Gender Differences in Recognition for Group Work

Heather Sarsons*

September 27, 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NOTE: Very preliminary – please do not cite without permission.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. 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 hiring[1] 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

[1] For 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. I use the data to show that the promotion gap appears when workers work in groups. I present a model of discrimination that allows for group work and show that the trends we see in the data are inconsistent with both statistical discrimination and taste-based discrimination, suggesting that some other form of bias is at play.

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

There are two main types of discrimination that economists consider: taste-based and statistical. Taste-based discrimination assumes that employers have a distaste for members of a certain group and therefore will not hire or promote them regardless of their skill. Statistical discrimination assumes that employers have priors over the average ability of each group. If they believe that one group is less skilled, they will hold individuals from that group to a higher standard and require them to invest in more skills or education. Anticipating this, workers from the discriminated-against group choose not make such a costly investment since their chance of being hired is low. This reinforced employers' beliefs.

In these models of discrimination, employees work alone and employers make decisions based on signals informative about one individual[2]. Allowing for group work complicates the employer's ability to infer an employee's ability from a signal. Figure 1 suggests that employers make different inferences about a man and a woman's type when employees work jointly. To test whether discrimination can explain Figure 1, I alter Coate and Loury's (1993) model of statistical discrimination to allow for group work and test its predictions with the data. Statistical discrimination models typically assume that workers are fully rational and know the hiring and promotion standards that employers hold them two. I consider 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. I will also describe the predictions that come out of a model of taste-based discrimination.

2.1 Basic Setup

The model begins with a set of workers who have already been hired. 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 the 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 p with associated cost c_p which is drawn from a 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, θ_c, about $j's$ probability of being a high type.

3. Worker i decides whether to work alone or collaborate. Collaborating reduces the cost of production, described in more detail below.

4. Workers complete their projects (either alone or with another worker) and send a signal to the employer, θ_e.

[2]For example, see Aigner and Cain (1977), Bjerk (2008), Coate and Loury (1993), and Fryer (2007).

5. Employers make promotion decisions.

2.1.1 Costs

Workers draw a cost associated with their project but the realized cost depends on whether they work independently or with a partner. Workers who work independently pay the full project cost, c_p. Collaborating lowers the cost of production to 0.

2.1.2 Signals

First, workers receive a signal from their potential collaborator. Recall that all workers are either high or low ability. Collaborators who are high types draw their signal from the distribution with CDF $\theta_c \sim F_H(\theta)$. Collaborators who are low types draw from $\theta_c \sim F_L(\theta)$. It is assumed that $F_H(\theta) \leq F_L(\theta)\ \forall\theta \in [0, 1]$ so that high types are more likely to draw high signals. Employers do not see θ_c.

After deciding whether to collaborate, the workers complete their projects and send a signal, θ_e, to the employer. This signal is drawn from the same distributions that θ_c is drawn from. A worker who works alone draws a signal θ_e from $F_H(\theta)$ if she is a high type and from $F_L(\theta)$ if she is a low type. If a worker chooses to collaborate, the signal she sends depends on her and her collaborator's types. If both are high types, they draw a signal from $F_H(\theta)$. If they are both low types, they draw from $F_L(\theta)$. If one is a low type and one is a high type, they draw from $F_H(\theta)$ with probability γ and from $F_L(\theta)$ with probability $1 - \gamma$.

To summarize, workers receive full information about the cost of the project, c_p, and a signal about the collaborator's type, θ_c. Workers decide whether to collaborate based on c_p, θ_c, and π_g. The employer then receives a signal, θ_e, 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_e, π, and the collaboration decision.

2.1.3 Payoffs

Workers who are promoted receive wage w while those who are not promoted receive 0. A worker who collaborates and is promoted has a total payoff of w while a collaborating worker who is not promoted has a total payoff of 0. A worker who works alone and is promoted has total payoff $w - c$ and one who is not promoted has total payoff $-c$.

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 Employer's Decision

The model is solved working backwards from the employer's promotion decision. The employer observes the worker's signal and collaboration choice and sets a cutoff rule that he uses to make promotion decisions.

2.2.1 Deciding whether to promote a solo worker

The employer wants to promote all high ability workers without promoting any low ability workers. When an employer sees a signal from a worker who works alone, he updates his beliefs about the worker's type according to Bayes' rule:

$$\beta_s(\theta_s) \equiv \mathbb{P}\left(a_i = H|\theta_e, \pi_g, S\right) = \frac{\pi_i f_H(\theta_e)\mathbb{P}(S|H)}{\pi_i f_H(\theta_e)\mathbb{P}(S|H) + (1-\pi_i) f_L(\theta_e)\mathbb{P}(S|L)}.$$

Here, $\mathbb{P}(S|H)$ is the probability that the worker would choose to work alone when she is a high type. This term is defined in the worker's program and depends on priors and the signals workers receive from one another.

