CENTRE DE RECHERCHE
SUR LES STRATÉGIES ÉCONOMIQUES

Gender Differences in Competitive Positions: Experimental Evidence on Job Promotion

Emmanuel Peterle and Holger A. Rau

January 2017

Working paper No. 2017–1

CRESE
30, avenue de l'Observatoire
25009 Besançon
France
http://crese.univ-fcomte.fr/

UNIVERSITÉ DE FRANCHE-COMTÉ

Gender Differences in Competitive Positions: Experimental Evidence on Job Promotion*

Emmanuel Peterlé[†1] and Holger A. Rau[‡2]

[1]Univ. Bourgogne Franche-Comté, CRESE EA3190, F-25000 Besançon, France
[2]University of Göttingen, Germany

January 2017

Abstract

This paper analyzes gender differences in access to competitive positions. We implement an experiment where workers can apply for a job promotion by sending a signal to their employer. We control for gender differences in anticipation of discrimination in a treatment where a computer randomly recruits. Discriminatory behavior by the employer is isolated in a treatment where workers cannot send signals. We find that gender disparity among promoted workers is highest when workers can apply for promotion and employers recruit. Strikingly, the gender composition in competitive position is balanced in the absence of a signaling institution. When signaling is possible, we observe that female workers who do not request a promotion are discriminated against.

JEL Classification numbers: C9, J24, J70.

Keywords: Experiment, Discrimination, Gender Differences, Real Effort.

*We thank Sabine Fischer, Kerstin Grosch, and Stephan Müller for helpful comments. We are also indebted to the audience of the 2016 ESA World Meeting in Jerusalem. We gratefully acknowledge financial support from the Fritz Thyssen Foundation.

†Univ. Bourgogne Franche-Comté, 45D Avenue de l'observatoire, 25030 Besançon (France), *E-mail: emmanuel.peterle@univ-fcomte.fr*

‡Corresponding author, Platz der Göttinger Sieben 3, 37073 Göttingen (Germany), *E-mail: holger.rau@uni-goettingen.de*

1 Introduction

The gender wage gap is a conspicuous phenomenon in labor markets. Though policymakers have carried out many attempts to attenuate the pay gap, it turns out that women still earn significantly less than men. In 2013, data of the OECD countries revealed an unadjusted gender wage gap of 15% (OECD, 2016). Many explanations exist for this prominent wage differential. There is evidence that a large fraction is caused by performance differences (e.g., Blau and Kahn, 1992). That is, women often work part time (Manning and Petrongolo, 2008), do not choose competitive high-income jobs (Daymont and Andrisani, 1984), and rarely work in executive positions (OECD, 2016).

The question remains: "why are women underrepresented in these areas?" One reason may be that gender differences in professional development exist. First, women often anticipate a shorter work life and therefore are less likely to advance their academic (Kahn, 1993) and professional careers (Blau and Kahn, 2000). Second, women may ask less often for job promotions (Babcock and Laschever, 2009).[1] In a field experiment Leibbrandt and List (2014) show that women are less likely than men to negotiate on wages. Recently, laboratory studies highlighted the role of gender differences in preferences (Croson and Gneezy, 2009) and offered behavioral explanations for womens' reluctance to apply for promotion. The experiments of Niederle and Vesterlund (2007) demonstrate that women tend to "shy away from competition." Whereas, Gneezy et al. (2003) show that women perform worse in competitive environments. These examples illustrate that attitudes toward competition may explain why women refrain from selecting competitive jobs or rarely apply for promotion. Although, these findings provide interesting insights, they neglect an important aspect of the labor market: the "demand side."

The adjusted gender wage gap is positive and suggests that supply-side-based differences cannot fully explain the occurrence of the pay differential. There is evidence that the demand-side matters, i.e., employers discriminate against equally qualified women in remuneration (Heinz et al., 2016) or hiring (Bohnet et al., 2015; Beaurain and Masclet, 2016) decisions. The presence of employers also crucially matters for the professional development of employees. First, it may be that few women are promoted because employers favor male applicants over female applicants. Second, the low existence of female executives may be a result of women not applying for promotion as they might anticipate the discriminatory behavior of employers. Analyzing the interplay of the supply and demand side in job-promotion decisions is promising, as it helps to understand how competitive

[1] Experimental evidence by Rigdon (2012) also suggests that women ask for less in modified demand ultimatum games. The author finds that this ultimately leads to a gender wage gap.

positions are achieved.

