NBER WORKING PAPER SERIES

GENDER PREFERENCES IN JOB VACANCIES AND WORKPLACE GENDER DIVERSITY

David Card
Fabrizio Colella
Rafael Lalive

Working Paper 29350
http://www.nber.org/papers/w29350

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
October 2021

We thank Ghazala Azmat, Omar Bamieh, Thomas Le Barbanchon, Pierre Boyer, Clément de Chaisemartin Armin Falk, Lucas Girard, Ingrid Haegele, Kerstin Holzheu, Patrick Kline, Michael Kosfeld, Francis Kramarz, Peter Kuhn, Marco Manacorda, Paolo Mangano, Steeve Marchand, Guy Michaels, Maria Guadalupe, Alexandra Roulet, Ity Shurtz, Benoit Schmutz, Camille Terrier, Andrea Weber, David Yanagizawa-Drott, Lennart Ziegler, Florian Zimmermann, and Josef Zweimüller for helpful comments. We thank Josef Zweimüller for providing us access to the data, and Damian Osterwalder for support in interpreting features of the data. Elke Lujansky-Lammer and Rainer Lichtfeld provided information on the institutional background. We also thank seminar participants at UC Santa Barbara, UC Berkeley, University of Lausanne, CREST, CERGE-EI, Insead, and Ben Gurion, as well as audience at the Brucchi Luchino Conference, the Ski and Labor Workshop, the Annual Meeting of the Italian Association of Labor Economics, IDSC of IZA Workshop, the Spring Meeting of Young Economist, and the Annual Congress of the Swiss Society of Economics and Statistics. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

Gender Preferences in Job Vacancies and Workplace Gender Diversity
David Card, Fabrizio Colella, and Rafael Lalive
NBER Working Paper No. 29350
October 2021
JEL No. J16,J63

ABSTRACT

In spring 2005, Austria launched a campaign to inform employers and newspapers that gender preferences in job advertisements were illegal. At the time over 40% of openings on the nation's largest job-board specified a preferred gender. Over the next year the fraction fell to under 5%. We merge data on filled vacancies to linked employer-employee data to study how the elimination of gender preferences affected hiring and job outcomes. Prior to the campaign, most stated preferences were concordant with the firm's existing gender composition, but a minority targeted the opposite gender - a subset we call non-stereotypical vacancies. Vacancies with a gender preference were very likely (>90%) to be filled by someone of that gender. We use pre-campaign vacancies to predict the probabilities of specifying preferences for females, males, or neither gender. We then conduct event studies of the effect of the campaign on the predicted preference groups. We find that the elimination of gender preferences led to a rise in the fraction of women hired for jobs that were likely to be targeted to men (and vice versa), increasing the diversity of hiring workplaces. Partially offsetting this effect, we find a reduction in the success of non-stereotypical vacancies in hiring the targeted gender, and indications of a decline in the efficiency of matching. For the much larger set of stereotypical vacancies, however, vacancy filling times, wages, and job durations were largely unaffected by the campaign, suggesting that the elimination of stated preferences had at most small consequences on overall job match efficiency.

David Card
Department of Economics
549 Evans Hall, #3880
University of California, Berkeley
Berkeley, CA 94720-3880
and NBER
card@econ.berkeley.edu

Fabrizio Colella
Department of Economics
HEC University of Lausanne
UNIL-Chamberonne
1015 Lausanne
Switzerland
fabrizio.colella@unil.ch

Rafael Lalive
Department of Economics
HEC University of Lausanne
1015 Lausanne
Switzerland
Rafael.Lalive@unil.ch

1 INTRODUCTION

Rules and informal prohibitions governing gender roles in the workforce have gradually diminished over the last century (e.g., Goldin, 2014). Nevertheless, men and women still tend to work at different firms.[1] How much of the remaining gender segregation in a modern labor market is due to skill or preference differences between males and females versus discriminatory employer preferences is hard to discern, particularly when employers are explicitly forbidden from expressing their preferences (Kuhn and Shen, 2013). Evidence from audit studies (e.g., Booth and Leigh, 2010, Kline *et al.* (2021)) suggests that employer preferences tend to reinforce gender stereotypes and segregation.[2] As noted by Cahuc *et al.* (2019), however, it may be difficult to infer the impacts of discriminatory preferences on hiring outcomes from the data that are collected in typical audit studies.

While the use of stated gender preferences (SGP's) in job recruiting was eliminated in the U.S. in the 1970s, they were still prevalent throughout Europe at the turn of the 21st century, including in Austria, where about 40% of job vacancies listed a preferred gender, despite a 1985 law banning SGP's. In mid 2004, an amendment to the Austrian Equal Treatment Act introduced financial penalties for including gender preferences in job ads, and in early 2005 the *Ombud for Equal Treatment* (OET) agency launched a campaign to inform the public about the law and these penalties. We use the changes after the information campaign to measure the effects of eliminating discriminatory signals at the earliest stage of the job matching process. Vacancy postings from the largest online job board in Austria contain explicit and unambiguous information on stated gender preference, and we combine these data with administrative data on the firms posting the vacancies and the recruits who filled the vacancies. The resulting data set allows us to document how the use of gender

[1] Hellerstein *et al.* (2008) report that a typical female employee of a larger establishment in the U.S. has about 60% female coworkers, while a typical male has only 40% female coworkers. For a broader sample of workers in the U.K., Mumford and Smith (2009), online Appendix Table A2, report corresponding rates of 70% female coworkers for women and 34% for men. Slightly larger gender gaps in exposure to female coworkers are estimated by Card *et al.* (2016) for Portugal and by Gerard *et al.* (2018) for Brazil.

[2] Azmat and Petrongolo (2014) and Baert (2017) present overviews of these studies. A number of authors have argued that the tendency to favor the gender group that stereotypically matches the job opening may arise from firms' intentions to reduce search time and improve expected outcomes, rather than from animus discrimination – see e.g., the discussion in Kuhn and Shen (2013).

preferences in the pre-campaign period varied across employers, and how actual hiring outcomes varied with the SGP's in the associated vacancies.

We show that prior to the campaign, firms with a mainly male workforce were likely to state a preference for males, and vice versa for firms with a mainly female workforce ("stereotypical" preferences). Nevertheless, a small fraction of employers were using SGP's to look for employees that were the opposite gender of the majority of their workforce – posting what we define as "non-stereotypical" vacancies. Vacancies with a gender preference were very likely (>90%) to be filled by a candidate of the preferred gender, even in the case of non-stereotypical preferences. Vacancies with gender preferences were also filled somewhat faster (especially those with non-stereotypical preferences), suggesting that SGP's provided highly informative signals when they were allowed. But would the elimination of these signals change actual hiring outcomes? Or would firms reach the same hiring decisions even after the elimination of SGP's?[3]

To answer these questions we use data on occupation, industry and gender composition of the firm's workplace to classify vacancies in the pre-campaign period into three groups: likely to specify a SGP for males, likely to specify a SGP for females, or likely to specify no preference. We then use the classification model to "tag" vacancies before and after the campaign and conduct simple event studies of hiring outcomes for different types of vacancies.

We find that the elimination of firms' abilities to advertise their gender preferences led to a significant rise (+2.5 percentage points) in the fraction of women hired to fill vacancies with a predicted male preference, and a smaller but still significant rise in the fraction of men hired to fill vacancies with a predicted female preference (+1 ppt). The vast majority of vacancies with predicted preferences were stereotypical, (e.g., concordant with the existing gender composition of the firm's workplace), and we show that hires after the campaign increased firm-based gender diversity. These findings suggest that some employers that were posting SGP's prior to the campaign were willing to hire workers of the opposite gender, and actually did so once the composition of their applicant

[3]A similar question arises for recent policies that forbid the collection of data on criminal records or credit histories at the job application phase (Agan and Starr, 2018; Doleac and Hansen, 2020; Bos *et al.*, 2018; Ballance *et al.*, 2020) or the collection of information on salary history (Agan *et al.*, 2020).

pool was broadened by the elimination of SGP's.

