Gender and Promotions: Evidence from Academic Economists in France *

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

July 8, 2016

Abstract

The promotion system for French academic economists provides a test for competing explanations of the promotion gap between men and women. In particular, promotions occur through national competitions for which we have information both on candidates and on those eligible to be candidates. We find that women are less likely to seek promotion than men, while gender has no significant effect on candidates' promotion rates. After eliminating a number of possible explanations for the gap in seeking promotion, our preferred interpretation is that women are less willing than men to participate in contests, in line with recent experimental evidence.

JEL Classification: J16, J7, I23

Key words: gender gaps, promotions, academic labour markets

*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, the editor of this journal and three referees, as well as seminar and conference participants. 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 SERC (LSE). 33 Boulevard du Port, 95011 Cergy-Pontoise Cedex, France. Email: clement.bosquet@u-cergy.fr. Website: https://sites.google.com/site/clementbosquet/

‡Aix-Marseille School of Economics and Sciences Po Paris. Also research fellow at the CEPR. 2 Rue de la Charité, 13236 Marseille cedex 02, France. Email: ppcombes@gmail.com. Website: http://www.vcharite.univ-mrs.fr/pp/combes/

§Aix-Marseille School of Economics. 2 Rue de la Charité, 13236 Marseille cedex 02, France. Also research fellow at CESifo. Email: cecilia.garcia-penalosa@univ-amu.fr. Website: http://www.vcharite.univ-mrs.fr/PP/penalosa/

1 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 percent of top executives in US firms and for France, Gobillon, Meurs and Roux (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.[1] 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, to look at various possible causes for the lower promotion rates of females and discriminate among them.

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.[2] 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.

Although the academic labour market is a very particular one, focusing on it has certain advantages. Unlike in many private sector jobs, where a promotion is associated with longer hours and a requirement for greater availability outside normal working 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 such as participating in

[1] For an overview of work on gender wage gaps see Blau and Kahn (2000). On one hand of the spectrum, Goldin and Rouse (2000) provide evidence of discrimination against women, on the other, Ichino and Moretti (2009) trace gender gaps to biological differences.

[2] See, for example, the analysis of auditions for orchestra membership and promotion in Goldin and Rouse (2000).

committees, seeking funding or performing editorial activities, these activities are not mandatory and not performing them would not imply that the individual is demoted. 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 several potential causes of the unexplained component of promotions gaps from associate to full professor. 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 against them. On the other, 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 differentiate to some extent between these hypotheses.

We use data for academic economists in France over the period 1991 to 2008, and find lower promotion rates for women, which are partly but not completely explained by the age structure and publication records. 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*. We find a negative but insignificant coefficient on gender on the later, even when restricting the sample to candidates close to admission threshold, thus finding no support for the presence of discrimination against women. In contrast, being a woman has a substantial negative impact on the likelihood to enter the *concours*.

We explore several potential explanations for this difference. First, we try to understand if the gap is the result of men applying too much or women too little, the evidence suggesting that the latter effect is in operation. Second, we find no evidence that the objective costs of undertaking the *concours* are the causes of the gender gap, implying that our results are consistent with both 'perceived' discrimination that would discourage women from applying as

well as with men and women having different attitudes towards contests.

The paper adds to a vast body of work on the different rates of promotion across genders; see, amongst others, Lazear and Rosen (1990) and Winter-Ebmer and Zweimuller (1997). In their highly 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 similar questions but use much more recent data, and find no evidence of discrimination for academics in the 1990s and 2000s.

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). Early on, empirical analyses identified both lower wages and lower promotion rates for female economists, while more recent work indicates that in the US salary gaps are explained by differences in academic rank, while promotions to tenure and to full professor rank are still affected by gender even after controlling for research output and demographic characteristics.[3] Evidence for the UK by Blackaby, Booth and Frank (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; see, for example, Lavy (2008), Bagues and Esteve-Volart (2010) and Bagues, Sylos-Labini and Zinovyeva (2015). The results of this literature generally do 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 work has addressed differences between male and female attitudes towards competition using mainly experimental evidence. This literature has proposed an alternative explanation for women's poorer performance in labour markets. It finds that women

[3]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, Singell and Ziliak (2001) and Ginther and Kahn (2004).

tend to be less likely to enter competitive situations and tend to perform worse in competitive situations than men. Since promotions to top jobs tend to be highly competitive, such behaviour could explain why we observe fewer women in these jobs. Our conclusions can be interpreted as being supportive of what the experimental literature has found. In line with the results in Niederle and Vesterlund (2007), Niederle and Vesterlund (2011) and Gupta, Poulsen and Villeval (2013) we find that once we control for observed productivity, including publications' quality, women are less likely than men to enter the *concours* and that this difference seems to be the result of women applying too little rather than of overconfident men applying too much.[4]

The paper is organised as follows. After describing the French academic system, section 2 examines the possible reasons why women are less likely to be promoted. Section 3 describes the data, an exhaustive panel of academic economists in France over the period 1991 to 2008. Our results are presented in section 4, while section 5 concludes.

2 Why are there so few female professors?

2.1 The French academic system

The French academic system has a number of features that we intend to exploit to test possible causes of the low rate of promotion of females in academia. 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.[5] Researchers in this category are hired by the CNRS, who pay their salaries, but are attached to a university and are hence a member of its economics department. Researchers have the possibility to undertake some teaching and they participate in department life in the same way as standard professors.

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* (DR) 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 and a much steeper slope for salaries over time. The salary scales are set by the Ministry of Higher Education and Research, and are identical for university

[4] See as well as the discussion in Bertrand (2011).

[5] CNRS stands for *Centre National de la Recherche Scientifique.*

professors and researchers as well as across departments.[6]

Promotions take place through a national contest, a *concours*, and are thus not decided by the department in which the individuals hold their current position.[7] Participation in this contest and the final ranking is 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. 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.

Members of the *concours*' committees are academic economists, drawn from various areas of expertise and universities, not necessarily the most prestigious ones. Because they have to be of the full professor rank and because of the age distribution of the population of academics, there is a strong male dominance in these committees. Committees change regularly, every two to four years depending on the particular instance.

The requirement of the *concours* differs across the two academic tracks. For university professors the contest is termed *concours d'agregation*. Over our sample period, the *agregation* was biannual and entailed four stages over, approximately, a 6-month period.[8] It included a research seminar, and three oral exams both in the candidate's field and in economics in general, one of which consisted of preparing in 24 hours a lecture on a topic randomly drawn by the applicant from a lengthy predetermined list. 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 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 terms of the costs involved in the two systems. 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 are usually not able to stay in the university where they held a rank B position.[9]

[6]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.

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

[8]The *agregation* was suspended in 2014 and a new promotion system has been put in place.

[9]In our sample, 81.8% of those promoted through the *agregation* have a new affiliation after their promotion. Only 3.4% of available positions were within commuting distance of Paris (urban area). On average, successful

After three years they are allowed to move to another university, if the latter wishes to recruit them, including their former university. For researchers, academics that are promoted can choose to stay at the university where they are or move to another department. The university does not need to have an open position for them since the researchers' salary is always paid by the CNRS.

