

Gender Differences in Recognition for Group Work

Heather Sarsons*

August 9, 2015

Abstract

Within academia, men are tenured at higher rates than women are in most quantitative fields, including economics. Researchers have attempted to identify the source of this disparity but find that nearly 30% of the gap remains unexplained even after controlling for family commitments and differences in productivity. Using data from academic economists' CVs, I test whether coauthored and solo-authored publications matter differently for tenure for men and women. While solo-authored papers send a clear signal about one's ability, coauthored papers are noisy in that they do not provide specific information about each contributor's skills. I find that men are tenured at roughly the same rate regardless of whether they coauthor or solo-author. Women, however, suffer a significant penalty when they coauthor. The results hold after controlling for the total number of papers published, quality of papers, field of study, tenure institution, tenure year, and the number of years it took an individual to go up for tenure. The result is most pronounced for women coauthoring with only men and is less pronounced the more women there are on a paper, suggesting that some gender bias is at play. I present a model in which bias enters when workers collaborate and test its predictions in the data.

NOTE: Very preliminary – please do not cite without permission. Comments welcome.

*Contact: sarsons@fas.harvard.edu. I thank Roland Fryer, Claudia Goldin, David Laibson, and Amanda Pallais for their guidance on this project. I also thank attendees of the Harvard Labour and Behavioural lunches, and the Seminar on Inequality and Social Policy for helpful feedback. This paper is intentionally solo-authored.

1 Introduction

In many industries, women and minorities are not only hired at lower rates than white men are, they are also promoted at lower rates. This phenomenon which is especially prominent in the STEM (science, technology, engineering, and math) fields, has been dubbed the “leaky pipeline”. Researchers have attempted to explain the leaky pipeline by looking at productivity differences between groups (Ginther and Kahn, 2004), differences in behaviour such as competitiveness and confidence (Niederle and Vesterlund, 2007), and the role that child-bearing plays for women (Ceci et al., 2014 or Ginther; Ginther and Kahn, 2004). Even after accounting for these factors, a significant portion of the gap remains unexplained. In academia, for example, over 30% of the observed gap in tenure rates can not be accounted for by observable productivity differences or family commitments (Ginther and Kahn, 2004).

Discrimination has been proposed as a factor contributing to the gap, but empirically testing for discrimination in promotion is difficult due to unobserved variables. The resume and audit studies typically used to test for discrimination in hiring¹ can not be used for promotion decisions. As such, most of the research on discrimination in promotion has been theoretical. Two notable examples are Fryer (2007) and Lehmann (2013) who extend Coate and Loury’s canonical model of statistical discrimination to include a promotion stage. Fryer demonstrates that minorities who are initially hired at a lower rate due to discrimination may be promoted at higher rates. This occurs if employers are “liberal” and believe that minorities who made it through the hiring stage must be truly exceptional. However, Lehmann argues that promotion may not increase if affirmative action skews employers’ views of minority hires. She develops a model in which employers are constrained in their hiring choices by a diversity policy but are unconstrained in their promotion decisions. Employers can differentially assign workers to tasks, some of which put workers in a better position for promotion. Employers who have a negative view of minority workers will place minority hires in “non-promotion tracks” and will consequently be promoted at a lower rate.

This paper proposes an alternative explanation for the promotion gap. I argue that bias can enter when workers can work in groups, a feature common in many workplaces today. While working with others lowers the cost of production, it gives the employer a noisy signal of each worker’s ability and he must make a judgment call as to who put in the most effort. I test this idea using data from

¹For examples, see Bertrand and Mullainathan (2004) or Pager (2003).

academic economists' CVs. Within academia, women hold a small fraction of full professorships across quantitative disciplines. Economics is no outlier. While women's representation among doctoral degree recipients has increased over time, there has not been a corresponding increase in their representation among tenured faculty.

Figure 1 motivates the paper. It shows the relationship between tenure and the fraction of an economist's papers that are solo-authored by tenure. The data behind the plot will be discussed in the body of the paper, but the figure provides evidence that women suffer a "coauthor penalty". While women who solo-author everything have roughly the same chance of receiving tenure as a man, women who coauthor most of their work have a significantly lower probability of receiving tenure. The penalty is not explained by coauthor selection and is robust to controlling for tenure institution, year of tenure, and field of study.

The remainder of the paper is organized as follows. Section 2 presents a model in which employers make promotion decisions based on group signals of productivity. The model provides a set of testable predictions that distinguish between several factors that could be driving the coauthor penalty, including bias, discrimination, and rational updating. Section 3 describes the data used to test the predictions of the model. The results are presented in Sections 4 and 5. Section 6 concludes.

2 Model

Existing models of statistical discrimination consist of employees working alone and sending signals about their abilities². Many workplaces today rely on group work, making inference about an employee's type more difficult. This difficulty is compounded when employees can choose to work alone or in groups, as is the case in academia. Figure 1 suggests that employers make different inferences about a man and a woman's type when employees work jointly. However, this could be due to statistical discrimination, taste-based discrimination, or bias. The model provides a set of predictions to distinguish between these channels. I also solve two cases: one in which workers are naive and do not consider that employers might infer something from their decision to collaborate and another where they are informed. These cases provide additional predictions that can be tested to see whether employees anticipate discrimination.

²For example, see Aigner and Cain (1977), Bjerck (2008), Coate and Loury (1993), and Fryer (2007),

2.1 Basic Setup

The model begins with hiring decisions already having been made. Workers belong to an identifiable group, men or women, denoted by $g \in \{M, W\}$. Nature assigns them a type (ability), $a \in \{L, H\}$, that is known to the worker but unobserved by other workers and the employer. Both employers and workers share the prior that a fraction π_w of female workers are high types and fraction π_m of male workers are high types, where $\pi_m > \pi_w$.

Workers must complete a project for their firm and can decide whether they want to work alone or with another worker. After receiving a signal from the worker, the employer makes a promotion decision. The exact sequence of events is as follows:

1. Worker i draws a project with associated cost c_i which is drawn from the distribution with CDF $G(c)$. At the same time, workers are randomly matched to another worker, j (“the collaborator”).
2. Collaborator j sends the worker a signal, α_j , about j 's probability of being a high type. Collaborators who are high types draw their signal from the distribution with CDF $\alpha_j \sim F_H(\alpha)$. Collaborators who are low types draw from $\alpha_j \sim F_L(\alpha)$. It is assumed that $F_H(\alpha) \leq F_L(\alpha) \forall \alpha \in [0, 1]$ so that high types are more likely to draw high signals. Employers do not see α_j .
3. Worker i decides whether to work alone or collaborate. Collaborating reduces the cost of production, described in greater detail below.
4. Workers complete their projects (either alone or with another worker) and send a signal to the employer, $\theta \in \{\theta_L, \theta_H\}$.
5. Employers make promotion decisions.

2.1.1 Costs and Signals

Workers draw a cost associated with their projects but the realized cost depends on whether they work independently or with a partner. Workers who work independently pay the full project cost, c_i . If two high ability workers collaborate, the cost of production falls, which is normalized to 0.

However, if a high ability worker collaborates with a low ability worker, or if two low ability workers collaborate, they each bear the full cost, c_i . For example, collaboration allows one to both minimize the time spent on a project and to raise the quality of the completed project. If the coworker on the project is unable to contribute much, the worker in question will have to spend as much time as she otherwise would to complete it. She could spend less time on the project but at the cost of a lower quality output since the coworker is not contributing to the quality.