To gain further insights we classify vacancies by the interaction between predicted gender preference and the gender composition of the firm's workplace. This allows us to compare stereotypical versus non-stereotypical vacancies. We find an *increase* in the hiring of women to fill vacancies with both stereotypical and non-stereotypical male predicted preferences (consistent with the objectives of the OET campaign), but a *decrease* in the hiring of women to fill jobs that were predicted to state a non-stereotypical preference for females. The latter impact was presumably an unintended consequence of the campaign, but is interpretable through the lens of a simple model where job seekers use a combination of general information about a firm (like its industry and gender composition), together with information in the SGP, in deciding where to apply. In such a model non-stereotypical SGP's can inform job searchers about preferences that are otherwise unexpected, allowing them to more easily find rare opportunities.

We also examine the impacts of the campaign on three key outcomes for filled vacancies: the time required to fill the vacancy; the wage on the newly established job; and the duration of the new job. We find that the average times required to fill vacancies with a predicted male preference (either stereotypical or non-stereotypical) were unaffected by the campaign, as was the time to fill predicted female vacancies at mainly female workplaces. Thus, for the vast majority of vacancies, matching speed remained constant, though we do find a significant rise in the time to fill vacancies with non-stereotypical female preferences. For wages we find small effects across the board once we control for compositional effects attributable to the gender gap in pay. For job durations we also find small effects, apart from a decline in the duration of jobs created from vacancies with a stereotypical female gender preference. We attribute the latter impact to a pre-existing trend toward shorter job durations for women in majority female workplaces, rather than to the OET's campaign.

Our findings contribute to the large literature on gender segregation at the workplace level, including Blau (1977), Groshen (1991), Petersen and Morgan (1995), Bayard *et al.* (2003), and Card *et al.* (2016). Specifically, we show the importance of direct employer signals – now outlawed in many higher-income countries but still in use in many lower-income countries – in reinforcing gender segregation.

Our findings also contribute to a growing literature concerned with policies that restrict the collection of information at early stages of the job matching process, including criminal records (Agan and Starr, 2018; Doleac and Hansen, 2020), credit histories (Bos *et al.*, 2018; Ballance *et al.*, 2020), and drug use (Wozniak, 2015). As in this literature, the key question in our paper is whether the prohibition of early-stage information affects ultimate hiring outcomes. Consistent with at least some of these studies, we find that it does. We also show that the quality of job matches does not seem to change much after early-stage information is disallowed.

Our work is closely related to studies of gender preferences in Chinese job matching markets, including Kuhn and Shen (2013), Helleseter *et al.* (2016), and Kuhn *et al.* (2020). Kuhn and Shen (2013) develop a signaling model of the decision by employers to state a gender preference that compares the cost of screening extra job applications to the expected benefits of being able to evaluate applicants of both genders. Implicit in this model is the assumption – validated by Kuhn *et al.* (2020) using employer call-backs to applicants – that SGP's are highly predictive of firm's preferences over candidates. We contribute to this literature by observing actual *hiring decisions* (rather than call-backs), by distinguishing between SGP's used by firms with higher and lower shares of female employees, and by studying wages and job durations for filled vacancies. Using hiring outcomes we also show how the elimination of stated gender preferences in job openings led to increases in gender diversity at hiring workplaces.

Most directly, our paper is related to Kuhn and Shen (2021), who study the effects of a decision by the job board in one Chinese city to eliminate SGP's from all posted vacancies. This instantaneous change allows Kuhn and Shen to compare applications to the same vacancy before and after the removal of SGP's, providing a credible design for measuring the reactions of applicants to SGP's. Consistent with our results, they find that the removal of SGP's leads to an increase in the number of applications from the non-preferred gender group. They also show that the share of call-backs to applicants from the non-preferred gender rises after the campaign – again, consistent with our findings on hiring outcomes. We view our results as highly complementary to those of Kuhn and

Shen (2021): they have information on applications and employer call-backs, whereas we have information on hiring and job outcomes, allowing us to assess changes in matching efficiency. The consistency of the findings across the two settings and between applications/call-backs versus hiring/job outcomes is highly reassuring, and suggests that the elimination of stereotypical SGP's can increase the diversity of hiring outcomes without large side effects on matching efficiency, albeit at some cost for employers that were previously trying to recruit against stereotype.

2 Background and Conceptual Framework

2.1 Background

Women and men tend to work at different workplaces (see e.g., the discussion of international evidence in Card *et al.* 2016). While some of the reasons for this gender segregation are widely understood and accepted (such as differences in training or field of study), a long-standing concern is that employers exert discretionary preferences for one gender or the other in a particular type of job or place of work, leading to diminished opportunities, particularly for women.[4] There are many steps in the hiring process at which gender can become an obstacle. Our focus, and the focus of the information campaign we study, is at the earliest stage, when employers create and publicize information on potential job openings. In our setting these vacancies appear on a job board (see below), providing information to job seekers and intermediaries that can affect subsequent decisions about which openings are pursued.

Employers can indicate a gender preference in a newly posted vacancy in several ways. Most directly, they can indicate that they prefer male or female applicants. Historically, such preferences were widespread (see Darity and Mason, 1998 for examples in the U.S. setting). The text in a vacancy notice can also provide clues as to which type of worker an employer is looking for, particularly in German, where the use of gendered occupational titles can indicate a preferred gender (e.g. *Bauarbeiterin* for female construction worker, or a *Bauarbeiter* for male construction worker).

[4] Bergmann (1974) showed how the forced segregation of women into certain occupations could lead to lower average wages for women and higher average wages for men.

In the U.S., the use of gender preferences in help-wanted notices was finally outlawed by the Supreme Court in 1973, nearly a decade after the 1964 Civil Rights Act banned gender discrimination in hiring.[5] Similar changes came later in other advanced countries. Austria adopted its first Equal Treatment Act (ETA) in 1979, and outlawed gender discrimination in job advertisements in a 1985 amendment, albeit without legal sanctions. Monetary sanctions for gender preference in ads by temporary help agencies were introduced in 1992, and a June 2004 amendment extended fines to the entire private sector. Title 1 of the 2004 amendment states that recruiters *"may not advertise a position publicly or within an enterprise (company) exclusively for men or women"*, and that the job advertisements may not contain any reference from which it could be inferred that members of one sex would be favored.[6] However, exceptions could be granted in very specific cases (e.g. female care worker in shelters for women).

In the second quarter of 2005, the *Ombud for Equal Treatment* (OET), an agency that offers advice on gender equality, analyzed 36,000 newspaper job advertisements published in all major newspapers of Austria.[7] The study found that only 68% were fully or partially gender neutral (Lujanski-Lammer, 2006). Over the next few months the OET conducted a major information campaign, contacting non-compliant firms and temporary help agencies, suggesting how to write gender neutral

[5] The case, Pittsburgh Press Co. v. Pittsburgh Commission on Human Relations, 413 U.S. 376 (1973), specifically focused on the practice of listing help wanted advertisements by gender in a newspaper. See Beller (1983) for an analysis of the effect of the 1964 Civil Rights Act in the U.S. on gender segregation across occupations. There are relatively few studies of the effect of the law on gender-related outcomes: there is a much richer literature on race-related outcomes (e.g., Brown, 1984; Hersch and Shinall, 2015).

[6] This meant that advertisements had to be made gender neutral by specifying both male and female forms of any occupation, e.g. *Bauarbeiter/Bauarbeiterin.* Title I of ETA also prohibits discrimination in promotion and pay, while Title II prohibits discrimination on the basis of ethnicity, age, religion, or sexual orientation. Some provisions required changes to state laws, which had to be in effect by December 2004. Recruiters who post an SGP face a fine of 360 Euros after one warning for private sector recruiters and with no warning for public sector recruiters. Job seekers can seek damages of one month's salary if they would have received the job but for gender, and 500 Euros in case a recruiter ignored an application because of gender.