In this paper we present an experiment on job-promotion decisions where workers and employers simulatenously decide. The experiment aims at investigating the impact of promotion opportunities for the filling of competitive positions. We study the role of gender differences in competitiveness when women can apply for a job promotion in the presence of employers. A major interest is whether female applicants anticipate that workers may discriminate against them. The experiment also investigates employers' reactions, i.e., whether employers discriminate against females when workers can apply for a promotion. Moreover, we focus on the treatment of employees who decide not to apply for a job promotion. Our experiments extend the real-effort framework of Niederle and Vesterlund (2007). In the main treatment employees first decide whether to apply for a competitive job or whether to work under a piece rate. In a next step, the employer is informed of the number of applicants for the job and their gender. Afterwards, the employer makes the promotion decision. Finally, the workers generate the firm revenue in their allocated positions (competitive or non-competitive job). We make use of two control treatments which isolate the impact of job-promotion opportunities on the behavior of employees and employers. The first control treatment tests whether the possibilities females have to apply for promotion affects employers' choices. Therefore, workers cannot apply for job promotion in this treatment. Another control treatment isolates whether women anticipate possible discrimination by the employers. Here, the promotion decision is not taken by an employer. Instead, the selection is made by the computer.

Our data find a substantial gender gap in competitive positions when employees can apply for these positions. We find that more than 70% of competitive positions are held by male employees, whereas only 29% of women advance to competitive jobs. A closer look indicates that females signal their willingness to compete significantly less often than males. Notably, this drives our gender gap in access to competitive positions in two aspects. First, males are more likely to be selected as they more frequently send signals. Second, employers prefer to promote male employees when male and female employees decide not to apply for promotion. We do not observe raw discrimination against female workers in the absence of the opportunity to send signals. In this case the gender pay gap shrinks, i.e., only 55% of men hold competitive positions.

Our findings contribute to a better understanding of females' access to competitive positions. The insights on the effects of job-promotion opportunities may help to improve the design of job environments to attenuate the gender wage gap.

2 Experimental design

2.1 Design overview

In our experiment, we use the real-effort task introduced by Niederle and Vesterlund (2007). Participants perform a simple arithmetic task where they add sets of five two-digits numbers for five minutes. Our experiment is divided into four games.

In the first game, all participants perform the task and are paid according to a piece-rate payment scheme. They receive €0.50 for every correctly solved problem. In game two, all participants enter a three-player, winner-takes-all tournament. They are matched with two counterparts, and performances are compared to determine payoffs (Niederle and Vesterlund 2007). The participant that achieves the best performance receives €1.50 for every correctly solved problem. The two losing participants do not earn payment for this game. Ties at the top position are solved with a random draw. The third game resembles an experimental *promotion game* to which participants take part either in the role of worker or the role of employer. We conduct three treatments which follow a between-subjects design that differs in the actions participants can undertake in game three. We describe this game in further detail in the next subsection. Finally, we elicit participants' risk attitude in game four of our experiment. We implement the lottery-choice task introduced by Eckel and Grossman (2002). Participants are given six gambles with two possible outcomes (low payoff/high payoff) each. Both events occur with a likelihood of 50%. The lotteries differ both in expected payoffs and in the variance of the outcome distribution (see table 1 below). Choices are numbered from 1 to 6 and are sorted from the safest to the most risky lotteries. Risk is measured as the standard deviations of payoffs (Eckel and Grossman 2002).

Choice	Low Payoff (€)	High Payoff (€)	Exp. payoff	Implied CRRA Range
1	5.60	5.60	5.60	$3.46 < r$
2	7.20	4.80	6.00	$1.16 < r < 3.45$
3	8.80	4.00	6.40	$0.71 < r < 1.16$
4	10.40	3.20	6.80	$0.50 < r < 0.71$
5	12.00	2.40	7.20	$0 < r < 0.50$
6	14.00	0.40	7.20	$r < 0$

Table 1: The gamble choices in game four.

2.2 The promotion game

In game three of the experiment, participants are randomly assigned either to the role of an employer or of a worker. Experimental firms are formed by matching together one employer and three workers. Those groups take part in an experimental setting that we refer to as the *promotion game*. In each experimental firm, one competitive position is to be filled. In this position the payment of the worker depends on the performance in a winner-takes-all tournament. Whereas, two workers are paid according to a piece rate. The task is exactly the same as the task performed in games two and three of the experiment. The worker who takes part in the tournament is competing against two workers from other firms, who are also taking part in the tournament. For workers, payoffs are similar to those previously described. Under a piece-rate scheme, the worker earns €0.50 for each correctly solved problem. Under the tournament scheme, the worker who performs best earns €1.50 for each problem solved, while the others earn nothing. The employer's payoff is entirely determined by the performance of the three workers in the firm. More precisely, the employer earns €0.20 for each problem correctly solved by workers performing under a piece rate. The employer receives €1 for each problem correctly solved by the worker performing in the tournament, if this worker wins the competition. The employer receives nothing for the problems solved by the worker performing in the tournament if she loses the tournament.