A worker who works alone perfectly reveals her type, sending signal θ_H if $a_i = H$ and θ_L if $a_i = L$. If a worker chooses to collaborate, the signal she sends depends on her and her collaborator's types. If both are high types, they send signal θ_H with probability 1. If they are both low types, they send signal θ_L with probability 1. If one is a low type and one is a high type, they send θ_H with probability γ and θ_L with probability $1 - \gamma$ where $\gamma \in (0, 1)$ is a known constant. In academia, this is akin to saying that two high types produce amazing papers but a low and a high type can only produce a good paper.

To summarize, workers receive full information about the cost of the project, c_i , and a signal about the collaborator's type, α_j . Workers decide whether to collaborate based on c_i , α_j , and π_g . The employer then receives a signal, θ , from the worker and also knows whether the worker collaborated and the cost of the project. The employer does not know the worker's true type, nor the type of her collaborator. He decides whether to promote the worker based on θ , π , and the decision to collaborate (or not).

2.1.2 Payoffs

Workers who are promoted receive wage w while those who are not promoted receive 0. A worker who is promoted thus receives the following ex-post payoffs, regardless of type:

- If worker i collaborates with j and $a_j = H$, i 's payoff is w .
- If worker i collaborates with j and $a_j = L$, i 's payoff is $w - c_i$.
- If worker i works alone, payoff is $w - c_i$.

A worker who is not promoted receives the following ex-post payoffs, regardless of type:

- If worker i collaborates with j and $a_j = H$, i 's payoff is 0.

- If worker i collaborates with j and $a_j = L$, i 's payoff is $-c_i$.
- If worker i works alone, payoff is $-c_i$.

Employers who promote a high ability worker receive payoff $\chi_H - w$ where $\chi_H > w$. Employers who promote a low ability worker receive payoff $\chi_L - w$ where $\chi_L < w$.

2.2 Naive Worker's Decision

For now, I assume that workers are naive and do not consider the fact that employers infer something from their decision to coauthor. Instead they assume that employers will promote workers who send a signal θ_H and not promote anyone who sends a signal θ_L , regardless of whether they collaborate or work alone. This assumption is relaxed in Section 2.4, but the data suggest that workers do not collaborate strategically and that employers do not uniformly take the decision to collaborate as an additional signal.

Recall that upon drawing the project, the worker is matched to another worker with whom she can collaborate. Collaborating is both beneficial and costly. Collaborating lowers the worker's expected cost of production but also lowers the probability that the worker will be promoted since there is some chance the other worker is a low type. Upon receiving a signal α_j from the potential collaborator, the worker updates her beliefs about the collaborator's ability and weighs the costs and benefits of collaborating.

Consider a high ability female worker, i , who draws a project and is matched to a male worker who sends signal α_j . Worker i updates her belief about j 's ability according to Bayes' rule:

$$\varphi(\alpha_j) \equiv \mathbb{P}(a_j = H | \pi_m, \alpha_j) = \frac{\pi_m f_h(\alpha_j)}{\pi_m f_h(\alpha_j) + (1 - \pi_m) f_l(\alpha_j)}.$$

Worker i will collaborate if the expected cost reduction from collaborating outweighs the possibility that j is a low type and the certainty of drawing a high signal if i works alone³. This is formalized in equation 1 below where the right-hand side is expected payoff of collaborating and the left-hand

³Allowing for risk aversion does not change the predictions of the model but makes high ability workers less likely to collaborate.

side is the payoff from working alone.

$$\begin{aligned}
w\mathbb{P}(a_j = H|\pi_m, \alpha_j) + \gamma w\mathbb{P}(a_j = L|\pi_m, \alpha_j) - c_i\mathbb{P}(a_j = L|\pi_m, \alpha_j) &\geq w - c_i \\
w\varphi + \gamma(1 - \varphi)w - (1 - \varphi)c_i &\geq w - c_i \\
w(\varphi + \gamma(1 - \varphi) - 1) &\geq (1 - \varphi)c_i - c_i \quad (1) \\
w \left[\frac{1 - \varphi - \gamma(1 - \varphi)}{\varphi} \right] &\leq c_i
\end{aligned}$$

Equation 1 shows how Worker i 's decision to collaborate changes with production costs and beliefs. The worker is more likely to coauthor as the cost of the project increases and as the probability that worker j is a high type increases. Relating this to the decision to coauthor, some projects, such as RCTs, might be so costly to complete on one's own that collaborating is attractive even if the coauthor may not turn out to be an ideal match. As the probability that the coauthor is a high type increases, the greater are the expected cost savings, making collaborating the optimal choice. This can occur if π_m , the belief over how many qualified men exist in the population, increases or if the man draws a high signal, α_j . The worker is less likely to collaborate as the wage increases since they are less willing to risk losing their promotion by being matched to a low ability coworker. High wages also begin to outweigh the cost of working alone, making the expected cost saving from collaborating less attractive.

Equation 1 defines a cutoff α^* for which worker i is indifferent between working alone or working with worker j . Through the same calculation, cutoff signals can be defined for all worker types:

1. A low ability group g worker receiving signal $\alpha_{j,g}$ from a group g worker will collaborate if

$$\begin{aligned}
\gamma\varphi_{j,g}w - (1 - \varphi_{j,g})c_i &\geq -c_i \\
\gamma\varphi_{j,g}w &\geq -\varphi_{j,g}c_i \\
-\gamma w &\leq c_i
\end{aligned}$$

which holds for any positive wage. Male and female low ability workers will coauthor regardless of the signal α they receive. They pay the cost of production, c_i , regardless of whether they collaborate but collaborating raises the probability that the employer will believe they are a high type.

2. A high ability group g worker receiving signal $\alpha_{j,g}$ from a group g worker will collaborate if

$$w \left[\frac{\varphi_{j,g} + \gamma(1 - \varphi_{j,g}) - 1}{\varphi_{j,g}} \right] \leq c_i$$

which implicitly defines a cutoff $\alpha_{j,g}^*(\pi_g, \gamma, c)$, below which the worker will choose to work alone. Note that for a given signal $\tilde{\alpha} = \alpha_{j,m} = \alpha_{j,w}$, since $\pi_m > \pi_w$, we will have that $\alpha_{j,w}^* > \alpha_{j,m}^*$. Because of people's prior that there are fewer qualified women than men, a woman with the same credentials as a man is less likely to be a high type. As such, both male and female workers will hold female workers to a higher standard than male workers.

Prediction 1: Both men and women who are high ability will require a higher signal, α^* , from women than from men in order to work with them. Women who collaborate will therefore be, on average, higher ability than men who collaborate.

2.3 Employer's Decision

The employer wants to promote all high ability workers without promoting any low ability workers. For a solo signal, the decision rule is simple: he will promote if $\theta = \theta_H$ and not promote if $\theta = \theta_L$. Because workers perfectly reveal their types, the employer has no reason to rely on his priors and will promote men and women at the same rate.

Prediction 2: Men and women of the same ability level who always work on their own will be promoted at the same rate.

Note that if employers were also naive and did not take the decision to collaborate as a signal, the decision rule would also be to promote any worker who has collaborated and sent a θ_H signal and not promote otherwise. I show that employers use the decision to collaborate as a signal for women in Section 4.1 and continue the model assuming that employers are not naive.

