



SciencesPo.

Department of Economics

Discussion paper 2013-17

**GENDER AND COMPETITION:
EVIDENCE FROM ACADEMIC
PROMOTIONS IN FRANCE**

**Clément Bosquet
Pierre-Philippe Combes
Cecilia Garcia-Peñalosa**

Sciences Po Economics Discussion Papers

Gender and Competition: Evidence from Academic Promotions in France *

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

October 24, 2013

Abstract

Differences in promotion across genders are still prevalent in many occupations. Recent work based on experimental evidence indicates that women participate less in or exert lower effort during contests. We exploit the unique features of the promotion system for French academics to look at women's attitudes towards competition in an actual labour market. Using data for academic economists over the period 1991-2008 we find that, conditional on entering the competition, there is no difference in promotions across the genders, which is difficult to reconcile with either discrimination or a poorer performance of women in contests. In contrast, women have a substantially lower probability than men to enter the promotion contest. Our data does not support that this gap is due to differences in costs or in preferences concerning department prestige, indicating that women are less willing than men to take part in contests.

JEL Classification: J16, J7, I23

Key words: gender gaps, promotions, academic labour markets

*This paper has benefited from the excellent research assistance of Philippe Donnay and Charles Laitong. We are grateful to Marco Manacorda, Esteban Aucejo, Javier Ortega, Felix Weinhardt, Stefano Bosi, Rodolphe Dos Santos Ferreira, Vélayoudom Marimoutou, Michel Mougeot, Boris Racine, as well as participants at seminars at the LSE and Greqam for helpful suggestions and assistance. This project has been supported by the Agence Nationale de la Recherche, grant ANR-11-BSH1-0014.

[†]London School of Economics and Political Science (SERC). Houghton Street, London WC2A 2AE, United-Kingdom. Email: c.bosquet@lse.ac.uk. Website: <https://sites.google.com/site/clementbosquet/>

[‡]Aix-Marseille School of Economics and Sciences Po. Also affiliated to 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. 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 particularly in seniority. The literature on gender wage gaps is vast, but differences across genders in promotions have received much less attention.¹ 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. Differences in characteristics account for a large proportion of the gap, with the remaining fraction being usually attributed to preferences or discrimination. A recent literature has proposed an alternative explanation, namely that women tend to be less likely to enter competitions and tend to perform worse in competitions than men². Since promotions to top jobs tend to be highly competitive, such behaviour could explain why we observe fewer women in these jobs. Existing evidence is based on experiments conducted in laboratory conditions, and although the results present strong support for gender differences in attitudes during and towards contests, it is not obvious that such attitudes are also prevalent in actual labour markets.

This paper uses the particular features of the French academic system, namely the fact that promotions occur through national contests, to look at the effect of gender on attitudes towards competition. In the US and the UK academic systems or in private sector jobs only information on actual promotions is available, while the list of applicants tends not to be known.³ In France, academics are public servants and promotion occurs through a national contest or *concours*, with the lists of applicants being publicly available at the time of the *concours*. We can hence examine how gender affects both the likelihood that an academic applies for promotion and the probability of promotion conditional on having applied. As a result, we can ask whether the lower observed promotion rate of women is due to lower success in competitions or to their unwillingness to enter them.

Using data for academics has important 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 committees, seeking funding or performing editorial activities, these activities are to a large extent voluntary. They are not ‘required’ by the promotion 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, unlike for most types of jobs, the key consideration in actual promotion deci-

¹For an overview of work on gender wage gaps see Blau and Kahn (2000). See Lazear and Rosen (1990) for a theoretical analysis of gender gaps in promotion and Winter-Ebmer and Zweimuller (1997) and Goldin and Rouse (2000) for empirical evidence.

²See Gneezy, Niederle and Rustichini (2003), Niederle and Vesterlund (2007) and Gupta, Poulsen and Villeval (2013)

³An exception is the analysis of auditions for orchestra membership and promotion in Goldin and Rouse (2000).

sions, an individual's productivity (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. On the one hand, women may be less likely to be promoted conditional on having applied, and this could be due either to discrimination or under-performance during the *concours*. 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, to their facing a different tradeoff between salaries and department prestige, or simply to an unwillingness to enter the contest. The special 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 test for 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*. The latter is unaffected by gender, thus rejecting the idea that females underperform in contests, and lending no support to the presence of discrimination against women. In contrast, being a woman has a substantial negative effect on the likelihood to enter the *concours*.

We explore several potential explanations for this difference. We find no evidence that either a different trade-off between income and department prestige nor the cost undertaking the *concours* and the mobility implications associated with it (see below) are the causes of the gender gap. The results seem to indicate that women are less willing than men to enter competitions for promotion, and this could be due either to different attitudes towards contests or to the expectation of being discriminated against. When we look at features of each year's *concours*, we find evidence consistent with both aspects.

Our contribution to the literature is twofold. First, we provide evidence on female attitudes towards competition in an actual labour market and hence add to the growing body that has addressed this question but which has so far been based on experimental evidence; see Gneezy et al. (2003), Niederle and Vesterlund (2007), Niederle and Vesterlund (2011) and Gupta et al. (2013). Our conclusions are partly supportive of what the experimental literature has found. In line with the results in Niederle and Vesterlund (2007), Niederle and Vesterlund (2011) and Gupta et al. (2013) we find that once we control for observed productivity women are less likely than men to enter the *concours* and that this difference does not depend on the visible costs associated with the contest. In contrast, we find no evidence that women are less likely to be promoted conditional on being a candidate, which can be interpreted as a lack of evidence for the under-performance in contests found by Gneezy et al. (2003).

The paper also contributes to the literature using data on academics to try to understand gender gaps in labour markets, which dates 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; see Johnson and Stafford (1974) and Farber (1977). More recent work, such as Ginther and Hayes (1999) and Ginther and Kahn (2004), 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. 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). 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, and their findings suggest that these are partly due to differences in the outside offers received by men and by women, 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. Our paper focuses on a different promotion system, and thus we can 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 organised as follows. Section 2 starts by examining the possible reasons why women are less likely to be promoted, we then describe the French academic system and explain how the different hypotheses can be tested. 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 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. These differences may in turn have three possible causes: discrimination, differences in costs of or rewards from promotion across genders, and different attitudes in and towards the promotion process itself. Let us consider them in turns.

The first possible explanation is simply that those making the promotion decisions discriminate against women. This discrimination may be due to tastes or statistical, and implies that a promotion committee would, for given research output and performance during the selection process, prefer to promote a man. Female candidates may also be less likely to succeed if they tend to under-perform in contests, as shown by Gneezy et al. (2003). As we will see below, the competitive nature of the promotion process for (most) academic promotions in France implies that if women ‘hold back’ during competitions they may perform worse than men and hence will be promoted less often. Both discrimination and underperformance during the promotion process will result in female candidates having a lower likelihood of promotion than male candidates.

Alternatively, women may differ in the objective costs of and rewards from promotion and this can make them less willing to apply, which in turn result in different observed promotion rates. There can be various causes of differences in the propensity to seek promotion. First, there may be costs associated with promotion, and if these are higher for women than for

men, they could result in the former applying for a promotion less often than men. In France, promotions tend to require moving to a different department and thus tend to imply moving to a different city (see below for details on the French academic system). If women's bargaining power in the household is lower, it may be harder for female than for male academics to impose the cost of moving on their families. As a result, the cost of a promotion is probably higher for women and hence they will be less likely to seek promotion.

Another reason why women may be less likely to apply for promotion is if the payoff is lower than for men. One possible cause of differences in payoff is that women may have different preferences over department prestige and income. Suppose that women have a lower marginal utility of income because, often, they are the second earner in the household. They may then be willing to be in a more prestigious department even if it implies lower wages. Since in the US and UK more prestigious departments also have tougher promotion thresholds (see McDowell et al. (2001)), this would result in a lower promotion probability for women.

Lastly, women may choose not to apply for promotion even if we observe no objective differences in the cost or payoff of promotion. This type of 'self-selection' out of the promotion race can be due to two reasons. On the one hand, female academics may expect to be discriminated against during the promotion process. Such a belief will reduce the expected payoff and hence the propensity to seek promotion. Note that this may happen even in the absence of actual discrimination. On the other hand, women may choose not to apply for promotion simply because they are less inclined to compete in tournaments than men, as indicated by experimental evidence; see Niederle and Vesterlund (2007), Niederle and Vesterlund (2011) and Gupta et al. (2013) as well as the discussion in Bertrand (2011). While certain aspects of an academic career, such as obtaining a PhD, depend only on individual performance, promotion typically entails a competition between various candidates in which only one gets the price. If women are less inclined to enter competitions, then they will apply less often for promotion even if costs and benefits are the same.