Our experimental setting aims at identifying two main channels through which female underrepresentation in competitive positions may occur. For that reason, we run three treatments that differ in the set of actions available to participants. We first present the *Baseline* treatment, which allows both workers and employers to take action. We then present the *No Selection* and the *No Signal* treatments.

The *Baseline* promotion game consists of three steps. In step one, participants in the role of a worker can state their willingness to enter competition. This statement takes the form of a costly signal that is displayed in the next step to the employer. Workers receive an endowment of €0.20 which can be employed to send a signal or it can be kept for themselves.[2] In step two, the employer selects which worker will get promoted to the competitive position. When employers decide on this they are provided with information regarding workers' demographics (gender, age, level of study).[3] Importantly,

[2]We employ a costly signal to integrate a monetary trade-off for the workers to apply or not to apply. Moreover, it can be argued that in real life applying for promotion is time-consuming and comes at opportunity costs.

[3]The information on demographics was collected at the beginning of the computer experiments. We

employers also receive information on whether workers have sent a signal to apply for promotion. In step three, workers are informed of the remuneration scheme allocated by the employer. Afterwards, workers perform in the task. Workers who have been promoted to the competitive position compete against workers from other firms who were also promoted by their employers. While workers perform the task in the final step of the promotion game, employers face survey questions regarding their beliefs on workers' performances. More precisely, employers are asked to rank the three workers of the firm from the best to the worst-performing in the mandatory tournament of game two.[4] We are particularly interested in the gender composition among promoted workers. An unbalanced composition in the *Baseline* treatment may be due to females "shying away from competition" (Niederle and Vesterlund, 2007) or because of employers selecting men over women (i.e., gender discrimination). To disentangle between these two potential channels, we implement two additional treatments.

In the *No Selection* treatment, the employer does not select a worker to be promoted. Instead, the selection process is automatic and computerized according to a fair and fully disclosed rule. If only one of the three workers sends a signal, she automatically accesses the competitive position. If more than one worker sends a signal, the computer randomly chooses who among them will access the competitive position. Finally, if no signal is sent the computer randomly picks one of the three workers to enter the tournament. In this treatment, employers do not have direct impact on the selection process. An imbalanced gender composition among promoted participants can only be attributed to gender differences in attitude toward competition.

The *No Signal* treatment differs from the *Baseline* in that workers cannot express their willingness to enter the competition. Workers cannot send a signal and employers are only shown demographics (gender, age, level of study) when selecting a worker to be promoted. In this treatment, gender differences in competitiveness have no direct impact on the selection process. Only the choice of the employers could explain a potential gender imbalance in access to promotion.

---

also present the employers with additional information beside gender. The reason is that we intend to avoid making workers' gender a focal point for employers.

[4]We did not incentivize this as we want to avoid strategic behavior motivated by hedging (Blanco et al. 2010). If the elicitation of the beliefs was incentivized, employers could hedge their outcome by stating that they believe that the promoted worker performed worst.

2.3 Experimental procedure

The experiment was programmed using z-Tree (Fischbacher 2007). In total, 248 participants were recruited with ORSEE (Greiner 2015) and took part in the experiment (see table 2 below). One hundred and twenty-three women and 125 men participated in the experiment. The subject pool consisted of students from the University of Göttingen from various fields. The sessions lasted approximately 45 minutes. Subjects' average earnings were €10.62, including a show-up fee of €5.

Treatment	Sessions	Participants
Baseline	5	96
No Selection	4	64
No Signal	5	88
Total	14	248

Table 2: Experimental procedure

3 Results

In this section we report our findings. We start with an overview of the overall result on the gender composition in competitive positions. As a next step we focus on the drivers for the emergence of the outcome we observe. We analyze workers' decisions to send a signal and employers' promotion choices in separate sections. When using non-parametric test methods, we always report two-sided $p - values$ throughout.

3.1 Gender composition in competitive positions

We first focus on the gender balance in competitive positions, i.e., the proportion of females among the promoted participants. Recall that in the *Baseline* treatment, the promotion process is affected by both the decision of the worker to send a signal and the choice of the employer. In our control treatments we isolate the drivers for the promotion outcome. First, in the *No Selection* treatment employers cannot select employees. Hence, only a worker's decision to send a signal affects the promotion outcome. Second, in the *No Signal* treatment employers select employees, but workers cannot send a signal. As a consequence, imbalanced gender compositions in this treatment can only be attributed to unequal treatment from the employers' side.

