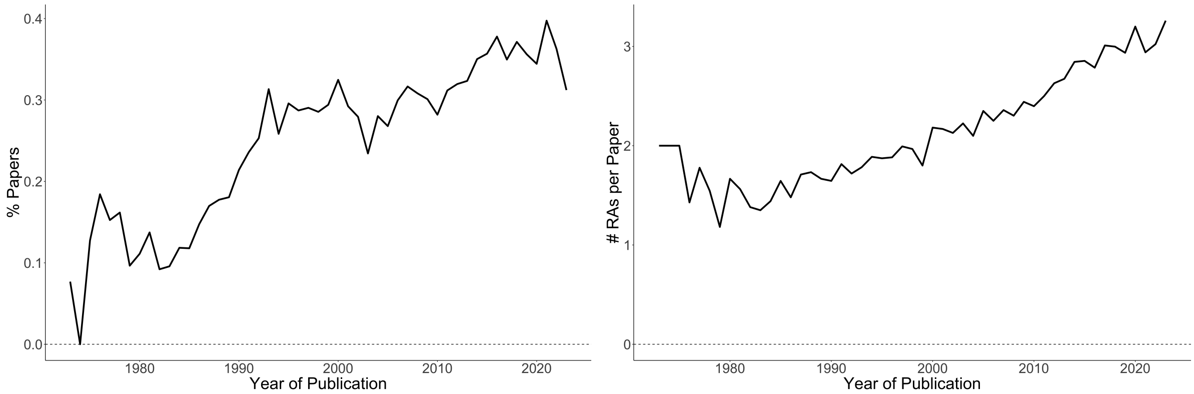

(a) Share Papers with RAs

(b) Average RA Team Size

Notes: Panel (a) shows the evolution of the share of papers mentioning at least one RA in the acknowledgments. Panel (b) shows the evolution of the average number of RAs mentioned in the acknowledgments of a paper conditional on the paper mentioning at least one RA. The underlying sample consists of all papers in the NBER Working Paper Series. The time period covered is 1973-2023.

mentioned research assistants in the acknowledgments, around 35% did so between 2013-2023. The increasing share of papers listing RAs might merely reflect an increase in the share of authors thanking their RAs and not an actual increase in the number of active RAs. However, this concern is alleviated by Figure 1b which shows a parallel increase in the average number of RAs mentioned on papers that list at least one RA, growing from roughly 1.5 to over 3 RAs per paper. Investigating the drivers behind this development is outside the scope of this paper, but likely causes are the growing popularity of empirical work and the recent rise of research teams in the economics profession (Jones, 2021).

Panel (a) and (b) of Figure 2 illustrate how the gender composition of RAs and PIs has changed over the same period. Even though men outnumber women in virtually all years across both RAs and PIs,[5] the share of women has steadily increased over time, going from an average of roughly 25% among RAs in the 1980s to about 35% in the most recent years, and from 0% to over 25% among PIs. Noticeably, the share of female RAs has expanded comparatively little over the same period that the share of female PIs has seen a large and consistent increase. Despite not being conclusive evidence, this trend questions the hypothesis that increasing the share of female faculty would by itself lead to a considerable increase in the number of women in more junior stages in the economics profession.

Analogously to panels (a) and (b), panels (c) and (d) of Figure 2 show the evolution of the racial composition of RAs and PIs. Both RAs and PIs are predominantly White throughout the whole period of observation, but the share of Asians, Blacks, and Hispanics has consistently increased over time for both groups. Unlike Asians, Hispanics and Blacks remain underrepresented and make up only around 8% and 1% of RAs in the 2020s (7%

[5] Female RAs outnumber male RAs in my data for 1973 and 1975. However, it is questionable whether these numbers are representative because I only observe a total of 4 and 12 RAs in those two years.

Figure 2: Composition of RAs and PIs over Time

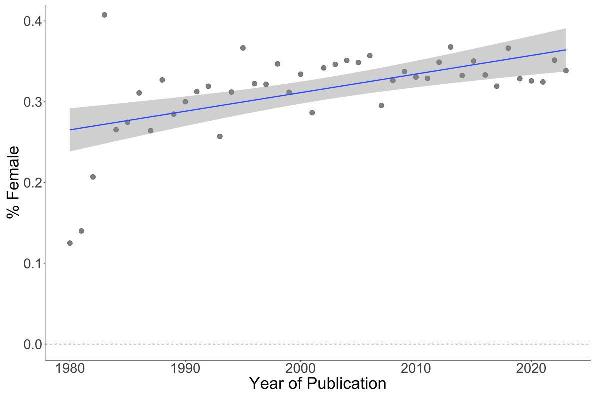

(a) Share Female RAs

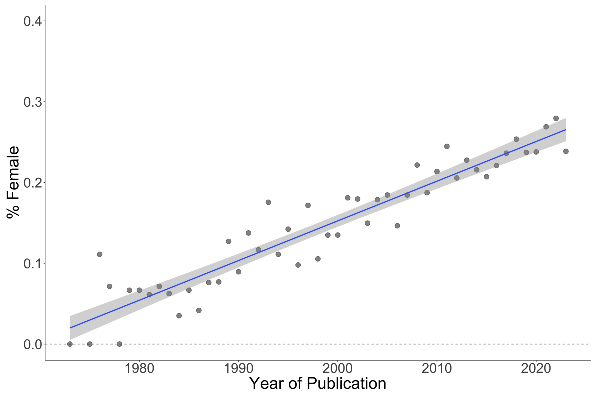

(b) Share Female PIs

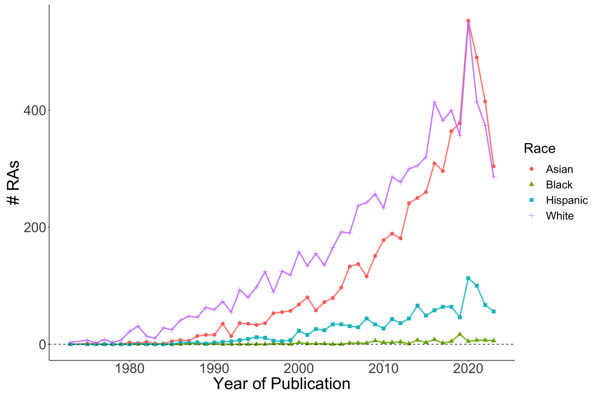

(c) Racial Composition of RAs

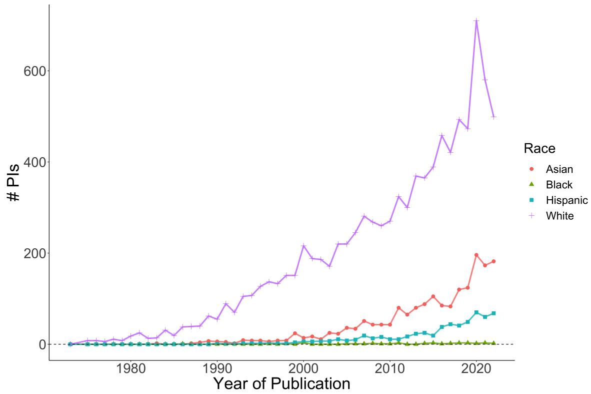

(d) Racial Composition of PIs

Notes: Panels (a) and (b) show the evolution of the share of female RAs and PIs over time, while panels (c) and (d) show the evolution of the racial and ethnic compositions of RAs and PIs over time. The blue line in panels (a) and (b) corresponds to a linear regression with the share of women on the RHS and year of publication on the LHS. The gray shaded area reflects the 95% confidence interval. The underlying sample consists of all papers in the NBER Working Paper Series. In panel (a), I drop the years prior to 1980, which suffer from a very small sample size of RAs (around 10 per year). Appendix Figure A1 replicates panel (a) including the years 1973-1979.

and less than 1% of PIs). Moreover, whereas Asian RAs slightly overtake White RAs in recent years, White PIs still constitute almost 75% of all observations in 2023. Although these results must be read carefully because of the imperfections of name-based inference for race and ethnicity, the overall conclusions from these figures are not threatened by moderate levels of measurement error.

4 Sorting Results

In this section, I examine the role of gender, race, and ethnicity in the formation of RA-PI links and show that there exist strong sorting patterns along those dimensions, even though my data does not allow me to determine the driver behind this sorting. As illustrated in Figure 3, a naive comparison of the share of female RAs by the gender of the PI reveals a clear gap between male and female PIs. The share of female RAs is visibly larger among

Figure 3: Evolution of Female RA Share by PI Gender

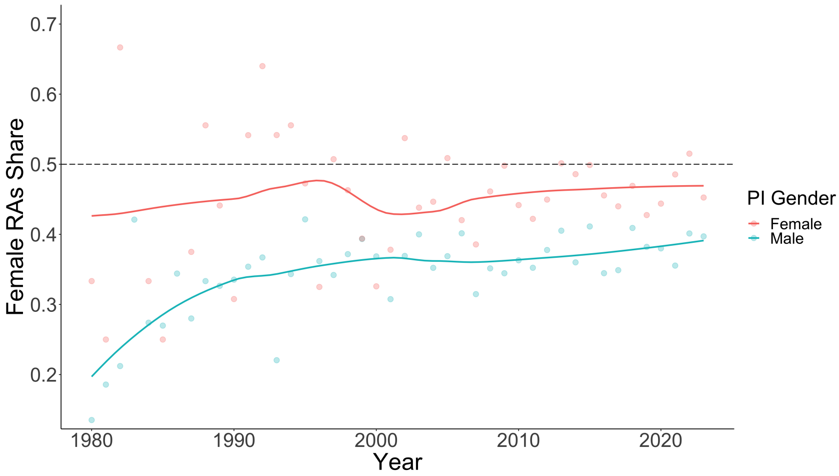

Notes: This figure shows the annual share of female RAs grouped by the gender of their PI. Data points reflecting RAs with female PIs are shown in red, and data points reflecting RAs with male PIs are shown in blue. Each point represents the share of female RAs in a given year. The two solid lines reflect local regressions (LOESS) using a smoothing parameter of 0.65. The underlying data is taken from the NBER Working Papers dataset described in Section 2. The observation period is restricted to 1980-2023 because of empty cell counts in earlier years.

female PIs than among male PIs throughout the whole period of observation, although the gap narrows over time, decreasing from an average difference of 13.9 percentage points for 1980-1989 to an average of 8.7 percentage points for 2014-2023. The compression of the gap is primarily driven by a considerable increase in the female share among RAs working with male PIs, which grows from 27.2% to 37.7%. For female PIs, the share of female RAs also increases although less strongly, growing from 41.1% to 46.4%. The sorting patterns in Figure 3 are suggestive, but they might conceal other factors that could play a role, such as the affiliation, seniority, or field of research of a given PI. The rest of this section presents analyses that explore and subsequently rule out a variety of alternative explanations for the sorting observed above.