It is obvious that promotion is very costly for university professors. There is a cost that in principle is gender-neutral, the cost of preparing the various exams. The *agregation* also involves 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.

2.2 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. Understanding these two processes would be of interest even if promotion rates did not differ across men and women, as it is possible for equal rates to be the result of differences in the two stages that compensate each other. 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 in and 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 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

candidates had to move 674 kilometers.

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

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

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_j are vectors of coefficients.

The first possible explanation for the promotion gap across genders is hence that those making the promotion decisions discriminate against women. Discrimination can be taste-based, statistical or implicit and may appear in different forms. Taste-based discrimination implies that female candidates are less likely to be promoted, i.e. $p_{j0}^f < 0$. A negative p_{j0}^f is also consistent with statistical discrimination if committees believe that for a given publication track women are on average less deserving. Alternatively, committees may value the research of women less, i.e. $p_{j1}^f < 0$. This can occur if women contribute less than their male coauthor to a joint paper and the committee engaged in statistical discrimination, or be due to implicit discrimination if the publications of females are perceived as being of lower quality.[10] Discrimination can then reduce female promotion rates directly as well as indirectly as a lower probability of success will reduce the likelihood that women apply. Note, also, that what is relevant for the decision to apply is not the actual probability of success in the *concours* but the expected probability. Consequently, if women believe that they will be discriminated against, i.e. that p_{j0}^f and/or p_{j1}^f are negative, then 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 $C_{ij} = c_j + c_j^f \delta_i$. For researchers, cost parameters are given by $c_r = c_0$ and $c_r^f = c_0^f$, where the terms c_0 and c_0^f 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

[10] See Bertrand, Chugh and Mullainathan (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.

to fear losing tournaments or contests more than men.[11] Higher subjective costs of competing hence may lead women to apply less often than men even if the objective costs and benefits are the same as those of men.

We suppose that these subjective costs are the same across tracks,[12] 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. In the US system, there is no reason for women to have different objective costs of promotion from men. In contrast, in the French system, there may be differences across the sexes: the opportunity cost of time could be higher for women if they undertook a disproportionate amount of domestic work, making it more costly for them to prepare the *agregation*, while 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.

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 and is equal across tracks.[13] 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 try to identify discrimination by considering the impact of being female on the conditional probability of promotion. Discrimination is captured by $p_{j0}^f < 0$ and $p_{j1}^f < 0$, and implies a negative impact of being female in the two tracks. 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 *agregation* are higher than those of seeking promotion as a researcher, we have $c_{u0}^f > c_{r0}^f$, 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,

[11] See, for instance, McGregor and Elliot (2005) and Buser (2016).

[12] 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'agregation,* 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 *agregation* solicits much greater interest from the academic community, with results at each stage being widely followed and discussed, which increases the preasure on candidates. It is hence difficult to argue that one or the other contest is more competitive.

[13] The fact that the promotion is decided at the national level implies we do not have the quality versus promotion tradeoff found in the US, where more prestigious departments also have tougher promotion thresholds; see McDowell et al. (2001).

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 $c_0^f > 0$.

2.3 Empirical specification

We consider the same specification for both the probability to apply $\Pr_j(A, i, t)$ and the probability of success conditional on applying, $\Pr_j(S, i, t)$. For an outcome O where $O = A, S$, we have

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

Either probability is a function of age Age_{it} and its square, whether or not individual i has published in Econlit-classed journals (i.e. whether individual i is a 'publisher', measured by the dummy Pub_{it}), the number of publications and the average quality of these publications, 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.

These last two variables have several interpretations. They could measure the positive effects that being in a more stimulating academic environment has, be an indicator of peer pressure to seek promotion, or simply capture higher unobserved ability since the best junior candidates will have been recruited by the departments with highest research quality or in a particularly attractive location, i.e. Paris and its commuter belt. There are as well externalities due to having colleagues who are also preparing the *agregation* since candidates often work together and share the burden of preparing lectures on the various topics. Our robustness tests will add further control variables relating to the role of coauthor networks in particular or having some controls interacted with the woman dummy δ_i.

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. Overall we are interested in whether the coefficients β_{j0}^f in the two regressions are significantly different

from zero, as well as in the differences between the coefficients estimated for the two tracks, at the university or CNRS. We assume that the probability is linear and use OLS in the regressions. This choice is made for two reasons. First, we use interacted terms to test different hypotheses and this leads to interpretation issues in nonlinear regressions (see Puhani (2012)). Second, we perform some Oaxaca-Blinder decompositions which would have been also imprecise with nonlinear specifications. Results are robust to the use of a logit or a probit model (results available upon request).

3 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 to 2008. For each individual we have information on age, rank, publications (see below) and department. We keep only individuals that are in departments larger than 4 full-time equivalent academics, which removes economists that are isolated in universities without real economics departments. We would have liked to control for characteristics of individuals' family life, such as whether they are married and the number and ages of children, but such data were not available. For the US, Ginther and Kahn (2004) find that having children has only a weak effect of the promotion probabilities of female economists in the US and none on their productivity. Moreover, having children and deciding to pass the *agregation* may be simultaneous decisions, as women may choose to postpone childbirth until after promotion or postpone the *agregation* because of maternity decisions, thus it is difficult to view family circumstances as an exogenous variable determining promotions.

We use the list of candidates that applied to and those who succeeded in becoming a rank A university professor or a rank A CNRS researcher (DR). 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.

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. Two issues are relevant. First, it is possible for CNRS researchers to take the *agregation*, 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 in our data no such individuals and

the promotions are de facto considered to be internal. We will hence consider both those that hold a rank B position as professors or at the CNRS as potential candidates for the *agregation* but only those with a CNRS rank B position as potential candidates to promotion to DR. The second question is how to deal with older candidates, who may have characteristics that may make them choose not to apply for promotion (poor publication record, unsuccessful past applications, low unobserved ability, etc.). In order to avoid having in our pool an increasing number of candidates unlikely to apply we introduce age limits. We hence consider as potential candidates only individuals aged between 32 and 55.[14]

We define two categories of department, 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 Bosquet et al. (2014) for details). 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.

3.1 Measuring research output

One of our key explanatory variable is an individual's research output. We measure individual output at date t by the cumulative publication record between the first year for which we have an observation for the individual and date t.[15] Publication records are measured as weighted sums of publications. All publications come from the EconLit database, which includes more than 560,000 papers published in more than 1200 journals –including 46 in French– between 1969 and 2008. Three dimensions enter the weighted scheme of publications: the quality of journals, the number of authors and the publication's relative number of pages.

We measure the quality of publications using the journals weighting scheme proposed by Combes and Linnemer (2010). They 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.[16] Two different degrees of convexity in the distribution of journals'

[14] Our results hold with different upper bounds, and the results without age limits are available at Bosquet, Combes and Garcia-Peñalosa (2014). A robustness excercise for younger individuals is performed below. 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. These figures indicate that women do no leave academia more often than men.

[15] 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.