The employer will promote a worker if the expected payoff from doing so is greater than the wage the employer must pay the worker. That is, the employer will promote the worker if

$$\begin{aligned} \chi_H \beta_s(\theta_e) + \chi_L(1-\beta_s(\theta_e)) &\geq w \\ \beta_s(\theta_e) &\geq \frac{w - \chi_L}{\chi_H - \chi_L} \end{aligned}$$

which defines a cutoff $\theta^*_{e,solo}$ at which the employer is indifferent between promoting and not promoting the worker. The cutoff depends on the priors, π_g and the belief that a high ability worker would choose to work alone, $\mathbb{P}(S|H)$. If the employer believes that a large fraction of workers from group g are high ability (π_g is close to 1), he will not need to set a high threshold so

$\theta^*_{e,solo}$ will fall. Similarly, if high ability workers are likely to work alone ($\mathbb{P}(S|H)$ is close to 1), the employer will believe that the worker is a high type regardless of the signal θ_e so the cutoff will fall.

The employer will promote any solo worker who sends a signal greater than $\theta^*_{e,solo}$ and not promote workers who send signals below this cutoff. Note that since $\pi_m > \pi_w$, women who work alone will be tenured at a lower rate than men who work alone. However, if the game is extended to multiple periods so that workers can send many signals, the signals start to outweigh the employer's prior. Each additional high solo signal will bring a woman's chance of promotion closer to that of a man.

Prediction 1: After one solo signal, women will be promoted at lower rates than men (provided that $\pi_m > \pi_w$) but additional "high" signals from women will start to close the promotion gap.

2.2.2 Collaborating workers

Signals from workers who collaborate contain more information. The employer considers the signal, the worker's decision to collaborate, and both the worker and the collaborator's group identity. I consider the case of a female worker who collaborates and a male worker who collaborates below.

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 types[3]. Upon receiving a signal, θ_e, 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[4]:

$$\begin{aligned}\beta_c(\theta_e) &\equiv \mathbb{P}\left(a_i = H|\theta_e,\, \pi_w, CA\right) \\ &= \frac{\pi_w \mathbb{P}(CA|H_i)\left[\pi_g \mathbb{P}(CA|H_j) f_H(\theta_e) + (1-\pi_g)\mathbb{P}(CA|L_j)\left(\gamma f_H(\theta_e) + (1-\gamma) f_L(\theta_e)\right)\right]}{Total\ Probability}\end{aligned}$$

[3]One 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.

[4]The total probability is $\pi_w \mathbb{P}(CA|H_i)\left[\pi_g \mathbb{P}(CA|H_j) f_H(\theta_e) + (1-\pi_g)\mathbb{P}(CA|L_j)\left(\gamma f_H(\theta_e) + (1-\gamma) f_L(\theta_e)\right)\right] + (1-\pi_w)\mathbb{P}(CA|L)\left[\pi_g \mathbb{P}(CA|H)(\gamma f_H(\theta_e) + (1-\gamma) f_L(\theta_e)) + (1-\pi_g)\mathbb{P}(CA|L) f_L(\theta_e)\right]$

The employer's belief depends directly and indirectly on his priors, π_m and π_w. As his views about women become more favourable, he is more likely to believe that a given woman is high ability. However, the priors also influence workers' willingness to work with one another ($\mathbb{P}(CA|H_i)$ and $\mathbb{P}(CA|H_j)$). If π_f is close to 0, both men and women are very unlikely to collaborate with women. A woman would have to send a very high signal (θ_c) in order to convince someone to work with her since the odds of her being a high type are so low. For example, in the case of coauthoring, if people that most women are "low types", women would have to prove their skills above and beyond what a man would have to do to attract a coauthor. While this may be frustrating for women, it has a positive effect on the employer's belief. If someone agrees to coauthor with a woman, it must have been because she sent a very high signal and is thus likely to be high ability. Therefore, while a low prior works against a woman in that she is expected to be low ability, it works for her if she is able to get a collaborator.

Recall that employers do not receive any information about the worker's collaborator except for the collaborator's gender. Because both men and women require a higher θ_c from women in order to work with him, women who work with other women are worse off when priors are low. This is shown graphically in Figure 2 where beliefs about men, π_m, are held constant and beliefs about women vary. For low π_w, workers will only work with women if they send a high θ_c (and are therefore likely to be high types). This means that the collaborator must have sent the worker a very high signal as well. The employer is actually less likely to believe that the female worker up for promotion is a high type when her collaborator is a woman rather than a man. This is because the male collaborator would not have had to send as high of a signal so his probability of being a high type is lower than a female collaborator's. As π_w increases, female collaborators do not have to send as high of signals and the female worker up for promotion begins to receive more credit.