Figure 1: Gender composition in competitive positions

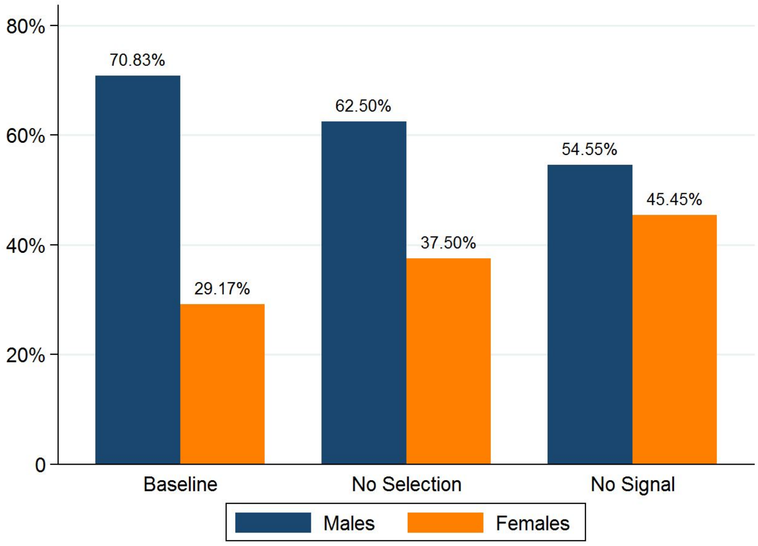

Figure 1 reports gender compositions in competitive positions across our three treatments. A conspicuous finding is that we observe the highest gender gap when employees can send signals and employers select. That is, we find that only 29.17% of promoted participants in *Baseline* are females. This gender composition is significantly different from a perfectly balanced distribution (proportion test, $p = 0.041$).[5] By contrast, more female workers are promoted when employers cannot select employees. In *No Selection*, it can be seen that the proportion of female workers increases to 37.50%. In this treatment the gender composition is not significantly different from a perfectly balanced distribution (proportion test, $p = 0.181$). The gender composition is the closest to a perfectly balanced composition in *No Signal*, where 45.45% of promoted participants are females (proportion test, $p = 0.777$). To summarize, our data suggest that a gender gap in competitive positions occurs when employees can send signals and employers decide on promotion.[6] The second finding indicates that parts of this differential may be driven by employers' selection process. The third finding emphasizes that the lion's share of the gender gap may be induced by employees' willingness to apply for promotion.

Result 1:
(a) The signaling institution entails a gender gap in promotion to competitive positions.
(b) The gender gap is mitigated when employers do not select. It is close to non-existent

[5] Proportion tests are run to compare the observed distributions with theoretical distributions drawn from the numbers of females and males among workers.

[6] All statistical analyses also hold for χ^2 tests (Baseline: $\chi^2(1) = 4.508,\ p = 0.037$; No Selection: $\chi^2(1) = 1.689\ p = 0.194$; No Signal: $\chi^2(1) = 0.242,\ p = 0.622$).

when workers cannot signal their willingness to compete and employers select.

In what follows, we further analyze the drivers of Result 1 in terms of employees' behavior and potential discriminatory behavior from the employers' side. The data emphasized that the gender gap is lowest when workers do not have the ability to signal their willingness to compete. Therefore, in the next section we will focus on workers' decisions to send a signal. Afterwards, we explore the role of employers in the promotion process, and how this may affect the probability of women being promoted.

3.2 Gender differences in applications to competitive positions

In this section we only focus on workers and study their inclination to apply for a competitive position. The analysis concentrates on the situations where workers have the opportunity to reveal their willingness to compete by sending a signal. We focus on the treatments where signals can be sent to the employer (*Baseline* treatment) or to the computer (*No Selection* treatment).

Our results are in line with previous experimental evidence on gender differences in the willingness to compete (e.g., Niederle and Vesterlund 2007; Croson and Gneezy 2009; Masclet et al. 2015; Heinz et al. 2016).[7] We find in both treatments that women send a signal significantly less often than men. In *Baseline*, 16.67% of female workers send a signal, whereas this holds for 44.44% of male workers ($\chi^2(1) = 6.546,\ p = 0.011$). In *No Selection*, a significantly lower share of female workers (19.23%) send a signal than male workers (45.45%; $\chi^2(1) = 3.814,\ p = 0.051$). Interestingly, the distribution of subjects that send a signal does not significantly differ between the treatment where employers can select the person to be promoted (*Baseline*) and the treatment where the computer decides (*No Selection*) ($\chi^2(1) = 0.007,\ p = 0.936$). Hence, one might at first conclude that the presence of employers does not influence workers' decision to send signals. However, in the *Baseline* treatment several channels may be at work. First, similar to the *No Selection* treatment, it may be that the decision to signal competitiveness is affected by individual risk preferences. Second, it is also likely that female and male workers anticipate that the employer will receive information on their demographics. Thus, it may be that female and male workers react differently to the presence of employers.