Case 1: Promotion of a female (Group W) worker who decides to collaborate

Like the workers, the employer holds the belief that π_w women are high types⁴. Upon receiving a signal, θ_i , from a female worker who collaborated with someone from group g , the employer will update his belief that the worker is a high type according to Bayes' rule:

$$\begin{aligned} \beta_{w,g}(\theta) &\equiv \mathbb{P}(a_i = H | \theta_H, \pi_w) = \frac{\mathbb{P}(\theta_H | a_i = H) \pi_w}{\mathbb{P}(\theta_H | a_i = H) \pi_w + \mathbb{P}(\theta_H | a_i = L) (1 - \pi_w)} \\ &= \frac{\pi_w (\beta_{j,g} (1 - F(\alpha_g^*)) (1 - F(\alpha_w^*)) + \gamma(1 - \beta_{j,g}) (1 - F(\alpha_g^*)))}{\pi_w (\beta_{j,g} (1 - F(\alpha_g^*)) (1 - F(\alpha_w^*)) + \gamma(1 - \beta_{j,g}) (1 - F(\alpha_g^*))) + (1 - \pi_w) (\gamma \beta_{j,g} (1 - F(\alpha_w^*)))} \\ &= \frac{\pi_w (1 - F(\alpha_g^*)) [\beta_{j,g} (1 - F(\alpha_w^*)) + \gamma(1 - \beta_{j,g})]}{\pi_w (1 - F(\alpha_g^*)) [\beta_{j,g} (1 - F(\alpha_w^*)) + \gamma(1 - \beta_{j,g})] + (1 - \pi_w) [\gamma \beta_{j,g} (1 - F(\alpha_w^*))]} \end{aligned}$$

In the above equation, $\beta_{w,g}(\theta)$ is the probability that the worker being considered for promotion is high ability and $\beta_{j,g}$ is the probability that the collaborator, j , from group g is high ability. α_g^* and α_w^* are the thresholds that i had set for j and that j had set for i . For simplicity, assume that the employer does not have information on the coworker so that $\beta_{j,g} = \pi_g$. This occurs when the employer is only considering worker i for promotion and has therefore not yet received a signal from worker j as he or she is not up for promotion. In the context of coauthoring, this amounts to assuming employers do not have information on how good a person's coauthors are⁵. Allowing employers to receive a signal from worker j changes the results insofar as it changes the employer's belief about whether j would agree to work with i . Given the complementarity in ability and output, the higher is β_j , the more likely it is that worker i is also a high type. This is further discussed in Section 2.4. For now, however, I set $\beta_{j,g} = \pi_g$, giving

$$\beta_{w,g}(\theta) = \frac{\pi_w (1 - F(\alpha_g^*)) [\pi_g (1 - F(\alpha_w^*)) + \gamma(1 - \pi_g)]}{\pi_w (1 - F(\alpha_g^*)) [\pi_g (1 - F(\alpha_w^*)) + \gamma(1 - \pi_g)] + (1 - \pi_w) [\gamma \pi_g (1 - F(\alpha_w^*))]} \quad (2)$$

Because $\beta(\theta) \in [0, 1]$, employers are, except for in extreme cases, uncertain about the worker's

⁴One might think that since hiring has already occurred, employers should set $\pi_m = \pi_w$ since they would try not to hire any low ability workers. If this is the case, workers will promote a man and a woman with the same signal θ with equal probability even when they have collaborated. I test for this in Section IV and show that employers treat men and women with the same signals differently.

⁵I relax this assumption when I empirically test the model's predictions.

type and therefore less likely to promote a high ability collaborating worker than a high ability independent worker.

Prediction 3: High ability men and women who collaborate are less likely to be promoted than high ability men and women who work alone.

The employer's belief, $\beta_{w,g}(\theta_i)$, depends directly and indirectly on his priors, π_m and π_w . As his views about women become more favourable, the more likely he is to believe a given woman is high ability. The priors also influence the cutoffs that workers set for one another which provides additional information about their types. Specifically, as the threshold that i sets for j ($\alpha_{j,g}^*$) increases, $\beta_{w,g}$ falls. To see this, recall that $\alpha_{j,g}^*$ is the standard that high ability workers set for their matches (low ability workers will always collaborate and so have no standards). If they set an impossibly high standard, they will never collaborate since no one can reach their standard. Thus, a high type would never collaborate to begin with so the worker in question must be a low type. As $\alpha_{j,g}^*$ decreases, $\beta_{w,g}$ increases to a value that is less than one. As more high types are willing to collaborate, a collaborating worker sending θ_H are more likely to be high types. Similarly, as worker j raises the standard for women (α_w^* increases in the expression above), $\beta_{w,g}$ increases. High ability workers are setting such a high standard for women that the only workers who will agree to collaborate with female workers are low types. Since low types cannot on their own send a high signal, worker i must be a high type.

Recall that employers do not update their beliefs about the worker's collaborator. The employer simply believes that the collaborator is a high type with probability π_g . Because of this, employers think that female collaborators are more likely to be low types than male collaborators. A woman collaborating with a woman is therefore more likely to be a high type than a woman collaborating with a man. Since at least one of the workers needs to be a high type in order to send θ_H , the employer is more likely to believe the collaborator was the high type if the collaborator is a man. We should therefore see employers promoting more women who collaborate with women than women who collaborate with men. This gap decreases as the prior about women approaches the prior about men.

Prediction 4: Both men and women who work with women are more likely to be promoted than those who work with men.

Case 2: Promotion of a male (Group M) worker who decides to collaborate An Employer updates his belief about a male worker according to

$$\beta_{m,g}(\theta) = \frac{\pi_m (1 - F(\alpha_g^*)) [\beta_{j,g} (1 - F(\alpha_m^*)) + \gamma(1 - \beta_{j,g})]}{\pi_m (1 - F(\alpha_g^*)) [\beta_{j,g} (1 - F(\alpha_m^*)) + \gamma(1 - \beta_{j,g})] + (1 - \pi_m) [\gamma\beta_{j,g} (1 - F(\alpha_m^*))]}$$

Following the same logic as in Case 1, employers believe that men coauthoring with women are more likely to be high types than men who coauthor with other men. The employer's beliefs again change directly and indirectly with π_g . However, because employers and workers share the belief that there are more qualified men than women, the model predicts that men who collaborate are more likely to be promoted than women who collaborate. The promotion gap grows as beliefs about women (men) become less (more) favourable.

Prediction 5: Between men and women who collaborate, a woman is less likely to be promoted than a man, even if they have the same α .

Prediction 6: The greater is the difference $\pi_m - \pi_w$, the fewer women will be promoted as β_{wm} falls. Case 1 is pictured graphically in Figure 2. In this figure, I set $\pi_m = 0.8$ and show how employers' beliefs change as π_w changes, which in turn affects α_w^* . Case 2 is similar except that priors over women would be held constant at a lower value, since $\pi_w < \pi_m$. This means that the employer has more positive beliefs about the male worker for any value of π_m .

An employer will choose to promote a worker who collaborates if

$$\begin{aligned} \chi_H \mathbb{P}(a_i = H | \theta_H, \text{collaborate}, \pi_w) - \chi_L (1 - \mathbb{P}(a_i = H | \theta_H, \text{collaborate}, \pi_w)) - w &\geq 0 \\ \chi_H \beta(\theta) - \chi_L (1 - \beta(\theta)) &\geq w \\ \beta^*(\theta) &\geq \frac{w + \chi_L}{(\chi_H + \chi_L)} \end{aligned}$$

which defines a threshold belief that depends on π and α_g^* . If workers are informed, this defines an equilibrium where workers use β^* to decide whether they should collaborate. This is further

explored in Section 2.4

2.4 Informed Worker's Decision

If workers know that employers take the decision to collaborate as an additional signal of ability, they should collaborate strategically. Specifically, the worker now chooses to coauthor if

$$w\mathbb{P}(a_j = H|\pi_m, \alpha_j) \beta_{ig} + \gamma w\mathbb{P}(a_j = L|\pi_m, \alpha_j) \beta_{ig} - c_i \mathbb{P}(a_j = L|\pi_m, \alpha_j) \geq w - c_i$$

$$w \left[\frac{1 - \beta_{ig}\varphi - \gamma(1 - \beta_{ig}\varphi)}{\varphi} \right] \leq c_i$$

Since there is now some probability that employers will attribute credit to the coworker ($\beta_{ig} < 1$), workers are less likely to collaborate than in the naive case. However, since the employer has more optimistic beliefs about a worker who collaborates with a woman than with a man ($\beta_{i,w} > \beta_{i,m}$), workers face a tradeoff between collaborating with a woman and working alone. On one hand, women are less likely to be high ability which lowers the chance of drawing θ_H . On the other hand, if workers do draw θ_H , more of that credit will be given to the worker being promoted. Men in particular benefit from working with women. Both male and female workers are therefore more likely to work with women than in the naive case. If this occurs, employers realize that workers are accepting lower draws of female coworkers and attribute even more of a high signal to the worker being promoted, especially if that worker is a man.