Prediction 2: For low π_w, women are more likely to be perceived as high types if they collaborate with men. As π_w increases, women are more likely to be high types if they collaborate with other women.

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

$$\beta_c(\theta_e) \equiv \mathbb{P}\left(a_i = H|\theta_e, \pi_m, CA\right)$$

$$= \frac{\pi_m \mathbb{P}(CA|H_i)\left[\pi_g \mathbb{P}(CA|H_j) f_H(\theta_e) + (1-\pi_g)\mathbb{P}(CA|L_j)\left(\gamma f_H(\theta_e) + (1-\gamma) f_L(\theta_e)\right)\right]}{Total\ Probability}$$

The employer's beliefs again change directly and indirectly with π_g. For low π_w, men are better off coauthoring with men. As π_w increases, though, employers' beliefs about women who collaborate will fall and men receive more credit if they coauthor with women.

Overall, an employer will choose to promote a worker who collaborates if

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

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

Since there is uncertainty over who had the idea, put in effort, and so on, the probability that a given collaborating worker is a high type is lower than the probability that a solo worker with the same θ_e is a high type, regardless of gender. Because of this additional uncertainty that is not present when workers work alone, employers are 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.

2.3 Naive Worker's Decision

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 sends a sufficiently high signal, θ_e, regardless of whether they collaborate or work alone. While the employer sets different standards for collaborating workers and for solo workers, naive workers assume that $\theta^*_{e,solo} = \theta^*_{e,collab} = \theta^*_e$. 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 θ_c 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 θ_c. Worker i updates her belief about j's ability according to Bayes' rule:

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

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 alone[5]. This is formalized in equation 1 below where the right-hand side is expected payoff of collaborating and the left-hand side is the payoff from working alone.

$$\begin{aligned}
\mathbb{E}\,(\text{payoff collaborate}) &\geq \mathbb{E}\,(\text{payoff solo}) \\
w\mathbb{P}\,(\theta_e \geq \theta_e^*|collab) &\geq w\mathbb{P}\,(\theta_e \geq \theta_e^*|solo) - c \\
w\left[(1 - F_H(\theta_e^*))\,\varphi + (1-\varphi)\left(\gamma(1 - F_H(\theta_e^*)) + (1-\gamma)(1 - F_L(\theta_e^*))\right)\right] &\geq w(1 - F_H(\theta_e^*)) - c \\
w(1-\gamma)(1-\varphi)\left(F_L(\theta_e^*) - F_H(\theta_e^*)\right) &\leq c \qquad (1)
\end{aligned}$$

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 (φ 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

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

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, θ_c. 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 θ_c^* 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

$$\begin{aligned} w\left[(1-\varphi)\left(1-F_L(\theta_e^*)\right)+\varphi\left(\gamma(1-F_H(\theta_e^*))+(1-\gamma)(1-F_L(\theta_e^*))\right)\right] &\geq w\left(1-F_L(\theta_e^*)\right)-c \\ w\varphi\gamma\left(F_H(\theta_e^*)-F_L(\theta_e^*)\right) &\leq c \end{aligned}$$

   which holds for all positive wages and costs. Low ability workers will therefore always be willing to collaborate. They will only work alone if no other workers are willing to work with them.

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

$$w(1-\gamma)(1-\varphi)\left(F_L(\theta_e^*)-F_H(\theta_e^*)\right)\leq c$$

   which implicitly defines a cutoff $\theta_c^*((\pi_g,\gamma,c))$, below which the high ability worker will choose to work alone. Note that for a given signal $\tilde{\theta}_c$, since $\pi_m > \pi_w$, the cutoff for women will be higher than the cutoff for men: $\theta_{c,w}^* > \theta_{c,m}^*$. 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 hold female workers to a higher standard than male workers.

Prediction 4: Both men and women who are high ability will require a higher signal, θ_c^*, 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. Because workers are naive, the worker's decision influences the employer's cutoff rule but the worker does not realize this. Specifically, the probability that a given worker would coauthor ($\mathbb{P}(CA|H)$ from the employer's problem) is defined as

$$\begin{aligned} \mathbb{P}(CA|H) &= \mathbb{P}\left(w(1-\gamma)(1-\varphi)(F_L(\theta_e^*) - F_H(\theta_e^*) \leq c\right) \\ \mathbb{P}(CA|L) &= 1 \end{aligned}$$

where all low types would like to collaborate and high types will collaborate under certain conditions.

