Gender and Promotions: Evidence from Academic Economists in France *

Clément Bosquet† Pierre-Philippe Combes‡
Cecilia García-Peñalosa§

February 15, 2018

Abstract

The promotion system for French academic economists provides an interesting environment to examine the promotion gap between men and women. Promotions occur through national competitions for which we have information both on candidates and on those eligible to be candidates. We can then examine the two stages of the process: application and success. Women are less likely to seek promotion and this accounts for up to 76% of the promotion gap. Being a woman also reduces the probability of promotion conditional on applying, although the gender difference is not statistically significant. Our results highlight the importance of the decision to apply.

JEL Classification: J16, J7, I23

Key words: gender gaps, professional advancement, competition

*We are grateful to Esteban Aucejo, Stefano Bosi, Rodolphe Dos Santos Ferreira, Florence Goffette-Nagot, Marco Manacorda, Vélayoudom Marimoutou, Michel Mougeot, Javier Ortega, Boris Racine, Marie-Claire Villeval, Felix Weinhardt, three referees, as well as seminar and conference participants. We thank Laurent Linnemer for earlier work on a smaller version of the data set for the *agrégation*. This project has been supported by the Agence Nationale de la Recherche, grant ANR-11-BSH1-0014.

†ThEMA, UMR 8184, Université Cergy-Pontoise, CNRS, F-95000 Cergy-Pontoise and CEP (Urban and Spatial programme, LSE). 33 Boulevard du Port, 95011 Cergy-Pontoise Cedex, France. Email: clement.bosquet@u-cergy.fr. Website: https://sites.google.com/site/clementbosquet/

‡Univ Lyon, CNRS, GATE Lyon Saint-Etienne UMR 5824, 93 chemin des Mouilles, F-69131 Ecully, France; Sciences Po, Department of Economics, 28, Rue des Saints-Pères, 75007 Paris, France. Also, research fellow at CEPR. ppcombes@gmail.com. Website: https://www.gate.cnrs.fr/ppcombes/

§Aix-Marseille University, CNRS, EHESS, Centrale Marseille, AMSE. 5 Bd Maurice Bourdet, 13001 Marseille, France. Also research fellow at CESifo. Email: cecilia.garcia-penalosa@univ-amu.fr. Website: http://www.vcharite.univ-mrs.fr/PP/penalosa/

I Introduction

Despite the rapid increase in female educational attainment over the last decades, the labour market outcomes of men and women still differ in terms of wages and seniority. The literature on gender wage gaps is vast, but differences across genders in promotions have received much less attention. Yet these gaps can be large. For example, Bertrand and Hallock (2001) find that women account for only 2.5% of top executives in US firms and for France, Gobillon et al. (2015) show that the gender gap in the probability of being an executive increases along the wage ladder from 9% to 50%. Existing work has shown that differences in characteristics account for a large proportion of the gap, with the remaining fraction being usually attributed to preferences or discrimination (for an overview of work on gender wage gaps see Blau and Kahn, 2000). Measuring the role played by these two aspects has proven difficult. This paper uses the particular features of the French academic system, namely the fact that promotions occur through national contests involving successive steps, to look at various possible causes for the lower promotion rates of females.

The observed promotion rates may be due to lower success rates of female candidates or to them being less likely than men to seek promotion. The argument that women do not seek promotion is often found in popular discussion without clear evidence of its relevance. Measuring the importance of this mechanism has proven difficult as in most contexts only information on applicants and their success rate tend to be known (see, for example Goldin and Rouse, 2000). In France, academic promotion occurs through a national contest or *concours*, with the list of applicants being publicly available at the time of the *concours*. Moreover, because academics are public servants, we have information on all those holding junior positions which constitute the pool of 'potential applicants'. We can hence examine not only whether gender affects a candidate's promotion probability, but also whether it impacts the decision to be a candidate.

Focusing on the academic labour market has certain advantages. Unlike in many private sector jobs, where a promotion is associated with longer hours, academics have similar obligations and constraints at all hierarchical levels. Even if more senior academics

tend to be involved in university administration and outside responsibilities, these activities are not mandatory. Female associate professors should thus not feel more constrained in terms of combining career and family duties by becoming full professors, and there is hence no obvious reason why they would prefer not to be promoted. Male and female academics are also likely to have rather homogeneous labour market attachment, as argued by Kahn (1995), removing one of the reasons often branded to justify lower promotion rates for females. A further advantage of these data is that a major consideration in actual promotion decisions, an individual's productivity in terms of number and quality of publications, can be observed and thus controlled for by the researcher.

In this context, we consider the two stages of promotion. On the one hand, women may be less likely to be promoted conditional on having applied, and a possible cause of this could be discrimination. On the other hand, female academics may have a lower propensity to apply for promotion than males, which could be explained by the requirements of the contest being more costly for women, their valuing the promotion less, the expectation of discrimination, or simply to an unwillingness to enter the contest. The particular features of the French academic system, such as a national salary scale and the existence of several categories of academics with different requirements during the contest and upon promotion, allow us to investigate these hypotheses.

We use exhaustive data for academic economists in France over the period 1991 to 2008, and find that the promotion rate of women is half as low as that of men. We then consider separately the determinants of the likelihood to enter a promotion *concours* and the probability of being promoted conditional on having entered the *concours*. Being a woman has a significant negative impact on the likelihood to enter the *concours*. The probability of being a candidate is 50% higher for men, and this gap is solely explained by gender. The conditional probability of promotion is also lower for women than for men, but the gap is not statistically significant. Interestingly, we have information about the various steps of the promotion decision. Candidates are first declared 'admissible' and then the committee decides whom to promote amongst these. We find that the main difference across the genders is in the first stage, with women being less likely to be

considered admissible, while in the final stage, candidates are more likely to be promoted if female. A possible interpretation of such results is that positive conscious discrimination occurs in the final stage while negative unconscious discrimination takes place early on when evaluating the worth of women's research output. Overall, the lower propensity to apply accounts for between 50% and 76% of the promotion gap.

The paper adds to a vast body of work on the different rates of promotion across genders; see, amongst others, Lazear and Rosen (1990), Winter-Ebmer and Zweimuller (1997), and Azmat and Ferrer (2017). In their influential paper, Goldin and Rouse (2000) use data from American orchestras to examine the gender gap in hiring and promotion. Their results indicate that discrimination against women was a crucial factor in the 1970s. We address a similar question but use much more recent data, and exploit our knowledge of who could have been a candidate to identify the effect of decisions taken at this stage on observed promotions.

Our analysis contributes to a recent literature that focuses on gender gaps in academia, dating back to the seminal work of Cole and Cole (1973). Recent work indicates that in the US salary gaps are explained by differences in academic rank, while promotions are affected by gender even after controlling for research output and demographic characteristics.[1] Evidence for the UK by Blackaby et al. (2005) indicates that there are both gaps in promotions and in within-rank pay across genders, while Sabatier (2010) documents the existence of a promotion gap in France. Most of this literature has considered the US and the UK, which have an academic labour market with much greater wage and promotion flexibility than those found in most European countries. We examine whether promotion gaps also exist in a labour market that operates in an entirely different way, with salaries being fixed at the country-wide level and promotions being decided by national committees and not by the department where the individual is employed.

The paper is also related to a number of recent analyses which focus on the effect of the sex of committee members on female promotions.[2] This literature generally does

[1] See Johnson and Stafford (1974), Farber (1977), Ginther and Hayes (1999) and Ginther and Kahn (2004). Some studies claim a decline in the promotion gap over time, while others find that it is large even in recent decades; see McDowell et al. (2001) and Ginther and Kahn (2004).

[2] See, for example, Lavy (2008), Bagues and Esteve-Volart (2010) and Bagues et al. (2017).

not support the idea that having more women in a committee increases the probability of success of female candidates. We focus on a different aspect of the promotion process, exploiting information on potential candidates to assess the importance of the decision of whether or not to apply for promotion. Lastly, a growing body of experimental work has addressed differences between male and female attitudes towards competition and argued that women tend to be less likely to enter competitive situations than men.[3] Our result that women have a lower probability than men to enter the *concours* is consistent with the experimental evidence.

The paper is organised as follows. After describing the French academic system, Section II examines the possible reasons why women are less likely to be promoted. Section III describes the data and our results are presented in Section IV, while Section V concludes.

II Why are there so few female professors?

The French academic system

The French academic system has a number of features that we intend to exploit. There are two types of academic positions in France. The most common are university positions, where the individual is a professor with a substantial teaching load. There exist also a number of public research instances, of which the largest is the CNRS, that have full research positions.[4] Researchers in this category are hired by the CNRS, who pay their salaries, but are attached to a university and participate in department life.

For all types of position there is an entry level category equivalent to assistant professor, termed 'Rank B', which includes the *maître de conférences* positions at the university and the *chargé de recherche* at the CNRS. The individual can then be promoted to 'Rank A', the equivalent to full professor, a position denoted *directeur de recherche* at the CNRS. Both rank A and rank B positions are tenured. The promotion from rank B to rank A entails a substantial salary increase, with the salary scales being set by the Ministry of Higher Education and Research and identical for university professors and

[3] See Niederle and Vesterlund (2007), Niederle and Vesterlund (2011) and Gupta et al. (2013).
[4] CNRS stands for *Centre National de la Recherche Scientifique*.

researchers as well as across departments.[5]

In several fields, notably economics, promotions take place through a national contest, a *concours*, with the decision taken by a committee consisting of academic economists of the full professor rank. Participation in this contest and the final ranking are public information. There is no limit to the number of times a candidate may apply for a rank A position. The fact that departments do not make promotion decisions is important for our purposes.[6] In a system in which there is a positive correlation between prestige and promotion threshold, women could choose to select into less prestigious department because promotion is easier in those, and hence the measured gap would underestimate the actual promotion gap.