When workers are informed, most of the predictions from Sections 2.2 and 2.3 hold but become more extreme: for most values of π_w , women require a higher α to be accepted as collaborators and men working with women are even more likely to be promoted than men working with men. While these predictions are not empirically testable, there is one prediction that can be tested using our data:

Prediction 7: High types in the informed case are less likely to collaborate than in the naive case. High ability women in particular are more likely to work alone.

3 Data

The sample consists of economists who went up for tenure between 1975 and 2014 in one of the top 30 PhD-granting universities⁶ in the United States. To account for people who went up for tenure, were denied it, and moved, I collected historical faculty lists from 16 of the 30 schools and locate over 90% of faculty who had ever gone up for tenure at these 16 institutions. To find individuals who had gone up for tenure at the remaining 14 schools, I looked at the top 75 U.S. institutions, the top 5 Canadian institutions, and the top 5 European institutions to locate anyone who went up for tenure at a top 30 U.S. school and then moved to another school. I also checked economists' CVs at the major Federal Reserve Boards in the U.S. This leaves us with a sample of 542 economists, suggesting the analyses suffers from a small sample selection problem. The thirty schools included have, on average, 19.6 tenured professors who had gone up for tenure between 1975 and 2014 and whose primary appointment is in the institution's economics department. Of these 535 economists, some went up for tenure at schools outside of the U.S. or in non-economics departments and are therefore excluded from the sample. For example, someone who received tenure at the London School of Economics and moved to the University of Michigan as a professor is not included in the sample. Nor is someone who was tenured at Chicago Booth but now works at Stanford. However, assuming that for every individual who was tenured outside of the schools in the dataset there is another individual who did not receive tenure and left for another school or the private sector, I should have at least 588 economists. The direction of this bias depends on whether men or women are more likely to leave academia for the private sector and whether those who leave are likely to have coauthored.

From an individual's CV, I code where and when he received his PhD, his employment and publication history, and his primary and secondary fields. To determine whether someone received tenure, I follow the guidelines on each school's website as to when tenure decisions are made. The majority of schools require faculty to apply for tenure after 7 years. I consider one year before and after the 7th year to account for people who go up for tenure early or late because of a leave of absence, for example. I put universities into bins of 3 based on their ranking and assume that an individual is denied tenure if that person moves to a lower-ranked university group after 6-8 years.

⁶Ranking is from <https://ideas.repec.org/top/top.usa.html> where only PhD-granting institutions are included. For example, the National Bureau of Economic Research is not included in the ranking even though it ranks second on the IDEAS list.

For example, a person who moves from Harvard to MIT after 6 years is not assumed to have been denied tenure since he moves within the same bin of schools. Someone who moves from Harvard to UCLA after 6 years is assumed to have been denied tenure since he moves to a lower group of schools. As another example, a person who moves 5 or fewer years after his initial appointment is not assumed to have been denied tenure since he moved before the tenure window (years 6 through 8 at an institution) starts .

I use the RePEc/IDEAS ranking of economics journals to control for the quality of a person's publications. I take the top 85 journals and give the top journal a score of 85. The lowest quality journal has a score of zero.

Table 1 presents summary statistics of the data. Approximately 70% of the full sample received tenure at the first institution they went up for tenure at but this masks a stark difference between men and women. Only 50% of women receive tenure while 77% of men do. There are small productivity differences between men and women: on average, men have published 0.6 more papers than women by the time they go up for tenure and they go up for tenure half a year earlier. The difference in publication counts is driven by coauthored papers, however, as there is no statistically significant difference between the number of solo-authored papers that men and women produce before tenure. There is also a small and significant difference in the journal rank that men and women publish in.

If women are tenured at lower rates because they are on average less productive than men, a woman who publishes the same amount of papers, or more, as a man who is tenured should also receive tenure. The remainder of the paper explores the tenure gap between men and women and tests the predictions from the model.

4 Empirical Strategy and Results

I begin by showing that the tenure gap is driven by the decision to coauthor. I then test the predictions surrounding decisions to collaborate and whom to collaborate with that were laid out in Sections 2.2 and 2.4.

4.1 Main Results

4.1.1 Paper type and tenure

Figure 3 plots the relationship between total publications and tenure. An additional paper is associated with a 4.5% increase in the probability of receiving tenure for both men and women but a constant gender gap between promotion rates persists. Women are on average 18% less likely to receive tenure than a man, even after controlling for productivity differences. The OLS regression lines in Figure 3 are plotted after controlling for average journal rank, the number of years it took the person to go up for tenure, tenure institution, year of tenure, and field, suggesting that productivity differences or field selection are not driving the gap.

As Figure 1 illustrated, the composition of papers matters for tenure, at least for women. Solo-authored papers are clear signals of a worker’s ability. In the model, employers start with different priors about men and women. After receiving a solo signal from both a man and a woman, the employer will update his beliefs upward. The employer continues to update his beliefs upward the more solo signals he receives until both the man and the woman are believed to be high types. The gap in tenure rates should therefore close the more solo-authored papers women produce. Figure 4 plots the relationship between solo-authored papers and tenure and shows that the tenure gap closes as women produce more solo-authored papers. The regression lines are estimated by running the regression

$$T_{ifst} = \beta_1 S_i + \beta_2 (fem_i \times S_i) + \beta_3 CA_i + \beta_4 (fem_i \times CA_i) + \beta_5 fem_i + \gamma' Z_i + \theta_f + \theta_s + \theta_t + \epsilon_{ifst} \quad (3)$$

I plot the coefficient on S_i for men and women controlling for an individual’s number of coauthored papers (CA_i), field (θ_f), tenure-granting institution (θ_s), year of tenure (θ_t), and a vector of individual-level characteristics (Z_i) including average journal ranking and time to tenure. Table 3 presents the full results from this estimation using a probit model.

The model predicts that individuals with mostly coauthored papers will be less likely to receive tenure than an individual whose papers are mostly solo-authored since the employer must now infer ability from the paper quality and from the decision to coauthor. Additionally, if women are believed to be lower ability, the “coauthor penalty” will be more pronounced for women than for men.

The tenure gap will therefore not close as women produce coauthored papers. Figure 5 plots the coefficient on CA_i from equation 3. While an additional coauthored paper increases the probability of receiving tenure, it helps a man more than it helps a woman. The tenure gap grows the more coauthored publications individuals have, conditional on the number of solo-authored papers they have. This is consistent with Prediction 5. However, looking at the size of the coefficients in Table 3, an additional coauthored paper is not less helpful to a man than an additional solo-authored paper as the model predicts. In fact, there is no difference in how coauthored and solo-authored papers correlate with tenure probability for men. Employers either think that all men are high types, which is unlikely given that some men do not receive tenure, or they take coauthoring as an additional signal about a woman’s ability but not a man’s.