[7] The OET had conducted several previous analyses of job advertisements between 1993 and 2005, concluding that compliance with the ETA was far from universal (Lujanski-Lammer, 2006). Gender preferences may be correlated with preferences over age, ethnicity, or sexual orientation, which were also outlawed in 2004. Age requirements are recorded in our data, but this information is very coarse, only distinguishing apprenticeships from openings for trained workers. Moreover, experience requirements were not outlawed in 2004, and could substitute for age requirements. Preferences for ethnicity and sexual orientation are not recorded in our data. The 2005 campaign of the Ombud for Equal Treatment did not address preferences for age, ethnicity, or sexual orientation, so based on the evidence reported below regarding gender preferences, we suspect that outlawing these preferences in 2004 probably had little or no impact.

job advertisements, and notifying firms about their breach of the law. The OET also informed the help wanted sections of newspapers about the requirement of gender neutrality, so newspapers could inform employers. After the information campaign, in the fourth quarter of 2005, the OET collected a new round of data on job ads in newspapers, and found that compliance with the law had increased to about 79%, substantially higher than in Spring 2005. Below we show that by mid-2006 the vast majority of postings in the Austrian Employment Service's online job board had eliminated gender preferences.

2.2 CONCEPTUAL FRAMEWORK

Next, we outline a simple conceptual framework, based on the model developed by Kuhn and Shen (2013) (hereafter, KS) that helps guide our empirical analysis. We assume that gender preferences are signals to job searchers (and those who assist job searchers) about a firm's preferences. Generalizing KS slightly, we assume that worker i applies to vacancy j if the subjective probability of being hired λ_{ij} exceeds some threshold τ_i.[8] We assume that λ_{ij} depends on observed information about the employer (such as the gender composition of its existing workforce), on the characteristics of the specific job, and on any stated gender preference. A stated preference can reinforce other prior information, or it can serve notice that the firm has a specific interest in a gender group that is a minority at its workplace, or in the occupation it is recruiting.

In the absence of any SGP the female share of applicants to a given job would be expected to vary with the fraction of female workers at a firm and the share of females in the relevant occupation, reflecting prior information available to searchers.[9] In the presence of an SGP, however, the female share of applicants will depend on the degree to which searchers believe that the firm will follow its stated preference. Although we do not see application flows, we observe the gender of hired workers, and as we show below, SGP's are very strong predictors of actual hiring outcomes, over-riding other

[8]KS assume that workers apply to all jobs with no SGP, and to jobs with an SGP that matches their own gender.

[9]For example, if $\tau_i \sim U[0,1]$ then the probability a specific worker applies is just λ_{ij}, so in a large population of searchers the female share of applicants would be $\mu_F\overline{\lambda}_{Fj}/(\mu_F\overline{\lambda}_{Fj} + (1-\mu_F)\overline{\lambda}_{Mj})$, where μ_F is the female share of all searchers and $\overline{\lambda}_{Gj}$ is the expected value of λ_{ij} for gender group G. One would expect $\overline{\lambda}_{Fj}$ to be higher for jobs at a firm with more female workers, or for more traditionally-female occupations.

information about the firm (such as the gender composition of its existing workforce) in predicting the gender of hired workers. Similarly, KS show that the specification of a gender preference in their setting almost completely eliminates applications from the other gender.

On the demand side, a key question is why a firm would ever purposefully limit its applicant pool by specifying a gender preference.[10] KS assume that applicant screening reveals (at some cost) a candidate's match value to the firm, which has a gender-specific mean and an idiosyncratic component. The match value includes a prior assessment of the likely quality of candidates of a given gender, as well as any discriminatory taste factors. If the mean match values for men and women are far apart relative to the dispersion in the idiosyncratic component, then it is optimal for the firm to limit screening to the higher-mean gender. If the gender-specific means are relatively similar, or the idiosyncratic components are more disperse, then it is optimal to screen applicants from both genders.[11]

This model suggests that if the firm's prior assessments of match quality are correct, or it has strong tastes for one gender, then the elimination of the ability to state SGP's will increase screening costs without necessarily inducing employers to hire workers from the previously excluded gender group. If, for example, an employer has a strong distaste for women and states a preference for males when SGP's are allowed, then it seems unlikely that employer will hire women when SGP's are eliminated, even if women apply for jobs at the firm. On the other hand, if some employers are using out-dated stereotypes to form their priors on match quality, they may be positively surprised by the quality of non-stereotypical candidates when SGP's are eliminated, and may end up filling a position with someone who would have been screened out prior to the change.

There is one case where the elimination of SGP's may lead to a *rise* rather than a fall in stereotypical hiring. That is the case of non-stereotypical SGP's. (For example, an engineering firm with only

[10] The same question can be asked about other qualifications specified in a vacancy, such as experience or education.

[11] Specifically KS assume that the match value of candidate i for a given opening j is a random variable $v_{ij} = v^{G(i)} + \beta\epsilon_{ij}$ where v^G is the mean for gender group G, $G(i)$ is an index function given $i's$ gender, β is a scaling factor, and ϵ_{ij} is an extreme value type 1 variate. They show that a firm's decision depends on $(v^M - v^F)/\beta$: when this term exceeds some threshold $c_U > 0$ it is efficient to screen only men, when if falls below another threshold $c_L < 0$ it is efficient to screen only women, and in the intermediate range it is efficient to screen both groups.

male employees may request female candidates). In such cases an SGP can override prior information that might otherwise prevent people of the non-stereotypical gender from applying. This type of SGP is similar to notices (widely used in the U.S.) that an employer is seeking a diverse application pool, but has more bite. Eliminating such SGP's may prevent job searchers or intermediaries from learning that an employer is specifically looking for a non-stereotypical candidate, and in the absence of this information they may not think it is worthwhile to apply.

3 EMPIRICAL ANALYSIS

3.1 DATA SOURCES

Our empirical analysis utilizes data from the job matching platform of the Austrian Employment Service, *Arbeitsmarktservice* (AMS). AMS administers Austria's income support programs (UI benefits, unemployment assistance, and related transfers) and also runs its active labor market programs. The agency has local offices in each of Austria's 104 labor market districts, at which people claiming UI have to register in person and meet regularly with staff in order to maintain benefit eligibility.

Since 1987 the AMS job-matching platform has gathered information on vacancies from firms. The platform includes almost 60% of all vacancies posted by Austrian firms, with higher coverage rates of openings in manufacturing and construction and lower rates in banking and finance (Kettemann *et al.*, 2018; Mueller *et al.*, 2019; Ziegler, 2021). Firms post job advertisements that list the characteristics of the open position and the desired qualifications of potential applicants. This information is entered in pre-configured fields, one of which is the preferred gender of applicants (which can be left blank).[12] Fortunately for our purposes, the preferred gender field remained in the system even after the 1984 law change because of the exemption granted by the law to certain job searches. The gender preference field allows us to examine how the 2005 information campaign, targeted to job ads, affected the use of gender preference statements in the AMS.

[12]Our data include the pre-configured fields, but not the actual text posted in the online job board. Thus, we cannot extract potential information on gender preferences from the text, contained for example in job titles.

AMS staff use the vacancy information in their pre-selection (*Vorauswahl*) service, which routes lists of job seekers registered with AMS to firms with openings. Firms can select interviewees from among those suggested by AMS staff; they can also consider direct applications from individual workers. Vacancy information is also used by caseworkers, who can search for openings that fit the profile of a client and suggest them to the client for follow-up. If an AMS client is selected for a job opening the identity of that person is recorded in the system - a feature we use below.

The AMS job board differs from some other systems (such as the one studied by Kuhn and Shen, 2013) in the direct role played by AMS staff, though many commercial posting services use algorithms to suggest list of candidates to employers, similar to the *Vorauswahl* service. Nevertheless, as we show below, the concordance rate between SGP's in the pre-campaign period in Austria and the gender of the hired worker was similar to the concordance rate between SGP's and job applicant genders in the job board studied by Kuhn and Shen (2013). This suggests that intermediation by AMS staff and caseworkers did not necessarily diminish or magnify the role of SGP's in the matching process relative to a process driven by worker-level choices.

3.2 Characteristics of Vacancies

We obtained AMS vacancy data for the period September 1997 to December 2013, providing multiple years of data before and after the 2004 and 2005 events. Our dataset includes the occupation sought by the prospective employer, education requirements for the job, whether the job is full- or part-time, and whether the contract is fixed-term or open-ended. It also includes the gender preference selected by the employer and another indicator for whether the position is targeted to people aged 20+ (who have normally completed training). Finally, it includes an initial posting date and the date the vacancy was filled or closed (if the vacancy was withdrawn without being filled). For vacancies that were filled by AMS clients there is also an identifier of the recruited employee. We refer to such vacancies as "AMS hires".