The requirements of the *concours* differ across the two tracks. For university professors the contest is termed *concours d'agrégation*. Over our sample period, the *agrégation* was biannual and entailed four stages over, approximately, a 6-month period. It included a research seminar and three oral exams, one of which consisted of preparing in 24 hours a lecture on a topic randomly drawn from a lengthy predetermined list. At the various stages of the competition, a number of candidates are eliminated, and after the penultimate stage a list of candidates that are *admissible* is provided by the jury. These candidates then undertake the final test, after which a ranking of all successful candidates is provided. The *concours* hence takes time and requires substantial preparation outside the candidate's field of expertise. In contrast, the effort involved in the CNRS promotion *concours* is minimal. The candidate simply declares him/herself a candidate and submits a vitae and a research project to the committee. This *concours* takes place annually and does not involve any meeting between the candidate and the committee.

There is a second difference in the costs. For university professors, a list of open positions is published and at the end of the *concours* candidates choose, sequentially and starting with the highest ranked, which department to join. When promoted, individuals

[5]Some departments pay, out of their own funds, an extra salary on top of the one paid by the university/CNRS. This practice is, however, restricted to only a few members in a handful of departments.

[6]Some internal promotions exist for individuals that have undertaken substantial administrative tasks at the university, but they are rare. See Combes et al. (2008) for more details.

are usually not able to stay in the university where they held a rank B position.[7] After three years they are allowed to move to another university, if the latter wishes to recruit them, including their former university. CNRS researchers that are promoted can choose to stay at the university where they are or move. The university does not need to have an open position for them since the researchers move with their position.

It is obvious that promotion is very costly for university professors. The *agrégation* involves an ex-ante cost in terms of preparing the exams and a substantial cost ex-post because of the geographical mobility involved in being promoted to full professor. Since candidates seeking to become full professor are typically between 30 and 40 years of age, the process occurs at a moment in the life cycle when family constraints are likely to be substantial. If women are less geographically mobile than men, then the cost is likely to be greater for them.

Discrimination, differences in payoffs or self-selection?

Although a substantial literature has examined the promotion gap across genders, a clear explanation is still lacking. Women may be less likely to be promoted either because they apply for promotion less often than men or because of a lower probability of being promoted conditional on being a candidate. In turn, gender differences in the underlying probabilities may have three possible causes: discrimination, differences across genders in the costs and benefits of promotion, and different attitudes towards the promotion process itself. The French system presents a number of features that will allow us to evaluate the various hypotheses.

To understand the possible mechanisms in operation consider the following model of promotion. Suppose that individual i is a potential candidate for promotion at time t and let the subscript $j = u, r$ denote whether the model applies to university professors or researchers. A candidate's promotion probability is the product of the probability of success conditional on having applied, $\Pr_j(S, i, t)$, and the likelihood that s/he applies

[7] In our sample, 81.8% of those promoted through the *agrégation* have a new affiliation after their promotion. Only 3.4% of available positions were within commuting distance of Paris (urban area). On average, successful candidates had to move 674 kilometers. This is large: would France be a disk, its radius would be around 420 kilometers.

for promotion $\Pr_j(A, i, t)$. The decision whether to be a candidate is in turn determined by a comparison between individual costs and benefits from entering the contest. An individual will be a candidate if and only if the expected cost, C_{ij}, is lower than the product of the expected probability of success and the value of being promoted, V_{ijt}. That is, if $C_{ij} < \Pr_j(S, i, t) \times V_{ijt}$.

The (conditional) probability of success of individual i at time t is assumed to be given by

$$\Pr_j(S, i, t) = p_{j0} + p^f_{j0}\ \delta_i + X_{it}\ p_{j1} + X_{it}\ p^f_{j1}\ \delta_i + Z_{it}\ p_{j2}\ , \tag{1}$$

where δ_i is a dummy taking the value 1 for females, X_{it} is a vector of measures of the individual's research output, Z_{it} is a vector of individual characteristics at date t, such as age, and the various p_{jn} and p^f_{jn} $(n = 0, 1, 2)$ are vectors of coefficients.

Taste-based, statistical and implicit discrimination all imply that female candidates are less likely to be promoted, i.e. $p^f_{j0} < 0$. Implicit discrimination may appear if committees believe that for a given publication track women are on average less deserving, for example if women are believed to contribute less than their male coauthor to a joint paper.[8] Discrimination can reduce female promotion rates directly as well as indirectly as a lower probability of success will reduce the likelihood that women apply. Also, the decision to apply is based on the expected probability of success. Hence, if women believe that they will be discriminated against, they will be less likely to apply for promotion than men irrespective of whether they are actually discriminated.

Women may also differ from men in the costs of and rewards from promotion and this can make them less willing to apply for promotion, causing 'self-selection' out of the promotion race. We suppose that costs are given by $C_{ij} = c_j + c^f_j \delta_i$. For researchers, cost parameters are given by $c_r = c_0$ and $c^f_r = c^f_0$, where the terms c_0 and c^f_0 capture any psychological or subjective costs incurred when taking part in a contest that may differ for men and women. The differences can be due to an unwillingness of women to participate in contests or to the fact that women tend to fear losing tournaments or contests more

[8] See Bertrand et al. (2005) for a discussion of implicit discrimination and Wenneras and Wold (1997) on differences in the way in which male and female research output is valued by committees.

than men.[9] Higher subjective costs of competing hence may lead women to apply less often than men.

We suppose that these subjective costs are the same across tracks,[10] and assume that for the university track the overall costs are given by $c_u = c_0 + c_1$ and $c_u^f = c_0^f + c_1^f$, where the terms c_1 and c_1^f are the objective costs due to the time required to prepare the contest and the expected probability of moving place of residence. Differences in intra-household bargaining power could make it harder for female than for male professors to impose the cost of moving on their families, thus making $c_1^f > 0$.

Turning to the value of being promoted, we suppose it takes the form $V_{it} = v_0 + (v_1 + v_1^f \ \delta_i) Income_t$. The intrinsic value of the promotion is given by v_0, which we assume is common across sexes, while the variable $Income_t$ captures the increase in income associated with promotion. Men and women may have different preferences over income if, for example, women have a lower marginal utility of income because they are the second earner in the household, in which case $v_1^f < 0$.

In this context, we consider three hypotheses. First, we examine whether there is a negative impact of being female on the conditional probability of promotion, captured by $p_{j0}^f < 0$. Second, we examine the determinants of the likelihood to apply for promotion. If women were less likely to apply, this might be due to them having higher objective costs of promotion. Because the costs of the *agrégation* are higher than those of seeking promotion as a researcher, we have $c_u^f > c_r^f$, i.e. we expect to find a more negative effect of being female on the probability of applying for promotion for professors than for researchers. Alternatively, women may apply less either because they expect to be discriminated against (i.e. they expect $p_{j0}^f < 0$), they have different preference for income, $v_1^f < 0$, or they are less willing to take part in contests because of a higher subjective cost of competing captured by $c_0^f > 0$.

[9] See, for instance Buser (2016).

[10] It is not obvious whether the degree of competition is greater for university professors or for researchers. On the one hand, a much smaller number of researchers are promoted each year in the CNRS than in the *concours d'agrégation*, but the number of potential candidates is also larger in the latter. On the other, although the list of candidates and rankings are available on internet for both, the *agrégation* solicits much greater interest from the academic community, with results at each stage being widely followed and discussed, which increases the preasure on candidates.

Empirical specification

Our empirical analysis considers five possible outcomes: the unconditional probability of being promoted, denoted $\Pr_j(P,i,t)$, the probability of applying, $\Pr_j(A,i,t)$, the probability of success conditional on applying, $\Pr_j(S,i,t)$, and, in the case of university professors, the probabilities of being declared admissible conditional on applying, $\Pr_j(Ad,i,t)$, and of succeeding conditional on being admissible, $\Pr_j(SAd,i,t)$.

We consider the same specification for all probabilities. For an outcome O where $O = P, A, S, Ad, SAd$, we have

$$\begin{aligned}\Pr_j(O,i,t) &= \beta_{j0} + \beta_{j0}^f\ \delta_i + \beta_{j1}\ Age_{it} + \beta_{j2}\ Age_{it}^2 + \beta_{j3}\ Pub_{it} + \beta_{j4}\ Pub_{it} \times Quantity_{it} \\ &+ \beta_{j5}\ Pub_{it} \times Quality_{it} + \beta_{j6}\ Int.\ Department_{it} + \beta_{j7}\ Ile\ de\ France_{it}. \quad (2)\end{aligned}$$

The outcome is a function of age Age_{it} and its square, whether or not individual i has published in Econlit-classed journals ('publisher' dummy Pub_{it}), the number of publications and their average quality, denoted respectively $Quantity_{it}$ and $Quality_{it}$, both in logs (see below for the exact measurement). With δ_i being a dummy for females, β_{j0}^f measures the differences in the probability for men and women with the same characteristics. The dummy $Int.\ Department_{it}$ has a value of 1 if the candidate is in a top department and that denoted $Ile\ de\ France_{it}$ has a value of 1 if the individual holds a position in Paris or within commuting distance of it (see below for the details).

These last two variables have several interpretations. They could measure the positive effects of a more stimulating academic environment, be an indicator of peer pressure to seek promotion, or capture higher unobserved ability since the best junior candidates will have been recruited by the departments with the highest research quality or in a particularly attractive location. There are also externalities due to having colleagues who are simultaneously preparing the *agrégation*.

Obviously, the qualities that lead to promotion conditional on applying are also those that make a potential candidate apply, hence it is not possible to run a selection model. We consider a linear probability model and use OLS in the regressions, although, as we

discuss below, a number of robustness analyses on the functional form are performed. Our central question is whether the coefficients β^f_{j0} in the regressions are significantly different from zero, implying that gender affects outcomes. We will also examine the differences between the coefficients estimated for the two tracks, at the university or CNRS. Recall that the cost of applying differs considerably across tracks, hence if high costs were preventing women from applying we would expect to find a higher value for β^f_{u0} than for β^f_{r0}.