Whether employers consider coauthoring to be a signal for women but not for men is partially testable by looking at how R^2 changes when different variables are included in equation 3. If tenure committees treat the composition of papers as a source of information for women but not for men, including the coauthor and solo-author variables in place of a “total papers” variable should explain a greater amount of the total variation in tenure decisions for women. That is, the R^2 should from estimating equation 4 be higher than the R^2 from estimating equation 5 below on the female sample. However, there should be little difference between the R^2 s when estimating the two equations on the male sample.

$$T_{ifst} = \beta_1(TotalPapers_i) + \gamma'Z_i + \theta_f + \theta_s + \theta_t + \epsilon_{ifst} \quad (4)$$

$$T_{ifst} = \beta_1S_i + \beta_2CA_i + \gamma'Z_i + \theta_f + \theta_s + \theta_t + \epsilon_{ifst} \quad (5)$$

The results are shown in Table 4. Using the male sample in Columns 1 and 3, the R^2 does not change after including the coauthor and solo author information separately. However, the R^2 for women increases drastically from 0.393 to 0.540.

While the results for women fit with the model, the results for men do not. This suggests that some bias could be at play: employers look only at the quality of a man’s work when evaluating him, regardless of whether he completed it on his own or in a group⁷. When women collaborate, however, how much and what the woman contributed comes into question.

⁷I show in Section 4.2.3 that employers do take coauthoring as a signal when junior men coauthor with senior men.

It could be, however, that employers practice taste-based discrimination and this leads to differential treatment of men and women. For example, employers might have a distaste for promoting women but, because of potential lawsuits, are unable to refuse tenure to women who have proven themselves capable by solo-authoring. Discriminating employers can make the case, though, that a woman who coauthors is not that good and was riding off of her coauthors' efforts. This is tested in the next section.

4.1.2 Taste-based discrimination: are women who coauthor with women better off than women who coauthor with men?

If employers statistically discriminate, women will be denied tenure regardless of whom they coauthor with (men or women). Conversely, the statistical discrimination model presented in this paper predicts that, as long as $\pi_w < \pi_m$, a woman who coauthors with other women is more likely to receive tenure than a woman who coauthors with men. Although female coauthors have higher standards set for them, they are also believed to be lower ability. If an employer receives a high signal from a woman who coauthored with another woman, the employer is more likely to believe that the woman is a high type since her female coauthor is more likely to be a low type than a man.

I rule out taste-based discrimination by estimating

$$\begin{aligned}
 T_{i\text{fst}} = & \beta_1 S_i + \beta_2 (\text{fem}_i \times S_i) + \beta_3 CA\text{fem}_i + \beta_4 (\text{fem}_i \times CA\text{fem}_i) + \beta_5 CA\text{male}_i & (6) \\
 & + \beta_6 (\text{fem} \times CA\text{male}_i) + \beta_7 CA\text{mix}_i + \beta_8 (\text{fem} \times CA\text{mix}_i) + \beta_9 \text{fem}_i \\
 & + \gamma' Z_i + \theta_f + \theta_s + \theta_t + \epsilon_{i\text{fst}}.
 \end{aligned}$$

As before, S_i is the number of solo-authored and coauthored papers individual i has. Individuals' coauthored papers are broken down into categories based on the gender of their coauthors. $CA\text{fem}_i$ is the number of coauthored papers and individual has in which all of the coauthors are female. Similarly, $CA\text{male}_i$ is the number of papers individual i has in which all of the coauthors are male. $CA\text{mix}_i$ is the number of papers an individual has in which the coauthors consist of men and women. The results in Table 5 show that most of the coauthoring penalty that women suffer is incurred when they coauthor only with men. Papers in which there is at least one other woman ($CA\text{mix}$) have a much smaller, and insignificant, negative association with tenure and papers with

only women have a positive, but also insignificant, relationship⁸.

Again, the results support the model’s predictions only for women. Men do not benefit more from coauthoring with women than a man: an additional paper with a woman is associated with a 6.1% increase in tenure probability while an additional paper with a man is associated with a 7.0% increase in tenure probability. The results provide further suggestive evidence of gender bias in promotion decisions.

4.2 Coauthor Selection

It is difficult to test the hypothesis that colleagues set standards for one another and select coauthors based on these standards. Individuals likely have different amounts of information about one another and coauthors may be selected for different reasons (providing data, for example). Job market papers, however, serve as an initial and widespread indicator of ability. Assuming that an individual’s job market paper is a good proxy of his or her ability, I first test whether high ability women anticipate discrimination and choose not to coauthor. Finding no evidence for this, I then test whether women are held to higher standards than men for coauthorship (Prediction 1). If so, women who coauthor should have better job market papers than men who coauthor.

4.2.1 Do women anticipate discrimination?

If women know that employers statistically discriminate in the face of a noisy signal, high ability women should choose to solo author since doing so clearly reveals their type. Testing whether women anticipate discrimination is difficult as there is no clear measure of a person’s ability. As a proxy for ability, I use the quality of journal that an individual’s job market paper was published to estimate

$$FracCA_{i_{fst}} = \beta_1 abil_i + \beta_2 (fem_i \times abil_i) + \beta_3 fem_i + \beta_4 TotPapers_i + \beta_5 T_i + \theta_f + \theta_s + \theta_t + \epsilon_{i_{fst}} \quad (7)$$

where $FracCA_{i_{fst}}$ is the fraction of person i ’s papers that are coauthored, $abil_i$ is person i ’s ability (job market paper rank), and T_i is a dummy variable for being tenured. If higher ability women predict that employers will discriminate, they will try to reveal their ability by solo-authoring a

⁸Unfortunately because there are so few papers with only female authors, this estimate is particularly noisy.

greater fraction of their pre-tenure publications. We would therefore expect $\beta_2 < 0$. However, ability seems to be uncorrelated with the fraction of papers that are solo-authored for both men and women (Table 6). While men coauthor more than women do overall, there is no evidence that women along the ability distribution act strategically in their choice to coauthor or solo-author. Column 4 of Table 6 asks whether high ability women coauthor more with women. This could also provide a clearer signal of ability if employers attribute more credit to men but divide credit evenly among female coauthors. However, I do not find evidence of this.

Overall, the results suggest that women either do not know that there is a coauthor penalty and therefore do not choose coauthors strategically, or that the benefit to coauthoring is sufficiently high such that women will take the coauthor penalty to produce a better paper. Another possibility is that they do not know their own ability and therefore coauthor as they think they are low ability.

4.2.2 Do women have to send higher signals?

While high ability women do not appear to be solo-authoring more than other women, workers could still use the job market paper signal as a screening mechanism for potential coauthors, as predicted in Section 2.2 of the model. The model predicts that women have to have a higher signal, α^* , in order to convince people to coauthor with them. I test this prediction by regressing the number of coauthors an individual has had at the time of tenure on ability, again measured by an individual’s job market paper. According to the model, assistant professors who have a well published job market paper should have an easier time gaining coauthors. Women in particular, though, will need a top job market paper to gain coauthors.

Figure 6 plots the regression line of

$$\text{NumCoauthors}_{ifst} = \beta_1 \text{abil}_i + \beta_2 \text{TotPapers}_{it} + \beta_3 \text{TotalCAs} + \theta_f + \theta_s + \theta_t + \epsilon_{ifst} \quad (8)$$

where the dependent variable, $\text{NumCoauthors}_{ifst}$, is the number of unique coauthors individual i has had at the time he or she goes up for tenure. I control for the total number of coauthors an individual has (TotalCAs_i) but the following results are robust using the total number of coauthors as the dependent variable.

The relationship between ability and number of coauthors is small ($\beta_1 = 0.01$) but significant at the 5% level for women. There is no significant relationship between ability and number of

coauthors for men, however. These results are again in line with the models predictions for women, who have more coauthors if they have a better published job market paper, but not for men.