2.4 Informed Worker's Decision

If workers know that employers take the decision to collaborate as an additional signal of ability, they know that $\theta^*_{e,solo} \neq \theta^*_{e,collab}$ and they will collaborate strategically. Specifically, the worker now chooses to coauthor if

$$\begin{aligned} w\mathbb{P}\left(\theta_{e,c} \geq \theta^*_{e,c}|collab\right) &\geq w\mathbb{P}\left(\theta_{e,s} \geq \theta^*_{e,s}|solo\right) - c_i \\ w\left[(1-\varphi)\left(1-F_L(\theta_e^*)\right) + \varphi\left(\gamma(1-F_H(\theta_e^*)) + (1-\gamma)(1-F_L(\theta_e^*))\right)\right] &\geq w\left(1-F_L(\theta_e^*)\right) - c \end{aligned}$$

where $\theta^*_{e,collab} > \theta^*_{e,solo}$.

Since workers now know there is some probability that employers will attribute credit to the coworker, workers are less likely to collaborate than in the naive case. They are held to a higher standard and are less likely to be promoted. High ability workers in particular are better off working alone than collaborating, leading to the following prediction:

Prediction 5: 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.

2.5 Taste-based Discrimination

The above model does not speak to taste-based discrimination but the predictions of such a model are straightforward. With taste-based discrimination, employers have a distaste in hiring or promoting workers from a particular group. In a simple world with taste-based discrimination, employers would not promote women regardless of how well they perform or whom they work with. However, if employers face potential lawsuits from failing to promote qualified women, they might promote all high-performing women who work alone and not promote any women who work in a group. Employers can not dispute the qualifications of a woman who works alone but they can argue that the output from women who work in a group is due to the other group members. In this case, women who collaborate should never be promoted, regardless of whom they work with and their output.

Prediction 6: Under taste-based discrimination, either no women will be promoted or all women who collaborate will not be promoted.

3 Data

The sample consists of economists who went up for tenure between 1975 and 2014 in one of the top 30 PhD-granting universities[6] in the United States. To account for people who went up for tenure, were denied it, and moved into industry, non-US schools, or non-top30 schools, 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 a sample of 552 economists.

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

[6]Ranking 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.

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. 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 80 journals and give the top journal a score of 80. 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 52% of women receive tenure while 77% of men do. There is no statistically significant difference in the number of papers that men and women produce although men do tend to publish in slightly better journals. If women are tenured at lower rates because of such productivity differences, controlling for the number and rank of publications should explain the tenure gap. The remainder of the paper explores the tenure gap and tests the predictions from the model.

4 Empirical Strategy and Results

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 5.7% 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 OLD regression lines in Figure 3 are plotted by estimating

$$T_{ifst} = \beta_1 TotPapers_i + \beta_2 fem_i + \gamma' Z_i + \theta_f + \theta_s + \theta_t + \epsilon_{ifst} \tag{2}$$

separately for men and women. The dependent variable, T_{ifst}, is the probability that individual i in field f at school s in year t receives tenure. $TotPapers_i$ is the number of papers individual i had at the time he or she went up for tenure and fem_i indicates gender. The vector of individual-level controls, Z_i, includes average journal rank and the number of years it took the person to go up for tenure. Finally, I include tenure institution, year of tenure, and field fixed effects as we might expect tenure standards to vary over time and by field and department.

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. Prediction 1 states that 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 using the estimates from

$$\begin{aligned} T_{ifst} \;=\; & \beta_1 S_i + \beta_2 (fem_i \times S_i) + \beta_3 CA_i + \beta_4 (fem_i \times CA_i) + \beta_5 fem_i \\ & + \gamma' Z_i + \theta_f + \theta_s + \theta_t + \epsilon_{ifst} \end{aligned} \tag{3}$$

The coefficient on S_i is plotted separately for men and women after controlling for an individual's number of coauthored papers (CA_i), and individual and school-level controls mentioned above. Table 2 presents the full results from this estimation using a probit model. The results are in line with the model's predictions: women with few solo-authored papers have a low chance of receiving tenure but the tenure gap narrows as the signal from the solo papers begins to outweigh the employer's prior.

The model also 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 as long as π_w is not too low. 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 the prediction that if π_w is sufficiently high but still lower than π_m, women will receive less credit for group work.

Looking at the size of the coefficients in Table 2, though, an additional coauthored paper for a man has the same effect on tenure as a solo-authored paper. An additional solo-authored paper is associated with a 7.3% increase in tenure probability and an additional coauthored paper is associated with an 8% increase. This is at odds with Prediction 3 which states that high ability men and women who collaborate should be promoted at lower rates than those who solo-author (holding productivity constant).