III The Data

Our sample consists of the entire population of French academic economists provided by the French Ministry of Higher Education and Research and the CNRS for the years 1991-2008. We have information on age, rank, publications and department. We keep only individuals that are in departments larger than 4 full-time equivalent academics, which removes universities without real economics departments. Data on individuals' family life, such as whether they are married and the number and ages of children, are not available. For the US, Ginther and Kahn (2004) find that having children has only a weak effect on the promotion probabilities of female economists in the US and none on their productivity. Moreover, having children and deciding to pass the *agrégation* may be simultaneous decisions, as women may choose to postpone childbirth until after promotion, thus it is difficult to view family circumstances as exogenous variables.

We use the list of candidates that applied to and those who succeeded in becoming a rank A professor or researcher. For professors we have data for the nine biannual *concours* taking place from 1992 to 2008, while for researchers we have the thirteen annual *concours* between 1996 and 2008. All the years are pooled together for each track.

Although all academic economists of rank B the year prior to the *concours* are eligible for promotion, it is important to consider how to define potential candidates. It is possible for CNRS researchers to take the *agrégation*, and some individuals in our dataset do so. In contrast, although it is in principle possible for rank B university professors to apply for a rank A CNRS position, there are no such individuals in our data. We will hence

consider all rank B individuals as potential candidates for the *agrégation* but only those with a CNRS rank B position as potential candidates to promotion to rank A at CNRS. A second question is how to deal with older candidates, who may have characteristics that may make them choose not to apply for promotion. In order to avoid having in our pool an increasing number of candidates unlikely to apply we introduce age limits. We hence consider only individuals aged between 32 and 55.[11]

We define two categories of departments, somewhat equivalent to the division in the US between the top-50 and other departments. France has a substantial number of national academic publications in French, and hence we define prestigious departments as those that have the largest research output in international journals (see the supplementary material). For this reason, we will term them 'international' departments and the rest 'national' departments. The international departments account for about one third of the academics each year.

Measuring research output

We measure individual output at date t by the cumulative publication record between the first year of publication and date t.[12] Publication records are measured as weighted sums of publications. All publications come from the EconLit database over the period 1969 and 2008. Three dimensions enter the weighted scheme of publications: the quality of journals, the number of authors and the relative number of pages. We measure the quality of publications using the journal ranking proposed by Combes and Linnemer (2010), divide each publication by the number of authors, and weight by the number of pages to capture the idea that longer articles contain more ideas, considering an article's length relative to the average length in that journal the same year.[13]

[11] Our results hold with different bounds, see the supplementary material, Bosquet et al. (2018). Sample attrition could also be important. In fact, 16.4% of men and 13.9% of women leave the sample before 2008 at age less than 55, indicating that women do no leave academia more often than men.

[12] As an alternative, we have computed degressive publication scores, with older publications having a smaller weight than recent ones. Measuring research output this way does not change our results (see the supplementary material).

[13] EconLit includes more than 560,000 papers published between 1969 and 2008 in more than 1200 journals, including 46 in French. Combes and Linnemer (2010) provide weights for all EconLit journals using various recursive impact factors built from Thomson Reuters Web of Knowledge impact factors and from Google Scholar citations. Two different degrees of convexity in the distribution of journals' weights

The output of individual i at date t is thus a weighted sum of individual i's articles published between the individual's first year of publication, t_{0i}, and date t, so that

$$y_{it} = \sum_{a_i \in [t_{0i}, t]} \frac{W(a_i)}{n(a_i)} \frac{p(a_i)}{\bar{p}_t}, \tag{3}$$

where $p(a_i)$ is the number of pages of article a_i written by i between t_{0i} and t, $\bar{p}_t$ the average number of pages of articles in the journal for the year of publication, $n(a_i)$ the number of authors of the article, and $W(a_i)$ the weighting scheme for journals. Each individual receives three scores: a dummy variable equal to 1 if the individual has at least one publication in an EconLit-listed journal between t_{0i} and t, the number of single-author-equivalent published articles ('quantity'), $\sum_{a_i \in [t_{0i}, t]} 1/n(a_i)$, and the average quality of these publications, defined as y_{it} divided by the quantity.

Descriptive statistics

In order to run regressions equivalent to those found in the literature on promotions in academia, where only outcomes are observed, we construct a sample that includes all rank A and rank B academics for each of the years that we will be using latter on. Table 1 gives the decomposition of this sample by gender.

Table 1: Observations and share of rank A by gender and track

rank A	Total	%	Women (25.6% of population)	%	Men (74.4% of population)	%
			Total sample			
0	11,322	64.8	3,686	82.3	7,636	58.8
1	6,145	35.2	792	17.7	5,353	41.2
			University professors (85% of population)			
0	9,797	66.0	3,212	82.8	6,585	60.0
1	5,052	34.0	666	17.2	4,386	40.0
			CNRS researchers (15% of population)			
0	1,525	58.3	474	79.0	1,051	52.1
1	1,093	41.7	126	21.0	967	47.9

are proposed and we use the most convex one (i.e. the one that most values quality), but our results are unchanged when we use the least convex one (see the supplementary material). Gibson et al. (2014) show that this is the journal ranking that best explains the salaries at the University of California economics departments.

There were between 1,165 and 2,143 academic economists in France depending on the year, and the vast majority of them are university professors, with researchers accounting for only 15.0% of the total. Women represent 25.6% of observations, and they are overrepresented amongst university professors and underrepresented amongst researchers. This difference could be due to the fact that obtaining a position as a researcher tends to require a stronger publication record than for university positions and, as we will see below, women tend to have a weaker research output than men. Slightly over a third of the population hold a rank A position, with the fraction being lower for university professors (34.0%) and higher for researchers (41.7%). Overall, the gender promotion gap is large, 23.5 percentage points on average, and is smaller for university professors than for researchers (22.8 and 26.9 points, respectively).

Table 2: Descriptive statistics of potential candidates

	University			CNRS		
	Women	Men	Diff.	Women	Men	Diff.
Prob. Candidate	0.060	0.090	-0.030^a	0.115	0.203	-0.088^a
	(0.005)	(0.004)	(0.006)	(0.017)	(0.014)	(0.024)
Prob. Promotion	0.011	0.020	-0.010^a	0.017	0.045	-0.028^b
	(0.002)	(0.002)	(0.003)	(0.007)	(0.007)	(0.012)
Age	41.0	43.1	-2.1^a	41.3	42.2	-0.9^c
	(0.1)	(0.1)	(0.2)	(0.3)	(0.3)	(0.4)
Publisher	0.559	0.590	-0.032^a	0.779	0.863	-0.084^a
	(0.010)	(0.007)	(0.012)	(0.022)	(0.012)	(0.024)
Quantity	1.08	1.42	-0.34^a	2.63	3.14	-0.51^a
	(0.03)	(0.03)	(0.05)	(0.15)	(0.11)	(0.19)
Quality	0.45	0.44	0.01	1.62	2.02	-0.41
	(0.06)	(0.04)	(0.07)	(0.27)	(0.21)	(0.36)
Int. Department	0.309	0.273	0.035^a	0.556	0.537	0.018
	(0.009)	(0.006)	(0.011)	(0.026)	(0.018)	(0.032)
Ile de France	0.349	0.272	0.077^a	0.469	0.426	0.043
	(0.009)	(0.006)	(0.011)	(0.026)	(0.018)	(0.032)

Notes: Standard errors in brackets. ^{a}, ^{b}, ^{c} Difference between men and women significant at the 1%, 5% and 10% level, respectively. Productivity measures (quantity and quality) are in levels. We take their logs in the regression analysis. There are 8,085 observations from 1,869 individuals and 1,132 observations from 191 individuals in the university and CNRS samples respectively. These correspond to 601 women (2,647 observations) and 1,268 men (5,438 observations) in the university sample and 61 women (358 observations) and 130 men (774 observations) in the CNRS sample.

Table 2 reports some descriptive statistics for the sample of potential candidates. We have 8,085 observations, including a total of 1,869 different academics, with the average

number of observations per individual being 4.3. The panel is unbalanced as individuals enter the pool of potential candidates and exit it either because they are promoted or leave academia/France, or reach the maximum age for us to consider them as potential candidates. Males are more likely than females to apply for promotion. The probability of being a candidate to the *agrégation* is 6.0% for women and 9.0% for men, while for researchers these figures are 11.5% and 20.3%. The (unconditional) probabilities of being promoted on a given year are small, 2% and 4.5% for men, and 1.1% and 1.4% for women, depending on the track. Again, the promotion probability is at least twice as large for men as for women. The probability of having published in EconLit journals is 59.0% and 55.9% for men and women professors, respectively. These figures are not large, but it is important to bear in mind that our publication criteria is stringent given the strong tradition in France to publish books and in national journals which are not necessarily in EconLit. Male professors publish more articles than female professors but we find no significant difference in the average quality of publications. In the CNRS, differences are more pronounced and indicate that men also have a greater quantity of publications. The aspect in which women fare better than men is affiliation: women are more often than men in international departments or in a department in or around Paris (*Ile de France*).[14]

Table 3 reports descriptive statistics for the sample of actual candidates to promotion and for those promoted. We indicate by a subscript next to the average for men whether the difference between the male and female averages is statistically significant. The probabilities of success in the *agrégation* are 18.1% for women and 21.6% for men, while for the CNRS contest these figures are 13.3% and 21.1%. Contrary to the probability of applying for promotion, these differences are not statistically significant. We find no significant difference in the quantity or quality of publications.