[16] Gibson, Anderson and Tressler (2014) show that this is the journal ranking that best explains the salaries at the University of California economics departments.

weights 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. We divide each publication by the number of authors, a standard practice in the literature, 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.

The output of individual i at date t is then 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, $\bar{p}_t$ is 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.

3.2 Descriptive statistics

Table 1: Percentage rank A by gender and track

rank A	Total	%	Women	%	Men	%
	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	3,212	82.8	6,585	60
1	5,052	34	666	17.2	4,386	40
	CNRS researchers (15% of population)					
0	1,525	58.3	474	79	1,051	52.1
1	1,093	41.7	126	21	967	47.9

All years are pooled together, every two years over 1992-2008 for university members and every year over 1996-2008 for the CNRS.

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 to estimate

the probability of being a candidate and of being promoted conditional of being a candidate. Table 1 gives the decomposition of this sample by gender in terms of institutional affiliation and rank. There were between 1,165 and 2,143 academic economists in France depending on the year, and the vast majority of the sample consists of university professors, with researchers accounting for only 15.0% of the total. Women account for 25.6% of observations, and they are over-represented amongst university professors and under-represented amongst researchers, where they account for only 22.9% of the population. 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%). Women do not seem to choose a career path that offers higher average promotion rates, which would counterbalance negative discrimination. 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 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, leave academia or the country over our sample period, or reach the maximum age for us to consider them as potential candidates. Males are much more likely than females to apply for promotion. The probability of being a candidate to the *agregation* 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 table indicates that some of these differences are likely to be explained by differences in observable characteristics. In our sample, 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, especially given the strong tradition in France to publish books and the large number of national journals, some of which are not in EconLit. Male professors publish more articles than female professors but we find no

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)
Network size	0.93	0.99	0.06	2.34	2.46	0.12
	(0.03)	(0.03)	(0.04)	(0.18)	(0.11)	(0.20)
% Men in Network	0.694	0.798	0.104^a	0.689	0.775	0.086^a
	(0.012)	(0.007)	(0.013)	(0.024)	(0.013)	(0.025)

All years are pooled together, every two years over 1992-2008 for university members and every year over 1996-2008 for the CNRS. Standard errors in brackets. ^{a}, ^{b}, ^{c} 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.

significant difference in the average quality of publications. In the CNRS, differences are more pronounced and indicate that men have a greater quantity of publications.

Table 2 indicates that the one aspect in which women fare better than men is affiliation: in the larger sample 30.9% of women are in international departments, while only 27.3% of men are; women are also more likely to be in a department in or around Paris (*Ile de France*).[17] We also construct for each individual two measures of research networks. Our measures are based on coauthorship, obviously an imperfect measure of actual networks, but one that is quantifiable with the EconLit data. Our first measure is the size of an individual's network, defined as the total number of different coauthors s/he has had over her/his publishing lifetime; the second is the fraction of network members that are men. Men do not have larger coauthor networks but have larger shares of other men in their networks.

Table 3 reports descriptive statistics for the sample of actual candidates to promotion (left

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

Table 3: Descriptive statistics of candidates and promoted candidates (age 32 to 55)

	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)
Journal quality	1.14	1.59	2.13	6.08^c	2.85	3.33	3.54	15.07
	(0.36)	(0.27)	(0.49)	(1.03)	(1.41)	(0.66)	(2.16)	(3.86)
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)
Network size	1.95	1.45^a	5.22	3.75^b	2.53	1.94	5.17	4.82
	(0.17)	(0.08)	(0.75)	(0.25)	(0.49)	(0.19)	(1.83)	(0.69)
% Men in Network	0.675	0.820^a	0.721	0.813^c	0.716	0.851^b	0.795	0.816
	(0.030)	(0.016)	(0.044)	(0.021)	(0.059)	(0.026)	(0.057)	(0.040)

See Table 2. 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.

panel) and for those promoted (right panel). 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 *agregation* 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 between men and women are not statistically significant. In general, we find no significant difference in the quantity or quality of publications. The differences between men and women among promoted candidates are even weaker, with no significant gap in output. Actual candidates have higher publication records, higher probability to be in international departments, larger coauthor networks and higher shares of men in their networks than potential candidates, and promoted candidates have even better characteristics than actual candidates.

We have seen that there are only small differences in the average research output of male and female potential candidates. It is however possible that there are significant differences

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

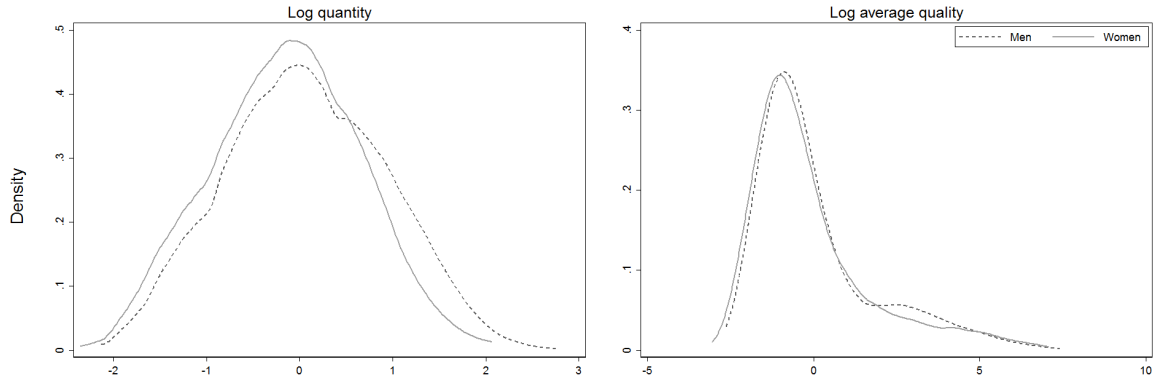

in the distribution of output and since promoted candidates are selected from the top tail, differences in the thickness of this tail could explain observed promotion differences. Figure 1 thus plots the distribution of the (log) quantity and (log) average quality of research output for men and women. The distribution of quantity for female professors is further to the left than that for males and specially has less mass at the top of the distribution, indicating that the most productive men publish more than the most productive women. In contrast, there is little difference in the distribution of quality, although we observe a mild bimodality for men that is absent for women.

A caveat is in order at this point. In what follows we will explain promotions as a function of research output, and we have seen that women are less productive in terms of quantity; our analysis will hence focus on whether women are less likely to be promoted than men conditional on research output. The latter is however not exogenous. In fact, if women expect not to be promoted (because they expect to be discriminated or because they believe they will not seek promotion in the future), they may have lower incentives to publish. As a result, their lower publication record could be the result of a lack of promotion rather than its cause. If this were the case, our results would be underestimating the negative impact of being a woman on promotions.

4 Results

4.1 The promotion of academic economists

We start by examining the determinants of *holding* a rank A position. The results are reported in Table 4, where all specifications include year fixed effects. In column (1), only time fixed effects

and gender are included in the model, with the dummy *Woman* being equal to 1 for women. The marginal effect on gender is significant at the 1% level and large at -0.196. Including age reduces the gender effect to -0.155, indicating that a large fraction of the difference in promotion is indeed due to the fact that women in the sample are younger than men. Column (3) includes our three measures of research output: whether or not the individual publishes, the quantity of publications and their quality. All three are highly significant and increase the probability of promotion.