While the results for women fit with a statistical discrimination model, the results for men do not. It appears that employers do not take the decision to coauthor as a signal for men but do for women. Employers could be practicing taste-based discrimination which leads to the 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

If employers have a distaste for promoting women, women will be denied tenure regardless of whom they coauthor with (men or women). To test for taste-based discrimination, I separate the number of coauthored papers an individual has with women and with men and estimate

$$\begin{aligned} T_{ifst} &= \beta_1 S_i + \beta_2 (fem_i \times S_i) + \beta_3 CAfem_i + \beta_4 (fem_i \times CAfem_i) + \beta_5 CAmale_i \\ &\quad + \beta_6 (fem \times CAmale_i) + \beta_7 CAmix_i + \beta_8 (fem \times CAmix_i) + \beta_9 fem_i \\ &\qquad + \gamma' Z_i + \theta_f + \theta_s + \theta_t + \epsilon_{ifst}. \end{aligned} \tag{4}$$

As before, S_i is the number of solo-authored and coauthored papers individual i has. $CAfem_i$ is the number of coauthored papers and individual has in which all of the coauthors are female. Similarly, $CAmale_i$ is the number of papers individual i has in which all of the coauthors are male and $CAmix_i$ is the number of papers an individual has in which the coauthors consist of men and women. The results in Table 3 show that the coauthoring penalty is almost entirely driven from coauthoring with men. An additional coauthored paper with a man has zero marginal effect on tenure. Papers in which there is at least one other woman ($CAmix$) have a smaller effect on tenure for women than for men (8% vs. 3.5%) but still have a positive marginal impact. Papers with only women are also positively associated with tenure and there is no statistical difference between the association for men and women, due in part to noise[7].

The results suggest that taste-based discrimination is not at play as women are treated differently based on their coauthors' genders. If an employer simply did not like women, no women who coauthor would be promoted which is not the case here.

Overall, the trends we see in the data are not in line with a model of statistical discrimination or taste-based discrimination. Some other form of bias could be at play. For example, 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 group[8]. When women collaborate, however, how much and what the woman contributed comes into question. It could also be that women select coauthors who have already established themselves, such as senior faculty, which leads the employer to believe the senior person put in the most effort or had the idea for the project. I now turn to some of these alternative explanations.

4.2 Coauthor Selection

4.2.1 Do women anticipate discrimination?

If women know that employers statistically discriminate in the face of a noisy signal, high ability women are better off solo-authoring as this clearly reveals their type. Testing whether women anticipate discrimination is difficult as there is no clear measure of a person's ability. I proxy for

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

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

ability using the quality of journal that an individual's job market paper was published in. If women anticipate discrimination, there should be a negative correlation between ability and the fraction of one's paper that are coauthored. To test this I estimate

$$FracCA_{ifst} = \beta_1 abil_i + \beta_2(fem_i \times abil_i) + \beta_3 fem_i + \beta_4 TotPapers_i + \beta_5 T_i + \theta_f + \theta_s + \theta_t + \epsilon_{ifst} \quad (5)$$

where $FracCA_{ifst}$ 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 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 4). 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. I also find no evidence that high ability women strategically coauthor with other women rather than men.

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 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 5 checks whether women are more likely to coauthor with senior professors. Each specification shows evidence that women are less likely to coauthor with senior faculty, although the difference between men and women in insignificant. Figure 6 plots the relationship between paper composition and tenure now controlling for the fraction of an individual's papers that have senior coauthors. The results do not substantially change. Coauthor selection along seniority lines therefore does not appears to be driving the results.

5 Clear signals: Testing against other coauthoring conventions

Employers may exhibit bias when evaluating women who send 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 250 sociologists in the sample and 40% are female. Table 6 presents sample statistics: tenure rates are comparable for men and women and men tend to produce more solo-authored papers than women.

I test whether men and women are treated differently when they coauthor papers in Table 7. 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 roughly 4% 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. Women do not seem to coauthor strategically and employers do not treat coauthored papers as noisy signals for men. The results are more in line with a model in which workers do not know their ability or do not anticipate employer discrimination, and where employers update on signals differently for men and women.