Table 2 presents estimation results for regressions that test whether the gender of the PI is predictive of the gender of the RA. If there was no gender-based sorting between RAs and PIs, the share of female RAs should be the same for both male and female PIs, and thus the gender of the PI should have no predictive power for the gender of the RA. Since the RA-PI relationship is non-symmetric, I perform the same analysis one more time, but with the roles of RAs and PIs switched. The level of observation in both cases is unique RA-PI *interactions*, where I define an interaction as a co-occurrence of a given RA and a given PI on the same paper. RA-PI links are identified as described in Section 2. For my main specification, I then estimate regressions of the form:

$$\text{Female RA}_i = \beta \text{Female PI}_i + \gamma X_i + \theta_i + \tau_i + \varepsilon_i, \tag{1}$$

Table 2: Sorting Between RAs and PIs by Gender

Panel A: RA Gender Conditional on PI Gender						
	All Interactions		First Interaction Only		# Interactions with	
	Female RA				Female RAs	Male RAs
	(1)	(2)	(3)	(4)	(5)	(6)
Female PI	0.096***	0.100***	0.092***	0.099***	0.009	0.021
	(0.009)	(0.011)	(0.009)	(0.012)	(0.013)	(0.015)
Controls		✓		✓	✓	✓
Affiliation PI	✓	✓	✓	✓	✓	✓
Year of Publication	✓	✓	✓	✓	✓	✓
Dep. Variable Mean	0.383	0.381	0.386	0.385	1.122	1.136
Observations	28,562	16,207	25,431	14,334	5,512	8,822
Panel B: PI Gender Conditional on RA Gender						
	All Interactions		First Interaction Only		# Interactions with	
	Female PI				Female PIs	Male PIs
	(1)	(2)	(3)	(4)	(5)	(6)
Female RA	0.060***	0.058***	0.058***	0.058***	-0.005	-0.016*
	(0.006)	(0.008)	(0.006)	(0.008)	(0.017)	(0.008)
Controls		✓		✓	✓	✓
Affiliation PI	✓	✓	✓	✓	✓	✓
Year of Publication	✓	✓	✓	✓	✓	✓
Dep. Variable Mean	0.208	0.210	0.210	0.213	1.104	1.128
Observations	28,562	16,763	25,431	14,930	3,187	11,743

Notes: This table shows estimation results for a set of regressions that test the predictive power of RA gender for PI gender and vice versa. All columns include fixed effects for the university affiliation of the PI and the year of publication. Columns (2) and (4) add controls for PI (RA) race and ethnicity as well as the seniority of the PI. Columns (1) and (2) use the full sample of RAs and PIs, columns (3) and (4) only consider the first interaction for a given RA-PI pair, and columns (5) and (6) are aggregated at the PI (RA) level. Standard errors clustered at the year level are shown in parentheses. * p < 0.10, ** p < 0.05, *** p < 0.01.

where i indicates a unique RA-PI interaction. Female RA$_i$ and Female PI$_i$ are dummy variables indicating whether the RA and PI for a given RA-PI interaction i are female. X_i is a vector of control variables and includes a set of race and ethnicity dummies as well as the seniority of the PI, measured as the number of years since PhD completion and grouped into 5-year bins. Finally, θ_i and τ_i are fixed effects for the university affiliation of the PI and the year of publication of the working paper. ε_i is the error term.

The first two columns in Table 2 report results based on Equation 1, without controls in column (1) and with controls in column (2). The coefficients for Female PI in panel A show that interactions involving a female PI are between 9.6 and 10.0 percentage points more likely to involve a female RA compared to interactions involving a male PI. This represents an increase of over 25% compared to a sample average of 38.3% (38.1%) of interactions with female RAs. Similarly, the results in panel B show that conditional on the RA being female, I am 6.0 (5.8) percentage points more likely to observe a female PI compared to when the RA is male, a 28.8% (27.6%) increase relative to the sample average. These results, all of which are statistically significant at the 1%-level, suggest that there exists strong gender-based sorting between RAs and PIs.

As pointed out above, I cannot determine what is driving the sorting because I observe neither the pool of RA applicants nor the pool of RA position openings and the PIs associated with them. I can, however, determine whether this sorting is driven by the extensive margin, e.g. a PI might be less likely to hire RAs of the opposite gender, or by the intensive margin, e.g. a given RA-PI pair might work together for longer if both are of the same gender than if they are of opposite genders.

Columns (3)-(6) report results for my analysis of the contribution of the extensive and intensive margin to the previously identified sorting. Columns (3) and (4) isolate the extensive margin by counting each unique RA-PI combination only once, thereby removing the contribution from repeated interactions between the same RA and PI. For the contribution pf the intensive margin in columns (5) and (6), I use the number of interactions per unique RA-PI pair as outcome variable and split the sample into interactions with female RAs (PIs) and male RAs (PIs). The coefficient for Female PI (Female RA) then reflects the difference between male and female PIs with respect to the average number of interactions with female (column (5)) and male RAs (column (6)).

The estimates from columns (3)-(6) show that virtually all of the sorting is driven by the extensive margin, i.e. RA-PI pairs are more likely to form within the same gender, but once a pair has been formed, there are no differences in the intensity of the collaboration. The coefficients for Female PI and Female RA in columns (3) and (4) remain virtually unchanged compared to columns (1) and (2). This suggests that the sorting is not driven by repeated interactions between the same RA and PI.

My results for the intensive margin support this interpretation. Female RAs seem to interact slightly less often with male PIs than male RAs, but this difference is small, corresponding to 1.4% of the average number of interactions, and is only significant at the 10% level. All estimates in columns (5) and (6) are of similarly small magnitude and imprecisely estimated. Overall, female PIs seem to have slightly more interactions with the same RA than male PIs, and female RAs seem to have slightly fewer interactions with the same PI than male RAs, but these trends hold irrespective of the other party's gender.

One weakness of the present analysis is my coarse measure of interaction intensity, i.e. the number of papers that a given RA and PI appear on together. Since paper publications, even working paper publications, are rather irregular events, this measure might conceal differences at a more granular level. This issue is also reflected in the fact that I observe most RA-PI pairs only once. My results indicate that there exist no large differences in the intensive margin. Future analysis could use a more granular intensity measure, e.g. months of collaboration, to investigate whether this result holds up at a higher level of detail.

One concern regarding the analysis presented above might be imperfections in the assignment algorithm that I use to identify RA-PI pairs. To alleviate this concern, Table A1 in the appendix replicates the analysis in Table 2, but restricts the sample to solo-authored papers which enables me to assign RA-PI links with certainty. I find that the sorting pattern is even stronger for solo-authored papers which suggests that the results shown above

Table 3: Sorting Between RAs and PIs by Race and Ethnicity

Panel A: RA Race/Ethnicity conditional on PI Race/Ethnicity				
	Asian RA (1)	Black RA (2)	Hispanic RA (3)	White RA (4)
Asian PI	0.247***	-0.006***	-0.029***	-0.214***
	(0.016)	(0.002)	(0.008)	(0.014)
Black PI	-0.089	0.049*	0.088	-0.050
	(0.070)	(0.029)	(0.061)	(0.077)
Hispanic PI	-0.034	-0.003	0.166***	-0.121***
	(0.023)	(0.004)	(0.024)	(0.026)
Controls	✓	✓	✓	✓
Year of Publication	✓	✓	✓	✓
Affiliation PI	✓	✓	✓	✓
Dep. Variable Mean	0.397	0.006	0.085	0.513
Observations	12,020	12,020	12,020	12,020

Panel B: PI Race/Ethnicity conditional on RA Race/Ethnicity				
	Asian PI (1)	Black PI (2)	Hispanic PI (3)	White PI (4)
Asian RA	0.107***	-0.0005	-0.001	-0.103***
	(0.008)	(0.001)	(0.004)	(0.009)
Black RA	-0.045	0.032	-0.004	0.014
	(0.034)	(0.023)	(0.027)	(0.046)
Hispanic RA	-0.012	0.004	0.087***	-0.073***
	(0.011)	(0.003)	(0.011)	(0.015)
Controls	✓	✓	✓	✓
Year of Publication	✓	✓	✓	✓
Affiliation PI	✓	✓	✓	✓
Dep. Variable Mean	0.142	0.004	0.046	0.811
Observations	10,289	10,289	10,289	10,289

Notes: This table shows estimation results for a set of regressions that test the predictive power of RA race/ethnicity for PI race/ethnicity and vice versa. All columns include fixed effects for the university affiliation of the PI and the year of publication. Columns (2) and (4) add controls for PI (RA) gender as well as the seniority of the PI. Columns (1) and (2) use the full sample of RAs and PIs, columns (3) and (4) only consider the first interaction for a given RA-PI pair, and columns (5) and (6) are aggregated at the PI (RA) level. Standard errors clustered at the year level are shown in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

might suffer from attenuation bias introduced through random measurement error in the PI assignment. This in turn would imply that the estimates reported in Table 2 are a lower bound for the true extent of gender sorting between RAs and PIs.

Table 3 replicates the analysis from above for race and ethnicity and also shows strong sorting patterns. For example, my estimates suggest that conditional on the PI being Asian, the RA is 24.7 percentage points more likely to be Asian as well compared to if the PI was White. This represents an increase of 62% over the sample mean. The sorting pattern holds for Black, Hispanic, as well as White RAs and PIs. Except for the estimates for Black PIs, which are based on a small number of observations, all coefficients are precisely estimated.