We link vacancies with an AMS hire to the ASSD using the (anonymized) identifier of the hired employee. ASSD covers all private sector employees in Austria – about 80% of the labor force – and provides information on employment spells, including days worked at each establishment in a year and total earnings, which we convert to an average daily wage[13] (see Zweimüller *et al.*, 2009 for more discussion of the ASSD).

In the period from September 1997 to December 2013 there were 5.2 million vacancies on the AMS system. Column 1 of Table 1 shows some characteristics of these vacancies, including average duration in the system (around 60 days), the share advertising for full time positions (77%) and for positions with an unlimited contract (79%), the share requiring at least upper secondary schooling (52%) and the share posted by firms with less than 4 employees (44%). Among all posted vacancies in the period up to July 2004, 27% stated a preference for males and 22% stated a preference for females (see panel B).

Unfortunately, the employer identifiers used in the AMS vacancy system cannot be linked to the ASSD. Our approach is to focus on vacancies filled by an AMS client (i.e., the AMS hires) for which we observe a person identifier that can be linked to the ASSD. For each AMS hire we search the ASSD for all employment spells involving that worker, and link the vacancy to the spell with the starting date closest to the closing date of the vacancy. We provide a full description of the matching procedure in the Appendix A.

A concern for our analysis is the representativeness of vacancies filled by AMS hires. Column 2 of Table 1 shows the characteristics of filled vacancies, which as shown in the bottom row of the table represent about 87% of posted vacancies. Relative to the characteristics in column 1 there are few major differences, suggesting that *filled* vacancies are similar to the broader population of filled or withdrawn vacancies. Next, in column 3, we present the characteristics of the 1.7 million filled vacancies with an AMS hire. This subset, which represent 26.6% of all filled vacancies, have a slightly lower mean duration (44 days vs. 58 days), are slightly more likely to advertise for a position with an unlimited contract, and are slightly less likely to require high-school level education (46% vs. 52%). In the pre-campaign period AMS-filled vacancies were also somewhat more likely to state a

[13]The ASSD database does not contain information on working hours, though we know the part-time status of a vacancy.

Table 1: Vacancy Characteristics

	Subsample			
	All Outflows (1)	**Hired** (2)	**Hired AMS** (3)	**Matched** (4)
Panel A: Vacancy Characteristics				
Avg Vacancy Duration (days)	52.72	51.86	40.97	40.91
	(30.23)	(29.89)	(28.45)	(28.37)
Full Time Position (share)	0.77	0.77	0.78	0.78
	(0.42)	(0.42)	(0.41)	(0.41)
Unlimited Contract (share)	0.79	0.78	0.81	0.82
	(0.41)	(0.41)	(0.39)	(0.39)
Small Firms (share)	0.44	0.44	0.44	0.43
	(0.50)	(0.50)	(0.50)	(0.50)
Upper Secondary Education (share)	0.52	0.52	0.46	0.47
	(0.50)	(0.50)	(0.50)	(0.50)
Panel B: SGP - pre-campaign period				
Preference for Men (share)	0.26	0.26	0.30	0.30
	(0.44)	(0.44)	(0.46)	(0.46)
Preference for Women (share)	0.22	0.22	0.23	0.23
	(0.41)	(0.42)	(0.42)	(0.42)
Observations	4,998,146	4,333,444	1,151,996	987,271
Observations - pre-campaign period	1,622,911	1,419,630	462,450	397,496

Notes: This table reports means and shares of selected characteristics of vacancies in four sub-samples. Period 1997-2013. Pre-campaign period refers to the period 1997-2004. The sub-sample "all" refers to all vacancies, "Hired" is for vacancies that are filled, "Hired AMS" is the sub-set of vacancies filled through AMS, and "Matched" is the sub-set of the AMS hires (in column 3) that we matched with a firm in the Austrian Social Security Database (ASSD). Vacancy Duration refers to "*Vacancy Closing Month (last day) - Vacancy Posting Date*" as the employment spell starting date is not available for non matched vacancies.

Source: AMS-ASSD data, own calculations.

preference for a male (31% versus 26% of filled vacancies), whereas the fraction stating a preference for a female is about the same as in the overall population of vacancies (23% versus 22%).

More information is presented in Appendix Tables A.2 and A.3, where we show 1-digit industry and occupation shares for the same sets of vacancies summarized in Table 1. Jobs in mining, manufacturing, construction, public administration and health are somewhat over-represented in the set of vacancies with an AMS hire, while jobs in hotels and restaurants, finance, real estate and business services, and health care are under-represented. At the occupation level, positions for clerks, operatives, and unskilled (elementary) occupations are over-represented among vacancies filled by AMS hires, while those for managers, professional and technical workers, and shop and sales workers are under-represented.

Finally, column 4 presents the characteristics of the subset of vacancies with an AMS hire that we successfully matched to an employer in the ASSD data. We lose about 14% of potential matches because of missing or incomplete data, or because there is no new employment spell in the ASSD that is an obvious match for the filled vacancy in the AMS data. Reassuringly, the characteristics in columns 3 and 4 are quite similar (as are the corresponding industry and occupation shares shown in Appendix tables A.2 and A.3), suggesting that the matching process is approximately random.

While AMS-filled vacancies do not have exactly the same characteristics as the more general population of filled vacancies, we believe the differences are modest. In the remainder of the paper we therefore use the sample described in column 4 of Table 1 as our main analysis sample. As a robustness check we use the characteristics of vacancies in column 1 to construct weights that represent the inverse probability that a vacancy was filled by an AMS client, and re-estimate our main models using these weights (see below). Reassuringly we find relatively small discrepancies from our unweighted models.

3.3 Use of Stated Gender Preferences Prior to 2005

We begin our analysis by looking at what happened to the fraction of vacancies with a stated preference after the introduction of fines for noncompliance with the Equal Treatment Act in June 2004, and the information campaign in early 2005, Figure 1 (Recall that the AMS system allowed employers to select a preferred gender even after they were declared illegal to accommodate exemptions from the law).

Figure 1: SGP in time

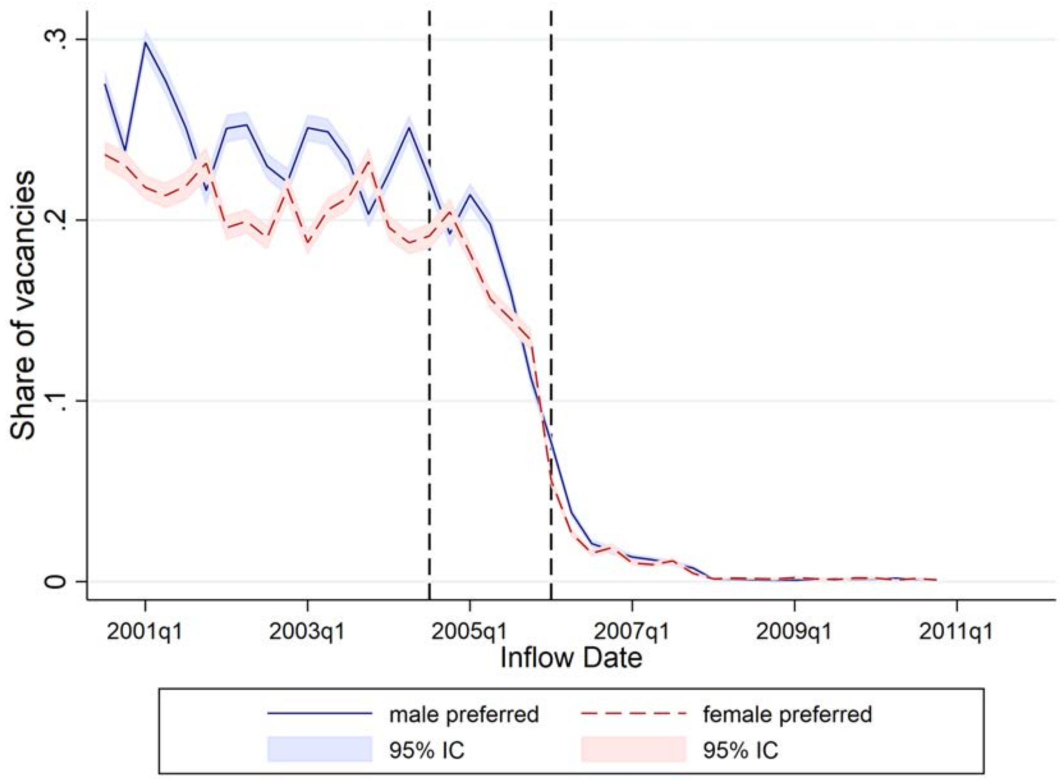

Notes: This figure shows that share of posted vacancies in the AMS job board system that specify a preference for females (dashed line), or males (solid line). The vertical dashed lines depict the date of a law change introducing sanctions for posting gender preferences (July 2004), and the timing of an information campaign to alert employers and newspapers about the law (Spring 2005). Since the law allows some exemptions, the system continued to allow employers to post preferences throughout the sample period.
Source: AMS-ASSD data, own calculations.