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] Bjerk, 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

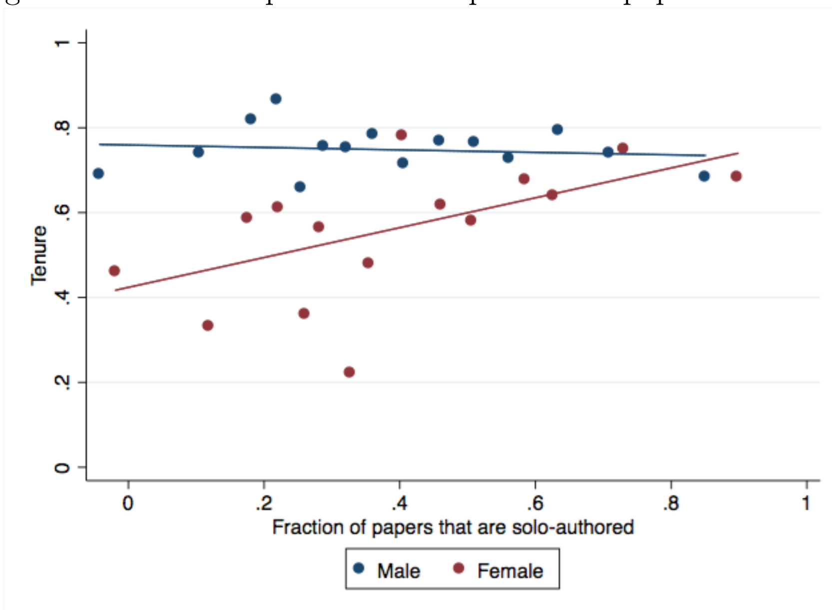

Figure 2: Employer's Updating about Women as π_w Changes

Figure 3: Relationship between Number of Publications and Tenure

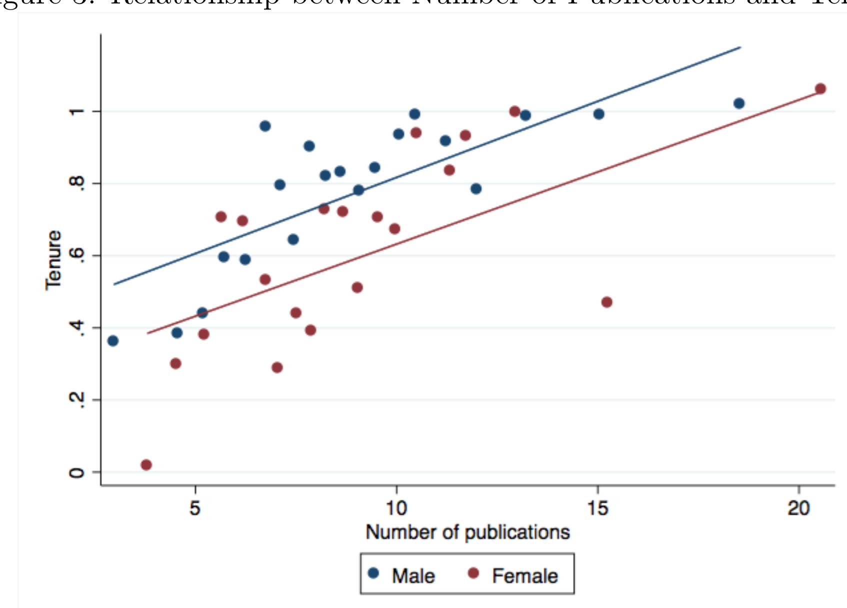

Figure 4: Relationship between Number of Solo-Authored Publications and Tenure

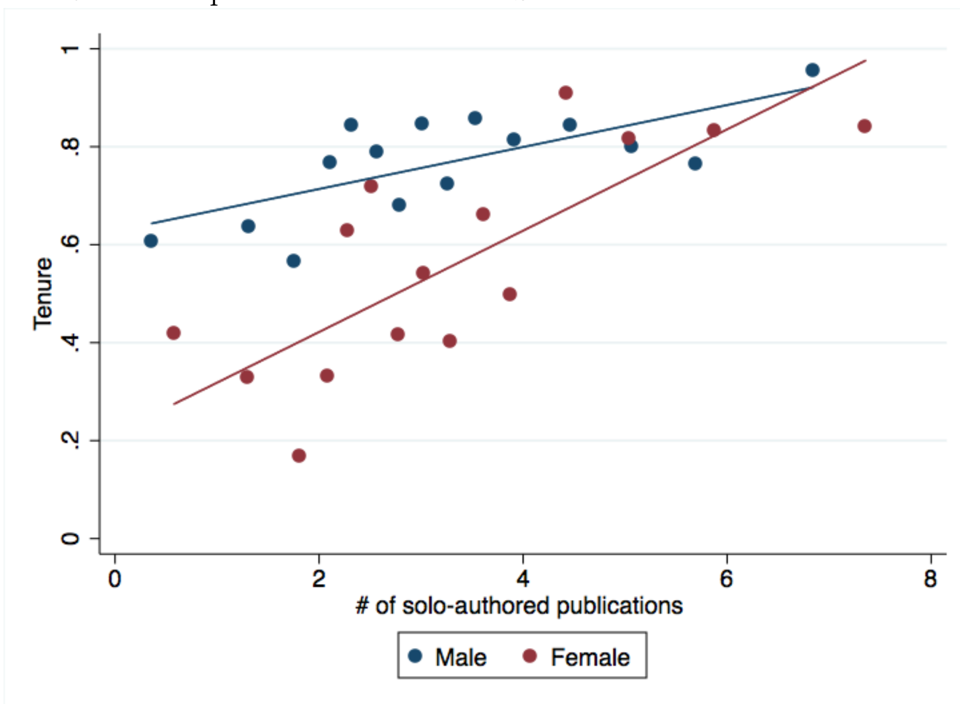

Figure 5: Relationship between number of coauthored publications and tenure

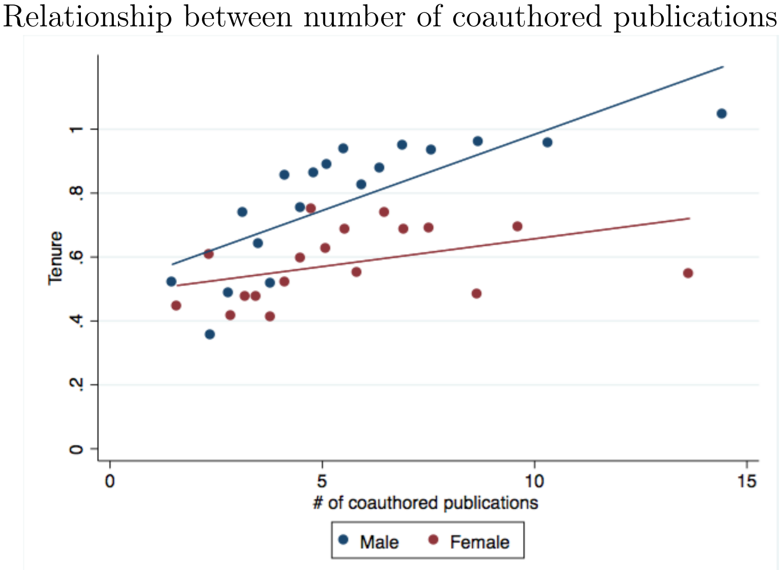

Figure 6: Controlling for Coauthor Seniority