To isolate these effects on the determinants of applying for promotion, we run probit regressions on the probability of sending a signal. We present a regression for each treatment where workers can send a signal (*Baseline*, *No Selection*).

[7]Note that with the exception of Heinz et al. (2016) these papers do not model employers.

Table 3: Determinants of applying for promotion in *Baseline* and *No Selection*.

	Probability to send a signal (Probit estimates)	
	Baseline	No Selection
Female	-0.745**	-0.519
	(0.348)	(0.440)
Risk tolerance	0.270**	0.285**
	(0.123)	(0.126)
Level of study	-0.017	-0.016
	(0.063)	(0.059)
Economist	0.214	-0.224
	(0.335)	(0.434)
Score in first tournament	0.044	0.127**
	(0.039)	(0.062)
Constant	-1.783**	-2.474***
	(0.781)	(0.893)
Observations	72	48
Mc Fadden's R^2	0.1581	0.1950

Robust standard errors in parentheses
*** $p < 0.01$ ** $p < 0.05$ * $p < 0.1$

The regressions control for demographics such as the gender of the worker (*female*), the number of semesters studied so far (*level of study).* We also add a dummy which is positive for either economic or business students (*economist*). The models control for subjects' risk preferences (*risk tolerance*), measured in part four of our experiment.[8] Finally, *score in first tournament* controls for participants' ability which is measured as the number of correctly solved puzzles in the mandatory tournament (part two of our experiment).

The first column confirms our previous finding, that women are less likely than men to send a signal in *Baseline*. The second regression suggests that this is not the case in *No Selection*. That is, the coefficient of *female* is negatively significant in *Baseline* and insignificant in *No Selection*. We find evidence that risk-averse individuals tend to shy away from competition. That is, *risk tolerance* is positively significant in both treatments. Thus, subjects with a more pronounced degree of risk tolerance are more likely to send a signal. Interestingly, Table 1 shows that the determinants of applying for promotion in

[8]Risk tolerance is a variable ranging from 1 to 6, which takes the value 1 if the individual has chosen the risk-free lottery and 6 if the individual has chosen the riskiest lottery.

No Selection contrast from the ones in *Baseline*. In *No Selection*, being a female does not appear to have a significant impact on the decision to send a signal. By contrast, *risk tolerance* and subject's ability significantly increase the probability of applying for promotion.

The regression results shed new light on the seemingly similar gender gap in competitiveness between the *Baseline* and the *No Selection* treatments. Although females in *No Selection* send a signal significantly less often than men, being a woman *per se* does not appear to affect this decision.[9] It turns out that the gender difference in risk tolerance is the most compelling argument to explain the observed gender-application gap in this treatment. Indeed, we find in the *No Selection* subsample that women are significantly more risk averse than men (Mann-Whitney test, $p = 0.013$).[10]

By contrast, in the *Baseline* treatment, the signal sent by candidates to the competitive position is considered by the employer, who also has access to information regarding age, level of study, and particularly gender. Women anticipating gender-based discrimination from employers may be reluctant to spend money on a signal that will anyway be disregarded. This is not the case in the *No Selection* treatment, where the computer automatically promotes the applying worker or makes a random draw if too many workers apply. The anticipation of discriminatory practices may therefore offer a possible explanation for our findings.[11] We summarize the findings in Result 2.

Result 2:
(a) In the presence and absence of employers, women send less signals than men.
(b) In the presence of employers, female and risk-averse workers send less signals. In the absence of employers, only risk aversion can explain why workers send less signals.

3.3 Employers' selection of applicants

In this section we turn to the analysis of employer behavior. Another possible explanation for gender differentials in access to promotion is the presence of discriminatory practices

[9]Results presented in table 3 hold when specifying a bootstrap estimation of standard errors.

[10]In the *Baseline* subsample, women and men do not differ in risk tolerance (Mann-Whitney test, $p = 0.322$). However, a significant difference can be found in *No Signal* (Mann-Whitney test, $p = 0.006$). The pooled data suggest that females are more risk averse than males (Mann-Whitney test, $p < 0.001$), which is in line with other experiments (e.g., Eckel and Grossman 2008; Charness and Gneezy 2012)

[11]An alternative explanation for our findings can be that women, compared to men, are more reluctant to display initiative. Evidence of this can be found in the bargaining literature which reports that: "women don't ask" in negotiations (Babcock and Laschever 2009). An interpretation may be motives of image concern, i.e., women could be perceived as acting unladylike if they behave too demanding.