In the years prior to 2004 the fractions of AMS-filled vacancies with a stated preference were fairly stable, apart from an obvious seasonal pattern. At the July 1 2004 effective date of the amendment introducing fines there is no obvious trend break in either rate, suggesting that the introduction of sanctions had little direct effect. In the second quarter 2005, however, the shares of vacancies with male or female preferences both begin to decline sharply, and by mid-2006 only a small share of vacancies express a gender preference. The trend break in 2005 coincides exactly with the information campaign conducted by the OET from the second to the fourth quarter of 2005. The coincidence of timing suggests that the OET's campaign effort was responsible for the near elimination of SGP's, though this effect has to be understood as coming after sanctions were introduced into the Equal Treatment Act, and in light of evolving attitudes toward gender.

Given the patterns in Figure 1, we consider the period up to and including 2004 as the "pre" period for our event studies and difference-in-differences models, and the period from 2006 onward as the "post" period. For our main models we drop 2005 (the year of transition), although in our event study graphs we retain these observations.

How did AMS-filled vacancies stating a male or female gender preference differ from those with no SGP? We investigate these differences in Table 2, which uses data from the 2000-2004 (pre-campaign) period. About 21% of vacancies stated a gender preference for females (column 1), 26% for males (column 3), and 53% stated no preference (column 2). During these years females made up about 46% of all people hired for AMS-filled vacancies (see column 4 of the table). The share of women hired for jobs with no SGP was quite similar to their overall share, but for vacancies with a female SGP, 96.8% were filled by a female, while for vacancies with a male SGP, only 2.6% were filled by a female. Thus, compliance with SGP's was extremely high.

We also report the mean and median duration of filled vacancies, the mean log daily wage associated with the newly created jobs, and the duration of these jobs (mean and median). We can see a couple of interesting patterns that will carry through our entire analysis. First, job openings with an SGP *for either gender* tend to be filled a little faster.[14] Second, jobs created by filling a vacancy with a female SGP – which are nearly all held by women – pay relatively low wages but have longer durations than than those created by filling vacancies with no SGP. (Some of this pay gap could be due to differences in hours of work). Symmetrically, jobs created from a vacancy with a male SGP – which are nearly all held by men – pay relatively high wages and have a shorter duration.

[14]Vacancy filling times reported in Table2 are based on the elapsed time from the day the vacancy was posted and the day the new employee started work. The times reported in Table 1, which pertain to all filled vacancies including those we cannot match to the ASSD, are based on the number of days between the posting date of the vacancy and the last day of the month in which it was removed from the AMS job board.

Table 2: Descriptive Statistics on New Jobs Associated with Filled Vacancies

	SGP			
	Female SGP (1)	**No SGP** (2)	**Male SGP** (3)	**Total** (4)
Share of vacancies (share)	0.213	0.529	0.259	
Panel A: Outcomes				
Share of women hired (share)	0.967	0.462	0.028	0.457
Vacancy filling time (mean)	30.70	34.53	30.65	32.71
Vacancy filling time (median)	20	23	19	21
Log wage of the hire (mean)	3.58	3.83	3.98	3.81
Job duration (mean)	425	397	338	388
Job duration (median)	195	177	144	171
Panel B: Context				
Share of women in the firm (share)	0.691	0.480	0.247	0.464
Share of women in occupation (share)	0.659	0.496	0.271	0.473

Notes: This table reports means, medians and shares of selected characteristics of vacancies in the matched sample by Stated Gender Preference (SGP) during the 2000-2004 (pre-campaign) period. Vacancy filling time refers to the difference in days between the starting date of the new job associated with the vacancy and the posting date of the vacancy.

Source: AMS-ASSD data, own calculations.

The key question motivating our work is the extent to which SGP's tend to perpetuate gender segregation across workplaces. Figure 2 shows histograms of the share of women in each establishment in our sample in 1999 (5 years before the campaign) and 2007 (3 years after). In both years the distribution of the share of women is inverse-U-shaped, with a relatively large share having <10% female workers and another relatively large group having 100% females.[15] Comparing histograms for the two years we can see clear evidence that the degree of gender segregation across firms in Austria fell somewhat between 1999 and 2007. The average share of women also increased over this period, however, making it hard to credit the change solely to the elimination of SGP's.

Figure 2: Gender Mix Within Firms

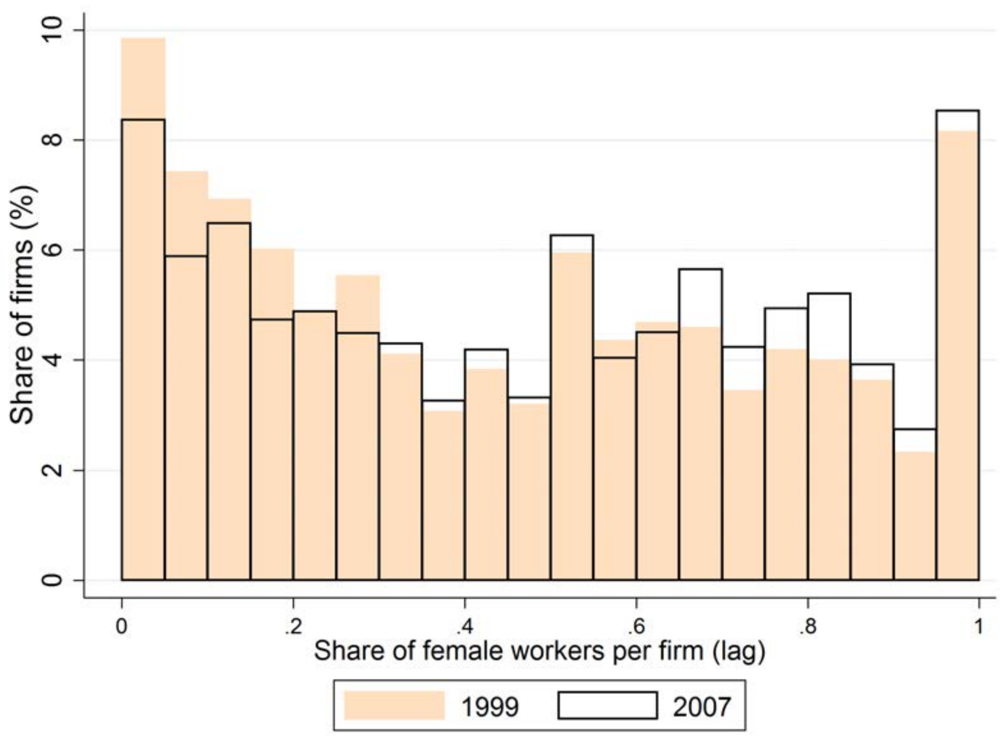

Notes: This figure displays the histogram of the share of women in firms for 1999, and 2007. Histograms are estimated using .05 bin size.
Source: ASSD

How did the use of stated gender preferences in the years prior to the campaign vary with the female share of employees? To provide some initial descriptive evidence, we construct a measure of

[15]The "spikes" in the distributions in Figure 2 are due in part to firms of size 1 or 2. See D'Haultfœuille and Rathelot (2017) for a discussion of small unit bias in measuring characteristics of firms.

the fraction of female employees at the firm that posted the vacancy in the year prior to the posting date. We then use data on vacancies in the pre-campaign period to estimate the average fraction of vacancies with either no SGP, or a male or female SGP, among workplaces in different ranges of the lagged female share. The results, presented in Figure 3, show the expected pattern. Workplaces with few women (toward the left of the figure) tended to post vacancies with either a *male* SGP (about 60% of the time) or *no* SGP (around 40% of the time). Symmetrically, workplaces with few men (toward the right of the figure) tended to either post vacancies with a *female* SGP (about 60% of the time) or *no* SGP (around 40% of the time). Thus, when they were allowed, SGP's tended to be used to reinforce existing gender segregation patterns. There was also some limited use of SGP's to recruit women at firms with <50% female share, and to recruit men at firms with <50% males. In our analysis below we refer to such preferences as non-stereotypical SGP's.