4.2.3 Senior coauthors

It could be that junior women select different types of coauthors than junior men do. If junior women coauthor more frequently with senior men while junior men coauthor with their male peers, the effect we see could be due to senior people being more established and therefore more likely to receive credit than junior faculty. Table 7 re-estimates equation 3 but includes a control for the number of senior coauthors an individual has. Looking at column 1, coauthoring with more senior researchers is associated with a 2.4% lower chance of receiving tenure but does not detract from the relationship between being female and coauthoring and receiving tenure. In Column 2, the senior coauthor term is interacted with the female dummy. Surprisingly, coauthoring with senior professors is not negatively correlated with tenure for women. In fact, it is correlated with a 1% increase in tenure probability while it is associated with a 4% decrease in tenure probability for men.

The fact that coauthoring with a senior professor hurts junior men suggests that tenure committees do take into account a person's decision to coauthor. However, employers seem to care about different aspects of who someone is coauthoring with when considering women and men for tenure.

5 Clear signals: Testing against other coauthoring conventions

That women are promoted at the same rate as men when they solo-author suggests that employers discriminate or have a bias that enters when they receive unclear signals, such as a coauthored paper. If this is true, we would expect the effect to diminish if individuals could truthfully signal their contribution. In sociology, authors are listed in order of contribution. Redoing the analysis using data from sociology provides a placebo check although it is imperfect given the different gender composition of faculty.

The sociology sample consists of randomly sampled faculty at the top 20 sociology PhD-granting schools in the U.S. There are 221 sociologists in the sample and 40% are female. Table 7 shows

that tenure rates are comparable for men and women as are most productivity measures.

I test whether men and women are treated differently when they coauthor papers in Table 8. I estimate equation 3 but include measures of the number of papers that researcher i is first author on. In column 1, I include the number of coauthored papers that a researcher is first author on as well as the female dummy interaction term. In column 2, I include the fraction of a researcher's coauthored papers that she is first author on and the interaction term.

Being first author on papers is strongly correlated with tenure for both men and women. It is associated with a 4.5% increase in tenure probability, regardless of gender. Importantly, women are not penalized for coauthoring. The coefficient on the female/total-coauthored papers interaction term is insignificant. Because of the small sample, however, the results are quite noisy. Future work will expand this sample and include other faculties that use different authorship conventions.

6 Conclusion

While the results presented in this paper are correlations, they provide suggestive evidence that gender bias exists in academic promotion decisions. The bias enters when workers send unclear signals (coauthored papers) that require some judgment on the part of the employer as to which worker made the greatest contribution. The data are not in line with a traditional model of statistical discrimination in which workers know their ability and anticipate employer discrimination, as this model predicts that high ability women will solo author in order to reveal their ability. The results are more in line with a model in which workers do not know their ability, do not anticipate employer discrimination, or where the benefits of coauthoring (costs of solo-authoring) are so high that workers of all abilities will choose to coauthor at some point.

Regardless of the reason, many occupations require group work. The tech industry, for example, prides itself on collaboration. In such male-dominated fields, however, group work in which a single output is produced could sustain the leaky pipeline if employers rely on stereotypes to attribute credit.

I also studied a profession in which individuals can choose to collaborate. If workers are put in teams and do not have the choice to work on their own, the model's predictions are amplified. Employers will rely primarily on their priors and women will be promoted at even lower rates. Bias, whether conscious or subconscious, can therefore have significant implications for the gender gap

in promotion decisions.

References

- [1] Aigner, Dennis J. and Glen C. Cain. 1977. "Statistical Theories of Discrimination in Labor Markets," *Industrial and Labor Relations Review*, 30(2): 175-187.
- [2] 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.
- [3] Bjerck, David. 2008. "Glass Ceilings or Sticky Floors? Statistical Discrimination in a Dynamic Model of Hiring and Promotion," *Economic Journal*, 118(530): 961-982.
- [4] Ceci, Stephen, Donna Ginther, Shulamit Kahn, and Wendy Williams. 2014. "Women in Academic Science: A Changing Landscape," *Psychological Science in the Public Interest*, 15(3): 75-141.
- [5] Coate, Stephen and Glenn Loury. 1993. "Will Affirmative Action Policies Eliminate Negative Stereotypes?" *American Economic Review*, 83(5): 1220-1240.
- [6] Fryer, Roland G. Jr. 2007. "Belief Flipping in a Dynamic Model of Statistical Discrimination." *Journal of Public Economics*, 91(5-6): 1151-1166.
- [7] 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.
- [8] Lehmann, Jee-Yeon. 2013. "Job Assignment and Promotion Under Statistical Discrimination: Evidence from the Early Careers of Lawyers." working paper.
- [9] Niederle, Muriele and Lise Vesterlund. 2007. "Do Women Shy Away From Competition? Do Men Compete Too Much?" *Quarterly Journal of Economics*, 122(3): 1067-1101.
- [10] Pager, Devah. 2003. "The Mark of a Criminal Record." *American Journal of Sociology*, 108(5): 937-975.