Table 4: Likelihood to hold a rank A position

	(1)	(2)	(3)	(4)	(5)	(6)
Woman	-0.196^a	-0.155^a	-0.042^b	-0.040^b	-0.038^b	-0.040^b
	(0.020)	(0.020)	(0.017)	(0.017)	(0.017)	(0.017)
Age		0.013^a	0.009^b	0.009^b	0.009^b	0.009^b
		(0.004)	(0.004)	(0.004)	(0.004)	(0.004)
Age^2		0.000	0.000^b	0.000^b	0.000^b	0.000^b
		(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
Publisher(Pub)			0.318^a	0.331^a	0.329^a	0.329^a
			(0.022)	(0.023)	(0.027)	(0.022)
Pub*Quantity			0.181^a	0.184^a	0.191^a	0.186^a
			(0.010)	(0.010)	(0.011)	(0.010)
Pub*Quality			0.039^a	0.041^a	0.041^a	0.041^a
			(0.005)	(0.005)	(0.005)	(0.005)
CNRS				-0.121^a	-0.124^a	-0.120^a
				(0.031)	(0.031)	(0.030)
Woman*CNRS				-0.035	-0.018	-0.036
				(0.057)	(0.059)	(0.057)
Woman*Pub					0.003	
					(0.048)	
Woman*Pub*Quantity					-0.039^c	
					(0.023)	
Woman*Pub*Quality					0.001	
					(0.011)	
Time FE	Yes	Yes	Yes	Yes	Yes	Yes
R^2	0.042	0.090	0.356	0.361	0.362	0.361
Observations	12,408	12,408	12,408	12,408	12,408	12,408

Linear probability model (OLS). Standard errors clustered by individuals in brackets. ^{a}, ^{b}, ^{c} Significant at the 1%, 5% and 10% level, respectively. In column (6), quantity and quality are replaced by scores putting more weights to recent publications (using a degressive function). Terms interacted with Woman are centered so that the coefficient associated with Woman is still the average effect of being a woman.

Once we control for research output, the effect of gender falls to -0.042, about a third of the one in column (2), indicating that the lower promotion rate for women is to a large extent due to them having published less. The gender effect is nevertheless still strong: being a woman reduces the probability of promotion almost as much as reducing by a quarter the individual's number of single-authored publications. The effects we obtain are lower but comparable in magnitude to those found for the US. Ginther and Kahn (2004) find a raw gender gap of -0.213,

which falls to -0.130 once age and publication records are included, indicating that differences in publications across genders explain a greater fraction of the gap in France than in the US.

Column (4) includes a dummy for being a CNRS researcher and its interaction with the gender dummy, and finds that the negative effect of being a woman on seniority is stronger for the CNRS, although this term is insignificant. Column (5) includes interactions between the gender dummy and publication scores and indicates that men and women do not have significantly different returns to publications.[18] Column (6) replaces quantity and quality with scores putting more weight on recent publications (using a degressive function) and shows that results are robust to these degressive publication scores.

4.2 Decomposing outcomes

The limitation of our analysis so far is that we do not know whether lower observed promotions are the result of a lower likelihood to apply for promotion or lower success in obtaining the promotions. We hence separately examine the two 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.

4.2.1 Likelihood to enter the *concours*

The results concerning the determinants of the likelihood to enter the *concours* are reported in Table 5. As expected, research output has a strong effect on the likelihood of being a candidate, with quality having a stronger effect for researchers.[19] 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 in the case of professors, where it tends to increase the probability of being a candidate, as expected. In the regression for the agregation we also include a dummy for whether the candidate holds a rank B CNRS position.

The unconditional probability of applying is lower for women than for men both for university professors (column 1) and researchers (column 3), and the negative impact of being female remains as large once we include individual characteristics, the marginal effect being -0.033 for professors and -0.071 for researchers. The results indicate that women have a lower likelihood

[18] Only *Quantity* has a marginally significant coefficient, at the 10% level. In what follows we will not consider different returns to publications as we found no significant effect for our other specifications.

[19] The results are robust to the inclusion of higher order terms to capture some non-linearity in the impact of publications' quality.

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

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.

to enter the contest for promotion. In terms of magnitude, the effect is large. Being a woman is equivalent to dividing the number of single-authored publications by around 2.

In order to try 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 costs of the *agregation* contest are substantially larger than those of applying for promotion at the CNRS, the coefficient on women should be higher for university professors than for researchers if differences across the genders 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.[20] The direct gender effect remains negative, while the interacted coefficient is negative but insignificant at the 10% level, implying that the effect of gender on the likelihood to enter the promotion contest is not higher for standard professors than in the research track. If anything, the larger coefficient for CNRS implies a stronger reluctance for women to enter this *concours*, not the university one.

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. In fact, the significant effect of the *Ile de France* dummy was entirely driven by its impact on women. To obtain the overall impact of gender, recall 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. This difference has several possible interpretations: it could be that the high concentration of economics departments in and around Paris makes it more likely that successful candidates can change department without moving their residence thus making women more willing to apply, 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 difference in the likelihood that men and women academics apply for promotion. The negative and significant coefficient is robust to different specifications of the model (results available upon request), such as the logit, the probit, and controlling for time since the first publication instead of age. We also run (not reported) a random effects model as well as a duration model where we estimate time to first application for promotion; both yielded equivalent results. The application gender gap has three possible causes. First, women may value promotions less, for example because they have a lower marginal utility of income; second, they may expect to be discriminated against and hence have a lower expected probability of success; third, they can have higher disutility costs from participating in contests. Below we will try to discriminate between these explanations.

[20] To save space, other interacted terms are not reported but are available upon request. Note that the specification in column (5) gives the same results as a weighted comparision between those in columns (2) and (3).

4.2.2 Gender differences in success in the *concours*

Table 6 reports results for the probability of being promoted conditional on having applied for promotion. The first four columns indicate 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. This is consistent with the fact that there is a stronger selection for those joining the CNRS than for those following the standard university track, implying that the latter are more likely to be non-publishers and have fewer publications. Being in an international department has a positive impact on the probability of passing the *agregation*, in line with our arguments above, while it has no impact on the probability of being promoted in the CNRS contest.