Figure 3: Share of Job Ads by Gender Preference

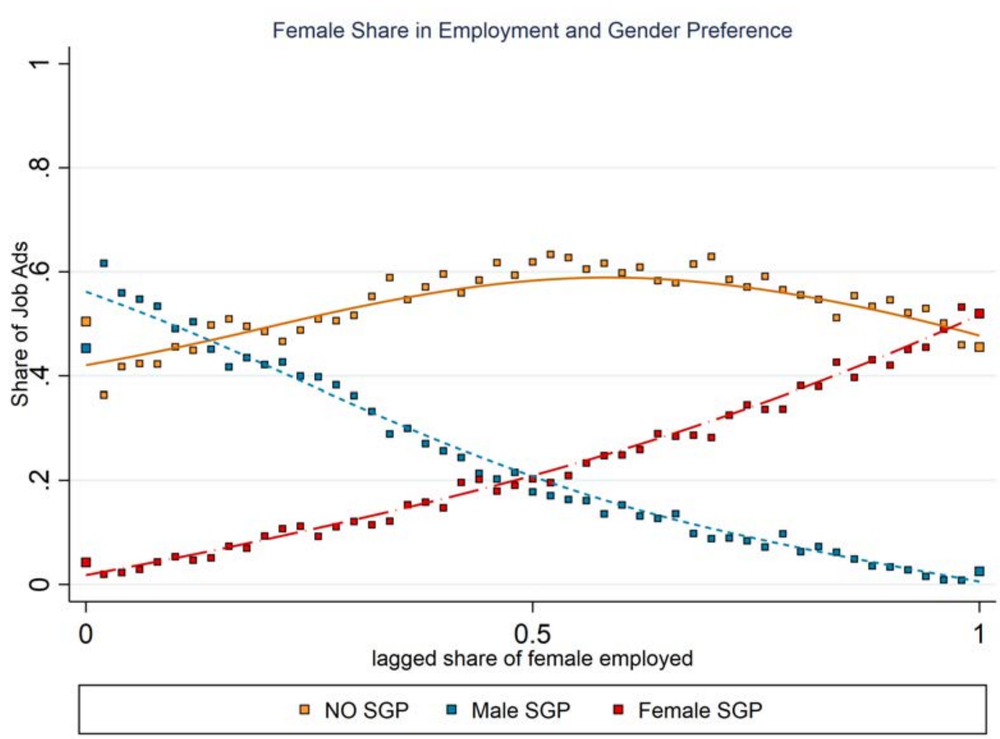

Notes: This Figure shows the share of vacancies that specify preferences for females (red line), males (blue line), or no gender preference (yellow line), by gender composition of the firm in the year prior to the job advertisement. Data are drawn from 2000-2004.

Source: AMS-ASSD data, own calculations.

Next we look at the relationship of SGP's to ***hiring outcomes***: Figure 4 shows the average share of women actually hired by firms with different (lagged) female employment shares, conditional on the type of SGP (male, female, or none). The figure shows stark differences in hiring outcomes for these three types of vacancies, even controlling for workplace composition. In the absence of any SGP (the yellow line in the figure), the fraction of females hired is approximately linear in the lagged female share, with an intercept that is just slightly above 0 and slope that is just slightly flatter than 1. For vacancies with a **male** SGP (blue line), however, the fraction of females hired rises only slightly with the workplace share of females, reaching about 20% even at firms that had all female workers in the previous year. In contrast, for vacancies with a **female** SGP (red line), the fraction of females hired starts at around 90% at workplaces with no women in the previous year, and rises slightly to essentially 100% at all-female workplaces.

Figure 4: Share of Female Hired by Gender Preference

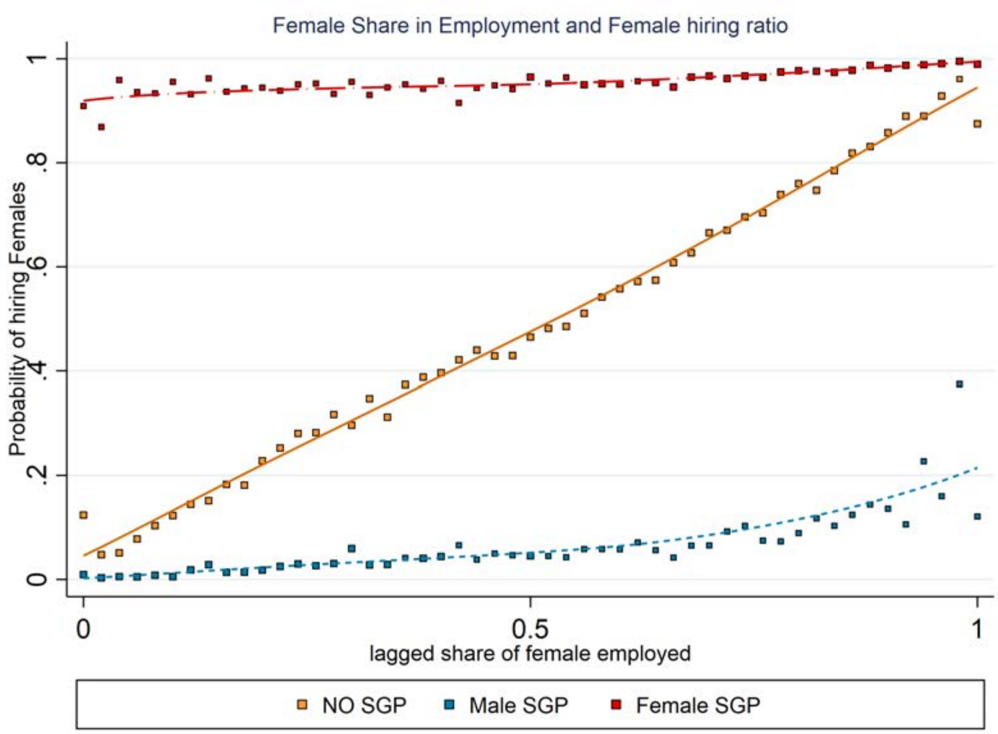

Note: This Figure shows shows the share of females hired by gender composition of the firm in the year prior to the hire, and by whether the vacancy stated a preference for females (red line), males (blue line), or neither. The data are from 2000-2004.

Source: AMS-ASSD data, own calculations.

The remarkable “Z” shape of Figure 4 suggests that, when SGP’s were allowed, they exerted a strong impact on hiring outcomes. Workers hired to fill vacancies with no SGP tended to perpetuate the status quo gender composition, with only a small probability of hiring a female at all-male firms (around 5%), or of hiring a male at all-female firms (also around 5%). But nearly all the new hires for vacancies with a female SGP were female (even at all-male workplaces) and nearly all the new hires for vacancies with a male SGP were male (even at all-female workplaces).

A potentially confounding factor in the interpretation of this Z pattern is the role of occupation. Many occupations are highly specific to one gender (Cortes and Pan, 2017), so the female share of a firm’s workforce is partly driven by occupation structure. We would expect new hires to obey the same structure, leading to a strong positive relationship between the lagged share of females and the average share of females in new hires. If employers tend to use a male SGP when they are recruiting for mainly male occupations and a female SGP for mainly female occupations, however, we would see a close match between SGP’s and the gender of hired workers, even if SGP’s have no effect on application or hiring decisions.