Existing work has had difficulty in testing these different hypotheses. Part of the problem is the fact that, in most countries, there exists no record of who applied for promotion. It is hence impossible to know whether the higher probability of men to hold a senior position is due to a lower conditional probability of success of women or to their lower propensity to apply. Moreover, even if data on applications were available, the tradeoff that exists in the Anglo-Saxon academic system between prestige of the department and promotion thresholds makes it difficult to assess the possible causes of a lower application propensity. In contrast, the French system presents a number of features that will allow us to evaluate the various options. We next turn to a description of these.

2.2 The French academic system

The French academic system has a number of features that we intend to exploit to test the above hypotheses. 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 pure

research positions.⁴ Researchers in this category are hired by the Cnrs, who pay their salaries, but are attached to a university and are hence located in its economics department just like the university professors are. 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* at the Cnrs. Both rank A and rank B positions are tenured, which implies that an individual who does not get a promotion can spend her/his entire career in a ‘junior’, i.e. rank B, position. Because of this possibility, we will term the two types of positions ‘rank A’ and ‘rank B’, rather than junior and senior. 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 identical for university professors and researchers and set by the Ministry of Higher Education and Research.⁵

Promotions take place through a national contest, a *concours*, and are thus not decided by the department in which the individual holds her/his current position.⁶ Participation in this contest is public information, and at the end of it a list with the ranking of those that have undertaken it is published. 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. A second feature of the French system is that there is no relationship between department prestige and salaries. Academic salaries are determined according to a national scale based on rank, and rank is decided at the national level.⁷ There is thus no reason to prefer being employed in a less prestigious department since salaries are the same across universities and the threshold for promotion is set nationally.

Members of the national committees that decide on promotion are academic economists, drawn from various universities in the country and areas of expertise. Committees tend to be large and represent a wide spectrum of universities, not necessarily the most prestigious. Because members of these committees 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

⁴Cnrs stands for *Centre National de la Recherche Scientifique*. There are other instances, such as the *Institut National de la Recherche Agronomique*, which accounts for about 5% of the academic economists in France. Unfortunately we have not been able to obtain the list of candidates for promotion for those and hence will not include them in our analysis.

⁵Within each broad rank, A and B, there are subranks that affect salary. Subrank promotion is also decided by a national committee, although the cost here is minimal, with application for a promotion requiring filling a form and submitting a vitae. Promotion to a higher subrank does not involve change of department, the number of promotions is not fixed ex ante, and the list of candidates is not public knowledge.

⁶Some 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.

⁷Some 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.

professors the contest, termed *concours d'agregation*, is biannual and entails four stages over, approximately, a 6-month period. It includes 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 candidates field of expertise. In contrast the effort involved in the Cnrs promotion *concours* is minimal. The candidate simply declares him/herself a candidate for a position as *directeur de recherche* (DR2) and submits a vitae and research proposal to the committee. This *concours* takes place annually.

There is a second difference in terms of the costs involved in the two systems. For university professors, departments publish before the start of the contest the positions that they have available. At the end of the *concours*, the candidates choose, sequentially and starting with the highest ranked, which university to join. When promoted, the individual is usually not able to stay in the university where s/he held a rank B position and has to move to a different department.⁸ After three years s/he is allowed to move to another university, if the latter wishes to recruit her/him, including the university where s/he held the rank B position. 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 both because of the geographical mobility involved in being promoted to full professor, but also because the individual faces considerable uncertainty about where s/he will be eventually recruited if s/he passes the contest since where s/he ends up depends on her/his ranking and hence on the entire pool of candidates. 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.

There are hence two differences between the *agregation* and the Cnrs *concours*. First, the former involves substantial costs that do not exist for the later, notably the fact that successful candidates need to change department. The second difference is that performance during the contest only matters in the *agregation*, as in the Cnrs the committee bases the decision on the vitae without meeting the candidates.

2.3 Testing the different hypotheses

We can use the features of the French academic system to confront the various hypotheses about lower promotion rates for women. In order to do so, we will examine both the determinants of the likelihood to apply for promotion and of the probability of obtaining a promotion conditional on being a candidate. We start by considering the gender gap in the probability of success conditional on being a candidate for promotion, and examine two possible explanations for a difference across the sexes:

⁸In our sample, 80% of those promoted through the *agregation* have a new affiliation after their promotion.

Under-performance in contests: If women tend to perform less well than men in contests, we expect to find a negative impact of being female on the conditional probability of promotion in the *agregation*. The absence of a gender gap in this probability can then be interpreted as no difference in the performance of women and men under competitive conditions.

Discrimination: Identifying discrimination is difficult. On the one hand, a negative impact of being female on the conditional probability of promotion can also be due to the presence of discrimination, and there is no feature of the data that would allow us to distinguish the effect of underperformance from that of discrimination. On the other, the absence of a gender effect could be the result of offsetting forces, with discrimination against women being compensated by better unobserved characteristics of women (not correlated to observed characteristics). As a result, it is not possible to rule out discrimination even in the absence of a significant gender effect.

Turning to the probability of being a candidate for promotion, we can test for two potential explanations for differences at this level by comparing the two tracks in the French academic system:

Higher cost of promotion: In the US system, it is not clear why women may have different costs of promotion from men. The nature of academia and the substantial flexibility associated with this type of jobs implies that jobs at the junior or senior level are equally likely to be compatible with family life. As a result, there is no reason why women would prefer not to be promoted. In contrast, in the French system, the mobility required if promoted (of professors but not researchers) implies that there is a very specific cost: the need to move to another university and hence, in most cases, to another city. It is possible that this cost is, on average, higher for women than for men, implying that the latter are less likely to apply for promotion. To test this hypothesis we can exploit the difference between the implications of promotion for standard professors and for researchers. If women were not trying to get promoted because of the cost of mobility, we would expect to find a more negative effect of being female on the probability of applying for promotion for professors than for researchers.

Differences in preferences about prestige and salary: We can also exploit the difference between promotion implications for professors and researchers to test the hypothesis of different preferences over salary and prestige across genders. Since for university professors being promoted to a senior position requires changing department, promotion brings a higher salary at the risk of having to move to a less good department. If women are less willing to trade-off prestige for income, then those who hold a junior position in a top department may not want to apply for promotion. In contrast, for junior researchers promotion does not require a change of department, and there is no reason for women not to seek promotion even if they are in a top department. If women have a stronger preference for department prestige and/or a lower one for income, those with a university position will be less willing to apply for promotion if they are in a top department than if they are not. Because promotion does not require changing department for researchers, we would expect to find no differences in the probability of applying for promotion for female researchers across departments.

Alternatively, the impact of gender on the probability of promotion may be common across the two tracks, implying that the differences in the propensity to seek promotion are due to either beliefs or attitudes towards contests. Testing for these effects is difficult. However, the

conditions under which the aggregation takes place can provide proxies for both a measure of expected discrimination and the competitive environment. We consider these in turns.

Expected discrimination: To examine whether women’s belief that they will be discriminated against affects their propensity to seek promotion, we proxy expected discrimination by the fraction of rank A professors that are women in a particular year, as a greater fraction could be interpreted by candidates as weaker discrimination. A positive effect of this variable interacted with being female would indicate that women’s decision to seek promotion responds to the observed promotion rate of women in the past.