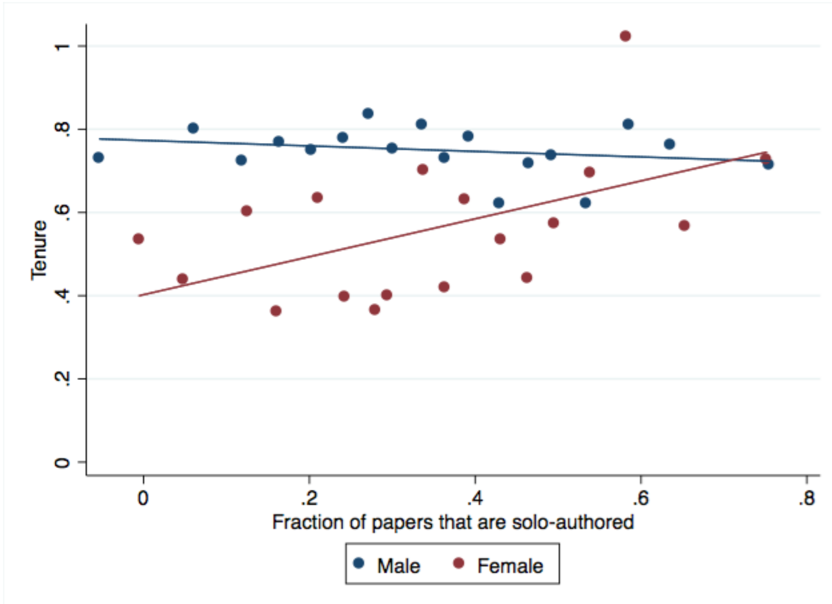

Tables

Table 1: Summary Statistics

	Full	Male	Female	p-value
Tenure	0.71	0.77	0.52	0.001
	(0.45)	(0.42)	(0.50)	
Total papers	8.6	8.7	8.2	0.164
	(4.0)	(4.1)	(3.4)	
Solo-authored	3.1	3.1	3.1	0.997
	(2.4)	(2.4)	(2.3)	
Coauthored	5.5	5.7	5.1	0.135
	(3.7)	(3.8)	(3.3)	
Years to tenure	6.7	6.6	7.1	0.021
	(1.9)	(1.9)	(1.9)	
Avg. Journal Rank:				
All Pubs.	45.5	46.3	42.6	0.048
	(18.8)	(19.1)	(17.3)	
Solo Pubs.	46.4	47.0	44.4	0.300
	(24.2)	(24.5)	(22.9)	
Coauthored Pubs.	45.6	46.6	42.2	0.054
	(22.3)	(22.8)	(20.1‘)	
Observations	552	422	130	