Table 6: Likelihood to be promoted conditional on applying

	University		CNRS		University: Selected candidates			
	All candidates		All candidates		"Admissibles"		Close to threshold	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Woman	-0.029	-0.046	-0.068	-0.063	0.004	-0.017	0.037	0.054
	(0.033)	(0.032)	(0.059)	(0.061)	(0.074)	(0.074)	(0.099)	(0.105)
Age		-0.032^b		-0.042		-0.062		-0.071
		(0.014)		(0.039)		(0.041)		(0.066)
Age^2		0.001		0.001		0.002		0.003
		(0.000)		(0.001)		(0.001)		(0.002)
Publisher(Pub)		0.157^a		0.028		-0.086		-0.100
		(0.048)		(0.137)		(0.105)		(0.152)
Pub*Quantity		0.074^a		0.075^c		0.038		0.018
		(0.020)		(0.041)		(0.048)		(0.071)
Pub*Quality		0.053^a		0.039^b		0.020		0.026
		(0.009)		(0.018)		(0.017)		(0.024)
Int. Department		0.093^a		-0.010		0.202^a		0.190^c
		(0.036)		(0.069)		(0.071)		(0.098)
Ile de France		0.023		0.067		-0.015		0.020
		(0.033)		(0.074)		(0.071)		(0.097)
Pos. other than univ.		0.001				0.072		0.032
		(0.036)				(0.088)		(0.128)
Time FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
R^2	0.030	0.164	0.043	0.197	0.048	0.114	0.050	0.112
Observations	781	781	198	198	236	236	145	145

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 for promotion, 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. Gender simply does not significantly matter.[21] We can interpret the insignificant coefficient as a lack of evidence of *ex-post* discrimination against women in the contest. This does not mean that discrimination does not matter, since expectations of discrimination may have an *ex-ante* effect on women's choices and hence on observed outcomes, as discussed above.

The absence of a gender gap in the conditional promotion probability in the two tracks can, however, be due to discrimination being offset by selection. Since fewer women are candidates than men, it is possible that they are drawn from the top of the distribution and hence that their unobserved ability is higher than that of men; see Petrongolo and Olivetti (2008) on selection and gender. The insignificant coefficient could then be the result of opposing effects canceling out: higher unobserved ability of women and discrimination against them.

To try to control for the role of selection, we perform two further tests exploiting the fact that for the *agregation* contest we have information on which individuals passed each stage of the competition, as well as the final rank. At the various stages of the competition, a number of candidates are eliminated, and after the penultimate stage a list of candidates that are *admissibles* is provided by the jury. These candidates then undertake the final test, after which a ranking of all *admissible* candidates is provided and, if p positions are available that year, the top p candidates are promoted. At this late stages, a substantial number of applicants have been eliminated, thus remaining candidates are presumably of similar observed and unobserved abilities. That is, focusing on the 'best' should eliminate the (presumably predominantly male) bottom tail of the distribution. 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 test focuses on *admissible* candidates only. Columns (5) and (6) of Table 6 present those regressions and support 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 international department matters. The regressions indicate that there is no significant difference between men and women in either the raw rate of passing the last stage of the contest or the one obtained after controlling for publications. To make our sample even more comparable, we construct from the final ranking of candidates a list of individuals close to the threshold. To do so we take, for each year, the n *admissible* candidates that did

[21] These results are robust to the use of a logit and a probit model.

not get promoted and the n lowest-ranked candidates that did get promoted and rerun the same regressions. The last two columns of Table 6 show that, if anything, women are *positively* discriminated, although the coefficient is not significant.

4.2.3 Interpreting the results

Our results indicate that being a woman does not significantly reduce the probability of being promoted conditional on entering the promotion contest, while they show that women are less likely than men to apply for promotion. In order to gauge the importance of women's lower propensity to apply, we perform a Oaxaca decomposition. For professors in our sample, the raw gender gap in the probability of being promoted in a given year to rank A amounts to 1 percentage point. Although this magnitude seems small, the probability of being promoted in a given year is only 2% for men and hence twice as large as that for women. As we have seen, this difference is largely driven by the decision to be a candidate, with an average male probability of being a candidate in a given year of 9% and that for women being 3 percentage points lower. In the CNRS, the annual raw probability of promotion is 37% lower for women than for men, and the raw probability of being a candidate 57% lower.

Table 7: Oaxaca decomposition of the likelihood to apply

	University		CNRS	
	$\Delta \Pr(C)$	%	$\Delta \Pr(C)$	%
Total gap	.029	100	.093	100
Woman	.033	115	.071	76
Characteristics	-.004	-15	.023	24
Age	-.015	-51	-.004	-5
Total outcome	.013	44	.028	30
Publisher	.01	33	.012	13
Quantity	.006	22	.02	22
Quality	-.003	-10	-.005	-5
International Department	0	-1	0	0
Ile de France	-.002	-6	0	0
CNRS	-.001	-2	.	.

To get rid of time fixed effects, the decomposition is made for each year separately before averaging over the full sample, with the number of observations per year as weights.

The first two columns of Table 7 present the Oaxaca decomposition for the probability of being a candidate to the *agregation*, based on column (2) of Table 5. Columns (3) and (4) present the same decomposition for the probability of being a candidate to a DR position at the CNRS, using column (4) of the same table. For both tracks, the gender gap in the

probability of being a candidate is mainly driven by the direct effect of gender. For professors, it amounts to over a 115% of the overall gap, the effect being offset by women being younger and their better department of origin. Men have a substantially better research output. Similar results are obtained for researchers; being a female accounts for 76% of the gap and a higher research output for 30%, with age favouring women. Both decompositions hence imply that the application gap comes mainly from the direct effect of being female, with women's lower publication record accounting for about a third of the difference.

Our decomposition indicates that the gap in seeking promotion is largely driven by a gap in the propensity to apply across the genders. Assessing whether women apply too little or men apply too much is not straightforward as we have no direct measure of which potential candidates should apply. Moreover, the question is particularly difficult to answer since it is possible that simultaneously women apply too little and men apply too much, with women who deserve promotion (i.e. at the top of the ability distribution) holding back and men at the bottom of the distribution applying when they have little chance of success.

To try to answer this question we consider three arguments. First, suppose that the gap is only due to men applying too much. Then the overconfident men would not be promoted and as a result their applications would be irrelevant implying that we should see no difference in the promotion rates of men and women conditional on age and research output. The fact that we do find gender differences in the probability of being rank A and these are not due to gender differences in the success rate of candidates can be interpreted as indicating that some women that should have seek promotion did not.

Second, if women were shying away from competition, we would expect to see female candidates being positively selected. Table 3 indicates that there are no significant differences across genders in terms of research output, while women are in better departments which could be an indication of positive selection. In fact, the gender gap in the probability of being in an international department or in the Paris region is substantially higher for candidates than for potential candidates (for international department the gap is three times as large for the former as for the latter, for the Paris region it is twice as large). Similarly, although there is no significant difference in network size amongst male and female potential candidates, amongst actual candidates women have networks that are on average a third larger than those of men. Overall, these figures indicate that female applicants are positively selected in terms of their department and network, a result that is consistent with the idea that women underapply.

Our third argument consists in looking at who is a candidate at different points of the distribution of the predicted probability of success.[22] We compute using the coefficients from equation (2) the value of $X_{it}\widehat{\beta}$ for each candidate, which is a measure of the 'quality of the potential candidate' that in turn determines the probability of success in the *agregation*, the latter depending obviously also on the number of positions available each year. We compute this measure of quality for all potential candidates and its distribution for each year. We then define the quartiles as well as top and bottom deciles of the (common across genders) distribution of quality. For each quartile/decile we then compute the share of women (men) in that part of the distribution that apply, i.e. the likelihood that a woman (man) in that part of the distribution seeks promotion. If women were applying too little, we would expect to see large gaps at the top of the distribution; if men were applying too much we would expect the gap to be mainly at the bottom of the distribution.