The sorting results presented in this section are robust to a variety of checks. For example, it might be the case that both RAs and PIs of a given gender tend to cluster in certain fields which could be driving the observed sorting. Figure 4 tests this hypothesis by estimating the regressions from Panel A of Table 2 for RA-PI interactions occurring within the same field of research. I identify field at the interaction level using the topics assigned to each NBER working paper. Figure 4 shows that gender sorting exists across all fields

Figure 4: Gender Sorting by Field

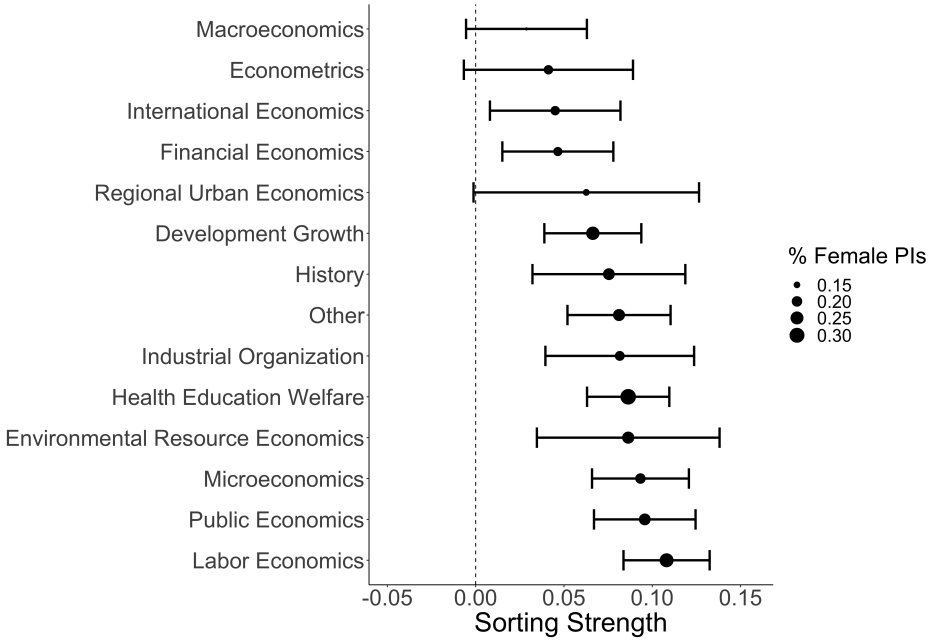

Notes: This figure shows the strength of sorting by field. Regression specifications are analogous to panel A of Table 2. Information on fields is taken from the "Topics" attribute associated with each NBER working paper. Black dots represent point estimates, and the size of each point reflects the share of female PIs within a given field. Bars reflect 95% confidence intervals for Standard errors clustered at the year level.

and that the strength of this sorting varies little between fields. Moreover, there is no clear correlation between the strength of sorting and the share of female PIs in a given field.

Appendix A.3 provides additional robustness checks that rule out other explanations for the sorting along gender, race, and ethnicity. Table A2 through Table A5 restrict the sample to PIs affiliated with US universities and universities whose economics department is ranked in the international top 20 according to RePEc to account for regional differences as well as differences between more and less prestigious universities. Table A6 through Table A9 restrict the sample to RA-PI interactions observed before 2000 or after 2010, respectively, to account for changes in sorting behavior over time. Figure A2 splits the sample into deciles by the share of undergraduate degrees in economics awarded to women at the university that the PI is affiliated with.[6] Finally, Table A10 and Table A11 replicate my sorting analysis for the subsample of graduate RAs, i.e. RAs that are classified as doctoral students at the time of their activity as RA. All of these robustness checks find similar sorting patterns along gender, race, and ethnicity that confirm the results from my main analysis.

[6]Data on the number of men and women majoring in economics at a given university in a given year is taken from IPEDS.

5 RA Experience and Career Outcomes

Table 4: Outcomes among PhD Graduates with and without RA Experience

	# Publications		Any Publication		# Publications (if any)	
	5Y Post-PhD (1)	10Y Post-PhD (2)	5Y Post-PhD (3)	10Y Post-PhD (4)	5Y Post-PhD (5)	10Y Post-PhD (6)
RA Experience	0.841***	1.56***	0.248***	0.260***	0.548***	1.00***
	(0.059)	(0.144)	(0.012)	(0.012)	(0.108)	(0.215)
Controls	✓	✓	✓	✓	✓	✓
Year of PhD Completion	✓	✓	✓	✓	✓	✓
PhD Institution	✓	✓	✓	✓	✓	✓
Dep. Variable Mean	0.334	0.634	0.116	0.139	1.81	3.46
Observations	34,292	28,680	34,292	28,680	6,345	5,260

Notes: This table shows estimation results for regressions that examine the correlation between academic career outcomes and RA experience. Columns (1) and (2) use the number of NBER publications within five and 10 years after PhD completion as outcome variable. Columns (3) and (4) use a binary indicator for having any NBER publications in the same time frame as outcome variable. Finally, columns (5) and (6) use the number of publications as outcome variable, but restrict the sample to only those PhD graduates with at least one NBER publication at any point in time. All columns include controls for gender and fixed effects for year and university of PhD completion. Standard errors clustered at the year level are shown in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

This section investigates the relationship between RA experience and career outcomes, specifically academic career outcomes. Past research and anecdotal evidence suggest that RA positions are particularly relevant for individuals pursuing careers in academia (Hansen, 1991; Krueger, 1991; Stock and Siegfried, 2014). The research skills accumulated during one's time as RA appear to play a major role in this context (Hansen, 1991; Krueger, 1991). In line with prior work, I present descriptive evidence that PhD graduates with RA experience have more NBER publications in the short, mid, and long term than their peers without RA experience. Moreover, I show that among PhD graduates with RA experience, those with a PI of the same gender exhibit better outcomes than those with a PI of the opposite gender.

Table 4 shows results for regressions that examine the hypothesis that PhD graduates with RA experience have better academic career outcomes than their peers without RA experience. For my main specification in columns (1) and (2), I regress the number of NBER publications on a dummy variable for RA experience, defined as being present in my NBER sample of RAs, while controlling for gender, year of PhD completion, and PhD-granting university. Column (1) considers the cumulative number of NBER publications within five years after PhD completion whereas column (2) considers all publications within ten years after PhD completion. For each regression, the samples are chosen to only include PhD graduates who I could have observed for at least five or ten years after PhD completion.

The coefficients for RA Experience suggest that a PhD graduate with RA experience will have 0.84 (1.56) more NBER publications within five (ten) years after finishing their degree than a PhD graduate without RA experience. For comparison, this is an increase of around 250% relative to the sample mean of 0.33 (0.63) among all PhD graduates. The coefficients in both columns are precisely estimated and highly significant.

The present analysis does not account for the potential of positive selection of PhD students into working as an RA. It appears likely that PhD students more interested and more proficient in academic research are going to work as RAs at higher rates than their peers who are e.g. planning to pursue a career in the public or private sector. My data does not allow me to accurately capture such selection or to control for it. I can, however, characterize the difference in outcomes between RAs and non-RAs further by separating the results into a part reflecting the extensive margin, i.e. do you publish at least one paper, and a part reflecting the intensive margin, i.e. given that you publish at least one paper, how many papers do you publish. While not answering the question about the fundamental driver behind the difference in outcomes, this analysis illuminates the proximate causes, namely whether former RAs are more productive or simply more likely to stay in academia.

For the analysis of the extensive margin, columns (3) and (4) replicate the regressions from columns (1) and (2), but instead of the number of publications, they use a binary indicator as outcome variable that is 1 if a PhD graduate has at least one NBER publication within five (ten) years after completing their degree and 0 otherwise. The results show that PhD graduates with RA experience are 24.8 (26.0) percentage points more likely to publish at least one NBER paper within five (ten) years after earning their PhD compared to their peers without RA experience, a 214% (187%) increase over the sample mean.

For the analysis of the intensive margin, columns (5) and (6) re-estimate the regressions from columns (1) and (2) using only PhD graduates who published at least one paper in the NBER Working Paper Series, although not necessarily within five (ten) years after PhD completion. Former RAs again outperform their peers without RA experience, publishing an average of 0.55 (1.00) more NBER papers within five (ten) years after finishing their PhD, a 30.4% (28.9%) increase over the sample mean.

Taken together, these results indicate that PhD graduates with RA experience are both more likely to pursue a career in academia and more successful at that pursuit. The differences are larger on the extensive margin than on the intensive margin, but both are sizeable and highly significant. Under the assumption that the results are not purely driven by selection, either because RAs are primarily composed of PhD students interested in an academic career or PhD students more skilled at research, this would suggest that RA experience provides PhD graduates with some form of human capital, e.g. in the form of research skills, or social capital, e.g. a stronger network of potential coauthors, that translates into better academic career outcomes. Determining the impact of selection on these results appears to be a worthwhile goal for future research.

Table 5: Former RAs, NBER Publications, and RA-PI Gender Alignment

	# Publications		Any Publications		# Publications (if any)	
	5Y Post-PhD (1)	10Y Post-PhD (2)	5Y Post-PhD (3)	10Y Post-PhD (4)	5Y Post-PhD (5)	10Y Post-PhD (6)
Same Gender	0.608***	1.04**	0.099**	0.114**	0.574*	0.904
	(0.197)	(0.429)	(0.041)	(0.052)	(0.333)	(0.749)
Controls	✓	✓	✓	✓	✓	✓
Year of PhD Completion	✓	✓	✓	✓	✓	✓
PhD Institution	✓	✓	✓	✓	✓	✓
Dep. Variable Mean	1.36	2.54	0.433	0.477	2.65	4.92
Observations	1,525	1,095	1,525	1,095	783	566

Notes: This table shows estimation results for regressions that examine the correlation between RA-PI gender alignment and career outcomes among RAs. Columns (1) and (2) use the number of NBER publications within five and 10 years after PhD completion as outcome variable. Columns (3) and (4) use a binary indicator for having any NBER publications in the same time frame as outcome variable. Finally, columns (5) and (6) use the number of publications as outcome variable, but restrict the sample to only those PhD graduates with at least one NBER publication at any point in time. All columns include controls for the gender of the RA and the PI as well as fixed effects for year and university of PhD completion. Robust standard errors are shown in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

The findings above suggest that RA experience might play an important role in the career success of PhD graduates. Together with my results in Section 4, which indicate that sorting between RAs and PIs might obstruct access to RA positions among minority groups, this raises the question of whether attributes like gender also factor into the impact of RA experience on subsequent career outcomes. The rest of this section therefore examines how gender and gender alignment between RA and PI interact with the apparent effect of RA experience on academic career outcomes. To answer this question, I estimate regressions of a similar form as those in Table 4:

$$\text{NBER Publications}_i = \beta_1 \text{Female RA}_i + \beta_2 \text{Female PI}_i + \beta_3 \text{Same Gender}_i + \delta_i + \tau_i + \varepsilon_i. \quad (2)$$

The dependent variable NBER Publications$_i$ captures the academic success of a former RA i measured either as a binary indicator that is 1 if i has any NBER publications within 5 (10) years after their PhD and 0 otherwise, or measured as the cumulative number of NBER publications within that same time frame. Female RA$_i$ and Female PI$_i$ are binary indicators for the gender of RA i and the gender of their PI(s). Note that because of how I construct the sample of graduate RAs (see Section 2), all PIs of a given RA i will be either male or female. δ_i and τ_i are fixed effects for the university that RA i earned their PhD from and the year of PhD completion. ε_i is the error term.