Table 2: Number of Papers and Tenure

	(1) Probit	(2) Probit	(3) Probit	(4) Probit
Total papers	0.057***	0.058***		
	(0.005)	(0.006)		
Solo-authored			0.073***	0.075***
			(0.011)	(0.010)
Fem x Solo			0.014	0.010
			(0.018)	(0.016)
Coauthored			0.080***	0.082***
			(0.008)	(0.009)
Fem x Coauthored			-0.055***	-0.059***
			(0.015)	(0.014)
Female	-0.183***	-0.175***	0.022	0.063
	(0.034)	(0.036)	(0.108)	(0.103)
School FE	Yes	Yes	Yes	Yes
Tenure Year FE	Yes	Yes	Yes	Yes
Field FE	No	Yes	No	Yes
Observations	547	544	547	544

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

Table 3: Coauthor gender and tenure

	(1)	
	Probit	
		x Female
Solo-authored	0.063***	0.009
	(0.008)	(0.015)
CA with only fem CAs	0.062***	0.024
	(0.017)	(0.027)
CA with only male CAs	0.068***	-0.068***
	(0.009)	(0.018)
CA with m and f CAs	0.080**	-0.045**
	(0.028)	(0.036)
Female	0.049	
	(0.099)	
Observations	542	

Regression controls for journal rank, time to tenure, school, year, and field FEs

Table 4: Anticipating Discrimination

	(1)	(2)	(3)
Ability	0.000	0.000	0.000
	(0.000)	(0.000)	(0.000)
Fem. x Ability	0.000	0.000	0.000
	(0.001)	(0.001)	(0.001)
Female	-0.039	-0.042	-0.046
	(0.047)	(0.048)	(0.050)
Total papers	0.012***	0.013***	0.014***
	(0.003)	(0.003)	(0.003)
Tenure			-0.026
			(0.033)
School FE	Yes	Yes	Yes
Tenure Year FE	Yes	Yes	Yes
Field FE	No	Yes	Yes
Observations	546	543	543

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 5: Number of Senior Coauthors

	(1)	(2)	(3)
Female	-0.140	-0.078	-0.049
	(0.269)	(0.238)	(0.229)
Total coauthors		0.198***	0.210*
		(0.043)	(0.094)
Years to tenure		-0.081	-0.061
		(0.053)	(0.078)
Coauthored		0.053	0.015
		(0.058)	(0.120)
Solo-authored		-0.174***	-0.136***
		(0.043)	(0.037)
School FE	No	No	Yes
Tenure Year FE	No	No	Yes
Observations	527	527	522

Sample: Individuals with at least one coauthor

Table 6: Comparison of Means

	Men	Women	p-value
Tenure	0.75	0.78	0.547
	(0.44)	(0.42)	
Total Papers	12.2	10.2	0.033
	(7.8)	(5.7)	
Total Coauthored	6.4	6.0	0.567
	(6.6)	(5.0)	
Total Solo	5.7	4.2	0.003
	(4.5)	(2.9)	
Length of Time to Tenure	7.6	7.5	0.686
	(1.6)	(1.7)	
Observations	150	100	

Table 7: Sociology: Publications and Tenure

	(1)	(2)	(3)
Total first author	0.050**		0.040*
	(0.017)		(0.016)
Fem. x First Author	0.026		0.006
	(0.040)		(0.028)
Fraction first author		0.403***	
		(0.043)	
Fem. x Frac. First Author		-0.042	
		(0.172)	
Solo papers	0.008	0.000	0.008
	(0.006)	(0.006)	(0.006)
Fem. x Total Solo	0.002	0.007	0.002
	(0.011)	(0.011)	(0.011)
Total coauthored	-0.010*	0.009	
	(0.004)	(0.007)	
Fem. x Total CA	-0.020	0.001	
	(0.017)	(0.015)	
Books	0.063*	0.058	0.063*
	(0.032)	(0.035)	(0.032)
Book chapters	0.007	0.005	0.007
	(0.013)	(0.012)	(0.013)
Female	0.026	0.010	0.026
	(0.114)	(0.163)	(0.114)
Observations	237	209	237

All regressions control for time to tenure.