Figures

Figure 1: Relationship between composition of papers and tenure

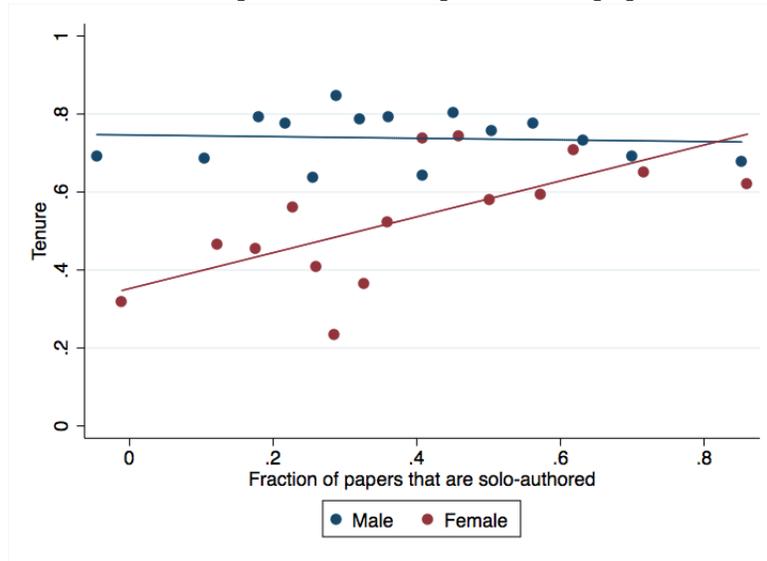


Figure 2: Employer Updating about Female Workers

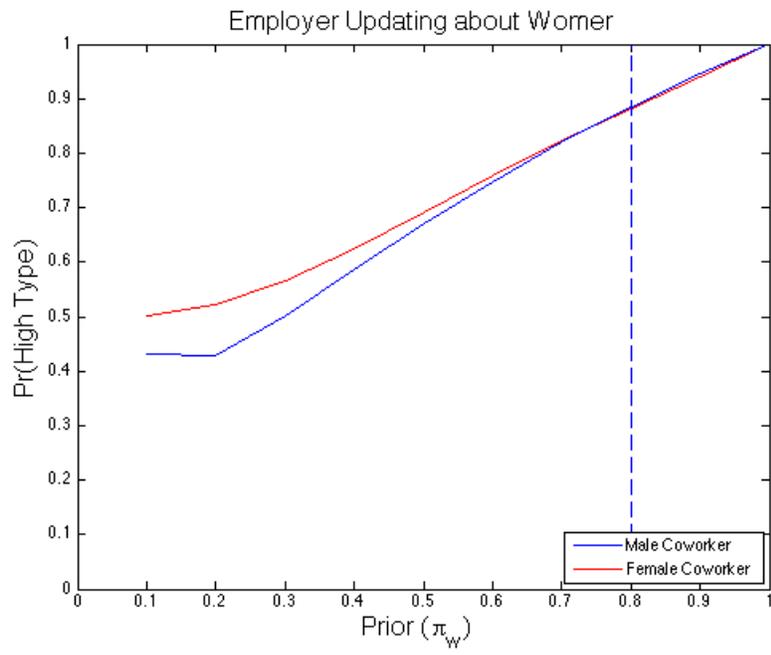


Figure 3: Relationship between Number of Publications and Tenure

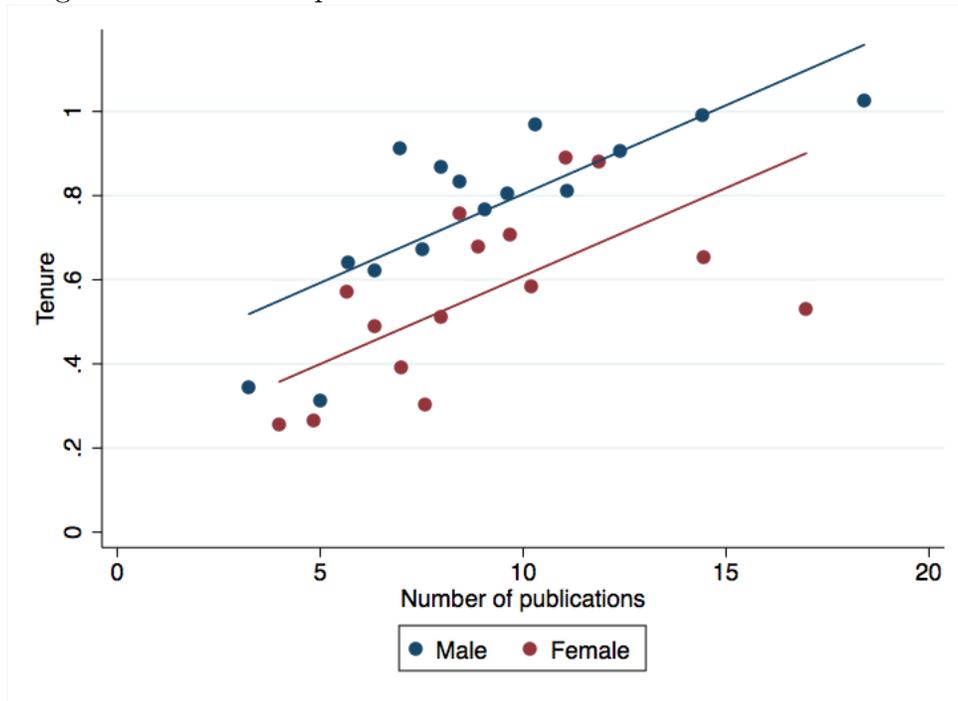


Figure 4: Relationship between Number of Solo-Authored Publications and Tenure

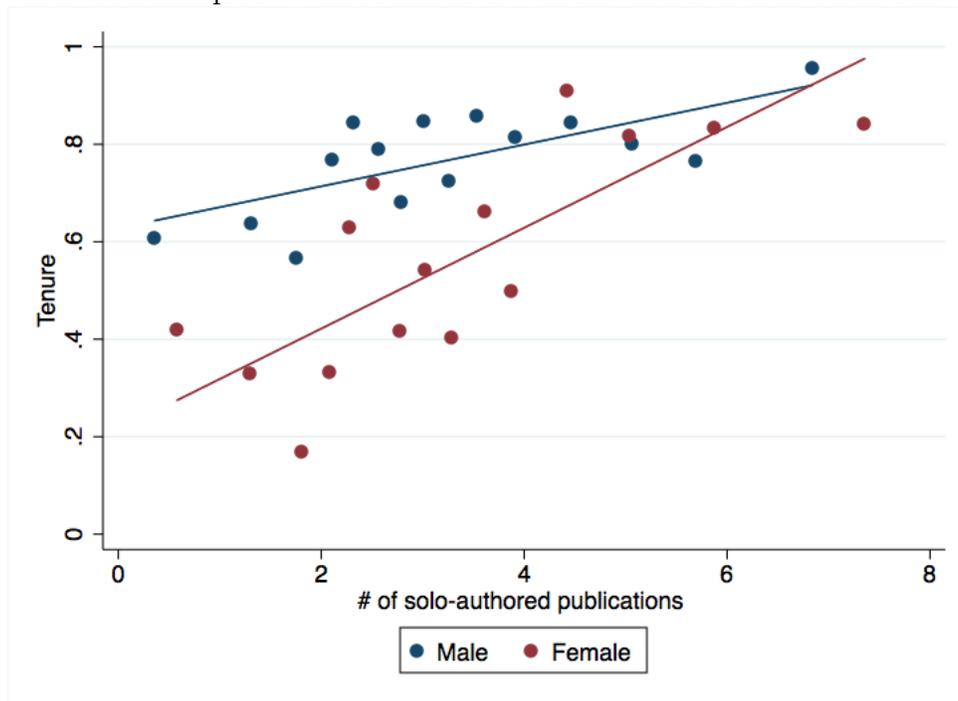


Figure 5: Relationship between number of coauthored publications and tenure

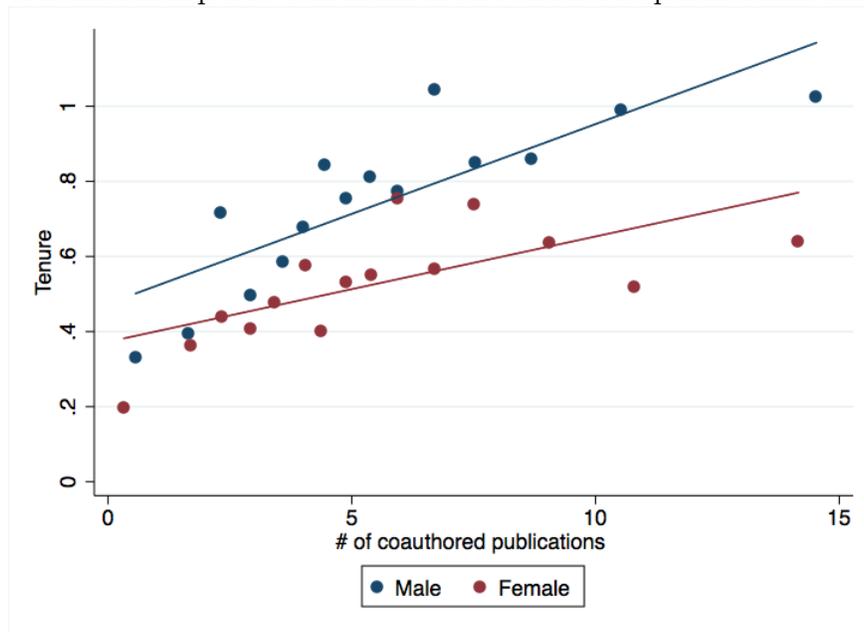
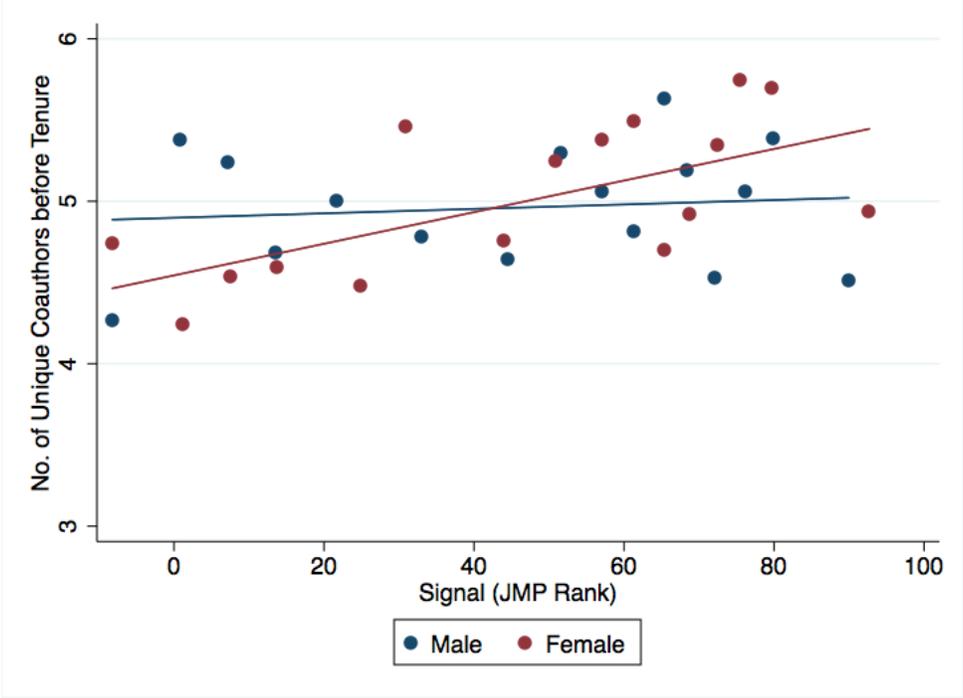