We have seen that there are differences in the average research output of male and female potential candidates. It is possible that there are also differences in its distribution, and since promoted candidates are selected from the top tail, differences in the thickness of this tail could explain observed promotion gaps. Figure 1 thus plots the distribution

[14]Women may be more likely to be in an international department for three reasons: positive discrimination, higher unobservable ability, or because of joint offers made to couples by top departments.

Table 3: Descriptive statistics of candidates and promoted candidates

	Candidates				Promoted candidates			
	University		CNRS		University		CNRS	
	Women	Men	Women	Men	Women	Men	Women	Men
Prob. Promotion	0.181	0.216	0.133	0.211	1.000	1.000	1.000	1.000
	(0.028)	(0.017)	(0.051)	(0.030)	(0.000)	(0.000)	(0.000)	(0.000)
Age	35.8	37.8[a]	42.8	43.6	34.9	35.9	42.5	40.8
	(0.3)	(0.2)	(0.7)	(0.4)	(0.4)	(0.3)	(2.3)	(0.7)
Publisher	0.824	0.809	0.978	0.930	0.882	0.852	1.000	0.949
	(0.028)	(0.016)	(0.022)	(0.019)	(0.056)	(0.032)	(0.000)	(0.036)
Quantity	2.14	2.37	5.50	4.58	2.39	3.18	5.58	6.54
	(0.15)	(0.10)	(0.59)	(0.25)	(0.37)	(0.28)	(1.18)	(0.65)
Quality	0.65	0.84	1.38	3.44	1.83	1.90	2.21	9.05
	(0.20)	(0.16)	(0.29)	(0.67)	(0.90)	(0.38)	(1.15)	(2.75)
Int. Department	0.399	0.287[a]	0.622	0.562	0.647	0.406[b]	0.667	0.692
	(0.036)	(0.019)	(0.073)	(0.037)	(0.083)	(0.044)	(0.211)	(0.075)
Ile de France	0.479	0.327[a]	0.289	0.443[c]	0.559	0.461	0.167	0.564[c]
	(0.037)	(0.019)	(0.068)	(0.037)	(0.086)	(0.044)	(0.167)	(0.080)

Notes: Standard errors in brackets. [a], [b], [c] Difference between men and women significant at the 1%, 5% and 10% level, respectively. Productivity measures (quantity and quality) are in levels. We take their logs in the regression analysis. There are 781 observations from 522 candidates and 230 observations from 93 candidates in the university and CNRS samples respectively. These correspond to 130 women (188 observations) and 392 men (593 observations) in the university sample and 16 women (45 observations) and 77 men (185 observations) in the CNRS sample. There are 162 and 45 promoted candidates in the university and CNRS samples respectively. These correspond to 34 women and 128 men in the university sample and 6 women and 39 men in the CNRS sample.

of the (log) quantity and (log) average quality of research output for men and women. The distribution of quantity for female professors is to the left of that for males and has less mass at the top of the distribution, indicating that the most productive men publish more than the most productive women. There is little difference in the distribution of quality, although we observe a mild bimodality for men that is absent for women.

IV Results

The promotion of academic economists

In order to compare our results with existing work, we consider the full sample of French academic economists and start by examining the determinants of the probability of *holding* a rank A position. The results are reported in the first three columns of Table 4, where

Figure 1: Smoothed densities of publication for men and women potential candidates

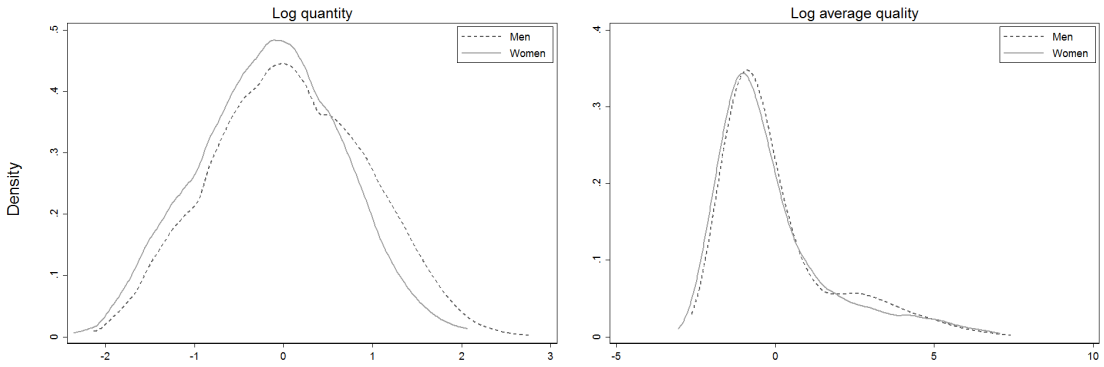

all specifications include year fixed effects. In column (1), only gender is included in the model, with the dummy *Woman* being equal to 1 for women, and exhibiting a negative and significant coefficient that implies that women are over 23.3 percentage points less likely than men to be rank A, a figure comparable to those found for the US.[15] Controlling for age and our measures of research output reduces the gender effect to -0.045, indicating that a large fraction of the difference in promotions is due to women being younger and having published less. Column (3) includes interactions between the gender dummy and publication scores and indicates that men and women do not have significantly different returns to publications.[16]

The next four columns of Table 4 consider our sample of potential candidates and report estimates of the probability of being promoted. There is a one percentage point raw gap across the genders for university professors, which amounts to women being half as likely to be promoted in a given year as men. This gap remains unchanged once we control for age, research output and department type (column 5). The last two columns concern the CNRS, where the raw gap of 0.029 falls slightly to 0.023 once we include controls, again amounting to about half the raw promotion probability for men (0.045). That is, our sample of potential candidates exhibits promotion gaps that are substantial and, contrary to what we find for the entire population of academic economists, are not reduced by controlling for age and productivity.

[15]For the US Ginther and Kahn (2004) find a raw gap of 21.3 percentage points.

[16]In what follows we will not consider different returns to publications as we found no significant difference for our other specifications.

Table 4: Likelihood to hold a rank A position and likelihood of promotion among potential candidates

	Likelihood to						
	hold a rank A position			be promoted (potential candidates)			
				University		CNRS	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Woman	-0.233^a	-0.045^a	-0.043^a	-0.009^a	-0.010^a	-0.029^a	-0.023^b
	(0.019)	(0.016)	(0.016)	(0.003)	(0.003)	(0.010)	(0.011)
Age		0.011^a	0.011^a		-0.009^a		0.020^a
		(0.002)	(0.002)		(0.001)		(0.004)
Age^2		0.000^b	0.000^b		0.000^a		-0.000^a
		(0.000)	(0.000)		(0.000)		(0.000)
Publisher(Pub)		0.316^a	0.317^a		0.041^a		0.032^b
		(0.021)	(0.026)		(0.007)		(0.015)
Pub*Quantity		0.179^a	0.182^a		0.015^a		0.037^a
		(0.009)	(0.010)		(0.003)		(0.008)
Pub*Quality		0.034^a	0.033^a		0.009^a		0.011^a
		(0.004)	(0.005)		(0.002)		(0.004)
CNRS		-0.114^a	-0.110^a		-0.038^a		
		(0.024)	(0.028)		(0.005)		
Woman*Pub			0.001				
			(0.045)				
Woman*Pub*Quantity			-0.024				
			(0.021)				
Woman*Pub*Quality			0.007				
			(0.011)				
Woman*CNRS			-0.013				
			(0.057)				
Int. Department					0.006		-0.002
					(0.004)		(0.011)
Ile de France					0.009^b		0.018
					(0.004)		(0.013)
Time FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
R^2	0.048	0.392	0.393	0.003	0.051	0.010	0.077
Observations	17,467	17,467	17,467	8,085	8,085	1,132	1,132

Notes: Linear probability model (OLS). Standard errors clustered by individuals in brackets. ^{a}, ^{b}, ^{c} Significant at the 1%, 5% and 10% level, respectively. Terms interacted with Woman are centered so that the coefficient associated with Woman is still the average effect of being a woman.

Decomposing outcomes

The limitation of our analysis so far is that we do not know whether observed promotions are the result of a lower likelihood to apply for promotion by women or lower success in obtaining the promotions. We hence examine separately the various steps. We start by considering what determines the decision to enter the contest, and then move to the determinants of success in the contest conditional on being a candidate.

The results on the determinants of the likelihood to enter the *concours* are reported in Table 5. As expected, research output has a strong effect, with quality having a stronger impact for researchers.[17] Being in an international department is insignificant in both tracks, while being at a department in Paris or within commuting distance of it (*Ile de France*) has a significant coefficient only for professors, as expected. In the regression for the *agrégation* we also include a dummy if the candidate holds a rank B CNRS position.

Table 5: Likelihood to apply for a promotion

	University		CNRS		Diff. in diff.	Diff. across dept. University	Diff. across dept. CNRS
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Woman	-0.030^a	-0.033^a	-0.093^b	-0.071^b	-0.028^a	-0.033^a	-0.071^b
	(0.008)	(0.007)	(0.040)	(0.036)	(0.008)	(0.008)	(0.036)
Age		-0.021^a		0.084^a	-0.008^a	-0.021^a	0.084^a
		(0.003)		(0.012)	(0.003)	(0.003)	(0.012)
Age^2		0.000^a		-0.002^a	0.000	0.000^a	-0.002^a
		(0.000)		(0.000)	(0.000)	(0.000)	(0.000)
Publisher(Pub)		0.114^a		0.118^a	0.124^a	0.114^a	0.117^a
		(0.015)		(0.045)	(0.015)	(0.015)	(0.045)
Pub*Quantity		0.046^a		0.108^a	0.046^a	0.046^a	0.107^a
		(0.007)		(0.028)	(0.007)	(0.007)	(0.028)
Pub*Quality		0.014^a		0.022^b	0.015^a	0.014^a	0.022^b
		(0.004)		(0.010)	(0.004)	(0.004)	(0.010)
Int. Department		0.004		0.020	0.007	0.005	0.016
		(0.010)		(0.048)	(0.010)	(0.010)	(0.050)
Ile de France		0.019^b		0.005	0.020^b	0.007	0.020
		(0.009)		(0.043)	(0.009)	(0.011)	(0.058)
CNRS		-0.127^a				-0.127^a	
		(0.011)				(0.011)	
Woman*CNRS					-0.056		
					(0.036)		
Woman*Ile de France						0.032^b	-0.041
						(0.016)	(0.077)
Interacted terms	No	No	No	No	Yes	No	No
Time FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
R^2	0.005	0.110	0.033	0.194	0.116	0.110	0.195
Observations	8,085	8,085	1,132	1,132	9,217	8,085	1,132

Notes: Linear probability model (OLS). Standard errors clustered by individuals in brackets. ^{a}, ^{b}, ^{c} Significant at the 1%, 5% and 10% level, respectively. Interacted terms: All variables interacted with applying to the CNRS *concours*. Ile de France has been centered in its interaction with Woman so that the coefficient associated with Woman is still the average effect of being a woman.