Table 8: Share of candidates by quartile/decile of the distribution of $X_{it}\widehat{\beta}$

	Women	Men	Diff.
Total sample	0.060	0.090	0.030^{a}
	(0.005)	(0.004)	(0.006)
1st decile	0.016	0.009	-0.007
	(0.008)	(0.004)	(0.008)
1st quartile	0.013	0.019	0.007
	(0.004)	(0.004)	(0.006)
2nd quartile	0.014	0.056	0.042^{a}
	(0.004)	(0.006)	(0.009)
3rd quartile	0.089	0.115	0.026^{c}
	(0.011)	(0.009)	(0.014)
Top quartile	0.130	0.168	0.038^{b}
	(0.014)	(0.010)	(0.018)
Top decile	0.139	0.190	0.051^{c}
	(0.022)	(0.017)	(0.029)

All years are pooled together, every two years over 1992-2008 for university members. 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 6.

The results are reported in Table 8. Although significance is not always high, the data indicates that the overall gap (reported in the first line) is mainly driven by differences at the top and middle of the distribution. For the first quartile and decile we find no gender gap, providing no support to the idea that men overapply. Differences are particularly large in the second and fourth quartiles, as well as in the top decile (tough the gap is only mildly significant in the latter case). These results indicate that even when their probability of success is high

[22] We consider only professors as there are too few observations for the CNRS.

women stay out of the competition, and are suggestive that they apply too little.

4.3 Robustness analysis

4.3.1 First applicants and young applicants

Table 9: Number of applications

	Actual candidates				Promoted candidates			
	University		CNRS		University		CNRS	
	Women	Men	Women	Men	Women	Men	Women	Men
Mean	1.46	1.64	2.82	2.53	1.36	1.41	3.33	2.35
	(0.05)	(0.04)	(0.35)	(0.23)	(0.08)	(0.05)	(0.71)	(0.27)
Median	1	1	3	2	1	1	3.5	2
Maximum	4	6	6	11	3	4	6	8

Standard errors in brackets.

To check the robustness of our results, we consider first applicants and young individuals. Men and women differ in the number of times they apply for promotion, as can be seen in Table 9. The differences in the average number of applications are not statistically significant, whether for candidates or for promoted candidates, and the median is equal across genders for professors and higher for women in the case of the CNRS. However, men show greater persistence, with the maximum number of applications being larger than for women both in the case of applicants and in that of successful applicants.

To consider these differences in more detail, Table 10 presents in the first four columns two subgroups: that of those who have never applied and that of those who have applied at least once. The regressions for *Apply 1st* hence give the probability that a potential candidate applies for promotion for the first time, while those for *Reapply* give the likelihood that an individual who has already applied but not succeeded applies again. For both tracks, women are less likely than men to be first applicants, whether in terms of the raw probability (not reported) or once we control for observables. The probability to reapply is also lower for women, but not significantly different across men and women, indicating that they differ in the decision to seek promotion but not in their persistence at trying to obtain it.[23]. The magnitude of the effects is large. For example, for professors, the probability that a man applies for the first time is 5.6% and being a woman reduces it by 2.4 percentage points (the corresponding figures for the CNRS are 9.2 and 3.6).

[23]To examine if past rejections discourage women, we also considered other specifications, for example with the number of times the individual has already seek promotion. These yielded equivalent results.

Table 10: First time applicants and young candidates

	Likelihood to apply				Young individuals, Univ.		
	Apply 1st		Reapply		Apply	Promoted	
	Univ. (1)	CNRS (2)	Univ. (3)	CNRS (4)	(5)	All (6)	Admiss. (7)
Woman	-0.024^a	-0.036^b	-0.032	-0.142	-0.053^c	0.092	0.176^b
	(0.006)	(0.018)	(0.031)	(0.121)	(0.028)	(0.071)	(0.088)
Age	-0.018^a	0.043^a	-0.076^a	-0.045	0.212	0.051	0.708^c
	(0.002)	(0.007)	(0.011)	(0.075)	(0.136)	(0.386)	(0.415)
Age^2	0.000^a	-0.001^a	0.001^a	0.001	-0.020	-0.009	-0.074^c
	(0.000)	(0.000)	(0.000)	(0.002)	(0.014)	(0.040)	(0.044)
Publisher(Pub)	0.082^a	0.048^c	0.164^a	-0.197	0.278^a	0.426^a	0.688^a
	(0.012)	(0.026)	(0.049)	(0.204)	(0.040)	(0.090)	(0.164)
Pub*Quantity	0.027^a	0.059^a	0.047^a	0.038	0.162^a	0.142^a	0.057
	(0.006)	(0.012)	(0.016)	(0.064)	(0.024)	(0.051)	(0.059)
Pub*Quality	0.010^a	0.011^b	0.029^a	0.051^b	0.028^a	0.065^a	0.048^a
	(0.003)	(0.005)	(0.010)	(0.023)	(0.009)	(0.017)	(0.017)
Int. Department	-0.004	0.011	0.019	0.016	0.030	0.060	0.040
	(0.007)	(0.021)	(0.032)	(0.125)	(0.042)	(0.099)	(0.103)
Ile de France	0.022^b	0.023	-0.064^c	-0.179	0.016	0.007	-0.089
	(0.009)	(0.026)	(0.036)	(0.140)	(0.055)	(0.088)	(0.096)
CNRS	-0.084^a		-0.266^a		-0.359^a		
	(0.008)		(0.073)		(0.039)		
Woman*Ile de France	0.007	-0.033	0.124^b	0.484^b	0.049		
	(0.013)	(0.034)	(0.060)	(0.182)	(0.058)		
Pos. other than univ.						-0.107	-0.162
						(0.085)	(0.121)
Time FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
R^2	0.083	0.144	0.221	0.216	0.243	0.292	0.340
Observations	6,803	951	1,282	181	746	183	100

See Table 6.

We next examine young individuals, so far excluded from our analysis. We consider individuals aged between 28 and 31 that are potential candidates to the *agregation* (we have no individuals in this age group applying for CNRS promotion). This sample is interesting because they are individuals that are at an early stage of their careers and hence have fewer publications; as a result seeking promotion early on can be seen as an indication of overconfidence. Table 10 reports regressions for the likelihood to enter the contest, as well as for the probability of being promoted conditional on having applied, both for all candidates and for the group of "admissibles"as in Table 6. The raw gap is 11.3 percentage points, as opposed to 3 in the core sample, and once we control for observables we find a coefficient which is about 60% larger than that in Table 6. Again, women are less likely to seek promotion than men, with the difference being larger in this age group than in our core sample but the coefficient being significant only at the 10% level, making it difficult to interpret these results. Columns (6) and (7) report regressions for the likelihood to be promoted conditional on having applied. As in our core sample, when

we consider all applicants we find an insignificant coefficient on *Woman* (column 6), while for the restricted sample the coefficient is positive and significant (column (7), the same result is obtained when we consider only individuals close to the threshold). That is, once the bottom tail of the distribution of applicants is eliminated, young female applicants are significantly more likely to be promoted than young male applicants. This could be the result of either positive discrimination or unobserved ability, our data not allowing us to distinguish between the two.