against women. Our data show that the underrepresentation of females in competitive positions narrows down in the *No Selection* treatment, where employers are not involved in the promotion process. This may be explained by an amplified gender gap induced by discriminatory behavior of the employers in the *Baseline* treatment. Our previous findings suggest that this may also be caused by women being more eager to apply for the competitive position when no discrimination is possible. We now investigate whether female workers suffer from an unequal treatment in our experiment. We therefore analyze the treatments (*Baseline* and *No Signal*) where employers have the opportunity to decide which worker should be promoted to the competitive position. Before analyzing possible discriminatory behavior, we focus on workers' performance to study whether favoritism of male workers may be justified by performance arguments. In all stages of our experiment we do not observe any gender gap in the task performance. This is supported by Mann-Whitney tests on gender differences in performance (piece rate: $p = 0.437$, tournament: $p = 0.658$, and the selection stage: $p = 0.448$). This confirms the findings in Niederle and Vesterlund (2007). As a consequence, if an employer intends to select the promoted worker solely on a performance criterion, gender should be regarded as irrelevant.

Table 4 reports estimates of probit regressions aiming at identifying the determinants of selection in the competitive positions. The models incorporate a dummy (*signal*) which is positive when employees send a signal to the employer. We also test for the impact of the information on workers' gender (*female*). Moreover, we control for the impact of the information regarding workers' demographics (*age* and *level of study*).

Model (1) shows that in *Baseline* workers who send a signal are significantly more often promoted. Interestingly, even when controlling for *signal*, females are less often selected for job promotion. That is, the coefficient of *female* is significantly negative. More information can be obtained by having a closer look at the data. When comparing women and men that send a signal, we do not observe any gender inequality in promotion decisions. In that case, 66.67% of women are selected for the tournament, compared to 56.25% of men (Mann-Whitney test, $p = 0.666$). However, significant gender inequality appears when comparing women and men that did not send a signal. In that case, only 10% of women are selected for job promotion as compared to 40% of men (Mann-Whitney test, $p = 0.013$). This finding is confirmed in Model (2) of table 2. In this regression we include a new interaction variable: *female* $\times$ *signal*. The coefficient associated with this variable is weakly significant, with the opposite sign, and has a similar amplitude to the female coefficient. This suggests that sending a signal cancels out the penalty of being a female on the probability of being selected.

Table 4: Access to promotion

	Probability to be selected for tournament (Probit estimates)			
	Baseline		No Selection	No Signal
	(1)	(2)	(3)	(4)
Signal	0.925**	0.449	1.744***	-
	(0.380)	(0.454)	(0.448)	
Female	-0.878**	-1.340***	-0.185	-0.300
	(0.359)	(0.440)	(0.426)	(0.351)
Female × *Signal*	-	1.379*	-	-
		(0.773)		
Age	-0.089**	-0.094*	-0.024	-0.153**
	(0.043)	(0.055)	(0.036)	(0.060)
Level of study	0.081	0.097	0.001	0.056
	(0.073)	(0.075)	(0.066)	(0.061)
Constant	1.372	1.626	-0.408	3.035**
	(1.101)	(1.375)	(0.931)	(1.451)
Observations	72	72	54	66
Mc Fadden's R^2	0.1916	0.2273	0.3611	0.0770

Robust standard errors in parentheses
$***p < 0.01, **p < 0.05, *p < 0.1$

Model (3) reports estimates in the *No Selection* treatment. Here, the selection process is carried out by the computer that uses only the signal sent as a determinant of choice. Our regression results in Model (3) validate this approach and highlight that the coefficient associated with *signal* is highly significant and positive. This demonstrates that only the signal mattered in the selection process. Model (4) reports estimates in the *No Signal* treatment. Here, employers select a worker to be promoted without having any information about the worker's willingness to compete. In this context, we find that the coefficient associated with *female* is insignificant. Hence, women do not appear to be generally discriminated against in the *No Signal* treatment.

Result 3:

(a) Workers who send a signal are significantly more often promoted.

(b) When signals are possible women are discriminated against, unless they send a signal.

Although we do not find evidence of discrimination in the *No Signal* treatment, where employers' decisions cannot be affected by workers' willingness to compete, our data suggest the presence of unequal treatment in the *Baseline* treatment. Women who do not display a willingness to apply for promotion are significantly less often promoted than men in the same position. An interpretation may be that employers in the *Baseline* treatment interpret a worker's application for promotion as a signal of a high-performing or motivated worker. This is probably the reason why workers who send a signal are most often promoted. Hence, female workers have the chance to outpace male competitors by sending a signal. The functioning of this process was demonstrated by the significantly positive effect of *female* $\times$ *signal* in Model (2). However, if workers do not send signals the employer has to hold a belief on workers' performance in the task. In this case, it is possible that employers in *Baseline* may hold the belief (wrongly)[12] that male workers achieve a better performance than women. As a consequence, it is possible for them to discriminate against women and show a preference for male non-applicants.