Unwillingness to participate in contests: Women may also have different attitudes from men concerning contests, and, in particular, be unwilling to take part in them. If this is the case, they are likely not to respond to changes in the payoff from being a candidate. We hence consider the effect of changes in the likelihood of promotion, which we measure by the number of available positions of rank A in a particular year. A greater number of positions should increase a potential candidate’s expected payoff and hence make him/her more likely to enter the competition. A positive coefficient on this variable hence indicates that candidates respond to this ‘objective’ measure of the likelihood of promotion, while no effect implies that individuals’ decision to be a candidate is determined by considerations other than the probability of promotion, which in the case of women could be an unwillingness to enter competitions.⁹

2.4 Empirical specification

In line with the existing literature, we suppose that the probability of individual i being rank R at year t is given by

$$\Pr(R, i, t) = \Phi(X'_{it}\beta_j) \quad (1)$$

where the two states R are being rank A or not, the subscript $j = u, r$ denotes whether we are looking at data for university professors or researchers, and $\Phi(\cdot)$ denotes the logistic density function. The term $X'_{it}\beta$ in equation (1) is

$$\begin{aligned} X'_{it}\beta_j = & \beta_{j0} + \beta_{j0}^f \delta_i + \beta_{j1} Age_{it} + \beta_{j2} Age_{it}^2 + \beta_{j3} Pub_{it} \\ & + \beta_{j4} Pub_{it} \times Quantity_{it} + \beta_{j5} Pub_{it} \times Quality_{it} \end{aligned} \quad (2)$$

implying that the probability of promotion is a function of age Age_{it} and its square, whether or not the individual has published in Econlit-classed journals (i.e. whether s/he 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 measured in logs (see below for the exact measurement). The variable δ_i is a dummy taking the value 1 for females, so that β_0^f measures the differences in promotion probability for men and women with the same characteristics. We would have liked to control for characteristics of the individual’s family life,

⁹A question that we have not addressed is 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*. 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 pressure on candidates. It is hence difficult to argue that one or the other contest is more competitive.

such as whether s/he is married and the number and ages of children, but such data were not available.¹⁰

The probability of holding a rank A position is obviously a combination of the probability of applying for promotion and that of obtaining a promotion conditional on having applied. We hence model separately the two process. Consider first the probability that individual i applies for promotion (i.e. takes the *concoirs*) at time t , which we assume to be

$$\Pr(C, i, t) = \Phi(X'_{it}\gamma_j) \quad (3)$$

where the two states C are being a candidate for promotion or not and

$$\begin{aligned} X'_{it}\gamma_j &= \gamma_{j0} + \gamma_{j0}^f \delta_i + \gamma_{j1} Age_{it} + \gamma_{j2} Age_{it}^2 + \gamma_{j3} Pub_{it} \\ &+ \gamma_{j4} Pub_{it} \times Quantity_{it} + \gamma_{j5} Pub_{it} \times Quality_{it} \\ &+ \gamma_{j6} Dept_{it} \end{aligned} \quad (4)$$

The dummy $Dept_{it}$ has a value of 1 if the candidate is in a top department (see below for the definition) and will capture whether individuals in top departments are less likely to apply for promotion because of the mobility costs involved.

Lastly, the probability of success in a *concoirs* conditional on being a candidate is given by

$$\Pr(S, i, t | C_{it} = 1) = \Phi(X'_{it}\alpha_j) \quad (5)$$

where the two states S are succeeding or failing in the *concoirs* (i.e. being promoted or not), C_{it} is a dummy that takes the value one if the individual applied for promotion, and

$$\begin{aligned} X'_{it}\alpha_j &= \alpha_{j0} + \alpha_{j0}^f \delta_i + \alpha_{j1} Age_{it} + \alpha_{j2} Age_{it}^2 + \alpha_{j3} Pub_{it} \\ &+ \alpha_{j4} Pub_{it} \times Quantity_{it} + \alpha_{j5} Pub_{it} \times Quality_{it} \\ &+ \alpha_{j6} Dept_{it} \end{aligned} \quad (6)$$

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.

Equations (1) and (2) capture the mechanism usually examined in the literature on academic promotions. Assuming that the estimated coefficient $\hat{\beta}_{j0}^f$ is different from zero, we can then estimate the two other models to see to what extent the difference in the probability of being rank A is due to women being less likely to enter the *concoirs* or to them having a lower success rate in the *concoirs* than men. That is, we will compare estimates of γ_{j0}^f and α_{j0}^f .

¹⁰In the light of existing evidence, this is likely not to be a major omission; for example, 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.

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, publication stock (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. Those (few) individuals for whom some of the individual characteristics (age or position for instance) are missing are also excluded.

In order to examine separately the determinants of the probability of applying for promotion and the likelihood to succeed in the *concours*, 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 (DR2). For professors we have the lists of candidates to promotion and actual promotions for the nine biannual *concours* taking place from 1992 to 2008, while for researchers we have the six annual *concours* between 2003 and 2008.

We define as potential candidates for promotion all academic economists of rank B the year prior to the *concours*. It is important to note that 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. This is probably due to the fact that few promotions are available each year (between 2 and 4) and they are perceived as being internal promotions for those already in the Cnrs. We hence consider as potential candidates for the *agregation* all academic economists in France of rank B, which yields 10,383 observations, while we have 681 observations for candidates susceptible to apply for Cnrs promotions.¹¹

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 Appendix A 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 population each year.

3.1 Measuring research output

Our key explanatory variable is an individual’s research output, and we measure the output of individual i at date t by her/his cumulative publication record between the first year for which we have an observation for that individual and date t .¹² We merge the data on publications with the list of French academics that includes individual characteristics (position, department, age, gender), merging by surname and initial, and then correct manually for those individuals having the same name and initial. Publication records are measured as weighted sums of publications.

¹¹The difference in the number of observations is mainly due to the fact that the Cnrs accounts for a small fraction of academics. We have, on average, 114 rank B researchers each year and 1,040 rank B university professors. Hence, there are each year 114 potential candidates for promotion at the Cnrs and 1,154 potential candidates to promotion through the *agregation*, which include both rank B professors and rank B researchers.

¹²As 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.

All publications come from the EconLit database, which includes more than 560,000 papers published in more than 1200 journals between 1969 and 2008. Three dimensions enter the weighted scheme of publications: the quality of journals, the number of authors and the relative number of pages.

We measure the quality of publications using the journals weighting scheme proposed by Combes and Linnemer (2010). Two different degrees of convexity in the distribution of journals' weights have been proposed by Combes and Linnemer, 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 also consider the number of authors of each publication, and assume constant returns to scale at the coauthor level, so that we divide each publication by the number of authors, a standard practice in the literature. Lastly, since many journals now publish both notes and longer papers, we weight by the number of pages to capture the idea that longer articles contain more ideas. We consider an article's length relative to the average length in that journal, so that the ratio of the number of pages of article a over the average number of pages of articles published in the journal is used as weight. Because editors' policies can vary over time, the mean number of pages is calculated for each year separately.

The output of a researcher i at date t is then a weighted sum of her/his articles a published between the first year in which s/he published an article, t_0 , and date t , implying that

$$y_{it} = \sum_{a \in [t_0, t]} \frac{W(a) p(a)}{n(a) \bar{p}} \quad (7)$$

where $p(a)$ is the number of pages of article a , \bar{p} is the annual average number of pages of articles in the journal, $n(a)$ the number of authors of the article, and $W(a)$ the weighting scheme for journals. Then, each individual receives three scores: a dummy equal to 1 if s/he has at least one publication in an EconLit-listed journal, the quantity of single-author-equivalent published articles, and the average quality of her/his articles, defined as y_{it} divided by quantity, i.e. $\sum_{a \in [t_0, t]} 1/n(a)$.

3.2 Descriptive statistics

Table 1: Percentage rank A by gender in 2008

rank A	Total	%	Women	%	Men	%
Total sample						
0	1425	67.7	512	83.0	913	61.4
1	680	32.3	105	17.0	575	38.6
University professors (91% of population in 2008)						
0	1321	69.0	477	83.5	844	62.8
1	593	31.0	94	16.5	499	37.2
Cnrs Researchers (9% of population in 2008)						
0	104	54.5	35	76.1	69	47.6
1	87	45.5	11	23.9	76	52.4

Table 1 gives the decomposition of our sample in terms of institutional affiliation and rank for the most recent year in our data, 2008. There were 2,105 (Cnrs and university) academic economists in France that year, and the vast majority of the population consists of university professors, with researchers accounting for only 9% of the total. Women account for 29% of observations, and they are over-represented amongst university professors and under-represented amongst researchers, where they account for only 24% 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 under a third of the population hold a rank A position, with the fraction being lower for university professors (31%) and higher for researchers (45.5%). Note that the data do not seem to indicate that women choose a career path that offers higher average promotion rates, which would counterbalance negative discrimination. Our data shows that feminisation is lower for researchers, which have a higher promotion rate, indicating that there is no selection of this type taking place. The gender promotion gap is large, 22 percentage points on average, and is smaller for university professors than for researchers (21 and 29 points, respectively).