The unconditional probability of applying is lower for women than for men in both

[17]The results are robust to the inclusion of higher order terms to capture some non-linearity in the impact of publications' quality (see the supplementary material).

tracks, and the negative impact of being female remains roughly as large once we include individual characteristics, the marginal effect being -0.033 for professors and -0.071 for researchers. The effect is large: being a woman is equivalent to dividing the number of single-authored publications by around 2, and implies that in both tracks the probability is about 50% higher for men than for women even when controlling for observables.

In order to understand what lies behind the gender gap in seeking promotion, we examine whether the effect of gender differs between the two types of positions. Since the objective costs of the *agrégation* contest are substantially larger than those in the CNRS, the coefficient on women should be higher for university professors than for researchers if differences in these costs where holding back potential female candidates. Column (5) of Table 5 hence runs a difference-in-differences regression where we have pooled together the data for the two types of *concours*. We interact all variables with a dummy taking the value 1 for the CNRS *concours* to allow for different impacts across the two tracks, our coefficient of interest being women interacted with this dummy.[18] The direct gender effect remains negative, while the interacted coefficient is negative but insignificant at the 10% level, implying that, if anything, the larger coefficient for CNRS implies a stronger reluctance for women to enter the CNRS *concours*. These results are not consistent with the hypothesis that the reason why women do not apply for promotion are the costs associated to seeking promotion in the university track.

The last two columns of Table 5 consider the role of location by interacting the female dummy with being in a department in or around Paris (we also considered interaction with being in an international department but the coefficient was never significant). For researchers we find an insignificant difference, while the coefficient is positive and significant for professors, indicating that the application gap between men and women in the *Ile de France* region is lower than for other women.[19] This difference has several possible interpretations: it could be that the high concentration of economics departments in and

[18]To save space, other interacted terms are not reported here but in the supplementary material. Note that the specification in column (5) gives the same results as a weighted comparison between those in columns (2) and (3).

[19]To obtain the overall impact of gender, note that the variable *Ile de France* has been centered. Its average value for women is 34.9%, implying that the total effect of gender is -0.012 for women in the *Ile de France* and -0.044 for the rest.

around Paris makes it more likely that successful candidates can change department without moving their residence, that being in Paris reduces the cost of preparing the *concours*, or it could simply capture better unobservables of women in the Paris region.

Our results thus indicate that there is a substantial gender gap in the likelihood to apply for promotion. The negative and significant coefficient is robust to different specifications of the model, such as logit, probit, and controlling for time since the first publication (see the supplementary material). We also run a random effects model as well as a duration model where we estimate time to first application for promotion; both yielded equivalent results.

Table 6: Share of candidates among potential candidates, reweighting

	University			CNRS		
	Women	Men	Diff.	Women	Men	Diff.
Prob. Candidate	0.060	0.090	-0.030^a	0.115	0.203	-0.088^a
	(0.005)	(0.004)	(0.006)	(0.017)	(0.014)	(0.024)
Wom. reweight logit	0.054	0.090	-0.036^a	0.145	0.203	-0.057^b
	(0.004)	(0.004)	(0.006)	(0.019)	(0.014)	(0.025)
Wom. reweight probit	0.054	0.090	-0.037^a	0.146	0.203	-0.057^b
	(0.004)	(0.004)	(0.006)	(0.019)	(0.014)	(0.025)
Men reweight logit	0.060	0.092	-0.032^a	0.115	0.186	-0.072^a
	(0.005)	(0.004)	(0.006)	(0.017)	(0.014)	(0.024)
Men reweight probit	0.060	0.092	-0.032^a	0.115	0.187	-0.072^a
	(0.005)	(0.004)	(0.006)	(0.017)	(0.014)	(0.024)

Notes: Standard errors in brackets. ^{a}, ^{b}, ^{c} Difference between men and women significant at the 1%, 5% and 10% level, respectively.

As further test of sensitivity to the functional form used, we estimate the application gap using propensity–score reweighting, which allows us to compute women's mean probability of applying if they had the same distribution of observables as men. We use the re-weighting method of DiNardo et al. (1996), which involves re-weighing females such that they correspond to the male distribution of observed characteristics (or vice-versa) and which they show is equivalent to non-parametric matching. Table 6 reports the estimated probability of applying using the raw data, as well as four re-weighting estimates that use all the control variables considered in our estimation, employ a logit or probit specification, and take either women or men as the reference group. The gap for university professors is stable across estimations. The results for the CNRS show

greater variation, with the estimates ranging between 0.057 and 0.088. These values are nevertheless large compared with the base male application probability (0.203) and highly significant, indicating that the application gap is not driven by differences in the distribution of observables nor by the particular functional form we employ.

We turn next to differences in the probability of being promoted conditional on having applied for promotion. Table 7 reports results for this probability and indicates that research output is the key determinant of the probability of success in both the university and CNRS *concours*. In the former, all three measures of research output have a significant coefficient, while for the CNRS promotion only quality matters, consistent with stronger selection in this track. Being in an international department has a positive impact on the probability of passing the *agrégation*, in line with our arguments above, while it has no impact on the probability of being promoted in the CNRS contest.

Table 7: Likelihood to be promoted conditional on applying

	University			CNRS		
	(1)	(2)	(3)	(4)	(5)	(6)
Woman	-0.029	-0.035	-0.046	-0.068	-0.079	-0.063
	(0.033)	(0.032)	(0.032)	(0.059)	(0.053)	(0.061)
Age		-0.032^{b}	-0.032^{b}		-0.048	-0.042
		(0.013)	(0.014)		(0.039)	(0.039)
Age^2		0.001	0.001		0.001	0.001
		(0.000)	(0.000)		(0.001)	(0.001)
Publisher(Pub)		0.175^{a}	0.157^{a}		0.004	0.028
		(0.046)	(0.048)		(0.124)	(0.137)
Pub*Quantity		0.074^{a}	0.074^{a}		0.063^{c}	0.075^{c}
		(0.020)	(0.020)		(0.038)	(0.041)
Pub*Quality		0.057^{a}	0.053^{a}		0.041^{a}	0.039^{b}
		(0.009)	(0.009)		(0.015)	(0.018)
Int. Department			0.093^{a}			-0.010
			(0.036)			(0.069)
Ile de France			0.023			0.067
			(0.033)			(0.074)
Pos. other than univ.			0.001			
			(0.036)			
Time FE	Yes	Yes	Yes	Yes	Yes	Yes
R^2	0.030	0.150	0.164	0.043	0.192	0.197
Observations	781	781	781	198	198	198

Notes: Linear probability model (OLS). Standard errors clustered by individuals in brackets. ^{a}, ^{b}, ^{c} Significant at the 1%, 5% and 10% level, respectively. Pos. other than univ.: assistant professor positions in economics outside France, as well as CNRS researchers, and assistant professors from other disciplines.

Turning to our coefficient of interest, the impact of being female on the probability of

success conditional on being a candidate, there is a negative but not significant difference across the genders. The coefficient on *Woman* is insignificant both for the raw probabilities and when we control for individual characteristics, for professors and researchers.[20] The coefficients are large, -0.046 for professors, implying a reduction in the probability of success of 21% compared to the male baseline (0.216), and -0.063 for the CNRS (a 30% gap). These results are not fully conclusive, as the lack of significance could be due to our sample being small (there are, respectively, only 188 and 41 women amongst the candidates) or could reflect the lack of *ex-post* discrimination against women.

To further understand the promotion outcome, we exploit the fact that for the *agrégation* we have information on which individuals passed each stage of the competition, as well as the final rank. After the penultimate stage a list of admissible candidates is provided and after the final test a ranking of all successful candidates is published. If p positions are available that year, the top p candidates are promoted. Our first test considers the probability that a candidate is declared admissible. We next examine the likelihood to be promoted for the sample of admissible candidates since this way we eliminate the (presumably predominantly male) bottom tail of the distribution of applicants. To make our sample even more comparable, we construct from the final ranking a list of individuals close to the threshold. We take, for each year, the n admissible candidates that did not get promoted and the n lowest-ranked candidates that were. If discrimination against women were taking place, it would be likely to appear in this reduced sample of candidates with similar observable and probably close unobservable characteristics.

The first two columns of Table 8 consider the probability that a candidate is declared admissible. Research output plays a major role, with the coefficients being even larger than in the previous table, as we would expect. The coefficient on *Woman* remains negative, large (amounting to about 20% of men's probability), and is significant at the 10% level. The next test focuses on admissible candidates only, and supports our hypothesis that candidates at this stage are more homogeneous than in the entire sample, as the quantity and quality of publications have no significant impact, and only being in an

[20]These results are robust to the use of a logit and a probit model (see the supplementary material).