Our experimental design allows us to get some input into the role of beliefs in the selection process. While workers were performing the task in game three, we asked employers a series of questions regarding their beliefs on workers' performance. More precisely, they were asked to rank the three workers of the firm from the best performing to the worst performing, according to their beliefs. We find that employer's ranking is directly related to their decision in the selection process. Females are considered on average as less performing than males in the *Baseline* treatment (Mann-Whitney test, $p = 0.002$) and to a lesser extent in the *No Signal* treatment (Mann-Whitney test, $p = 0.073$). When comparing only women and men who sent a signal in the *Baseline* treatment, no gender difference is found (Mann-Whitney test, $p = 0.666$). This support our interpretation that signaling willingness to compete cancels out the pre-existing prejudice against females.

4 Discussion

Several precautions should be taken when interpreting and extrapolating our findings. First, we acknowledge that a contextualized lab experiment is unlikely to capture the whole complexity of existing promotion processes. Our experiment aims at identifying specific behavioral mechanisms that cannot be observed in non-experimental data. We do believe that laboratory experiments do not compete with other empirical approaches,

[12]Recall, that no gender differences in performance exist. Thus, believing that male workers achieve a better outcome is incorrect.

but are rather complementary. Second, promotions in our experimental setting are based on the experimental design of Niederle and Vesterlund (2007). We therefore apply a tournament-based remuneration scheme where the expected value does not exceed the alternative piece-rate scheme. Although this may narrow down the scope of our study on promotion to *competitive* positions, it allows direct comparisons to existing experimental studies.

We find that women apply less frequently than men to competitive positions, i.e., they send signals significantly less often. This finding echoes previous experimental evidence of women behaving less competitive than men (e.g., Niederle and Vesterlund 2007; Niederle et al. 2013; Buser et al. 2014; Wozniak et al. 2014). This gender gap in competitiveness can partly be explained by women being more risk averse than men.[13]

Extending the Niederle and Vesterlund (2007) framework by adding the demand-side represented by an employer, offers interesting new insights. Our regression analysis suggests that being a female is relevant in the decision to apply for the competitive position when an employer has decision power. This is, however, no longer the case when the employer does not take part in the selection process (*No Selection*). This observation points to women anticipating discriminatory behavior from employers in *Baseline*. However, this does not mean that taking discretionary power away from employers is sufficient to achieve gender equality in the willingness to compete. Other determinants, such as risk aversion, persist even in the absence of an employer and prevent the gender gap from significantly decreasing.[14]

Ironically, the only evidence of gender prejudice that we observe is directed toward women who do not exhibit a willingness to compete. When workers have the ability to send a signal to the employer, we observe that women who ask for a promotion are not disadvantaged compared to men. More than a claim on one's willingness to compete, the signal may be interpreted by the employer as a proxy for productivity and/or motivation. In our experiment employers, on average, perceive female workers as less productive than male workers. However, there is no gender difference in performance. Interestingly, the (inaccurate) stereotype on performance is not held for women who send a signal. By contrast, in the *No Signal* treatment, workers do not have the opportunity to express their willingness to compete. In this context, we do not observe any discrimination against

[13] There are two sources of risk associated with the decision to send a signal in our experiment: i) Payoffs in the tournament are conditional on workers' relative performance; ii) the costly signal does not ensure tournament entry.

[14] Other potential determinants that we do not directly account for include self-confidence or social preferences.

women. In this treatment the gender composition in competitive positions is close to being perfectly balanced. The “signaling institution” may therefore be considered at the heart of the issues leading to gender inequality in access to promotion. One should, however, question the relevance of a working environment where individuals cannot express their preferences. So far, our analysis has pertained to the attainment of gender equality in promotion. Reaching perfect gender equality may, however, be undesirable, if it implies that women are assigned to positions that they are not aiming at.

5 Conclusion

In the current paper, we investigate gender differences in access to competitive positions when *both* workers and employers are active in the promotion process. We implement a simple experiment where workers can signal their willingness to be promoted to the competitive position and employers have a final say on the matter. We find that women are significantly less likely to be promoted. This is a consequence of the interplay between the supply-side (gender difference in competitiveness) and the demand-side (employer discrimination) in our experimental labor market. To disentangle between both channels, we implement two control treatments that either limit workers expressing their preferences (*No Selection*) or limit an employer's role in the promotion process (*No Signal*).