Table 2: Descriptive statistics of potential candidates

	Univ', panel 1992-2008				Cnrs, panel 2003-2008			
	Min.	Max.	Mean	Std. err.	Min.	Max.	Mean	St. err.
Women								
Prob. Candidate	0	1	0.062	0.241	0	1	0.103	0.305
Prob. Promotion	0	1	0.020	0.140	0	1	0.018	0.133
Age	27	68	41.6	8.9	29	64	42.8	9.1
Publisher	0	1	0.55	0.50	0	1	0.83	0.37
Quantity	0	15.0	1.02	1.55	0	15.0	3.14	3.22
Quality	0	62.5	0.42	2.83	0	16.6	1.01	2.31
Prob. Int. Dept.	0	1	0.32	0.47	0	1	0.58	0.49
Men								
Prob. Candidate	0	1	0.084	0.277	0	1	0.186	0.389
Prob. Promotion	0	1	0.033	0.179	0	1	0.044	0.205
Age	25	70	45.2	9.4	28	65	43.5	10.0
Publisher	0	1	0.57	0.50	0	1	0.91	0.29
Quantity	0	21.2	1.33	2.05	0	16.6	3.70	3.03
Quality	0	116.8	0.46	3.18	0	86.5	2.42	6.27
Prob. Int. Dept.	0	1	0.28	0.45	0	1	0.58	0.49

Int. Dept. = international department; see Appendix A for details. Productivity measures (quantity and quality) are in levels. We take their logs in the regression analysis. There are 10,383 observations from 2,136 individuals and 681 observations from 158 individual in the university and Cnrs samples respectively. These correspond to 687 women (3,328 observations) and 1,449 men (7,055 observations) in the university sample and 54 women (223 observations) and 104 men (458 observations) in the Cnrs sample.

Table 2 reports some descriptive statistics for the sample of potential candidates. We have 10,383 observations, including a total of 2,136 academics. The panel is unbalanced as individuals enter the pool of potential candidates (usually when they get their PhD) and exit it either because they are promoted or because they leave academia or the country over our sample period, with the average number of observations per individual being 4.9 years. The probability

of being a candidate to the *aggregation* is 6% for women and 8% for men, while for researchers these figures are 10% and 19%, implying a much larger gap. The (unconditional) probabilities of being promoted on a given year are small, 3% to 4% for men, and 2% for women. The table indicates that some of these differences are likely to be explained by differences in observable characteristics. Women are on average 3 years younger than men in the aggregation sample (though not in the Cnrs). The main difference lies in publication records. In our larger sample, the probability of publishing in EconLit journals is 57% and 55% for men and women, respectively. This figure is 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. The quality of publications is somewhat higher for men (0.46 compared to 0.42), while the quantity of publications is 30% higher for men. For the sample of researchers, we find a higher probability of being a publisher but a greater gender gap (91% for men but only 83% for women), and a particularly large difference in the quality of publications.¹³

The definition and construction of our two categories of departments is detailed in Appendix A and Table 10 gives the list of ‘international departments’. The last lines in the two panels of Table 2 indicate that the one aspect in which women seem to fare better than men is affiliation: in the larger sample 32% of women are in international departments, while only 28% of men are. When we focus only on researchers we find no difference in affiliation across the genders.

4 Results

4.1 The promotion of academic economists

We start by examining the determinants of being rank A for the entire population, which are reported in Table 3. 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. This gives us a sample of 16,077 observations. We can thus estimate the probability of *holding* a rank A position. All specifications include year fixed effects. In column (1), only time fixed effects and gender are included in the logit model, with the gender dummy being equal to 1 for women. The marginal effect on gender is significant at the 1% level and large, implying an odds ratio of 0.31.¹⁴ Including age reduces the effect on gender from -0.233 to -0.165, indicating that a large fraction of the difference in promotion is indeed due to the fact that the sample of women is younger than that of men. Column (3) includes our three measures of research

¹³Obviously some of the differences in research output are due to the age structure of the two populations, but Bosquet and Combes (2013*b*) show that even when controlling for age women have worse publication records. Women are sometimes also found to be cited less than men. This is the case in political sciences, as shown by Maliniak, Powers and Walter (2013), while for publications in biomedical and exact sciences, Kelchtermans and Veugelers (2013) find that women have a lower probability of reaching top citations but once they do they are as likely as men to stay there. In contrast, the data we use exhibits no gender gap in citations once research output is controlled for; see Bosquet and Combes (2013*a*).

¹⁴This values is obtained from $\exp(-1.182)$, -1.182 being the coefficient associated to the -0.233 marginal effect.

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 3: Likelihood to hold a rank A position: panel 1992-2008, marginal effects

	(1) logit	(2) logit	(3) logit	(4) logit	(5) probit	(6) OLS
Women	-0.233 ^a (0.006)	-0.165 ^a (0.007)	-0.053 ^a (0.007)	-0.049 ^a (0.007)	-0.048 ^a (0.007)	-0.044 ^a (0.007)
Age		0.015 ^a (0.002)	0.010 ^a (0.001)	0.010 ^a (0.001)	0.010 ^a (0.001)	0.007 ^a (0.001)
Age ²		-0.000 (0.000)	0.000 ^a (0.000)	0.000 ^a (0.000)	0.000 ^a (0.000)	0.000 ^a (0.000)
Publisher (Pub)			0.325 ^a (0.012)	0.331 ^a (0.012)	0.329 ^a (0.012)	0.341 ^a (0.010)
Pub*Quantity			0.148 ^a (0.004)	0.151 ^a (0.004)	0.152 ^a (0.004)	0.178 ^a (0.004)
Pub*Quality			0.040 ^a (0.002)	0.041 ^a (0.002)	0.041 ^a (0.002)	0.038 ^a (0.002)
Cnrs				-0.089 ^a (0.014)	-0.081 ^a (0.014)	-0.093 ^a (0.017)
Women*Cnrs				-0.042 ^c (0.025)	-0.039 (0.025)	-0.038 (0.026)
Time FE	Yes	Yes	Yes	Yes	Yes	Yes
pseudo-R ²	0.040	0.114	0.367	0.369	0.368	
Observations	16077	16077	16077	16077	16077	16077
log-likelihood	-9963	-9195	-6573	-6548	-6559	-6833

Standard error between brackets. ^a, ^b, ^c Significant at the 1%, 5% and 10% level, respectively.

Once we control for research output, the effect of gender falls to -0.053, about a quarter of the initial one, indicating that the lower promotion rate for women is to a large extent due to them having published less. The effect is nevertheless still strong: being a woman reduces the probability of promotion almost as much as having one rather than two single-authored publications. The effects we obtain are comparable 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 researcher and its interaction with Women, and indicates that the negative effect of being a woman on seniority is stronger for the Cnrs, although this term is significant only at the 10% level. Columns (5) and (6) report the estimations of a probit and an OLS model, respectively, and indicate that the result that women are less likely to be promoted are robust to the use of those specifications. There is however no robust evidence of a difference between the two tracks, with the coefficient on Women*Cnrs being insignificant.

The result that women have lower promotion rates than men contrasts with evidence on affiliation. Table 11 in Appendix B examines the likelihood for an individual to be in an international department, using the same explanatory variables as above. Women are more likely to be in a prestigious department than men, and the effect of being a woman on the probability of being in an international department is large.¹⁵ These results are surprising. As

¹⁵The positive effect of being female on the likelihood of being in an international department may have three

argued by Lazear and Rosen (1990), it is hard to understand why women would be discriminated against in promotions but not in other labour market experiences. In fact, in many instances the evidence tends to indicate that women have a lower probability of both being promoted and being hired, as for example in the case of top US orchestras; see Goldin and Rouse (2000). This does not seem to be the case for academic economists, and raises the question of whether the low promotion rate of female academics has a cause other than discrimination.

4.2 Decomposing outcomes

4.2.1 Gender differences in success in the *concours*

The limitation of our analysis so far is that it uses only outcomes and hence 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 examine the two separate steps. We start by considering what determines success in the contest, conditional on being a candidate, and then move to the determinants of the decision to enter the contest.

The determinants of success in the promotion contests are assumed to be gender and the three variables measuring publications. We also include a dummy that captures whether the individual was, at the time of application, in an international department. This variable could measure unobserved ability –e.g. being a good teacher– or the positive effects that being in a more stimulating academic environment could have on the preparation of the *concours*. 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. Lastly, in the regression for the *agregation* we include a dummy for whether the candidate holds a position other than rank B university professor in economics (i.e. the candidate is in the Cnrs, in another field such as mathematics, comes from abroad, etc.).