To isolate the effects of SGP’s while controlling for occupation, we used data from the EU Standard of Income and Living Conditions (SILC) data base, **excluding Austria**, to assign the share of females in each of 100 occupation groups (International Standard Classification of Occupations, ISCO, 2-digit major sub-groups). We then linked these gender shares to the vacancies in our data set, and divided vacancies in the pre-campaign period into 3 groups: those for mainly male occupations (with less than 30% female workers on average); those for mainly female occupations (with more than 70% female on average); and those for mixed-gender (or “gender neutral”) occupations. Finally, we constructed separate versions of Figure 4 for each of the three groups of occupations.

As shown in Figure 5, all three types of vacancies exhibit a Z pattern, though the precise shape of the Z – particularly the diagonal element representing vacancies with no SGP – varies by occupation group. In the absence of an SGP, the average fraction of women hired to fill *female* occupations starts at around 25% at all-male workplaces and rises to 100% at all-female workplaces. Symmetrically, the average fraction hired to fill *male* occupations starts at 0 at all-male workplaces and rises to around

Figure 5: Role of Female Share in Occupation

(a) Female Occupations

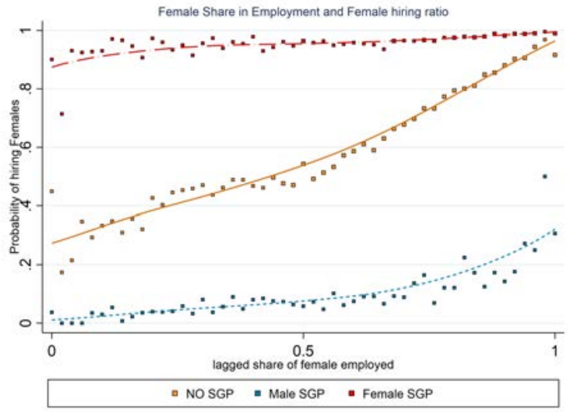

(b) Gender Neutral Occupations

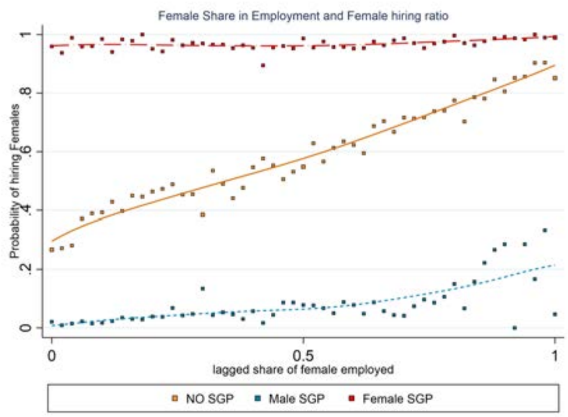

(c) Male Occupations

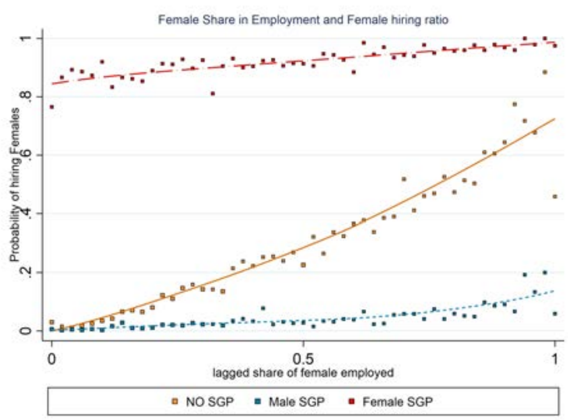

Note: This Figure shows shows the share of females hired by gender composition of the firm in the year prior to the hire, and by whether the vacancy stated a preference for females (red line), males (blue line), or neither (yellow line), for three groups of vacancies: those advertising for mainly female occupations (panel a), mainly male occupations (panel c) or occupations with a mixed gender composition (panel b), The data are from 2000-2004.

Source: AMS-ASSD data, own calculations.

75% at all-female workplaces. The average fraction hired to fill *gender neutral* occupations lies between these extremes, starting at around 25% and increasing to around 85%. Thus, for vacancies with no SGP, controlling for occupation flattens the relationship between the lagged female share and the probability of hiring a female. For vacancies with male or female SGP's, however, most of the new hires match the stated gender preference, regardless of whether we control for the gender mix of the occupation or not, and regardless of the gender mix of the workplace.

Stated gender preferences appear to influence hiring outcomes, but is there any evidence that they affect hiring efficiency (as is assumed in the KS model)? To assess this we look at vacancy filling times. Columns 1-4 of Table 3 present a series of simple models that relate the days required to fill a vacancy to the use of SGP's. To help provide a context for these models we also present a parallel set of models for the event that a female was hired to fill the vacancy (columns 5-8). We look at all vacancies (columns 1 and 5), and separately at vacancies for mainly male occupations (columns 2 and 6), for occupations with a mix of male and female workers (columns 3 and 7), and for mainly female occupations (columns 4 and 8).

We classify workplaces as "female" (more than 50% females in the previous year, with label "F workplace") or "male" (more than 50% males in the previous year, , with label "M workplace") and include as the main variables of interest indicators for a male or female SGP, interacted with dummies for mainly male or mainly female workplaces. All the models include fixed effects for occupation and industry, time effects, and dummies for 51 different intervals of the lagged share of females at the workplace.

Looking at the vacancy filling time results we see two broad patterns. First, the use of SGP's tends to (if anything) reduce vacancy filling times. Second, SGP's for the gender that is the opposite of the workplace majority tend to have the largest negative effects. The large negative effects of a female SGP in filling an opening for mainly male occupation at a mainly male workplace (-2.406 days, in the 4th row of column 2) and of a male SGP in recruiting for a mainly female occupation at a mainly female workplace (-6.517 days, in the first row of column 4) are particularly interesting. In both cases, the employer is stating a preference for a non-stereotypical gender (different than the majority of the occupation and of the existing workforce), and yet the time to fill the vacancy is reduced. As shown in the corresponding models for the gender of the newly hired worker (columns 6 and 8), such preferences have large effects on the probability of hiring the preferred gender, consistent with the patterns in Figure 5. Thus, it appears that employers with non-stereotypical preferences could easily find workers to match those preferences. Taken together, the results on filling times and gender outcomes suggest that in the pre-campaign period there were more females looking for jobs in mainly male occupations, and more males looking for jobs in mainly female occupations, than were demanded by the relatively small share of employers who were seeking to recruit them.

Table 3: Vacancy Filling Time and Female Hiring

	OLS Estimation							
Dependent Variable:	*Vacancy Filling Time - in days*				*Female hire*			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
M SGP in F workplace	-2.140***	-2.960***	2.038	-5.684***	-0.378***	-0.306***	-0.348***	-0.443***
[Non-stereotypical Male PGP]	(0.446)	(0.836)	(1.504)	(0.752)	(0.005)	(0.010)	(0.017)	(0.009)
M SGP in M workplace	-1.212***	-0.394	-0.739	-4.637***	-0.139***	-0.081***	-0.231***	-0.277***
[Stereotypical Male PGP]	(0.289)	(0.375)	(0.831)	(1.019)	(0.002)	(0.003)	(0.009)	(0.011)
F SGP in F workplace	-3.042***	-2.662***	-2.099**	-3.148***	0.240***	0.447***	0.182***	0.227***
[Stereotypical Female PGP]	(0.296)	(1.022)	(1.053)	(0.356)	(0.003)	(0.012)	(0.011)	(0.004)
F SGP in M workplace	-3.274***	-4.440***	-4.572***	-2.196***	0.469***	0.673***	0.335***	0.423***
[Non-stereotypical Female PGP]	(0.405)	(0.807)	(0.931)	(0.770)	(0.004)	(0.009)	(0.011)	(0.008)
All Vacancies	✓				✓			
Male Occupations		✓				✓		
Neutral Occupations			✓				✓	
Female Occupations				✓				✓
Observations	255,969	102,630	48,046	105,293	255,969	102,630	48,046	105,293

Notes: OLS estimation of the effect of the SGP on vacancy filling time and hiring gender. Period 2000-2004 (pre-campaign). Controls include occupation, industry and firm gender composition fixed effects as well as year FE. Beta coefficients reported and robust standard errors in parentheses. Columns (1) and (5) concerns all observations. In column (2) and (6) only job openings in female occupations are used, in column (3) and (7) in neutral occupations and in column (4) and (8) in male occupations only. *** significant at 1 percent, ** significant at 5 percent, * significant at 10 percent.