We report three main findings. First, gender disparity in competitive positions is highest when both employer and workers interact within the promotion process. It significantly decreases when the employer does not have discretionary power. In stark contrast, the gender composition among promoted participants is close to being perfectly balanced in a situation where employers select with no information on the willingness to compete. Second, women apply for the competitive position significantly less often than men. Although gender differences in risk tolerance appear to be the main driver of this gender gap, our data suggest that women may anticipate discriminatory behavior from the employer. Third, we do not find evidence of discriminatory behavior from employers when they are not informed of workers' willingness to compete. In an environment where workers have the ability to signal themselves, women who do not send a signal are significantly less likely to be promoted compared to men in the same situation. However, women signaling themselves do not suffer from unequal treatment. While employer discrimination is still at the heart of current policy debates, our findings rather point to the role of workers' preferences. If gender parity is a desirable outcome, the supply-side of the labor market would be the relevant policy target.

References

L. Babcock and S. Laschever. *Women don't ask: Negotiation and the gender divide.* Princeton University Press, 2009.

G. Beaurain and D. Masclet. Does affirmative action reduce gender discrimination and enhance efficiency? new experimental evidence. *European Economic Review*, 90:350–362, 2016.

M. Blanco, D. Engelmann, A. K. Koch, and H.-T. Normann. Belief elicitation in experiments: is there a hedging problem? *Experimental Economics*, 13(4):412–438, 2010.

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

F. D. Blau and L. M. Kahn. The gender earnings gap: learning from international comparisons. *The American Economic Review*, 82(2):533–538, 1992.

I. Bohnet, A. Van Geen, and M. Bazerman. When performance trumps gender bias: Joint vs. separate evaluation. *Management Science*, 62(5):1225–1234, 2015.

T. Buser, M. Niederle, H. Oosterbeek, et al. Gender, competitiveness and career choices. *The Quarterly Journal of Economics*, 129(3):1409–1447, 2014.

G. Charness and U. Gneezy. Strong evidence for gender differences in risk taking. *Journal of Economic Behavior & Organization*, 83(1):50–58, 2012.

R. Croson and U. Gneezy. Gender differences in preferences. *Journal of Economic literature*, 47(2):448–474, 2009.

T. N. Daymont and P. J. Andrisani. Job preferences, college major, and the gender gap in earnings. *Journal of Human Resources*, pages 408–428, 1984.

C. C. Eckel and P. J. Grossman. Sex differences and statistical stereotyping in attitudes toward financial risk. *Evolution and human behavior*, 23(4):281–295, 2002.

C. C. Eckel and P. J. Grossman. Forecasting risk attitudes: An experimental study using actual and forecast gamble choices. *Journal of Economic Behavior & Organization*, 68 (1):1–17, 2008.

U. Fischbacher. z-tree: Zurich toolbox for ready-made economic experiments. *Experimental economics*, 10(2):171–178, 2007.

U. Gneezy, M. Niederle, A. Rustichini, et al. Performance in competitive environments: Gender differences. *The Quarterly Journal of Economics*, 118(3):1049–1074, 2003.

B. Greiner. Subject pool recruitment procedures: organizing experiments with orsee. *Journal of the Economic Science Association*, 1(1):114–125, 2015.

M. Heinz, H.-T. Normann, and H. A. Rau. How competitiveness may cause a gender wage gap: Experimental evidence. *European Economic Review*, 90:336–349, 2016.

S. Kahn. Gender differences in academic career paths of economists. *The American Economic Review*, 83(2):52–56, 1993.

A. Leibbrandt and J. A. List. Do women avoid salary negotiations? evidence from a large-scale natural field experiment. *Management Science*, 61(9):2016–2024, 2014.

A. Manning and B. Petrongolo. The part-time pay penalty for women in britain. *The Economic Journal*, 118(526):F28–F51, 2008.

D. Masclet, E. Peterle, and S. Larribeau. Gender differences in tournament and flat-wage schemes: An experimental study. *Journal of Economic Psychology*, 47:103–115, 2015.

M. Niederle and L. Vesterlund. Do women shy away from competition? do men compete too much? *The Quarterly Journal of Economics*, (3):1067–1101, 2007.

M. Niederle, C. Segal, and L. Vesterlund. How costly is diversity? affirmative action in light of gender differences in competitiveness. *Management Science*, 59(1):1–16, 2013.

OECD. Oecd employment outlook 2016. 2016.

M. L. Rigdon. An experimental investigation of gender differences in wage negotiations. *Working Paper*, 2012.

D. Wozniak, W. T. Harbaugh, and U. Mayr. The menstrual cycle and performance feedback alter gender differences in competitive choices. *Journal of Labor Economics*, 32(1):161–198, 2014.