The results, reported in the first four columns of Table 4, indicate that research output is the key determinant of the probability of success both in the *agregation* and the 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 former are more homogeneous in terms of research output. Being in an international department has a positive impact on the probability of passing the *agregation*, in line with our arguments before, while it has no impact on the probability of being promoted in the Cnrs contest.

Turning to our coefficient of interest, the impact of being female on the probability of success conditional on being a candidate for promotion, we find no significant difference across the genders. In fact, the probability of success in the *agregation* is 24.8% for men and 24.1% for women, and hence not significantly different. The coefficient is insignificant both for the raw probabilities and also when we control for individual characteristics, for professors and

causes. The first is positive discrimination when departments consciously try to increase their female faculty. A second is unobservable abilities, such as organisational skills or teaching ability, which could be, on average, higher amongst female than male academics, specially since fewer women than men that enter academia (see Petrongolo and Olivetti (2008) on selection and gender). Lastly, a strong female presence could be the result of joint offers made to couples by top departments.

Table 4: Likelihood to be promoted conditional on applying, marginal effects

	Agreg' 1992-2008		Cnrs 2003-2008		Agreg': selected candidates 1992-2008			
	All candidates		All candidates		"Admissibles"		Close to threshold	
	(1) logit	(2) logit	(3) logit	(4) logit	(5) logit	(6) logit	(7) logit	(8) logit
Women	-0.007 (0.032)	-0.018 (0.028)	-0.051 (0.078)	0.044 (0.110)	0.044 (0.056)	0.040 (0.057)	0.097 (0.082)	0.103 (0.086)
Age		-0.018 (0.011)		-0.071 ^b (0.031)		-0.033 (0.026)		-0.042 (0.039)
Age ²		-0.000 (0.000)		0.002 ^b (0.001)		0.001 (0.001)		0.002 (0.002)
Publisher (Pub)		0.155 ^a (0.047)		0.077 (0.209)		0.112 (0.086)		0.142 (0.130)
Pub*Quantity		0.091 ^a (0.019)		-0.009 (0.065)		0.032 (0.039)		-0.024 (0.060)
Pub*Quality		0.047 ^a (0.006)		0.046 ^b (0.022)		0.031 ^b (0.014)		0.024 (0.022)
Int. Dept.		0.090 ^a (0.030)		0.019 (0.088)		0.130 ^a (0.046)		0.117 (0.080)
Other		-0.030 (0.032)				0.008 (0.069)		0.008 (0.108)
Time FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
pseudo-R ²	0.015	0.183	0.020	0.184	0.015	0.056	0.005	0.025
Observations	970	970	115	115	336	336	198	198
log-likelihood	-531	-441	-52	-43	-201	-192	-137	-134

Standard error between brackets. ^a, ^b, ^c Significant at the 1%, 5% and 10% level, respectively. Int. Dept. = international department; see Appendix A for details. Other: non-French Assistant Professors positions in economics, as well as Cnrs and Inra researchers, and assistant professors from other disciplines.

researchers. Gender simply does not matter.

To interpret this result consider first the Cnrs *concours*. Recall that the *concours* consists of submitting a vitae and research proposal. The candidate does not ‘perform’ in front of the committee taking the promotion decision, hence the only reason why we would expect to find a negative coefficient on women would be discrimination. We can then 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.

In the case of the *agregation*, interpreting the coefficient is more complicated as it captures both potential biases in the committee and the actual performance of the candidate during the oral exams that we cannot observe. A negative and significant coefficient could then be due to either discrimination or under-performance by women, as in the experiments performed by Gneezy et al. (2003). Our results find no evidence of either discrimination against or under-performance by women. It is also possible that discrimination is offset by over-performance or that, more generally, selection is an issue. Since fewer women are candidates than men, as we will see below, 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. 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 this, 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.

The first test focuses on admissible candidates only. Since at this stage a substantial number of candidates have been eliminated, if selection implies that the unobserved quality of male candidates is lower than that of female candidates, focusing on the ‘best’ should eliminate the (predominantly male) bottom tail of the distribution. Columns (5) and (6) of Table 4 present those regressions and indicate that there is no difference between men and women in either the raw rate of passing the last step of the context or the one obtained after controlling for output. Not surprisingly, only the quality of publications matters at this stage. 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 not get promoted and the n lowest-ranked candidates that did get promoted and rerun the same regression equations. 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 last two columns of Table 4 show that, if anything, women are *positively* discriminated although the coefficient is not significant. In fact, none of the variables we consider at this stage are significant.

4.2.2 Likelihood to enter the *concours*

We turn now to the determinants of the likelihood to enter the *concours*. The results are reported in Table 5. The unconditional probability of applying is lower for women than for men both for university professors and researchers (column 1), and the negative impact of being female is even larger once we include individual characteristics, the marginal effect being 0.025 for professors and 0.092 for researchers. As expected, research output has a strong effect on the likelihood of being a candidate in either of the *concours*, with quality having a stronger effect for researchers. Being in an international department is significant only in the case of professors, where it tends to increase the probability of being a candidate. This may be due to the fact that, as we saw earlier, being in a top department has a positive effect on the probability of success; potential candidates may anticipate this and hence be more likely to apply if they are in those departments. An additional effect may come from peer pressure to pass the *concours*. Removing this variable from the analysis has no impact on the other estimated coefficients.

The results indicate that women have a lower likelihood to enter the contest for promotion. In terms of magnitude, being a woman is equivalent to decreasing the number of single-authored publications by around 0.6 for the university professors and 2 for the Cnrs researchers,¹⁶ or

¹⁶The regression coefficients on gender and quantity in the regressions for the likelihood to apply for the *agregation* and to apply for a Cnrs promotion are, respectively, -0.455 , 0.740 , -1.007 and 0.501 .

Table 5: Likelihood to apply for a promotion: marginal effects

	Agreg', panel 1992-2008				Cnrs, panel 2003-2008			
	(1) logit	(2) logit	(3) probit	(4) OLS	(5) logit	(6) logit	(7) probit	(8) OLS
Women	-0.021 ^a (0.005)	-0.025 ^a (0.004)	-0.027 ^a (0.004)	-0.033 ^a (0.005)	-0.085 ^a (0.022)	-0.092 ^a (0.021)	-0.087 ^a (0.022)	-0.073 ^a (0.028)
Age		0.001 (0.002)	-0.001 (0.002)	-0.016 ^a (0.001)		0.100 ^a (0.011)	0.097 ^a (0.010)	0.064 ^a (0.007)
Age ²		-0.000 ^a (0.000)	-0.000 ^a (0.000)	0.000 ^a (0.000)		-0.002 ^a (0.000)	-0.002 ^a (0.000)	-0.001 ^a (0.000)
Publisher (Pub)		0.094 ^a (0.011)	0.094 ^a (0.011)	0.119 ^a (0.009)		0.033 (0.054)	0.029 (0.051)	0.075 (0.047)
Pub*Quantity		0.043 ^a (0.004)	0.043 ^a (0.004)	0.050 ^a (0.004)		0.050 ^b (0.020)	0.042 ^b (0.018)	0.044 ^b (0.018)
Pub*Quality		0.010 ^a (0.001)	0.010 ^a (0.002)	0.016 ^a (0.002)		0.040 ^a (0.007)	0.040 ^a (0.007)	0.036 ^a (0.007)
Int. Dept.		0.013 ^b (0.006)	0.012 ^b (0.006)	0.014 ^b (0.006)		-0.011 (0.028)	-0.006 (0.028)	-0.008 (0.029)
Cnrs		-0.087 ^a (0.002)	-0.086 ^a (0.002)	-0.130 ^a (0.009)				
Time FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
pseudo-R ²	0.011	0.262	0.258		0.030	0.277	0.274	
Observations	10383	10383	10383	10383	681	681	681	681
log-likelihood	-2779	-2075	-2087	-297	-289	-215	-216	-213

Standard error between brackets. ^a, ^b, ^c Significant at the 1%, 5% and 10% level, respectively. Int. Dept. = international department; see Appendix A for details.

to decreasing the quality of publications by 87% and 81% of its standard deviation, for the university professors and the Cnrs researchers, respectively.¹⁷

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 were holding back potential female candidates. The first column of Table 6 hence runs a difference-in-differences regression where we have pooled together the data for the two *concours*. We interact all variables with a dummy taking the value 1 for the Cnrs *concours* to allow for different impacts across the two tracks, our coefficient of interest being women interacted with this dummy.¹⁸ The coefficient is negative and significant at the 10% level, implying that the effect of gender on the likelihood to enter the promotion contest is twice as high in the research track. This result indicates that the time-cost of preparing the *agregation* and the implications in terms of mobility are not the main reason why women are less likely than men to apply for promotion.