Table 5 reports estimation results based on the regression equation in Equation 2 where the sample consists of all graduate RAs identified in my NBER data, i.e. RAs who I determine to have been PhD students during their time as RA. Columns (1) and (2) show results for my main specification with the number of NBER publications within five and ten years after PhD completion as outcome variable. The estimates indicate that an RA with a PI of the same gender publishes an average of 0.61 (1.04) more NBER papers within five (ten) years of completing their PhD than their peers with a PI of the opposite gender.

Similar to Table 4, I again split this result into a part reflecting the extensive margin and a part reflecting the intensive margin. The results for the extensive margin in columns (3) and (4) show that an RA with a PI of the same gender is 9.9 (11.4) percentage points more likely to have at least one NBER publication within five (ten) years after their PhD than their RA peers who had a PI of the opposite gender. This suggests that a PI of the same gender might function as a role model which in turn motivates PhD graduates to pursue a career in academia.

Columns (5) and (6) report similar, albeit less precise results for the intensive margin, i.e. when considering only former RAs who publish at least one NBER paper. Within five (ten) years after PhD completion, an RA with a PI of the same gender will have published an average of 0.57 (0.90) more papers than their peer with a PI of the opposite gender. The estimate for the five year window is significant at the 10% level whereas the estimate for the years window is not significant at the conventional levels. The decline in precision for the intensive margin is not least due to the reduced sample size when restricting the sample to graduate RAs who published at least one NBER paper. Overall, the coefficients confirm, however, the difference in outcomes between RAs who worked with a PI of the same gender and those who worked with a PI of the opposite gender.

The regressions in Table 5 use Same Gender as a common indicator for both female and male RAs with PIs of the same gender. I choose this specification over an alternative one that allows for differential gender alignment effects for female and male RAs because of the small sample size and the coarse outcome measure. To investigate the question of asymmetric gender alignment effects and to alleviate concerns about one group of RAs driving these results, Table A12 and Table A13 in the appendix report results when estimating the regression on split samples containing male and female RAs only. Overall, I find similar patterns as in Table 5. The effects of gender alignment appear to be somewhat larger for female than for male RAs, especially on the extensive margin. In line with prior research (Avilova and Goldin, 2018; Becker, Rouse and Chen, 2016; Boustan and Langan, 2019; Carrell, Page and West, 2010), this could suggest that having a female mentor plays an important role for female PhD graduates in particular when deciding whether to stay in academia or to go into the private or public sector. For most specifications, however, there are no statistically distinguishable differences between male and female RAs.

Figure 5 extends the analysis in Table 5 by separating NBER publications by graduate RAs into those that were coauthored with former PIs and those with no former PIs involved. Panel (a) and (b) plot the average number of coauthored and non-coauthored NBER publications, residualized by year of PhD completion and degree-granting university, for a given year t relative to the year of PhD completion. Each line corresponds to one RA-PI gender combination with red (blue) lines representing female (male) RAs and solid (dotted) lines representing female (male) PIs.

In line with the results from Table 5, RA-PI pairs of the same gender have on average

Figure 5: NBER Publications by RA-PI Gender Over Time

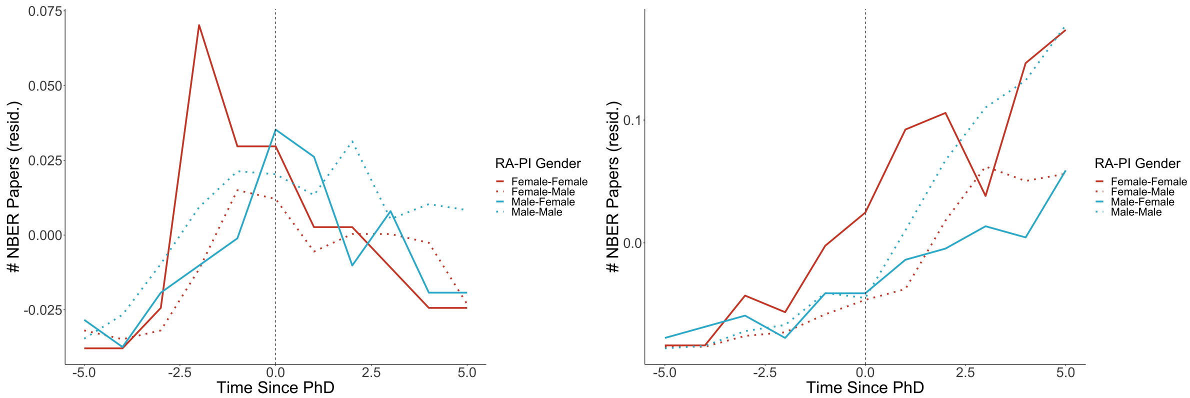

(a) Coauthored with PI (b) Not Coauthored with PI

Notes: Panel (a) plots the average residualized number of NBER papers that RAs coauthored with their former PI in a given year relative to the RA's year of PhD completion. Similarly, panel (b) plots the number of NBER papers not coauthored with a former PI. RAs are grouped by their own gender and the gender of their PI. Red (blue) lines represent female (male) RAs and solid (dotted) lines represent female (male) PIs. The number of NBER publications in a given year is residualized by year and university of PhD completion. The underlying sample only includes graduate RAs who I could observe for at least five years after PhD completion.

more publications than RA-PI pairs of opposite genders. Figure 5 shows, however, that this difference is primarily driven by publications that were not coauthored with former PIs. Although panel (a) indicates that female RAs with female PIs have more coauthored papers than female RAs with male PIs in the years -2 through 0 relative to the year of PhD completion, this difference is small in absolute terms. After PhD completion, both same-gender and opposite-gender RA-PI pairs coauthor at similar low rates. For papers not coauthored with former PIs, the gap between same-gender and opposite-gender RA-PI pairs opens up around the time of PhD completion, somewhat earlier for female RAs than for male RAs, and persists from there on. These results suggest that the previously observed gender alignment effect is driven by the achievements of the RAs themselves and not due to PIs artificially inflating the academic success of former RAs.

6 Discussion

The growing importance of research teams (Jones, 2021) and empirical research within the economics profession is well-known. I have presented new data that documents a parallel rise in the importance of RAs. As RA positions become a more integral part of an economist's career path, it is crucial to gain a better understanding of this form of academic apprenticeship. I contribute to this goal by offering large-scale evidence on the prevalence and composition of RAs and by presenting suggestive evidence concerning the importance of RA experience for later career outcomes. My results reveal potential conflicts between promoting the representation of minority groups in the economics profession on the one hand and promoting their success on the other. Finally, I highlight new questions that demand further research.

Despite the share of minority groups consistently trending upward, women, racial, and ethnic minorities remain underrepresented among RAs.[7] This finding matches prior research on female representation at the undergraduate (Bayer and Wilcox, 2019; Goldin, 2015), graduate (Hale and Regev, 2014; Lundberg and Stearns, 2019), and post-graduate stage (Chari and Goldsmith-Pinkham, 2017; Ceci et al., 2014; Ginther and Kahn, 2004; Meade, Starr and Bansak, 2021). While the reasons for this underrepresentation are most likely multi-faceted, my findings suggest two channels that could have an impact in this context. First, although my observational data does not allow me to infer whether RAs, PIs, or both sides are causing the sorting documented in Section 4, my results suggest that improving the representation of minority groups among faculty would also improve representation among RAs. That is, regardless of the cause or direction underlying the observed sorting patterns, a higher share of e.g. female PIs should entail a higher share of female RAs.

The extent of the improvement that could be achieved through this channel appears to be limited, however. A naive back-of-the-envelope calculation based on my results in Table 2 suggests that even given gender balance among PIs, the share of female RAs in 2023 would only increase from about 40.5% to roughly 42.9%. Thus, in the absence of additional changes in e.g. the sorting preferences of RAs and PIs or the number of female students choosing to major in economics, gender parity among faculty will not suffice to achieve gender parity among RAs.

The second channel one might turn to is intervening in the sorting itself. Since my results cannot make any claims regarding the causes behind the sorting, it remains unclear what the implementation of such an intervention should look like. If RAs are driving the sorting, e.g. female applicants might be less likely to apply to male PIs, a more diverse applicant pool might first require a more diverse pool of faculty as suggested by Hale and Regev (2014), which brings us back to the first channel discussed above. If PIs are driving the sorting, e.g. male PIs might have a preference for male RAs, breaking up the sorting would be theoretically possible. However, the legal environment universities operate in and potential pushback from faculty against interference with their hiring decisions call the practical feasibility of such an intervention into question.

Moreover, even if breaking up the sorting proved feasible, my results indicate that this could have undesirable consequences. I find that female RAs working with female PIs exhibit better academic career outcomes than their peers working with male PIs and vice versa. Thus, breaking up the sorting between RAs and PIs might come at the expense of worse career outcomes for RAs who are matched with PIs of the opposite gender. This suggests that there could be a conflict between promoting the representation of minority groups in the economics profession on the one hand and promoting their success on the other hand.

[7]The only exception to this pattern are Asian RAs whose numbers have recently caught up with those of White RAs.