4.3.2 Selection, networks and the competitive environment

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 try to assess the robustness of the impact of gender on the likelihood to seek promotion by considering the impact of selection. Our first exercise consists in controlling for unobservables by including a dummy for whether the individual has already applied for promotion (*Already app.*). Our data does not capture all aspects that could be important in promotion decisions, such as teaching ability, involvement in administrative duties or how innovative the candidate's work is considered by the academic community. As a result, it is conceivable that if we were to control for these unobservables the coefficient on being a woman would become insignificant. To palliate this problem, we include the dummy *Already app.*.

Columns (1) and (2) of Table 11 report regressions including this dummy. Both for professors and for the CNRS the coefficient on this variable is positive and highly significant. Nevertheless, its interaction with Woman is not significant and introducing this control does not change the sign or the significance of the coefficient on Woman. For both tracks, this coefficient is smaller in absolute value than in Table 5, falling from -0.034 for professors (-0.072 for the CNRS) to -0.024 (-0.065, respectively). Note also that the R^2 increase considerably: by almost 50% for professors while it more than doubles for researchers. This indicates that the *Already app.* dummy partly controls for unobservables, as expected, but does not erode the impact of being a woman on the probability of seeking promotion.

Second, we investigate some further possible explanations of the application gender gap. One of them is that some individual variables that we have so far ignored have an impact on the actual (or perceived) probability of success. In particular, we consider how a potential candidate's decision is affected by her/his research network or group of coauthors. The idea that networks are important in obtaining jobs and achieving promotions is widespread in the

literature, and the issue has been addressed for promotions in academia; see McDowell and Smith (1992) for the US, Combes et al. (2008) for France and Zinovyeva and Bagues (2015) for Spain. Coauthor networks have been shown to differ across genders, with females having fewer coauthors and a lower fraction of male coauthors; see McDowell and Smith (1992) and Boschini and Sjögren (2007). If women have smaller or less efficient networks, then this may affect the expected outcome and hence the payoff from entering the competition for promotion. Potential reasons for this effect are that candidates who have a member of their network in a promotion committee have a higher likelihood of success, but also that with a larger network the candidates' work may be better known and cited, or that this could provide extra information about how to best prepare for a *concours*. As a result, the gap that we find between men and women could be due to differences in coauthors and networks, and hence including these variables would have an effect on the coefficient on gender.

Recall that our measures are based on coauthorship, with the size of an individual's network being defined as the total number of different coauthors s/he has had over her/his publishing lifetime. More than in other contexts, academic networks are highly endogenous, with research output and affiliation being both causes and consequences of an individual's network size, and patterns of network formation have been shown to differ across the genders; see Boschini and Sjögren (2007). A caveat is in order. Networks are not exogenous and being a candidate for promotion can increase an individual's network, especially in a system where applicants are often candidates several times, raising questions about which way causality runs. As a result any effect we find has to be interpreted with care.

Table 11 reports in columns (3) and (4) the regressions for the likelihood of entering the promotion contest for the two tracks, to which we have added network variables. We find a positive and significant effect of the proportion of men and the size of the network in, respectively, the regression for professors and for researchers, indicating that networks matter. However, for both tracks, the coefficient on gender and the R^2 barely change as compared to the specifications in the previous two columns. That is, gender differences in networks do not seem to explain the lower propensity of women to seek promotion.

Our final test consists of examining whether the environment under which the competition takes place affects men and women differently. Since the effects of these variables can be identified only over the time dimension of the data we consider only the concours d'agregation for which we have a much larger number of observations than for the CNRS. In line with

Table 11: Networks and the competition's environment: Likelihood to apply for a promotion

	Selection		Networks		The environment		
	Univ.	CNRS	Univ.	CNRS	University		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Woman	-0.024^a	-0.063^a	-0.023^a	-0.060^a	-0.021^a	-0.021^a	-0.021^a
	(0.006)	(0.023)	(0.006)	(0.022)	(0.006)	(0.006)	(0.006)
Already applied	0.164^a	0.621^a	0.164^a	0.614^a	0.166^a	0.166^a	0.166^a
	(0.016)	(0.052)	(0.016)	(0.051)	(0.016)	(0.016)	(0.016)
Woman*Already app.	-0.006	-0.202^c	-0.006	-0.203^c	-0.014	-0.014	-0.015
	(0.029)	(0.121)	(0.029)	(0.115)	(0.029)	(0.029)	(0.029)
Pub*Network size			-0.002	0.034^c	-0.002	-0.002	-0.001
			(0.010)	(0.019)	(0.010)	(0.010)	(0.010)
% Men in Network			0.022^b	0.029	0.023^b	0.023^b	0.023^b
			(0.011)	(0.024)	(0.011)	(0.011)	(0.011)
% Women rank A							0.019
							(0.015)
Woman*% Wom. rank A					0.009^a	0.008^b	-0.010
					(0.003)	(0.004)	(0.021)
Avail.Pos.							0.002^a
							(0.001)
Woman*Avail.Pos.						-0.000	-0.000
						(0.001)	(0.001)
Time trend							-0.013^b
							(0.006)
Woman*Time trend							0.007
							(0.008)
Time FE	Yes	Yes	Yes	Yes	Yes	Yes	No
Usual controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes
R^2	0.152	0.439	0.153	0.444	0.154	0.154	0.153
Observations	8,085	1,132	8,085	1,132	8,085	8,085	8,085

Linear probability model (OLS). Standard errors clustered by individuals in brackets. ^{a}, ^{b}, ^{c} Significant at the 1%, 5% and 10% level, respectively. Usual controls include age, age squared, the three publications' variables, the dummies for being in an international department, in Ile de France and in the CNRS and the interaction between Woman and being in Ile de France.

existing literature (Bagues et al., 2015), we considered the gender composition of the national committees that decide on promotion. We run regressions that include a dummy equal to 1 if there was at least one woman in the committee (more than one is rare) and found no significant impact (results not reported).

Table 11 reports results for two alternative variables that affect the competitions environment. First, we consider the proportion of women amongst rank A professors in the year of the competition, which can act as a proxy for 'perceived discrimination' or be seen as a measure of role models for young women. This variable has increased steadily over the period, going from 8.9% in 1992 to 15.2% in 2008. The other variable is the number of available positions as professor in a particular year, which is usually known before the decision to be a candidate is taken, and fluctuates substantially, ranging between 15 and 33. Obviously, when more positions

are available the probability of promotion is higher. If women were more risk-averse, we would expect a stronger reaction than for men in years in which rank A positions were plentiful.