As we discussed earlier, if women have a stronger relative preference for department prestige, they may choose not to apply for promotion in order to stay in an international department.

¹⁷The coefficients on gender and quality in the *agregation* and the Cnrs promotion regressions are, respectively, -0.455 , 0.170 , -1.007 and 0.406 , and 3.07 is the standard deviation of the average quality of publications in our sample.

¹⁸To save space, other interacted terms are not reported but are available upon request.

Since promotion implies mobility for university professors, if females cared more about department quality they would be less willing to move –and hence to pass the *agregation*– whenever they are in an international than in a national department. In contrast, for researchers promotion does not require changing department and we expect to find no effect of department of origin on the likelihood of seeking a promotion.

We hence run again the regressions for the likelihood to apply for promotion and include an interaction between being a woman and being in a top department. If the tradeoff between rank (i.e. income) and department quality differed across the genders, this would be captured by a negative coefficient on being in a top department for female professors but not for female researchers.

Table 6: Differences across tracks and departments: likelihood to apply for a promotion, marginal effects

	All (1) logit	Agreg ⁷ , panel 1992-2008 (2) logit	Cnrs, panel 2003-2008 (3) logit
Women	-0.024 ^a (0.005)	-0.035 ^a (0.005)	-0.085 ^a (0.031)
Age	0.004 ^c (0.002)	0.001 (0.002)	0.100 ^a (0.011)
Age ²	-0.000 ^a (0.000)	-0.000 ^a (0.000)	-0.002 ^a (0.000)
Publisher (Pub)	0.102 ^a (0.012)	0.094 ^a (0.011)	0.031 (0.054)
Pub*Quantity	0.046 ^a (0.004)	0.043 ^a (0.004)	0.050 ^b (0.020)
Pub*Quality	0.011 ^a (0.001)	0.010 ^a (0.001)	0.040 ^a (0.007)
Int. Dept.	0.015 ^b (0.006)	0.004 (0.006)	-0.006 (0.032)
DR2 Cnrs*Women	-0.025 ^c (0.013)		
Women*Int. Dept.		0.031 ^b (0.014)	-0.016 (0.055)
Cnrs		-0.087 ^a (0.002)	
Interacted terms	Yes	No	No
Time FE	Yes	Yes	Yes
pseudo-R ²	0.260	0.263	0.277
Observations	11064	10383	681
log-likelihood	-2318	-2072	-215

Standard error between brackets. ^a, ^b, ^c Significant at the 1%, 5% and 10% level, respectively. Int. Dept. = international department; see Appendix A for details. Interacted terms: all variables interacted with applying to the DR2 Cnrs contest.

The marginal effects obtained from these regressions are reported in the second and third columns of Table 6. Being in an international department does not affect the likelihood of female researchers applying for promotion, while it has a positive effect on whether females enter the *agregation*. For male university professors the department of origin seems not to matter. This indicates, first, that the effect on being in an international department we obtained earlier is driven by the effect on women in top departments. Second, note that when we consider together

the two effects, that of women and that of women interacted with international department they cancel out. This seems to indicate that, for university professors, women in top departments are as likely as men to enter the *concours*. In contrast, women in national departments and those with a research position have a lower probability of being candidates for promotion.

The evidence so far indicates that there is a substantial difference in the likelihood that men and women academics apply for promotion. This gap appears both for professors and researchers despite the differences in the costs and implications of promotion, indicating that neither the direct costs associated with the *concours* nor differences in the way in which men and women trade-off income and the quality of their department are the causes of the gender gap in seeking promotion.

In order to gauge the importance of women’s lower propensity to apply for promotion, we perform a Oaxaca decomposition. Given the difficulty of interpreting decomposition stemming from logit regressions, we perform the decomposition using the OLS regression in Table 5 column (4). The raw gender gap in the probability of being promoted to rank A in our sample amounts to 1.2 percentage points. Although this magnitude seems small, the probability of being promoted in a given year is only 3.3% for men, implying that being a woman reduces it by one third. As we have seen, this difference is driven by the decision to be a candidate, with an average male probability of being a candidate of 8.4% and that for women being 2.2 percentage points lower, i.e. gender reduces the likelihood to apply for promotion by one quarter.

Table 7: Oaxaca decomposition of the likelihood to apply: Univ. only, OLS

	% of $\Delta \Pr(C)$ explained by
<u>Characteristics</u>	-86.8
Age	-115.1
Total outcome	39.0
<i>Publisher</i>	13.1
<i>Quantity</i>	28.9
<i>Quality</i>	-3.0
International Department	-3.5
Cnrs	-7.2
<u>Coefficient</u>	186.8
<u>Total gap</u>	100.0

Table 7 presents the Oaxaca decomposition for the probability of being a candidate, $\Pr(C)$, based on column (4) of Table 5, where the only coefficient that differs across the genders is the constant. The gender gap in the probability of being a candidate is driven by the direct effect of gender, which is partly offset by the fact that females have better characteristics in terms of age and affiliation. The fact of being younger has the strongest impact, with a magnitude slightly higher than that of the overall gap observed. Being in a international department also favors women, though to a much smaller extent. In contrast, their lower research output has a negative impact on their probability of being a candidate which amounts to 39% of the overall effect.¹⁹

¹⁹Interestingly, when we look at the three components of research output we see that women are less likely to

The effect of research output is however dwarfed by the direct impact of being female, which is almost five times larger. Without the advantage in characteristics, the gender gap would be almost twice as large as the one we observe, as indicated by the fact that the coefficient on women explains 187% of the overall gap. That is, the direct effect of being female implies that the probability of applying for promotion for women is about half that of men, although this effect is partly offset by women being younger and in better departments.

4.3 Networks and the competitive environment

4.3.1 Network effects

One possible reason for differences in the likelihood of seeking a promotion is that there is some individual variable that we have so far ignored and which has an impact on the actual (or perceived) probability of success. A candidate explanation are research networks or groups 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 Zinovyevay and Bagues (2012) 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 a candidate who has a member of her/his network in a promotion committee has a higher likelihood of success, but also that with a larger network the candidate's 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.

In order to test this hypothesis, we construct for each individual two measures of 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 the researcher has had over his/her publishing lifetime; the second is the fraction of network members that are men. 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). Moreover, being a candidate for promotion can increase networks, especially in a system where individuals are often candidates several years in a row, since other candidates that are met during sessions aimed at the preparation of the *agregation concours* may become coauthors. As a result, any resulting effect has to be interpreted with care, indicating correlation rather than as having causal implications.

Table 8 reports the regressions for the likelihood of entering the promotion contest for the two tracks, to which we have added network variables. Both variables have positive coefficients although only one of them is significant, the proportion of men in the case of university professors be publishers and, especially, have fewer publications, but they exhibit a higher quality of publications than men.