Rather than claiming that my findings are conclusive evidence in one direction or another, I present them to call attention to aspects that have so far remained understudied in the discussion about women and minority groups in academia. Improving our understanding of the fundamental causes underlying the observational findings documented in this paper presents itself as a promising and important avenue for future research. Evidence on questions like which side of the matching process is driving the sorting between RAs and PIs could enable universities to design policies that reconcile the promotion of both representation and success of minority groups in the economics profession.

7 Conclusion

In this paper, I have presented new descriptive evidence on the prevalence and composition of research assistants (RAs) in the economics profession throughout the last 50 years. Moreover, I have shown that there exist strong sorting patterns between RAs and their supervisors (PIs) along gender, race, and ethnicity. I discussed and subsequently ruled out several alternative explanations for these sorting patterns, such as the clustering of certain groups by field of research or by the geographic location of a university. Finally, I have provided suggestive evidence that an RA working with a PI of the same gender experiences better academic career outcomes than RAs working with PIs of the opposite gender, as reflected in higher rates of publication in the NBER Working Paper Series.

My results complement earlier research on the evolution and contemporary state of the representation and success of minority groups in academia and the economics profession in particular (Bayer and Rouse, 2016; Ceci et al., 2014; Lundberg and Stearns, 2019; Spoon et al., 2023). Despite recent progress, women as well as racial and ethnic minorities remain underrepresented among RAs. Considering the crucial role of RA positions as a stepping stone for students at various levels to gain research skills, build professional networks, and get exposed to the academic labor market, this finding adds to concerns that minority groups have less access and face higher barriers to careers in the economics profession. My findings point to sorting among RAs and PIs as one potential factor contributing to this problem. At the same time, I find that RAs working with PIs of the same gender experience better academic career outcomes, which in turn suggests that breaking up the sorting between RAs and PIs might have undesirable consequences.

My results raise several questions about the causes underlying the observational findings documented in this paper. A solid understanding of these causes is an essential requirement for the formulation of adequate policy responses to the persistent issue of bias against minority groups in academia. Further investigation of the questions raised in this paper therefore presents itself as a promising and important avenue for future research.

References

Avilova, Tatyana and Claudia Goldin. 2018. What can UWE do for economics? In *AEA papers and proceedings*. Vol. 108 American Economic Association 2014 Broadway, Suite 305, Nashville, TN 37203 pp. 186–190.

Bayer, Amanda and Cecilia Elena Rouse. 2016. "Diversity in the Economics Profession: A New Attack on an Old Problem." *Journal of Economic Perspectives* 30(4):221–242.

Bayer, Amanda and David W. Wilcox. 2019. "The unequal distribution of economic education: A report on the race, ethnicity, and gender of economics majors at U.S. colleges and universities." *The Journal of Economic Education* 50(3):299–320.

Becker, Charles M., Cecilia Elena Rouse and Mingyu Chen. 2016. "Can a summer make a difference? The impact of the American Economic Association Summer Program on minority student outcomes." *Economics of Education Review* 53:46–71.

Bertrand, Marianne and Sendhil Mullainathan. 2004. "Are Emily and Greg More Employable Than Lakisha and Jamal? A Field Experiment on Labor Market Discrimination." *American Economic Review* 94(4):991–1013.

Boustan, Leah and Andrew Langan. 2019. "Variation in Women's Success across PhD Programs in Economics." *Journal of Economic Perspectives* 33(1):23–42.

Bryan, Kevin A. 2019. "Young "Stars" in Economics: What They Do and Where They Go." *Economic Inquiry* 57(3):1392–1407.

Carrell, Scott E., Marianne E. Page and James E. West. 2010. "Sex and Science: How Professor Gender Perpetuates the Gender Gap." *Quarterly Journal of Economics* 125(3):1101–1144.

Ceci, Stephen J., Donna K. Ginther, Shulamit Kahn and Wendy M. Williams. 2014. "Women in Academic Science: A Changing Landscape." *Psychological Science in the Public Interest* 15(3):75–141.

Chari, Anusha and Paul Goldsmith-Pinkham. 2017. "Gender Representation in Economics Across Topics and Time - Evidence from the NBER Summer Institute.pdf.".

Colander, David and Arjo Klamer. 1987. "The Making of an Economist." *Journal of Economic Perspectives* 1(2):95–111.

Davies, Benjamin. 2022. "Gender sorting among economists: Evidence from the NBER." *Economics Letters* 217:110640.

Dupas, Pascaline, Alicia Sasser Modestino, Muriel Niederle, Justin Wolfers and The Seminar Dynamics Collective. 2021. Gender and the Dynamics of Economics Seminars. Technical Report w28494 National Bureau of Economic Research Cambridge, MA: .

Economist, The. 2020. "What it takes to become an academic economist is changing.pdf." *The Economist* .

Fortin, Nicole, Thomas Lemieux and Marit Rehavi. 2021. "Gender Differences in Fields of Specialization and Placement Outcomes among PhDs in Economics." *AEA Papers and Proceedings* 111:74–79.

Galiani, Sebastián. and Ugo Panizza. 2020. *Publishing and measuring success in economics.* London, UK: CEPR Press. OCLC: 1374831110.

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):193–214.

Goldin, Claudia. 2015. "Gender and the Undergraduate Economics Major: Notes on the Undergraduate Economics Major at a Highly Selective Liberal Arts College.".

Hale, Galina and Tali Regev. 2014. "Gender ratios at top PhD programs in economics." *Economics of Education Review* 41:55–70.

Hansen, W Lee. 1991. "The education and training of economics doctorates: Major findings of the executive secretary of the American Economic Association's commission on graduate education in economics." *Journal of Economic Literature* 29(3):1054–1087. Publisher: JSTOR.

Hilmer, Christiana and Michael Hilmer. 2007. "Women Helping Women, Men Helping Women? Same-Gender Mentoring, Initial Job Placements, and Early Career Publishing Success for Economics PhDs." *American Economic Review* 97(2):422–426.

Jones, Benjamin F. 2021. "The Rise of Research Teams: Benefits and Costs in Economics." *Journal of Economic Perspectives* 35(2):191–216.

Krueger, Anne O. 1991. "Report of the Commission on Graduate Education in Economics." *Journal of Economic Literature* 29(3):1035–1053. Publisher: American Economic Association.

Lundberg, Shelly and Jenna Stearns. 2019. "Women in Economics: Stalled Progress." *Journal of Economic Perspectives* 33(1):3–22.

Marschke, Gerald, Allison Nunez, Bruce A. Weinberg and Huifeng Yu. 2018. "Last Place? The Intersection of Ethnicity, Gender, and Race in Biomedical Authorship." *AEA Papers and Proceedings* 108:222–227.

Meade, Ellen E., Martha Starr and Cynthia Bansak. 2021. "Changes in Women's Representation in Economics: New Data from the AEA Papers and Proceedings." *FEDS Notes* 2021(2961).

Neumark, David and Rosella Gardecki. 1998. "Women helping women?" *The Journal of Human Resources* 33(1):220–246. Place: Madison Publisher: University of Wisconsin Press.

Rieke, Aaron, Vincent Southerland, Dan Svirsky and Mingwei Hsu. 2022. Imperfect Inferences: A Practical Assessment. In *2022 ACM Conference on Fairness, Accountability, and Transparency.* Seoul Republic of Korea: ACM pp. 767–777.

Sarsons, Heather, Klarita Gërxhani, Ernesto Reuben and Arthur Schram. 2021. "Gender Differences in Recognition for Group Work." *Journal of Political Economy* 129(1):101–147.

Spoon, Katie, Nicholas LaBerge, K. Hunter Wapman, Sam Zhang, Allison C. Morgan, Mirta Galesic, Bailey K. Fosdick, Daniel B. Larremore and Aaron Clauset. 2023. "Gender and retention patterns among U.S. faculty." *Science Advances* 9(42):eadi2205.

Stock, Wendy A. and John J. Siegfried. 2014. "Fifteen Years of Research on Graduate Education in Economics: What Have we Learned?" *The Journal of Economic Education* 45(4):287–303.

Stock, Wendy A. and John J. Siegfried. 2015. "The Undergraduate Origins of PhD Economists Revisited." *The Journal of Economic Education* 46(2):150–165.

Stock, Wendy A, John J Siegfried and T. Aldrich Finegan. 2011. "Completion Rates and Time-to-Degree in Economics PhD Programs." *American Economic Review* 101(3):176–187.

A Appendix

A.1 Extracting RAs from Acknowledgments

I classify every person as an RA that is mentioned in the acknowledgments of a paper as a *research assistant* or as having provided *research assistance*. I exclude all other kinds of being mentioned in the acknowledgments including people being thanked for their comments, discussions, or suggestions, administrative assistants, and data providers. This approach is chosen to exclusively extract people in "academic apprenticeship positions" as precisely as possible. Note, however, that the dataset of RAs compiled this way will contain undergraduate students, pre-doctoral and postgraduate fellows, as well as graduate students. Using data on PhD graduates, I am able to distinguish graduate student RAs from predocs and undergraduate RAs with a reasonable degree of accuracy.

A.2 Determining RA-PI Relationships

In order to assign each RA at least one PI, I use the following algorithm:

- **Solo-authored papers:** RAs in solo-authored papers can always be assumed to be RAs for the respective solo author[8]
- **Coauthored papers, all authors with same attributes**: In this case, it might not be possible to perfectly determine RA-PI relationships in a one-to-one correspondence. However, I am interested in RA-PI relationships with respect to personal attributes such as gender and race. Therefore, if all authors have the same attribute, identification of one-to-one correspondences is not required since the analysis is invariant to which author I assign as PI.
- **Coauthored papers, authors with different attributes:** This case is the only case that poses a problem, since it can make a difference which author(s) I classify as PI(s). To overcome this issue, I make a few basic and (hopefully) uncontroversial assumptions and subsequently assign PIs and RAs using a simple algorithm. The assumptions and algorithm are outlined below

Assumption #1 Every RA has at least one PI

Assumption #2 For every paper an RA is acknowledged on, at least one of the paper's authors is a/the PI of that RA

Assumption #3 The more papers a given RA and a given author appear on together, the more likely it is that that author is a/the PI of that RA

Algorithm This algorithm assigns a PI for each RA such that for any paper at least one author is classified as PI of that RA:

- **Step #1** For each RA, calculate the number of times that that RA and a given author appear together on a paper
- **Step #2** For each RA, calculate the total number of papers that that RA appears on
- **Step #3** For each paper for each RA on that paper, assign the author(s) with the highest ratio of $\frac{\#\text{ Papers with both RA and PI}}{\#\text{ Papers with RA}}$ as a/the PI of that RA.
- **Step #4** Where possible, I break ties by defaulting to the PI(s) affiliated with the same university that the RA received their PhD from.