Table 8: Conditional likelihood to be admissible and promoted, University

	Likelihood to be admissible		Likelihood to be promoted			
			Admissibles		Close to threshold	
	(1)	(2)	(3)	(4)	(5)	(6)
Woman	-0.041	-0.060^{c}	0.004	-0.017	0.037	0.054
	(0.037)	(0.035)	(0.074)	(0.074)	(0.099)	(0.105)
Age		-0.028^{c}		-0.062		-0.071
		(0.015)		(0.041)		(0.066)
Age^2		0.000		0.002		0.003
		(0.000)		(0.001)		(0.002)
Publisher(Pub)		0.235^{a}		-0.086		-0.100
		(0.049)		(0.105)		(0.152)
Pub*Quantity		0.097^{a}		0.038		0.018
		(0.021)		(0.048)		(0.071)
Pub*Quality		0.066^{a}		0.020		0.026
		(0.009)		(0.017)		(0.024)
Int. Department		0.039		0.202^{a}		0.190^{c}
		(0.039)		(0.071)		(0.098)
Ile de France		0.047		-0.015		0.020
		(0.037)		(0.071)		(0.097)
Pos. other than univ.		-0.033		0.072		0.032
		(0.039)		(0.088)		(0.128)
Time FE	Yes	Yes	Yes	Yes	Yes	Yes
R^2	0.035	0.204	0.048	0.114	0.050	0.112
Observations	781	781	236	236	145	145

Notes: The likelihood to be admissible is conditional on applying. The likelihood of promotion is conditional on being admissible. Linear probability model (OLS). Standard errors clustered by individuals in brackets. ^{a}, ^{b}, ^{c} Significant at the 1%, 5% and 10% level, respectively. Pos. other than univ.: assistant professor positions in economics outside France, as well as CNRS researchers, and assistant professors from other disciplines.

international department matters. The regressions indicate that there is a small, negative and insignificant effect of being a woman, which amounts to only 2% of the probability of success of admissible men (69%). The last two columns consider the $2n$ candidates close to the threshold, and report a positive and not significant coefficient, which is nevertheless of considerable size (10% of the probability of success which is 50% by construction).

Interpreting the results

Our results find a large but not significant gender gap in the probability to be promoted conditional on applying. A possible interpretation is that this gap is zero and observed promotions are entirely due to the application gap. Alternatively, we can take the opposite extreme view and decompose the overall probability of being promoted into the two

steps, ignoring differences in parameter significance. Denote by $Pr(A, m)$ the probability that men apply and by $Pr(S, m)$ that of succeeding conditional on applying, and so on. The gap in promotions between men and women can be expressed as $Pr(A, m) \times Pr(S, m) - Pr(A, f) \times Pr(S, f)$, which can in turn be written as $Pr(S) \times \Delta \Pr(A) + Pr(A) \times \Delta \Pr(S)$, where $\Delta \Pr(.) \equiv Pr(., m) - Pr(., f)$ is the gap in probabilities and $Pr(.)$ is the (unweighted) average $Pr(.) \equiv (Pr(., m) + Pr(., f))/2$. The first term denotes the differences due to the gap in application propensity, while the second term captures the gap due to differences in success conditional on being a candidate.

For both tracks, we decompose the overall gap into the probability of applying and that of being promoted conditional on applying. In the case of professors, the latter is the product of the likelihood of being admissible and that of success conditional on being admissible, hence we further decompose into these two stages.[21] We next use the coefficients reported in Tables 5, 7, and 8 (columns 2, 3, and 4, respectively), to compute which share of these contributions is explained by observables other than gender and how much is due to gender or 'unexplained'.

The top panel in Table 9 reports the results for professors, giving the probabilities for the two groups, the gap, and the contribution of each stage to the overall promotion gap, with the last two columns presenting the explained and unexplained shares. Almost 70% of the raw gender gap for professors is due to the application gap, and its contribution is even larger (76%) when we control for characteristics. The share of the gap that is explained by characteristics is small (16.4%) and favours women. For the CNRS we find that the gap in application propensity explains 55.3% of the raw gap, most of which (44.3 percentage points) is not explained by characteristics. Although not as strong as in the university sample, these results confirm that the application stage is crucial, and while men now have better characteristics than women, the share explained by characteristics remains moderate. Three clear results emerge. First, the application stage plays an important role as it explains between half and three-quarters of the overall promotion gap. Second, the gender gap in the probability of being a candidate is mainly driven by

[21] The details of the decomposition are provided in the appendix.

the direct effect of gender rather than by differences in publications. Third, for the sample for which we have data, our results indicate a gender gap in the admissibility stage of the contest but virtually none latter on.

Table 9: Decomposition of the unconditional probability of promotion

	Women	Men	Diff.	Contr. (%)	Expl. (%)	Unexpl. (%)
			University			
Applying	0.060	0.090	-0.030	69.6	-6.9	76.5
Promotion/Apply	0.181	0.216	-0.035	30.4	-9.5	39.9
Admiss./Apply	0.266	0.314	-0.048	28.3	-7.1	35.4
Promotion/Admiss.	0.680	0.688	-0.008	2.1	-2.1	4.2
Promotion	0.011	0.020	-0.009	100.0		
			CNRS			
Applying	0.115	0.203	-0.088	55.3	11.0	44.3
Promotion/Apply	0.133	0.211	-0.077	44.7	7.8	36.4
Promotion	0.015	0.043	-0.027	100.0		

Notes: Contr.: contribution of each probability to the overall promotion gap. Expl. (resp. Unexpl.): shares of the overall promotion gap that are due to the explained (rep. unexplained) term in each of the probabilities. The probability of promotion conditional on applying (Promotion/Apply) is the product of the probability of being admissible conditional on applying and of being promoted conditional on being admissible.

Given the importance of the gap in the propensity to apply, we would like to assess whether women apply too little or men apply too much. This is not straightforward since we have no direct measure of which potential candidates should apply. To try to answer this question we look at who is a candidate at different points of the distribution of the predicted probability of success. We compute a measure of $X_{it}\widehat{\beta}$ for each potential candidate using the coefficients obtained in the equation for the probability of being admissible (column 2 in Table 8) and we interpret it as a measure of the 'quality of the potential candidate'.[22] We compute this measure of quality for all potential candidates and its distribution. Figure 2 plots for men and women the non-parametric functions fitted to the probability of being a candidate as a function of $X_{it}\widehat{\beta}$ (centered around its mean) as well as the 95% confidence intervals. In order to increase the readability of the

[22]We consider only professors as there are too few observations for the CNRS. The results are similar if we use the coefficients for the probability of promotion. We choose the probability of being admissible as this equation has a higher R^2, indicating that we explain better this stage than the promotion stage. This difference is partly due to the fact that promotions depend on the candidate's quality but also on the number of positions available.

graph, we removed the top 5% of the $X_{it}\widehat{\beta}$ distribution.[23] Table 10 reports the gaps for the various quartiles of the distribution, as well as for the top 5%.

Figure 2: Graph of the non-parametric function f of the probability of being a candidate $P(C) = f(X_{it}\widehat{\beta})$ of admissibility, trimming top 5% of $X\widehat{\beta}$

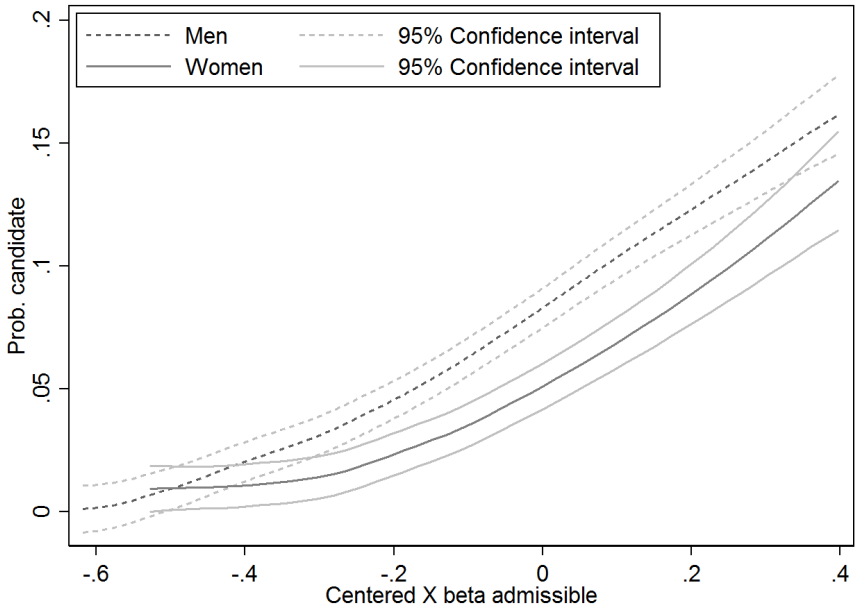

If men were applying too much we would expect the gap to be mainly at the bottom of the distribution, while if women were applying too little, we would expect to see large gaps at the top where the probability of promotion is substantial. At the bottom of the distribution, men do not seem to overapply as the gap is small for low values of $X_{it}\widehat{\beta}$. It then increases rapidly and remains roughly constant at 3 percentage points up to the 95th percentile. For the top 5% of the distribution, the gap increases dramatically: 25% of men and 15.8% of women apply, implying a 9 point gap. The relative gap is bell-shaped. It is particularly large in the second quartile where the probability of promotion is very low, it then declines steadily to 18% for the group between the 75th and 95th percentiles, but rises sharply again at the very top (a gap of 37%). The probability of promotion is three times as high for the top 5% than for the 75-95 percentile, yet the share of women that are candidates increases by only 10% while that for men rises almost by 44%. It is difficult to conclude from these results whether men overapply or women underapply. What we find is that men are very likely to apply in the second quartile where the probability of

[23]The range of values for this group is large (from 0.398 to 1.11 after centering), which made it difficult to appreciate visually the gap for the rest of the distribution.

success is low, and that the best women (those at the top 5%) are not increasing their likelihood to apply in a way that correspond to the increase in their probability of success when compared to next group down the distribution.