Table 8: Network effects: likelihood to apply for a promotion, marginal effects

	Agreg', panel 1992-2008			Cnrs, panel 2003-2008		
	(1) logit	(2) probit	(3) OLS	(4) logit	(5) probit	(6) OLS
Women	-0.024 ^a (0.004)	-0.026 ^a (0.004)	-0.032 ^a (0.005)	-0.091 ^a (0.021)	-0.088 ^a (0.022)	-0.066 ^b (0.028)
Age	0.000 (0.002)	-0.002 (0.002)	-0.017 ^a (0.001)	0.097 ^a (0.011)	0.095 ^a (0.010)	0.061 ^a (0.007)
Age ²	-0.000 ^a (0.000)	-0.000 ^a (0.000)	0.000 ^a (0.000)	-0.002 ^a (0.000)	-0.002 ^a (0.000)	-0.001 ^a (0.000)
Publisher (Pub)	0.086 ^a (0.012)	0.084 ^a (0.012)	0.104 ^a (0.010)	-0.048 (0.060)	-0.057 (0.057)	0.013 (0.052)
Pub*Quantity	0.042 ^a (0.004)	0.042 ^a (0.004)	0.051 ^a (0.004)	0.041 ^b (0.020)	0.034 ^c (0.019)	0.030 (0.019)
Pub*Quality	0.009 ^a (0.001)	0.009 ^a (0.002)	0.015 ^a (0.002)	0.032 ^a (0.007)	0.031 ^a (0.008)	0.030 ^a (0.008)
Int. Dept.	0.012 ^b (0.006)	0.012 ^b (0.006)	0.013 ^b (0.006)	-0.010 (0.027)	-0.007 (0.027)	-0.010 (0.029)
Cnrs	-0.087 ^a (0.002)	-0.086 ^a (0.002)	-0.131 ^a (0.009)			
Pub*Network size	0.004 (0.005)	0.006 (0.005)	0.001 (0.007)	0.039 ^b (0.018)	0.040 ^b (0.018)	0.045 ^b (0.022)
% Men in Network	0.014 ^b (0.006)	0.015 ^b (0.007)	0.024 ^a (0.008)	0.057 (0.044)	0.058 (0.043)	0.046 (0.041)
Time FE	Yes	Yes	Yes	Yes	Yes	Yes
pseudo-R ²	0.263	0.259		0.291	0.290	
Observations	10383	10383	10383	681	681	681
log-likelihood	-2072	-2083	-291	-211	-211	-209

Standard error between brackets. ^a, ^b, ^c Significant at the 1%, 5% and 10% level, respectively. Int. Dept. = international department; see Appendix A for details.

and the network's size for researchers. Despite the significance of these effects, the coefficients on gender barely change as compared to our earlier specifications. The only exception is the OLS regression for researchers where the coefficient increases from -0.073 (see Table 5) to -0.066. That is, gender gaps in networks do not seem to explain the lower propensity of women to seek promotion. The reason is probably that differences in network size across the genders are minor in France. In the larger sample, the average network size is 0.88 for men and 0.82 for women, while researchers have a large number of coauthors, averaging 2.9 for women and 3.2 for men. These differences across the genders are not statistically significant. There is a greater gap in the fraction of men in the network, which are 70% for females in the *agregation* sample and 80% for males, and although it is significant the effect of this variable is moderate and hence its inclusion barely affects the coefficient on gender.

4.3.2 The competition's environment

Our final test consists of examining whether the environment under which the competition takes place affects men and women differently. We consider two variables. The first is the proportion of women amongst rank A professors in the year of the competition, which can act as a proxy for 'perceived discrimination'. 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, which is known before the decision to be a candidate is taken, and ranges between 15 and 33, fluctuating substantially from year to year and showing no clear trend.

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 the Cnrs we have a shorter time span, and hence little change in the fraction of women DRs (which goes from 10.6 to 10.9), while there is too little variability in the number of positions, as either 2, 3 or 4 are available every year. 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 on men. Alternatively, we substitute the fixed effects by a time trend and estimate the impact of these variables for both men and women.

Table 9: The competition's environment: likelihood to apply for a promotion, *agregation*, panel 1992-2008, marginal effects

	(1) logit	(2) logit	(3) logit	(4) logit
Women	-0.093 ^a (0.030)	-0.082 ^a (0.031)	-0.090 ^a (0.031)	-0.079 ^b (0.032)
Age	0.000 (0.002)	0.000 (0.002)	0.000 (0.002)	0.000 (0.002)
Age ²	-0.000 ^a (0.000)	-0.000 ^a (0.000)	-0.000 ^a (0.000)	-0.000 ^a (0.000)
Publisher (Pub)	0.093 ^a (0.011)	0.093 ^a (0.011)	0.093 ^a (0.011)	0.093 ^a (0.011)
Pub*Quantity	0.043 ^a (0.004)	0.043 ^a (0.004)	0.043 ^a (0.004)	0.043 ^a (0.004)
Pub*Quality	0.010 ^a (0.001)	0.010 ^a (0.001)	0.010 ^a (0.001)	0.010 ^a (0.001)
Int. Dept.	0.013 ^b (0.006)	0.013 ^b (0.006)	0.006 (0.006)	0.006 (0.006)
Cnrs	-0.087 ^a (0.002)	-0.087 ^a (0.002)	-0.087 ^a (0.002)	-0.087 ^a (0.002)
Time trend		-0.012 ^a (0.004)		-0.012 ^a (0.004)
% Women rank A		0.016 (0.011)		0.016 (0.011)
Women*% Women rank A	0.007 ^c (0.004)	0.006 ^c (0.003)	0.006 ^c (0.004)	0.006 ^c (0.003)
Men*Avail.Pos.		0.002 ^a (0.001)		0.002 ^a (0.001)
Women*Avail.Pos.	-0.000 (0.001)	0.002 ^c (0.001)	-0.001 (0.001)	0.001 (0.001)
Women*Int.Dept*Avail.Pos.			0.001 ^b (0.000)	0.001 ^b (0.000)
Time FE	Yes	No	Yes	No
pseudo-R ²	0.263	0.262	0.264	0.263
Observations	10383	10383	10383	10383
log-likelihood	-2072	-2075	-2070	-2073

Standard error between brackets. ^a, ^b, ^c Significant at the 1%, 5% and 10% level, respectively. Int. Dept. = international department; see Appendix A for details.

Table 9 presents the estimated regressions. Column (1) indicates that the fraction of women amongst rank A professors has a greater effect on women than on men, and our specification with a time trend implies that this variable has no significant effect on men. This is consistent with the idea that perceived discrimination is a potential cause for females' lower propensity to apply for promotion. Nevertheless, the coefficient on being a woman remains negative and highly significant.

The number of available positions has no differential effect across the genders, as seen in column (1), and the coefficient on being female remains negative, large and highly significant. In column (2) we try to identify the effect for both men and women. For ease of interpretation, we interact this variable with a dummy that takes the value 1 for males (Men*Avail.Pos) as well as with the female dummy (Women*Avail.Pos), implying that the coefficient on this last variable captures the overall impact of the number of positions on women's decisions. We can see that a greater number of available positions increases the likelihood of being a candidate, as expected, capturing the idea that more positions imply a greater probability of being promoted and hence encourages individuals to apply. The coefficient is the same across the genders, although the standard error is much greater for females, with the coefficient being significant only at the 10% level. A possible interpretation is that women are, on average, influenced by the number of positions just as much as men, but that there is greater variability in their reactions.

To further explore this aspect we ask whether the large variability in this effect stems from the fact that not all groups of women behave in the same way. A possible division is between women who are in a national and those in an international department. Women in different types of departments may react differently towards competition because of unobserved characteristics; in particular, some of them may be more 'driven' which makes them both manage to get hired by a better department and be more responsive to the competition's environment. Columns (3) and (4) indicate that there is indeed a difference in the coefficients on the number of available positions for the two groups. Women in international departments respond to the number of positions available, with the coefficient being of the same magnitude as for men (since the overall coefficient for this group is the sum of those on Women*Avail.Pos. and $\text{Women*Int.Dept.*Avail.Pos.}$). In contrast, for women in national departments the decision to apply for promotion seems to be unaffected by how competitive the promotion contest is likely to be. A possible interpretation of this effect is that, contrary to their international department counterparts, women in national departments find the idea of taking part in the contest so unattractive that variations in the expected success rate do not affect their decision.

5 Conclusions

This paper has used data for promotions amongst academic economists in France to look at the attitudes of women during and towards contests. We have information both on who obtained and who was a candidate for promotion, which allows us to test different hypotheses about the reasons that cause women's lower promotion rates.

On the one hand, women may be less likely to be promoted conditional on having applied, and this could be due either to discrimination or under-performance during the promotion contest. On the other, female academics may have a lower propensity to apply for promotion

than males. This could in turn be due to the requirements of the contest being more costly for women, to their facing a different tradeoff between salaries and department prestige, or to an unwillingness to enter the contest. 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 test for these hypotheses.

We find that there is no difference between men and women in their probability of success, but that the later are less likely than the former to enter contests. Moreover, our results indicate that neither the differential cost of promotion nor the tradeoff between income and department quality are behind these differences.

There are two possible explanations for our findings. One is that although women are not discriminated against during the contest they believe they will be, and hence decide not to enter the competition for promotion. The data points at this as a partial explanation of gender differences, with women being more likely to apply the more females there are amongst top-rank academics. The alternative explanation is that women are less willing than men to enter contests, in line with the experimental evidence provided by see Gneezy et al. (2003), amongst others. When we try to assess the impact of the competition's environment, we find that, unlike for men and women in top departments, for women in other departments the likelihood to apply for promotion is unaffected by the number of available positions in a particular year.