Figure 6: Relationship between Ability Signal and Number of Coauthors



Tables

Table 1: Summary Statistics

	Full	Male	Female	p-value
Tenure	0.70 (0.46)	0.77 (0.42)	0.50 (0.50)	0.001
Total papers	8.6 (4.0)	8.6 (4.1)	8.0 (3.5)	0.061
Solo-authored	3.1 (2.4)	3.1 (2.4)	3.0 (2.3)	0.758
Coauthored	5.5 (3.7)	5.7 (3.8)	5.0 (3.3)	0.066
Years to tenure	6.7 (1.8)	6.6 (1.8)	7.0 (1.9)	0.022
<i>Avg. Journal Rank:</i>				
All Pubs.	47.7 (19.4)	48.6 (19.7)	44.7 18.1	0.050
Solo Pubs.	46.8 (24.1)	47.1 (24.7)	45.6 (22.5)	0.543
Coauthored Pubs.	45.5 (22.5)	46.8 (22.9)	41.3 (20.6)	0.018
Observations	542	415	127	

Table 2: Comparison of Means

	Male	Female	p-value
Tenure	0.760 (0.428)	0.487 (0.502)	0.001
Total papers	8.8 (4.2)	8.0 (3.6)	0.029
Coauthored	5.7 (3.9)	5.0 (3.3)	0.075
Solo-authored	3.1 (2.4)	3.0 (2.3)	0.651
Years to tenure	6.7 (1.7)	7.0 (1.9)	0.065
<i>Avg. Journal Rank:</i>			
All Pubs.	49.3 (19.2)	45.3 (17.9)	0.031
Solo Pubs.	47.9 (24.9)	45.3 (22.7)	0.327
Coauthored Pubs.	47.5 (22.7)	41.9 (20.7)	0.017
Observations	363	123	

Table 3: Number of Papers and Tenure

	(1)	(2)	(3)
	Probit	Probit	Probit
Total papers	0.060*** (0.006)		
Solo-authored		0.072*** (0.012)	0.074*** (0.011)
Fem x Solo		0.019 (0.019)	0.016 (0.017)
Coauthored		0.083*** (0.008)	0.085*** (0.009)
Fem x Coauthored		-0.055*** (0.016)	-0.057*** (0.015)
Female	-0.179*** (0.037)	-0.000 (0.111)	0.026 (0.107)
Observations	534	537	534

All regressions control for avg. journal rank and time to tenure

Table 4: Information from Paper Type

	(1)	(2)	(3)	(4)
	Male	Female	Male	Female
Total papers	0.056*** (0.008)	0.095*** (0.024)		
Solo-authored			0.052*** (0.011)	0.271*** (0.065)
Coauthored			0.057*** (0.008)	0.077** (0.025)
Years to tenure	-0.029** (0.010)	-0.092** (0.035)	-0.029** (0.010)	-0.091* (0.036)
Observations	408	114	408	114
r-squared	0.378	0.336	0.377	0.480

All regressions control for avg. journal rank

Table 5: Coauthor gender and tenure

	(1) Probit	
		x Female
Solo-authored	0.061*** (0.008)	0.016 (0.014)
CA with only fem CAs	0.061*** (0.015)	0.029 (0.026)
CA with only male CAs	0.070*** (0.009)	-0.067*** (0.019)
CA with male and fem CAs	0.091** (0.028)	-0.046 (0.035)
Female	0.014 (0.098)	
Observations	533	

All regressions control for avg. journal rank, time to tenure, school and year FEs

Table 6: Anticipating Discrimination

	(1)	(2)	(3)	(4)
Ability	0.000 (0.000)	0.000 (0.000)	0.000 (0.000)	0.236 (0.233)
Fem. x Ability	0.000 (0.001)	0.000 (0.001)	0.000 (0.001)	-0.202 (0.235)
Female	-0.037 (0.048)	-0.041 (0.050)	-0.044 (0.052)	2.615 (6.520)
Total papers	0.013*** (0.003)	0.014*** (0.003)	0.015*** (0.003)	-3.464 (3.401)
Tenure			-0.019 (0.034)	19.086 (20.132)
School FE	Yes	Yes	Yes	Yes
Tenure Year FE	Yes	Yes	Yes	Yes
Field FE	No	Yes	Yes	Yes
Observations	537	534	534	533

We proxy for ability using the journal ranking that an individual's job market paper was published in. Depending variable is the fraction of an individual's papers that are coauthored.

Table 7: Senior Coauthors and Tenure

	(1)	(2)
Solo-authored	0.067*** (0.010)	0.062*** (0.010)
Fem x Solo	0.012 (0.018)	0.019 (0.016)
Coauthored	0.089*** (0.009)	0.094*** (0.009)
Fem x Coauthored	-0.058*** (0.015)	-0.075*** (0.013)
No. of Sr. CAs	-0.024* (0.010)	-0.039*** (0.009)
Fem. x Sr. CAs		0.049** (0.018)
Female	0.061 (0.118)	0.051 (0.108)
Observations	511	511

All regressions control for journal rank. Senior coauthors are any coauthors who are full professors at time of tenure.

Table 8: Comparison of Means

	Men	Women	p-value
Tenure	0.75 (0.435)	0.77 (0.425)	0.633
Total Papers	12.2 (8.1)	10.5 (5.7)	0.089
Total Coauthored	6.6 (6.6)	6.1 (5.0)	0.594
Total Solo	5.6 (4.5)	4.4 (2.9)	0.022
Length of Time to Tenure	7.5 (1.6)	7.6 (1.7)	0.764
Observations	133	88	

Table 9: Sociology: Publications and Tenure

	(1)	(2)
Total first author	0.045*	
	(0.021)	
Fem. x First Author	0.018	
	(0.043)	
Fraction first author		0.451***
		(0.081)
Fem. x Frac. First Author		-0.178
		(0.181)
Solo papers	0.007	-0.000
	(0.006)	(0.007)
Fem. x Total Solo	0.000	0.002
	(0.010)	(0.009)
Total coauthored	-0.009	0.011
	(0.005)	(0.008)
Fem. x Total CA	-0.013	-0.001
	(0.018)	(0.015)
Books	0.052	0.040
	(0.046)	(0.054)
Book chapters	0.014	0.011
	(0.014)	(0.012)
Female	0.019	0.134
	(0.140)	(0.162)
Observations	205	180

All regressions control for time to tenure.