[8] There might be exceptions, e.g. in the case of PhD students getting assistance from RAs of their supervisor. However, I expect such cases to be (i) rather rare considering the sample of papers and (ii) potentially somewhat informative as those RAs might have been hired in coordination with the respective PhD students.

A.3 Additional Results

Figures

Figure A1: Gender Composition of RAs over Time - With Outliers

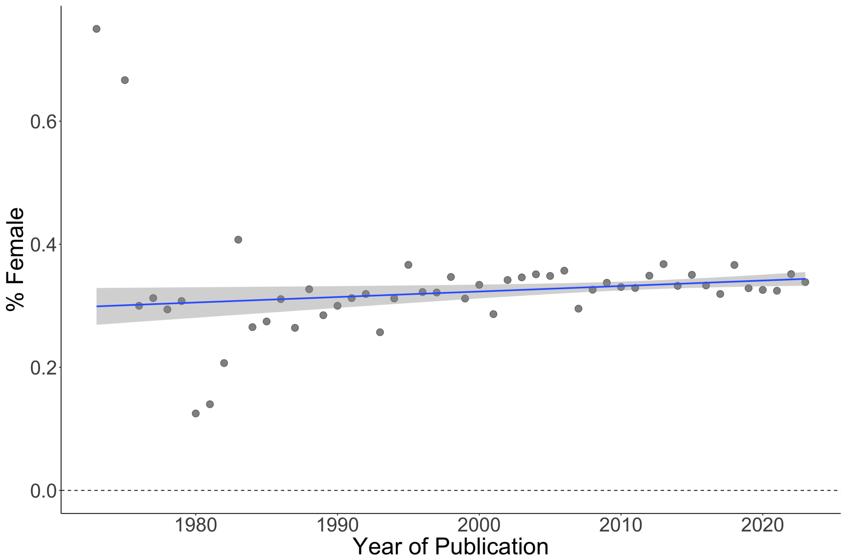

Notes: This figure shows the evolution of the share of female RAs by year of publication including the outlier years 1973-1979. The blue line corresponds to a linear regression with the share of women on the RHS and year of publication on the LHS. The gray shaded area reflects the 95% confidence interval. The underlying sample consists of all papers in the NBER Working Paper Series. The time period covered is 1973-2023.

Figure A2: Gender Sorting Share of Female Economics Majors

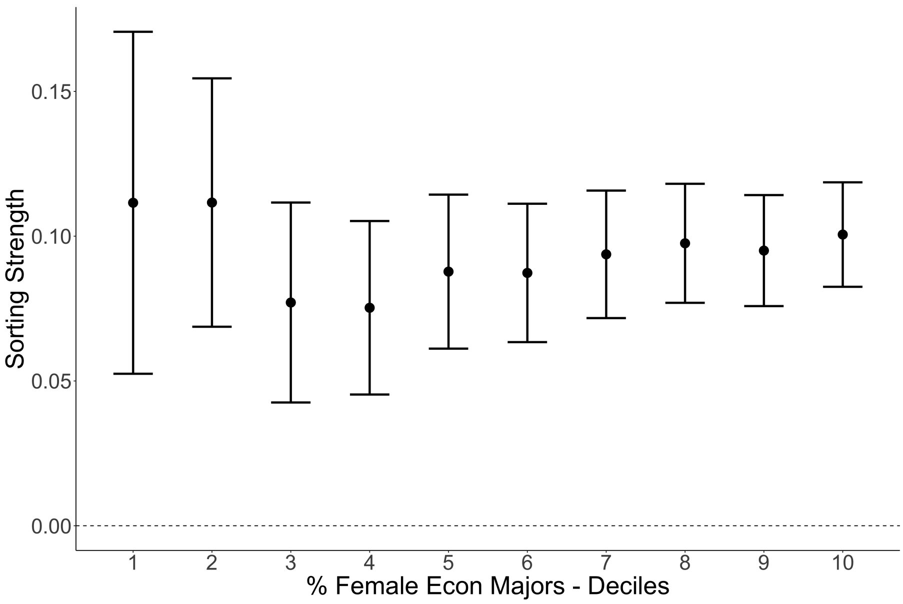

Notes: This figure shows the strength of sorting by the share of women majoring in economics. Regression specifications are analogous to panel A of Table 2. Information on the share of women majoring in economics is taken from IPEDS and matched to RA-PI pairs based on the university affiliation of the PI and the year of publication of the associated NBER working paper. Black dots represent point estimates. Bars reflect 95% confidence intervals for standard errors clustered at the year level.

Tables

Table A1: Sorting by Gender - Solo-Authored Papers

Panel A: RA Gender Conditional on PI Gender						
	All Interactions		First Interaction Only		# Interactions with	
	Female RA				Female RAs	Male RAs
	(1)	(2)	(3)	(4)	(5)	(6)
Female PI	0.179***	0.207***	0.172***	0.185***	0.176**	0.127
	(0.034)	(0.051)	(0.034)	(0.052)	(0.084)	(0.103)
Controls		✓		✓	✓	✓
Affiliation PI	✓	✓	✓	✓	✓	✓
Year of Publication	✓	✓	✓	✓	✓	✓
Dep. Variable Mean	0.360	0.369	0.353	0.363	1.199	1.169
Observations	2,598	1,455	2,234	1,233	448	785
Panel B: PI Gender Conditional on RA Gender						
	All Interactions		First Interaction Only		# Interactions with	
	Female PI				Female PIs	Male PIs
	(1)	(2)	(3)	(4)	(5)	(6)
Female RA	0.088***	0.073***	0.087***	0.079***	-0.102	0.012
	(0.018)	(0.022)	(0.019)	(0.023)	(0.100)	(0.031)
Controls		✓		✓	✓	✓
Affiliation PI	✓	✓	✓	✓	✓	✓
Year of Publication	✓	✓	✓	✓	✓	✓
Dep. Variable Mean	0.165	0.154	0.161	0.152	1.174	1.157
Observations	2,598	1,485	2,234	1,281	195	1,086

Notes: This table shows estimation results for regressions mirroring those in Table 2, but using only observations derived from solo-authored papers. All columns include fixed effects for the university affiliation of the PI and the year of publication. Columns (2) and (4) add controls for PI (RA) race and ethnicity as well as the seniority of the PI. Columns (1) and (2) use the full sample of RAs and PIs, columns (3) and (4) only consider the first interaction for a given RA-PI pair, and columns (5) and (6) are aggregated at the PI (RA) level. Standard errors clustered at the year level are shown in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table A2: Sorting by Gender - Top 20 Universities

Panel A: RA Gender Conditional on PI Gender						
	All Interactions		First Interaction Only		# Interactions with	
	Female RA				Female RAs	Male RAs
	(1)	(2)	(3)	(4)	(5)	(6)
Female PI	0.098***	0.102***	0.095***	0.105***	-0.031	-0.013
	(0.012)	(0.014)	(0.012)	(0.013)	(0.019)	(0.024)
Controls		✓		✓	✓	✓
Affiliation PI	✓	✓	✓	✓	✓	✓
Year of Publication	✓	✓	✓	✓	✓	✓
Dep. Variable Mean	0.372	0.376	0.377	0.378	1.12	1.13
Observations	14,702	8,322	13,155	7,403	2,796	4,607
Panel B: PI Gender Conditional on RA Gender						
	All Interactions		First Interaction Only		# Interactions with	
	Female PI				Female PIs	Male PIs
	(1)	(2)	(3)	(4)	(5)	(6)
Female RA	0.063***	0.066***	0.062***	0.065***	0.007	-0.017
	(0.009)	(0.012)	(0.009)	(0.013)	(0.022)	(0.011)
Controls		✓		✓	✓	✓
Affiliation PI	✓	✓	✓	✓	✓	✓
Year of Publication	✓	✓	✓	✓	✓	✓
Dep. Variable Mean	0.190	0.194	0.194	0.199	1.09	1.13
Observations	14,702	8,662	13,155	7,735	1,541	6,194

Notes: This table shows estimation results for regressions mirroring those in Table 2, but using only RA-PI pairs where the PI was affiliated with a university ranked among the international top 20 according to RePEc. All columns include fixed effects for the university affiliation of the PI and the year of publication. Columns (2) and (4) add controls for PI (RA) race and ethnicity as well as the seniority of the PI. Columns (1) and (2) use the full sample of RAs and PIs, columns (3) and (4) only consider the first interaction for a given RA-PI pair, and columns (5) and (6) are aggregated at the PI (RA) level. Standard errors clustered at the year level are shown in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table A3: Sorting by Race and Ethnicity - Top 20 Universities

Panel A: RA Race/Ethnicity conditional on PI Race/Ethnicity				
	Asian RA (1)	Black RA (2)	Hispanic RA (3)	White RA (4)
Asian PI	0.248***	-0.007**	-0.025**	-0.219***
	(0.020)	(0.003)	(0.011)	(0.018)
Black PI	-0.067	0.089**	0.036	-0.060
	(0.108)	(0.038)	(0.077)	(0.110)
Hispanic PI	-0.074**	-0.002	0.201***	-0.119***
	(0.028)	(0.005)	(0.031)	(0.030)
Controls	✓	✓	✓	✓
Year of Publication	✓	✓	✓	✓
Affiliation PI	✓	✓	✓	✓
Dep. Variable Mean	0.395	0.007	0.091	0.508
Observations	6,032	6,032	6,032	6,032
Panel B: PI Race/Ethnicity conditional on RA Race/Ethnicity				
	Asian PI (1)	Black PI (2)	Hispanic PI (3)	White PI (4)
Asian RA	0.122***	0.0003	-0.006	-0.114***
	(0.011)	(0.002)	(0.006)	(0.012)
Black RA	-0.082**	0.050	0.002	0.025
	(0.040)	(0.036)	(0.038)	(0.060)
Hispanic RA	-0.005	0.002	0.112***	-0.098***
	(0.014)	(0.003)	(0.016)	(0.020)
Controls	✓	✓	✓	✓
Year of Publication	✓	✓	✓	✓
Affiliation PI	✓	✓	✓	✓
Dep. Variable Mean	0.139	0.004	0.049	0.813
Observations	5,243	5,243	5,243	5,243