Our results have two main implications. On the one hand, they provide a link between laboratory evidence and behaviour observed in actual markets, and indicate how patterns of behaviour, namely women's lower propensity to take part in contests, can be observed in both contexts and what their consequences are for observed labour market outcomes. On the other, 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 differences across genders in preferences for taking part in competitions, which can be the result of 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 incite more women to apply. Alternatively, creating a system in which the default is that an individual will be considered for promotion after x years and s/he has to opt out instead of opting in, could also be a way of overcoming the differences in confidence across the genders.

Appendix

A Definition of international departments

Academia in France is organised around ‘research centers’, with a university potentially having several research centers in economics. We have defined ‘departments’ either by an economics department when it is the single affiliation where economists are found in a given university (which corresponds to the majority of cases), or by the aggregation of all research centres where there are economists in the university. We performed robustness checks by looking at research centres rather than departments, and obtained consistent results (available upon request). However, our notion of slightly aggregated economics departments better matches the reality of French academic research and hence we preferred to focus our discussion of the results on this concept.

We define ‘international departments’ as departments with the highest research output, measured by both the total stock and average stock of publications per member of the department in EconLit journals. Also, we wish to take into account both “normal” quality and top quality research, two measures that we denote CLm and CLh, respectively, following Bosquet and Combes (2013*b*). In order to do so, we apply the following procedure at each point in time. We first calculate the total research output of a department (in terms of the CLm and CLh measures) as the sum of research outputs of its members. We also calculate the two scores per member of the department. Hence, we obtain 4 scores per department: CLm total, CLh total, average CLm and average CLh. Departments are then ranked with respect to these 4 scores and their final score is the average of these rankings.

The cutoff point for top departments is arbitrary. We choose our cutoff at each point in time so as to have 35% of the total population in the top department category. Since departments may change their status over time, we considered as international departments for the entire sample period those departments which are more often than not in the top department category.

Table 10: List of international departments

Cepremap	1991	2004
Crest	1993	2008
Ecole Polytechnique	1991	2008
Ehess-Ens (Delta)	1991	2004
Enpc (Ceras, Latts)	1991	2004
Hec	1995	2008
University Aix Marseille 2-3	1991	2008
University Cergy	2005	2008
University Cergy-Paris 10	1991	2004
University Clermont 1	1991	2008
University Nancy 2-Strasbourg	2006	2008
University Paris 1	1991	2004
Paris School of Economics	2005	2008
University Paris 10	2005	2008
University Paris 9	1991	2008
Toulouse School of Economics	1991	2008

The last two columns give the first and last data at which the departments are observed in the sample, respectively. Note that some departments have been merged, as for instance the Cepremap, the Ehess-Ens and the University Paris 1 to form the Paris School of Economics from 2005.

B Likelihood to be in an international department and to hold a rank A position by category

Table 11: Likelihood to be in an international department and to hold a rank A position by category, 1992-2008, marginal effects

	International department					Rank A	
	All (1) logit	All (2) logit	All (3) logit	Univ. (4) logit	Cnrs (5) logit	Univ. (6) logit	Cnrs (7) logit
Women	0.020 ^b (0.008)	0.022 ^b (0.009)	0.068 ^a (0.009)	0.071 ^a (0.009)	0.079 ^b (0.034)	-0.064 ^a (0.007)	-0.120 ^a (0.023)
Age		-0.010 ^a (0.002)	-0.009 ^a (0.002)	-0.008 ^a (0.002)	-0.026 ^a (0.007)	0.011 ^a (0.002)	0.059 ^a (0.007)
Age ²		0.000 ^a (0.000)	0.000 ^a (0.000)	0.000 ^a (0.000)	0.001 ^a (0.000)	0.000 ^c (0.000)	-0.001 ^a (0.000)
Publisher (Pub)			0.192 ^a (0.012)	0.196 ^a (0.013)	0.041 (0.054)	0.324 ^a (0.013)	0.036 (0.039)
Pub*Quantity			0.051 ^a (0.005)	0.049 ^a (0.005)	0.083 ^a (0.017)	0.147 ^a (0.004)	0.145 ^a (0.012)
Pub*Quality			0.048 ^a (0.002)	0.047 ^a (0.002)	0.058 ^a (0.007)	0.041 ^a (0.002)	0.047 ^a (0.005)
Cnrs			0.185 ^a (0.024)				
Women*Cnrs			0.018 (0.031)				
Time FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
pseudo-R ²	0.012	0.014	0.082	0.056	0.118	0.347	0.487
Observations	16077	16077	16077	14848	1229	14848	1229
log-likelihood	-10022	-9999	-9315	-8578	-721	-6212	-432

Standard error between brackets. ^a, ^b, ^c Significant at the 1%, 5% and 10% level, respectively.

References

- Bertrand, M. (2011), New Perspectives on Gender, *in* O. Ashenfelter and D. Card, eds, ‘Handbook of Labor Economics’, Elsevier, pp. 1543–1590.
- 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. and Combes, P.-P. (2013a), ‘Are academics who publish more also more cited? Individual determinants of publication and citation records’, *Scientometrics* (DOI 10.1007/s11192-013-0996-6).
- Bosquet, C. and Combes, P.-P. (2013b), ‘Do Large Departments make Academics more Productive? Agglomeration and Peer Effects in Research’, *CEPR Discussion paper* (9401).
- 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*.
- 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.
- Ginther, D. K. and Hayes, K. J. (1999), ‘Gender Differences in Salary and Promotion for Faculty in the Humanities’, *American Economic Review Papers and Proceedings* **89**(2), 397–402.
- Ginther, D. K. and Kahn, S. (2004), ‘Women in Economics: Moving Up or Falling Off the Academic Career Ladder?’, *Journal of Economic Perspectives* **18**(3), 193–214.
- Gneezy, U., Niederle, M. and Rustichini, A. (2003), ‘Performance in Competitive Environments: Gender Differences’, *Quarterly Journal of Economics* **118**(3), 1049–1074.
- Goldin, C. and Rouse, C. (2000), ‘Orchestrating Impartiality: The Impact of “Blind” Auditions on Female Musicians’, *American Economic Review* **90**(4), 715–742.
- Gupta, N. D., Poulsen, A. and Villeval, M. C. (2013), ‘Gender Matching and Competitiveness: Experimental Evidence’, *Economic Inquiry* **51**(1), 816–835.

- Johnson, G. E. and Stafford, F. P. (1974), ‘Lifetime Earnings in a Professional Labor Market: Academic Economists’, *Journal of Political Economy* **82**(3), 549–69.
- Kahn, S. (1995), ‘Gender Differences in Academic Career Paths of Economists’, *Journal of Economic Perspectives* **9**(4), 193–204.
- Kelchtermans, S. and Veugelers, R. (2013), ‘Top Research Productivity and Its Persistence: Gender as a Double-Edged Sword’, *The Review of Economics and Statistics* **95**(1), 273–285.
- Lazear, E. P. and Rosen, S. (1990), ‘Male-Female Wage Differentials in Job Ladders’, *Journal of Labor Economics* **8**(1), S106–S123.
- Maliniak, D., Powers, R. and Walter, B. F. (2013), ‘The Gender Citation Gap in International Relations’, *International Organization*, (Available on CJO 2013 doi:10.1017/S0020818313000209).
- 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.
- Niederle, M. and Vesterlund, L. (2007), ‘Do Women Shy Away from Competition? Do Men Compete Too Much?’, *Quarterly Journal of Economics* **122**(3), 1067–1101.
- Niederle, M. and Vesterlund, L. (2011), ‘Gender and Competition’, *Annual Review of Economics* **3**(1), 601–630.
- 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.
- Sabatier, M. (2010), ‘Do Female Researchers Face a Glass Ceiling in France? A Hazard Model of Promotions’, *Applied Economics* **42**(16), 2053–2062.
- Winter-Ebmer, R. and Zweimuller, J. (1997), ‘Unequal Assignments and Unequal Promotion in Job Ladders’, *Journal of Labor Economics* **15**(1), 43–71.
- Zinovyevay, N. and Bagues, M. (2012), ‘The Role of Connections in Academic Promotions’. mimeo.