Because we are exploiting the time dimension of the data we cannot include time fixed effects as well as the variables of interest for the entire population. We hence consider two specifications. One includes time fixed effects and our two variables of interest interacted with the female dummy, which allows us to see whether effects differ across the genders but not to identify the impact for each gender separately. Alternatively, we substitute the fixed effects by a time trend and estimate the impact of these variables for both men and women.

Table 11 column (5) indicates that a larger fraction of women amongst rank A professors increase the likelihood that women seek promotion, consistent with both the importance of role models and the idea that perceived discrimination is a potential cause for females' lower propensity to apply for promotion. Nevertheless, when we include a female specific time trend (column 7) neither the share of female professors nor the trend (both interacted with the female dummy) have significant coefficients, indicating that it is difficult to distinguish between the way in which changes in female presence in this particular occupation and wider trends of female presence in professional life affect young female academics. As a result, we cannot either confirm or reject a role for role models or perceived discrimination. In either case, the direct impact of being a woman falls but remains significantly negative.

When we consider the number of available positions we find that a greater number of available positions increases the likelihood of being a candidate, capturing the idea that more positions imply a greater probability of being promoted and hence encourages individuals to apply, but the term interacted with the female dummy is never significant implying there are no differences across genders in the reaction to a greater probability of being promoted.

4.3.3 Assessing the bias from selection on unobservables

A further way of assessing the likelihood that the estimates are biased by unobservables is to follow the strategy proposed by Altonji, Elder and Taber (2005) and Oster (2015). Altonji et al. (2005) provide a measure of the relative importance of selection on unobservables. In particular, their measure captures how much stronger selection on unobservables must be relative to selection on observables in order to explain the effect of our variable of interest that has been estimated. The measure consists of comparing the coefficients 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 importance of unobservables is given by the ratio $\widehat{\rho} = \widehat{\beta}_F / \left(\widehat{\beta}_R - \widehat{\beta}_F\right)$. The numerator measures the effect that needs to be explained, 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 (2015) proposes a modified version of this ratio which takes into account the explanatory power of the observables through changes in the R^2 of the regressions. In particular she proposes to compute the ratio $\widehat{\delta} = \widehat{\rho}\left(R_F^2 - R_R^2\right) / \left(R_{\max}^2 - R_F^2\right)$, 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 (2015) and use 1.3 times R_F^2 as the latter.

Table 12: Selection on unobservables

Full from	Table 5 cols. (2) & (4)		Table 11 cols. (1) & (2)		Table 11 cols. (3) & (4)	
	$\widehat{\rho}$	$\widehat{\delta}$	$\widehat{\rho}$	$\widehat{\delta}$	$\widehat{\rho}$	$\widehat{\delta}$
University	−9.3	−29.5	4.4	14.2	3.8	12.2
CNRS	3.2	8.8	2.5	7.5	2.1	6.4

Full specifications are run without interacted terms.

Table 12 considers two sets of regressions. 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 three specifications: the regressions in columns (2) and (4) of Table 5 that include age, research output, and department of origin, as well as the regressions reported in Table 11 which add the dummy for whether the individual has already applied for promotion (columns (1) and (2)) and those that also consider networks (columns (3) and (4)). We compute both the basic ratio $\widehat{\rho}$ as well as $\widehat{\delta}$. Of the 12 ratios reported in Table 12, none is less than one. Their absolute value ranges from 2.1 to over 29.5, with an average of 8.7, 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, over 8 times as strong. This are large values, well above the commonly used threshold of a ratio of 1. This makes us conclude that the coefficient on gender is unlikely to be fully explained by selection on unobservables.

5 Conclusions

This paper uses data for promotions amongst 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 there is no significant difference between men and women in their probability of success conditional on being a candidate, but that the latter are less likely than the former to enter promotion contests. Moreover, our results indicate that the differential cost of promotion (the mobility associated with it) for women is important but does not fully explain the gender gap. The gender gap could have two causes: women applying too little or men applying too much, the latter being consistent with experimental evidence that finds male overconfidence in certain situations. Although measuring exactly who should apply for promotion is not possible with our data, a number of results point towards the gap being driven by too few female applicants.

There are various possible explanations for our findings. First, women could derive lower utility (monetary or otherwise) from promotion and this would reduce their likelihood to seek it. Second, although women are not discriminated against during the contest they may believe they will be, and hence decide not to enter the competition for promotion. The alternative explanation is that women are less willing than men to enter contests, in line with the experimental evidence provided by Gneezy, Niederle and Rustichini (2003), amongst others. 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 (all that is required is to submit a cv) 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 provide a link between laboratory

evidence and behaviour observed in actual markets. They indicate that patterns of behaviour, namely women's lower propensity to take part in contests, can be observed in actual labour markets and that they 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 seem to be partly driven by women being less confident than men. Building confidence is a process that is difficult 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 provide more incentives for women 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.

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.

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. (2015), 'Does the Gender Composition of Scientific Committees Matter?', *IZA Discussion Paper* (9199).

Bertrand, M. (2011), New Perspectives on Gender, *in* O. Ashenfelter and D. Card, eds, 'Handbook of Labor Economics', Elsevier, pp. 1543–1590.

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.

Boschini, A. and Sjögren, A. (2007), 'Is Team Formation Gender Neutral? Evidence from Coauthorship Patterns', *Journal of Labor Economics* **25**(2), 325–365.

Bosquet, C., Combes, P.-P. and Garcia-Peñalosa, C. (2014), 'Gender and Promotions: Evidence from Academic Economists in France', *LIEPP Working Paper* (29).

Buser, T. (2016), 'The Impact of Losing in a Competition on the Willingness to Seek Further Challenges', *Management Science* (forthcoming).

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 Working Paper* (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.

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', *The 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.

Ichino, A. and Moretti, E. (2009), 'Biological Gender Differences, Absenteeism, and the Earnings Gap', *American Economic Journal: Applied Economics* **1**(1), 183–218.

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.

McDowell, J. M. and Smith, J. K. (1992), 'The Effect of Gender-sorting on Propensity to Coauthor: Implications for Academic Promotion', *Economic Inquiry* **30**(1), 68–82.

McGregor, H. A. and Elliot, A. J. (2005), 'The Shame of Failure: Examining the Link between Fear of Failure and Shame', *Personality and Social Psychology Bulletin* **31**(2), 218–231.

Niederle, M. and Vesterlund, L. (2007), 'Do Women Shy Away from Competition? Do Men Compete Too Much?', *The 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. (2015), 'Unobservable Selection and Coefficient Stability: Theory and Evidence', *mimeo* .

Petrongolo, B. and Olivetti, C. (2008), 'Unequal Pay or Unequal Employment? A Cross-Country Analysis of Gender Gaps', *Journal of Labor Economics* **26**(4), 621–654.

Puhani, P. A. (2012), 'The Treatment Effect, the Cross Difference, and the Interaction Term in Nonlinear "Difference-in-Differences" Models', *Economics Letters* **115**(1), 85–87.

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.

Zinovyeva, N. and Bagues, M. (2015), 'The Role of Connections in Academic Promotions', *American Economic Journal: Applied Economics* **7**(2), 264–292.