Notes: This table shows estimation results for regressions mirroring those in Table 3, but using only RA-PI pairs where the PI was affiliated with a university ranked among the international top 20 according to RePEc. All columns include fixed effects for the university affiliation of the PI and the year of publication. Columns (2) and (4) add controls for PI (RA) gender as well as the seniority of the PI. Columns (1) and (2) use the full sample of RAs and PIs, columns (3) and (4) only consider the first interaction for a given RA-PI pair, and columns (5) and (6) are aggregated at the PI (RA) level. Standard errors clustered at the year level are shown in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table A4: Sorting by Gender - US Universities

Panel A: RA Gender Conditional on PI Gender						
	All Interactions		First Interaction Only		# Interactions with	
	Female RA				Female RAs	Male RAs
	(1)	(2)	(3)	(4)	(5)	(6)
Female PI	0.097***	0.101***	0.092***	0.100***	0.015	0.020
	(0.010)	(0.011)	(0.010)	(0.012)	(0.012)	(0.017)
Controls		✓		✓	✓	✓
Affiliation PI	✓	✓	✓	✓	✓	✓
Year of Publication	✓	✓	✓	✓	✓	✓
Dep. Variable Mean	0.380	0.379	0.383	0.381	1.13	1.14
Observations	22,850	13,055	20,253	11,529	4,392	7,137
Panel B: PI Gender Conditional on RA Gender						
	All Interactions		First Interaction Only		# Interactions with	
	Female PI				Female PIs	Male PIs
	(1)	(2)	(3)	(4)	(5)	(6)
Female RA	0.065***	0.060***	0.063***	0.059***	-0.006	-0.018*
	(0.007)	(0.010)	(0.008)	(0.010)	(0.019)	(0.010)
Controls		✓		✓	✓	✓
Affiliation PI	✓	✓	✓	✓	✓	✓
Year of Publication	✓	✓	✓	✓	✓	✓
Dep. Variable Mean	0.213	0.215	0.215	0.219	1.11	1.13
Observations	22,850	13,483	20,253	11,964	2,616	9,348

Notes: This table shows estimation results for regressions mirroring those in Table 2, but using only RA-PI pairs where the PI was affiliated with a university located in the USA. All columns include fixed effects for the university affiliation of the PI and the year of publication. Columns (2) and (4) add controls for PI (RA) race and ethnicity as well as the seniority of the PI. Columns (1) and (2) use the full sample of RAs and PIs, columns (3) and (4) only consider the first interaction for a given RA-PI pair, and columns (5) and (6) are aggregated at the PI (RA) level. Standard errors clustered at the year level are shown in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table A5: Sorting by Race and Ethnicity - US Universities

Panel A: RA Race/Ethnicity conditional on PI Race/Ethnicity				
	Asian RA	Black RA	Hispanic RA	White RA
	(1)	(2)	(3)	(4)
Asian PI	0.263***	-0.006**	-0.029***	-0.231***
	(0.015)	(0.002)	(0.008)	(0.014)
Black PI	-0.067	0.056*	0.017	-0.007
	(0.072)	(0.031)	(0.054)	(0.080)
Hispanic PI	-0.042*	-0.004	0.183***	-0.128***
	(0.024)	(0.004)	(0.024)	(0.025)
Controls	✓	✓	✓	✓
Year of Publication	✓	✓	✓	✓
Affiliation PI	✓	✓	✓	✓
Dep. Variable Mean	0.395	0.007	0.088	0.512
Observations	9,619	9,619	9,619	9,619
Panel B: PI Race/Ethnicity conditional on RA Race/Ethnicity				
	Asian PI	Black PI	Hispanic PI	White PI
	(1)	(2)	(3)	(4)
Asian RA	0.122***	-0.0005	-0.003	-0.117***
	(0.009)	(0.001)	(0.005)	(0.010)
Black RA	-0.043	0.036	-0.013	0.017
	(0.037)	(0.026)	(0.028)	(0.049)
Hispanic RA	-0.008	0.0006	0.099***	-0.084***
	(0.011)	(0.002)	(0.012)	(0.016)
Controls	✓	✓	✓	✓
Year of Publication	✓	✓	✓	✓
Affiliation PI	✓	✓	✓	✓
Dep. Variable Mean	0.137	0.004	0.044	0.819
Observations	8,288	8,288	8,288	8,288

Notes: This table shows estimation results for regressions mirroring those in Table 3, but using only RA-PI pairs where the PI was affiliated with a university located in the USA. All columns include fixed effects for the university affiliation of the PI and the year of publication. Columns (2) and (4) add controls for PI (RA) gender as well as the seniority of the PI. Columns (1) and (2) use the full sample of RAs and PIs, columns (3) and (4) only consider the first interaction for a given RA-PI pair, and columns (5) and (6) are aggregated at the PI (RA) level. Standard errors clustered at the year level are shown in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table A6: Sorting by Gender - Papers Published Before 2000

Panel A: RA Gender Conditional on PI Gender						
	All Interactions		First Interaction Only		# Interactions with	
	Female RA				Female RAs	Male RAs
	(1)	(2)	(3)	(4)	(5)	(6)
Female PI	0.100***	0.100**	0.089***	0.085*	0.001	-0.072
	(0.025)	(0.038)	(0.024)	(0.042)	(0.042)	(0.056)
Controls		✓		✓	✓	✓
Affiliation PI	✓	✓	✓	✓	✓	✓
Year of Publication	✓	✓	✓	✓	✓	✓
Dep. Variable Mean	0.349	0.349	0.350	0.353	1.16	1.18
Observations	4,182	2,377	3,565	2,034	718	1,316
Panel B: PI Gender Conditional on RA Gender						
	All Interactions		First Interaction Only		# Interactions with	
	Female PI				Female PIs	Male PIs
	(1)	(2)	(3)	(4)	(5)	(6)
Female RA	0.041***	0.025*	0.038***	0.022	-0.037	-0.009
	(0.011)	(0.012)	(0.011)	(0.013)	(0.026)	(0.023)
Controls		✓		✓	✓	✓
Affiliation PI	✓	✓	✓	✓	✓	✓
Year of Publication	✓	✓	✓	✓	✓	✓
Dep. Variable Mean	0.120	0.115	0.127	0.122	1.11	1.19
Observations	4,182	2,315	3,565	1,964	240	1,724

Notes: This table shows estimation results for regressions mirroring those in Table 2, but using only RA-PI pairs derived from NBER working papers that were published prior to or in the year 2000. All columns include fixed effects for the university affiliation of the PI and the year of publication. Columns (2) and (4) add controls for PI (RA) race and ethnicity as well as the seniority of the PI. Columns (1) and (2) use the full sample of RAs and PIs, columns (3) and (4) only consider the first interaction for a given RA-PI pair, and columns (5) and (6) are aggregated at the PI (RA) level. Standard errors clustered at the year level are shown in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table A7: Sorting by Race and Ethnicity - Papers Published Before 2000

Panel A: RA Race/Ethnicity conditional on PI Race/Ethnicity				
	Asian RA (1)	Black RA (2)	Hispanic RA (3)	White RA (4)
Asian PI	0.247***	-0.001	-0.014	-0.232***
	(0.068)	(0.002)	(0.030)	(0.073)
Black PI	0.143	-0.003***	-0.118***	-0.021
	(0.101)	(0.0008)	(0.027)	(0.104)
Hispanic PI	-0.224***	-0.003	0.262**	-0.036
	(0.077)	(0.003)	(0.122)	(0.123)
Controls	✓	✓	✓	✓
Year of Publication	✓	✓	✓	✓
Affiliation PI	✓	✓	✓	✓
Dep. Variable Mean	0.269	0.0006	0.049	0.682
Observations	1,626	1,626	1,626	1,626
Panel B: PI Race/Ethnicity conditional on RA Race/Ethnicity				
	Asian PI (1)	Black PI (2)	Hispanic PI (3)	White PI (4)
Asian RA	0.064***	0.004	-0.007	-0.061***
	(0.019)	(0.005)	(0.006)	(0.020)
Hispanic RA	-0.005	-0.008	0.049*	-0.036
	(0.024)	(0.005)	(0.028)	(0.036)
Controls	✓	✓	✓	✓
Year of Publication	✓	✓	✓	✓
Affiliation PI	✓	✓	✓	✓
Dep. Variable Mean	0.057	0.005	0.009	0.929
Observations	1,476	1,476	1,476	1,476

Notes: This table shows estimation results for regressions mirroring those in Table 3, but using only RA-PI pairs derived from NBER working papers that were published prior to or in the year 2000. All columns include fixed effects for the university affiliation of the PI and the year of publication. Columns (2) and (4) add controls for PI (RA) gender as well as the seniority of the PI. Columns (1) and (2) use the full sample of RAs and PIs, columns (3) and (4) only consider the first interaction for a given RA-PI pair, and columns (5) and (6) are aggregated at the PI (RA) level. Standard errors clustered at the year level are shown in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table A8: Sorting by Gender - Papers Published After 2010

Panel A: RA Gender Conditional on PI Gender						
	All Interactions		First Interaction Only		# Interactions with	
	Female RA				Female RAs	Male RAs
	(1)	(2)	(3)	(4)	(5)	(6)
Female PI	0.090***	0.096***	0.088***	0.098***	0.007	0.020
	(0.011)	(0.015)	(0.012)	(0.016)	(0.011)	(0.020)
Controls		✓		✓	✓	✓
Affiliation PI	✓	✓	✓	✓	✓	✓
Year of Publication	✓	✓	✓	✓	✓	✓
Dep. Variable Mean	0.394	0.390	0.396	0.393	1.09	1.11
Observations	17,349	9,876	15,951	8,983	3,529	5,454
Panel B: PI Gender Conditional on RA Gender						
	All Interactions		First Interaction Only		# Interactions with	
	Female PI				Female PIs	Male PIs
	(1)	(2)	(3)	(4)	(5)	(6)
Female RA	0.060***	0.067***	0.059***	0.067***	0.005	-0.015
	(0.008)	(0.012)	(0.008)	(0.012)	(0.019)	(0.009)
Controls		✓		✓	✓	✓
Affiliation PI	✓	✓	✓	✓	✓	✓
Year of Publication	✓	✓	✓	✓	✓	✓
Dep. Variable Mean	0.238	0.238	0.237	0.239	1.08	1.09
Observations	17,349	10,396	15,951	9,564	2,285	7,279