Table 10: Share of candidates by quartiles of the distribution of $X_{it}\widehat{\beta}$ of admissibility

	Women	Men	Diff.	in % / men	Prob. prom.
Total sample	0.060	0.090	-0.030^{a}	-34	0.017
	(0.005)	(0.004)	(0.006)		(0.001)
1st quartile	0.011	0.014	-0.002	-17	0.001
	(0.004)	(0.003)	(0.005)		(0.001)
2nd quartile	0.017	0.051	-0.034^{a}	-66	0.001
	(0.005)	(0.006)	(0.009)		(0.001)
3rd quartile	0.071	0.105	-0.033^{b}	-32	0.010
	(0.010)	(0.008)	(0.013)		(0.002)
Top quartile	0.146	0.190	-0.045^{b}	-23	0.053
	(0.014)	(0.010)	(0.018)		(0.005)
75-95 percentiles	0.143	0.174	-0.031	-18	0.038
	(0.016)	(0.011)	(0.020)		(0.005)
Top 5%	0.158	0.250	-0.092^{c}	-37	0.099
	(0.037)	(0.025)	(0.048)		(0.013)

Notes: Standard errors in brackets. ^{a}, ^{b}, ^{c} Significant at the 1%, 5% and 10% level, respectively. The values of $X_{it}\widehat{\beta}$ are calculated using all coefficients but the one associated with Woman from column (2) of Table 8.

As we saw above, women also have a lower conditional probability of being promoted, although the coefficient is not significant. This seems to be the result of a lower probability of being admissible and a positive effect of being female among borderline admissible candidates. Given how homogeneous candidates are in the last stage of the *agrégation*, we would expect that conscious discrimination would appear more among admissible candidates, a stage where it is difficult for someone outside the committee to assess differences between candidates. We find that, if anything, women are *positively* discriminated.

To further assess these differences, we compute the probability of a candidate being admissible and of an admissible candidate being promoted as a function of their $X_{it}\widehat{\beta}$ for the admissible stage. Figure 3 depicts these probabilities. The probability of being admissible is similar at the bottom of the distribution but a considerable gap appears for high-quality candidates. The promotion probability amongst admissibles shows an interesting pattern, as at the top of the distribution the probability of success of women is higher than that of men and close to one. These two apparently contradictory re-

Figure 3: Graphs of the distributions of the conditional likelihood to be admissible and promoted at a function of $X_{it}\widehat{\beta}$ of admissibility

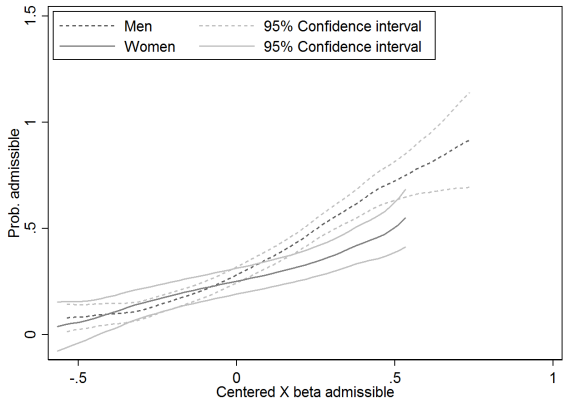

Probability of being admissible

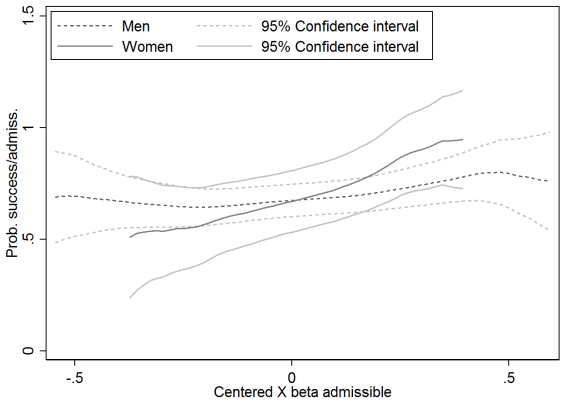

Probability of success among admissible

sults have a possible explanation. The negative effect at the admissibility stage could be the result of either men being better at 'promoting' their work (either because men overestimate their contribution or women underestimate it) or of women being subject to unconscious discrimination by committee members (the committee underestimate the contribution of a woman). In contrast, at the last stage of the competition, conscious positive discrimination could be in operation, favouring high productivity women.

Assessing the bias from selection on unobservables

Given the central role of the probability to apply, we performed a number of robustness tests of our results, looking at different subsamples, such as first-time applicants, and including additional variables such as coauthor networks, all of which confirmed the robustness of the effect we have identified.[24] A further concern is that despite our attempts to control for observable factors, our estimates could be biased by unobservable factors correlated with both research output and selection into seeking promotion. In this section we assess the likelihood that the estimates are biased by unobservables by following the strategy proposed by Altonji et al. (2005) who provide a measure of the relative importance of selection on unobservables. It evaluates how much stronger selection on unobservables must be relative to selection on observables in order to explain the effect of our variable of interest. The measure consists of comparing the coefficients

[24]Results are available in the supplementary material.

of two regressions, one with a restricted set of control variables $\widehat{\beta}_R$, and one with a full set of controls, $\widehat{\beta}_F$. The measure of the maximum importance of unobservables is given by the ratio $\widehat{\rho} = \widehat{\beta}_F / \left(\widehat{\beta}_R - \widehat{\beta}_F\right)$. The numerator is the effect the robustness of which is tested, while the denominator captures the fact that the smaller the gap between the two coefficients is, the less the estimate is affected by selection on observables, implying that selection on unobservables needs to be relatively stronger to explain the effect estimated in the full regression. Oster (2016) proposes a modified version of this ratio to take into account the explanatory power of the observables through changes in the R^2 of the regressions. She proposes to compute the ratio $\widehat{\delta} = \widehat{\rho}\,(R_F^2 - R_R^2) \,/\, (R_{\max}^2 - R_F^2)$, where R_F^2 and R_R^2 are the R-squared of the full and restricted regressions, while $R_{\max}^2$ is a measure of the maximum R^2 expected in the data. We follow Oster (2016) and use $R_{\max}^2 = 1.3 \times R_F^2$.[25] Both papers propose a rule of thumb such that when $\widehat{\rho}$ or $\widehat{\delta}$ are above 1 the probability that unobservables can drive the value estimated is low.

Table 11: Selection on unobservables

Likelihood to:	apply		be promoted	
	$\widehat{\rho}$	$\widehat{\delta}$	$\widehat{\rho}$	$\widehat{\delta}$
University	−9.3	−29.5	−2.7	−7.5
CNRS	3.2	8.8	13.9	36.4

Notes: $\widehat{\rho} = \widehat{\beta}_F/(\widehat{\beta}_R - \widehat{\beta}_F)$, $\widehat{\delta} = \widehat{\rho}(R_F^2 - R_R^2)/(0.3\ R_F^2)$. $\widehat{\beta}_R$ and R_R^2 are the gender coefficients and R-squared of columns (1) and (3) of Table 5 (likelihood to apply) and of columns (1) and (4) of Table 7 (likelihood to be promoted conditional on applying). $\widehat{\beta}_F$ and R_F^2 are the gender coefficients and R-squared of columns (2) and (4) of Table 5 (apply) and of columns (3) and (6) of Table 7 (be promoted).

Table 11 considers two sets of regressions. For the likelihood to apply, our restricted regressions are those in columns (1) and (3) of Table 5, with only gender as an explanatory variable. As full regressions we consider columns (2) and (4) of Table 5 that include age, research output, and department of origin. For the conditional probability of being promoted, we use as restricted regressions columns (1) and (4) of Table 7, and columns (3) and (6) as full regressions. None of the ratios is less than one. It is even always negative

[25] Oster proposes the value $R_{\max}^2 = 1.3 \times R_F^2$ following an analysis of papers in top economics journals performing randomised experiments. She chooses the cut-off so that the results in 90% of the published papers are robust. Nevertheless, we performed robustness exercises by looking at the cases of $R_{\max}^2 = 1.5 \times R_F^2$ and $R_{\max}^2 = 1.7 \times R_F^2$. Both yielded values for $\widehat{\delta}$ above 3 in absolute value (see the supplementary material).

for the University. Their absolute value ranges from 2.7 to 29.5 for the University and from 3.2 to 36.4 for the CNRS, with averages of 12.3 and 15.6 respectively, implying that in order to attribute the entire value of the coefficient on gender to selection effects, selection on unobservables would have to be at least twice that on observables and, on average, 14 times as strong. These are large values, which make us conclude that the coefficient on gender is unlikely to be fully explained by selection on unobservables.

V Conclusions

This paper uses data for promotions among academic economists in France to look at competing explanations for the gap in promotions between men and women, exploiting the fact that we have information both on who was a candidate for promotion and who could have been a candidate. On the one hand, women may be less likely to be promoted conditional on having applied, potentially due to discrimination; on the other, female academics may have a lower propensity to apply for promotion than males. The features of the French academic system, such as a national salary scale, the need to go through a national contest in order to be promoted, and the existence of several categories of academics with different requirements during the contest and upon promotion, allow us to examine various hypotheses for these differences.

We find that women are less likely than men to enter promotion contests, and that this gap accounts for between 50 and 76% of the promotion gap. Although we are unable to assess whether women underapply or men overapply, we find a particularly large gap at the top of the research-quality distribution, raising the question of whether some of the best women are self-censoring themselves. The conditional probability of promotion is also lower for women, but the coefficient is not significant. Our detailed analysis of the various stages of the competition indicates a negative effect of being a women in the early stages but a positive one amongst similar candidates in the last stage.