Notes: This table shows estimation results for regressions mirroring those in Table 2, but using only RA-PI pairs derived from NBER working papers that were published after the year 2000. All columns include fixed effects for the university affiliation of the PI and the year of publication. Columns (2) and (4) add controls for PI (RA) race and ethnicity as well as the seniority of the PI. Columns (1) and (2) use the full sample of RAs and PIs, columns (3) and (4) only consider the first interaction for a given RA-PI pair, and columns (5) and (6) are aggregated at the PI (RA) level. Standard errors clustered at the year level are shown in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table A9: Sorting by Race and Ethnicity - Papers Published After 2010

Panel A: RA Race/Ethnicity conditional on PI Race/Ethnicity				
	Asian RA (1)	Black RA (2)	Hispanic RA (3)	White RA (4)
Asian PI	0.261***	-0.006**	-0.035***	-0.221***
	(0.020)	(0.003)	(0.009)	(0.019)
Black PI	-0.107	0.070*	0.132	-0.097
	(0.084)	(0.035)	(0.084)	(0.087)
Hispanic PI	-0.036	-0.007*	0.150***	-0.106***
	(0.025)	(0.004)	(0.024)	(0.028)
Controls	✓	✓	✓	✓
Year of Publication	✓	✓	✓	✓
Affiliation PI	✓	✓	✓	✓
Dep. Variable Mean	0.451	0.009	0.086	0.456
Observations	7,675	7,675	7,675	7,675
Panel B: PI Race/Ethnicity conditional on RA Race/Ethnicity				
	Asian PI (1)	Black PI (2)	Hispanic PI (3)	White PI (4)
Asian RA	0.123***	-0.0007	-0.001	-0.119***
	(0.011)	(0.002)	(0.006)	(0.012)
Black RA	-0.032	0.036	-0.027	0.019
	(0.039)	(0.026)	(0.024)	(0.048)
Hispanic RA	-0.014	0.006*	0.095***	-0.078***
	(0.014)	(0.004)	(0.015)	(0.019)
Controls	✓	✓	✓	✓
Year of Publication	✓	✓	✓	✓
Affiliation PI	✓	✓	✓	✓
Dep. Variable Mean	0.176	0.004	0.060	0.764
Observations	6,398	6,398	6,398	6,398

Notes: This table shows estimation results for regressions mirroring those in Table 3, but using only RA-PI pairs derived from NBER working papers that were published after the year 2000. All columns include fixed effects for the university affiliation of the PI and the year of publication. Columns (2) and (4) add controls for PI (RA) gender as well as the seniority of the PI. Columns (1) and (2) use the full sample of RAs and PIs, columns (3) and (4) only consider the first interaction for a given RA-PI pair, and columns (5) and (6) are aggregated at the PI (RA) level. Standard errors clustered at the year level are shown in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table A10: Sorting by Gender - Graduate Student RAs

Panel A: RA Gender Conditional on PI Gender						
	All Interactions		First Interaction Only		# Interactions with	
	Female RA				Female RAs	Male RAs
	(1)	(2)	(3)	(4)	(5)	(6)
Female PI	0.090***	0.048	0.096***	0.072**	-0.170***	0.005
	(0.024)	(0.030)	(0.026)	(0.032)	(0.062)	(0.037)
Controls		✓		✓	✓	✓
Affiliation PI	✓	✓	✓	✓	✓	✓
Year of Publication	✓	✓	✓	✓	✓	✓
Dep. Variable Mean	0.285	0.284	0.285	0.281	1.211	1.194
Observations	4,941	2,812	4,172	2,346	659	1,687
Panel A: PI Gender Conditional on RA Gender						
	All Interactions		First Interaction Only		# Interactions with	
	Female PI				Female PIs	Male PIs
	(1)	(2)	(3)	(4)	(5)	(6)
Female RA	0.059***	0.037*	0.065***	0.044**	-0.020	0.008
	(0.016)	(0.019)	(0.018)	(0.021)	(0.041)	(0.033)
Controls		✓		✓	✓	✓
Affiliation PI	✓	✓	✓	✓	✓	✓
Year of Publication	✓	✓	✓	✓	✓	✓
Dep. Variable Mean	0.176	0.181	0.183	0.188	1.143	1.196
Observations	4,941	3,004	4,172	2,533	475	2,058

Notes: This table shows estimation results for regressions mirroring those in Table 2, but using only observations referring to graduate RAs and their respective PIs as defined in Section 2.1. All columns include fixed effects for the university affiliation of the PI and the year of publication. Columns (2) and (4) add controls for PI (RA) race and ethnicity as well as the seniority of the PI. Columns (1) and (2) use the full sample of RAs and PIs, columns (3) and (4) only consider the first interaction for a given RA-PI pair, and columns (5) and (6) are aggregated at the PI (RA) level. Standard errors clustered at the year level are shown in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table A11: Sorting by Race and Ethnicity - Graduate Student RAs

Panel A: RA Race/Ethnicity conditional on PI Race/Ethnicity				
	Asian RA (1)	Black RA (2)	Hispanic RA (3)	White RA (4)
Asian PI	0.211***	-0.006	-0.044**	-0.162***
	(0.036)	(0.009)	(0.022)	(0.040)
Black PI	0.048	0.219	-0.109***	-0.156
	(0.308)	(0.217)	(0.035)	(0.245)
Hispanic PI	0.033	-0.012**	0.050*	-0.072
	(0.044)	(0.006)	(0.026)	(0.050)
Controls	✓	✓	✓	✓
Year of Publication	✓	✓	✓	✓
Affiliation PI	✓	✓	✓	✓
Dep. Variable Mean	0.410	0.006	0.086	0.499
Observations	2,171	2,171	2,171	2,171
Panel B: PI Race/Ethnicity conditional on RA Race/Ethnicity				
	Asian PI (1)	Black PI (2)	Hispanic PI (3)	White PI (4)
Asian RA	0.088***	0.002	0.007	-0.088***
	(0.020)	(0.004)	(0.011)	(0.022)
Black RA	0.007	0.115	-0.058	-0.078
	(0.142)	(0.112)	(0.038)	(0.164)
Hispanic RA	-0.021	-0.002	0.029	-0.005
	(0.021)	(0.002)	(0.022)	(0.029)
Controls	✓	✓	✓	✓
Year of Publication	✓	✓	✓	✓
Affiliation PI	✓	✓	✓	✓
Dep. Variable Mean	0.126	0.003	0.046	0.829
Observations	1,847	1,847	1,847	1,847

Notes: This table shows estimation results for regressions mirroring those in Table 3, but using only observations referring to graduate RAs and their respective PIs as defined in Section 2.1. All columns include fixed effects for the university affiliation of the PI and the year of publication. Columns (2) and (4) add controls for PI (RA) gender as well as the seniority of the PI. Columns (1) and (2) use the full sample of RAs and PIs, columns (3) and (4) only consider the first interaction for a given RA-PI pair, and columns (5) and (6) are aggregated at the PI (RA) level. Standard errors clustered at the year level are shown in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table A12: Former RAs, NBER Publications, and RA-PI Gender Alignment - Female RAs

	# Publications		Any Publications		# Publications (if any)	
	5Y Post-PhD (1)	10Y Post-PhD (2)	5Y Post-PhD (3)	10Y Post-PhD (4)	5Y Post-PhD (5)	10Y Post-PhD (6)
Same Gender	0.759*	1.01	0.133*	0.122	1.04	1.18
	(0.413)	(0.806)	(0.076)	(0.102)	(0.715)	(1.35)
Year of PhD Completion	✓	✓	✓	✓	✓	✓
PhD Institution	✓	✓	✓	✓	✓	✓
Dep. Variable Mean	1.08	1.85	0.414	0.456	2.27	3.86
Observations	415	298	415	298	198	143

Notes: This table shows estimation results for regressions mirroring those in Table 5, but using only observations corresponding to female graduate RAs. Columns (1) and (2) use the number of NBER publications within five and 10 years after PhD completion as outcome variable. Columns (3) and (4) use a binary indicator for having any NBER publications in the same time frame as outcome variable. Finally, columns (5) and (6) use the number of publications as outcome variable, but restrict the sample to only those PhD graduates with at least one NBER publication at any point in time. All columns include controls for the gender of the RA and the PI as well as fixed effects for year and university of PhD completion. Robust standard errors are shown in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table A13: Former RAs, NBER Publications, and RA-PI Gender Alignment - Male RAs

	# Publications		Any Publications		# Publications (if any)	
	5Y Post-PhD (1)	10Y Post-PhD (2)	5Y Post-PhD (3)	10Y Post-PhD (4)	5Y Post-PhD (5)	10Y Post-PhD (6)
Same Gender	0.586***	1.02*	0.052	0.093	0.418	0.364
	(0.206)	(0.574)	(0.055)	(0.070)	(0.391)	(1.16)
Year of PhD Completion	✓	✓	✓	✓	✓	✓
PhD Institution	✓	✓	✓	✓	✓	✓
Dep. Variable Mean	1.46	2.80	0.440	0.484	2.78	5.28
Observations	1,110	797	1,110	797	585	423

Notes: This table shows estimation results for regressions mirroring those in Table 5, but using only observations corresponding to male graduate RAs. Columns (1) and (2) use the number of NBER publications within five and 10 years after PhD completion as outcome variable. Columns (3) and (4) use a binary indicator for having any NBER publications in the same time frame as outcome variable. Finally, columns (5) and (6) use the number of publications as outcome variable, but restrict the sample to only those PhD graduates with at least one NBER publication at any point in time. All columns include controls for the gender of the RA and the PI as well as fixed effects for year and university of PhD completion. Robust standard errors are shown in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.