There are various possible explanations for our findings. Women could derive lower utility from promotion and this would reduce their likelihood to seek it. Also, although women are not discriminated against during the contest they may believe they will be, and

hence decide not to enter the competition. The alternative explanation is that women are less willing than men to enter contests, in line with experimental evidence. It is difficult to distinguish between these hypotheses with the data available. However, the fact that the objective cost to seek promotion as a researcher is virtually null implies that even if the benefits from promotion and the expected probability of success were low, all individuals should apply for promotion. Hence not doing so indicates an unwillingness to enter the contest per se, which is therefore stronger for women.

Our results have two main implications. First, they indicate that women's lower propensity to take part in contests is observed in actual labour markets and that it can have important consequences for observed outcomes. Second, they raise the question of what type of policy intervention can help increase female promotion rates. The evidence in Gneezy et al. (2003) indicates that differences in contest participation are driven by women being less confident than men, while Azmat and Ferrer (2017) highlight the importance of aspirations in explaining gender gaps in performance. Building confidence and changing aspirations is a dificult process and probably starts in early childhood, but changes in the way in which deciding whether to enter a competition occurs could facilitate women's probabilities of climbing up the rank ladder. For example, a system of mentoring whereby junior faculty are assigned a mentor that 'proposes' them as candidates for promotion may increase women's propensity to apply. Alternatively, creating a system in which the default is that an individual will be considered for promotion after x years and the individual has to opt out instead of opting in, could also be a way of overcoming the differences across genders.

Appendix

A Decomposition

Denote by $Pr(A,m)$ (resp. $Pr(A,f)$) the probability that men (resp. women) apply, by $Pr(S,m)$ (resp. $Pr(S,f)$) that of succeeding conditional on applying and by $\Delta\Pr(.) \equiv Pr(.,m) - Pr(.,f)$ the gap in any probabilities between men and women. The gap in promotions (P) between men and women can then be expressed as:

$$\begin{aligned}\Delta\Pr(P) &= Pr(P,m) - Pr(P,f) = Pr(A,m) \times Pr(S,m) - Pr(A,f) \times Pr(S,f) \\ &= \frac{Pr(S,m) + Pr(S,f)}{2}\,\Delta\Pr(A) + \frac{Pr(A,m) + Pr(A,f)}{2}\,\Delta\Pr(S). \quad \text{(A.1)}\end{aligned}$$

The first term denotes the differences due to the gap in application propensity, while the second term captures the gap due to differences in success conditional on being a candidate.

Adding the admissible stage for the University track, with $Pr(Ad,.)$ the probability of being admissible conditional on applying and with $Pr(S,.)$ now that of succeeding conditionnal on being admissible, the gap in promotions can be expressed as:

$$\begin{aligned}\Delta\Pr(P) &= Pr(A,m) \times Pr(Ad,m) \times Pr(S,m) - Pr(A,f) \times Pr(Ad,f) \times Pr(S,f) \\ &= \frac{Pr(Ad,m)\ Pr(S,m) + Pr(Ad,f)\ Pr(S,f)}{2}\,\Delta\Pr(A) \\ &\quad + \frac{Pr(A,m) + Pr(A,f)}{2}\ \frac{Pr(S,m) + Pr(S,f)}{2}\,\Delta\Pr(Ad) \\ &\quad + \frac{Pr(A,m) + Pr(A,f)}{2}\ \frac{Pr(Ad,m) + Pr(Ad,f)}{2}\,\Delta\Pr(S). \quad \text{(A.2)}\end{aligned}$$

The explained part 'Expl.' of the overall promotion gap is

$$\frac{Pr(S,m) + Pr(S,f)}{2}\ (\bar{X_A^f} - \bar{X_A^f})\ \widehat{\beta_A} + \frac{Pr(A,m) + Pr(A,f)}{2}\ (\bar{X_S^m} - \bar{X_S^f})\ \widehat{\beta_S}, \quad \text{(A.3)}$$

and the unexplained part 'Unexpl.' is then $\frac{Pr(S,m)+Pr(S,f)}{2}\ \widehat{\beta_A^f} + \frac{Pr(A,m)+Pr(A,f)}{2}\ \widehat{\beta_S^f}$.

References

Altonji, J. G., Elder, T. E., and Taber, C. R. (2005). Selection on Observed and Unobserved Variables: Assessing the Effectiveness of Catholic Schools. *Journal of Political Economy*, 113(1):151–184.

Azmat, G. and Ferrer, R. (2017). Gender Gaps in Performance: Evidence from Young Lawyers. *Journal of Political Economy*, 125(5):1306–1355.

Bagues, M. and Esteve-Volart, B. (2010). Can Gender Parity Break the Glass Ceiling? Evidence from a Repeated Randomized Experiment. *Review of Economic Studies*, 77(4):1301–1328.

Bagues, M., Sylos-Labini, M., and Zinovyeva, N. (2017). Does the Gender Composition of Scientific Committees Matter? *American Economic Review*, 107(4):1207–1238.

Bertrand, M., Chugh, D., and Mullainathan, S. (2005). Implicit Discrimination. *American Economic Review*, 95(2):94–98.

Bertrand, M. and Hallock, K. F. (2001). The Gender Gap in Top Corporate Jobs. *Industrial and Labor Relations Review*, 55(1):3–21.

Blackaby, D., Booth, A. L., and Frank, J. (2005). Outside Offers and the Gender Pay Gap: Empirical Evidence from the UK Academic Labour Market. *Economic Journal*, 115(501):81–107.

Blau, F. and Kahn, L. M. (2000). Gender Differences in Pay. *Journal of Economic Perspectives*, 14(4):75–99.

Bosquet, C., Combes, P.-P., and Garcia-Peñalosa, C. (2018). Supplementary material to 'Gender and Promotions: Evidence from Academic Economists in France'. Available on https://sites. google.com/site/clementbosquet/.

Buser, T. (2016). The Impact of Losing in a Competition on the Willingness to Seek Further Challenges. *Management Science*, 62(12):3439–3449.

Cole, J. and Cole, S. (1973). *Social Stratification in Science*. Chicago: The University of Chicago Press.

Combes, P.-P. and Linnemer, L. (2010). Inferring Missing Citations. A Quantitative Multi-Criteria Ranking of all Journals in Economics. *GREQAM WP 2010-25*.

Combes, P.-P., Linnemer, L., and Visser, M. (2008). Publish or Peer-Rich? The Role of Skills and Networks in Hiring Economics Professors. *Labour Economics*, 15(3):423–441.

DiNardo, J., Fortin, N. M., and Lemieux, T. (1996). Labor Market Institutions and the Distribution of Wages, 1973-1992: A Semiparametric Approach. *Econometrica*, 64(5):1001–1044.

Farber, S. (1977). The Earnings and Promotion of Women Faculty: Comment. *American Economic Review*, 67(2):199–206.

Gibson, J., Anderson, D. L., and Tressler, J. (2014). Which Journal Rankings Best Explain Academic Salaries? Evidence from the University of California. *Economic Inquiry*, 52(4):1322–1340.

Ginther, D. K. and Hayes, K. J. (1999). Gender Differences in Salary and Promotion for Faculty in the Humanities. *American Economic Review Papers and Proceedings*, 89(2):397–402.

Ginther, D. K. and Kahn, S. (2004). Women in Economics: Moving Up or Falling Off the Academic Career Ladder? *Journal of Economic Perspectives*, 18(3):193–214.

Gneezy, U., Niederle, M., and Rustichini, A. (2003). Performance in Competitive Environments: Gender Differences. *Quarterly Journal of Economics*, 118(3):1049–1074.

Gobillon, L., Meurs, D., and Roux, S. (2015). Estimating Gender Differences in Access to Jobs. *Journal of Labor Economics*, 33(2):317–363.

Goldin, C. and Rouse, C. (2000). Orchestrating Impartiality: The Impact of "Blind" Auditions on Female Musicians. *American Economic Review*, 90(4):715–742.

Gupta, N. D., Poulsen, A., and Villeval, M. C. (2013). Gender Matching and Competitiveness: Experimental Evidence. *Economic Inquiry*, 51(1):816–835.

Johnson, G. E. and Stafford, F. P. (1974). Lifetime Earnings in a Professional Labor Market: Academic Economists. *Journal of Political Economy*, 82(3):549–69.

Kahn, S. (1995). Gender Differences in Academic Career Paths of Economists. *Journal of Economic Perspectives*, 9(4):193–204.

Lavy, V. (2008). Do Gender Stereotypes Reduce Girls' Human Capital Outcomes? Evidence from a Natural Experiment. *Journal of Public Economics*, 92(10-11):2083–2105.

Lazear, E. P. and Rosen, S. (1990). Male-Female Wage Differentials in Job Ladders. *Journal of Labor Economics*, 8(1):S106–S123.

McDowell, J. M., Singell, L. D., and Ziliak, J. P. (2001). Gender and Promotion in the Economics Profession. *Industrial and Labor Relations Review*, 54(2):224–244.

Niederle, M. and Vesterlund, L. (2007). Do Women Shy Away from Competition? Do Men Compete Too Much? *Quarterly Journal of Economics*, 122(3):1067–1101.

Niederle, M. and Vesterlund, L. (2011). Gender and Competition. *Annual Review of Economics*, 3(1):601–630.

Oster, E. (2016). Unobservable Selection and Coefficient Stability: Theory and Evidence. *Journal of Business and Economic Statistics*, forthcoming.

Sabatier, M. (2010). Do Female Researchers Face a Glass Ceiling in France? A Hazard Model of Promotions. *Applied Economics*, 42(16):2053–2062.

Wenneras, C. and Wold, A. (1997). Nepotism and Sexism in Peer-Review. *Nature*, 387(22):341–343.

Winter-Ebmer, R. and Zweimuller, J. (1997). Unequal Assignments and Unequal Promotion in Job Ladders. *Journal of Labor Economics*, 15(1):43–71.