To summarize: in the years prior to the 2005 campaign, nearly one-half of vacancies were posted with a stated gender preference. Most employers were stating stereotypical gender preferences that tended to reinforce gender segregation, though a minority were using non-stereotypical preferences to recruit against type. Moreover, firms that stated a gender preference were very likely to recruit a worker of the preferred gender, and were able to fill the vacancy somewhat faster, particularly in cases where the stated preference was for a gender different from the majority of the occupation and the firm's existing workforce. What we cannot tell from the observational data is the extent to which firms that advertised for a particular gender were willing to hire people of the opposite gender, and if so whether the resulting matches would be much different from the ones initiated prior to the campaign. These are the questions to which we now turn.

4 Effects of the Campaign on Stated Gender Preferences

In the post-campaign years we cannot see the gender preferences that employers would have stated in the absence of the campaign.[16] Instead, we develop a prediction model that identifies vacancies that would have been likely to have a male SGP, a female SGP, or no SGP in the pre-campaign period. We refer to these as "predicted gender preferences", or PGP's. Then we use this model to classify vacancies *before and after the campaign* into the three groups, and examine changes in outcomes for the three sets of vacancies after the campaign.[17]

4.1 Classifying Vacancies by Predicted Gender Preference

We use data on AMS-filled vacancies in the 2000-2004 period to estimate prediction models for the use of gender preferences. Specifically, we use information on the industry I_j (in 86 categories) and occupation O_j (in 100 categories) associated with the j^{th}vacancy, together with data on the lagged female share at the posting establishment F_j (in 51 intervals). In our first approach, we define a variable S_j for each vacancy equal to 1 if it has a female SGP, equal to -1 if it has a male SGP, and equal to 0 otherwise. (Note that $S_j = S_j^f - S_j^m$, where S_j^f is a dummy for a stated female preference and S_j^m is a dummy for a stated male preference). We then calculate the leave-out mean (LOM)

[16] Kuhn and Shen (2021) are able to compare applications to vacancies that were initially posted before gender preferences were removed from the job board they study, and remained posted after their removal. This design holds constant all features of the vacancy except the SGP's.

[17] This approach is related to designs that classify subgroups based on the likelihood of being affected by a treatment, as in the "gap design" for studying minimum wages (e.g., Card, 1992). Botosaru and Gutierrez (2018) discuss identification and estimation in related studies using repeated cross sectional data in which treatment status is only observed after treatment.

value $\overline{S}_{\sim j}$ for vacancies in each (I, O, F) cell and assign this back to vacancies in the pre-campaign and post-campaign periods.[18] Finally, we classify a vacancy as likely to have a female SGP (indicated by a dummy D_j^f) if $\overline{S}_{\sim j}$ is above some threshold:

$$D_j^f = 1[\overline{S}_{\sim j} > c_1]$$

and likely to have a male SGP (identified by D_j^m) if $\overline{S}_{\sim j}$ is below some other threshold:

$$D_j^m = 1[\overline{S}_{\sim j} < c_2].$$

Vacancies with neither condition true are classified as likely to have no SGP. We select the cutoffs c_1 and c_2 so that the predicted shares of vacancies with a female or male SGP match the actual shares with these preferences in the pre-campaign period. We refer to the dummies D_j^f and D_j^m as ***predicted gender preferences*** (PGP's). Note that we cannot assign leave out means to (I, O, F) cells with only 1 vacancy. Such vacancies are omitted from the sample used in all subsequent analysis.

This classification procedure is equivalent to calculating the leave-out mean fractions of vacancies in each (I, O, F) cell with female and male SGP's, $\overline{S}^f_{\sim j}$ and $\overline{S}^m_{\sim j}$, respectively, then classifying a vacancy based on whether $\overline{S}^f_{\sim j} - \overline{S}^m_{\sim j}$ is above some threshold (for a female SGP) or below some other threshold (for a male SGP). Thus, holding constant $\overline{S}^f_{\sim j}$, vacancies in a cell with a higher value of $\overline{S}^m_{\sim j}$ are less likely to be classified as a female PGP. Likewise, holding constant $\overline{S}^m_{\sim j}$, vacancies in a cell with a higher value of $\overline{S}^f_{\sim j}$ are less likely to be classified as a male PGP. We believe this ordered structure makes sense: an employer may be close to indifferent between posting a job with a particular gender preference or no gender preference, but would hardly ever be indifferent between posting a job with a male SGP or a female SGP.

In our second approach, we estimate a Random Forest (RF) model using 4-fold cross validation to classify SGP's using the same categorical I, O, F data for the 2000 to 2004 period.[19] We construct

[18]Since the post-campaign vacancies are never used in calculating the SGP shares, the LOM assigned to these vacancies includes all the data.

[19]Note that this approach can in principal calculate predicted preferences even for (I, O, F) cells with only 1 vacancy.

100 trees and take the majority vote for the classification of each vacancy as our final classification.

We show the prediction success rates for our two different approaches in Appendix Figure C.1. We define the success rate in a given year as the fraction of vacancies in that year with a given SGP type that were predicted to have that SGP type (i.e., $E[D_j^f|S_j^f = 1]$ and $E[D_j^m|S_j^m = 1]$). The data for 1999-2004 are "in-sample" predictions, though by using a leave-out mean (in our first approach) or k-fold cross validation (in our second) we reduce the risk of in-sample over-fitting. The data for 2005 and 2006 are "out-of-sample" predictions, since data from these years are not used in our models. We stop after 2006 because by 2007 most employers were no longer using SGP's. For comparative purposes we also show predictions that use the entire sample in either approach in the two right hand panels of the figure.

The upper left panel of the figure shows that the fraction of SGP's correctly predicted by our first approach in the in-sample period is just over 60% for vacancies with male SGP's and no gender preference, and around 55% for those with female SGP's. All three rates evolve fairly smoothly between the in-sample and out-of-sample period. The lower left panel shows that the RF classifier has somewhat higher success rates for male SGP's (around 65%) and no SGP (close to 70%), but a lower success rate for female SGP's (around 41%). In the out-of-sample period the RF's prediction success rate also falls more abruptly, suggesting some tendency for over-fitting despite the use of cross validation in building the classifier.

The patterns in two right hand panels illustrate the importance of not using data from a given vacancy in forming a prediction of its SGP. For both our classifiers, using the own-vacancy data leads to higher prediction success in the in-sample period but a noticeable degradation in performance in the out-of-sample period, indicative of over-fitting bias.

4.2 Measuring Effects on Hiring Outcomes

We evaluate the effects of the OET campaign on the use of stated gender preferences using a simple difference in differences approach that compares post-campaign outcomes to pre-campaign out-

However, for comparison with the LOM approach we drop these from the sample.

comes for three groups of vacancies: those with a predicted female preference ($D_j^f = 1$), those with a predicted male preference ($D_j^m = 1$) and those with no predicted preference ($D_j^f = D_j^m = 0$). Let y_j represent some outcome associated with the vacancy (e.g., the gender of the hired candidate or the wage paid to the new worker). Then we fit models of the form:

$$y_j = \beta_0 + \beta_1 D_j^f + \beta_2 D_j^m + \lambda_1 D_j^f Post_j + \lambda_2 D_j^m Post_j + X_j \delta + \varepsilon_j \tag{1}$$

where $Post_j$ is an indicator for vacancy j being listed in 2006 or later, and X_j is a set of control variables, including time effects (which absorb the main effect of $Post_j$), and dummies for industry, occupation and the female share of employees at the posting employer in the previous year. Recognizing that D_j^f and D_j^m are based on a first-stage prediction model, we bootstrap the standard errors for the estimated coefficients from this model.

The coefficients λ_1 and λ_2 in equation (1) measure the changes in the outcome y between the pre-campaign and post-campaign period for vacancies that were predicted to have a female or male SGP, respectively, *relative to the trend for vacancies that were predicted to have no expressed gender preference*. Since the law did not directly effect recruiters that would have posted vacancies with no gender preferences, the outcomes for their vacancies are a natural comparison group. We note, however, that general equilibrium effects could potentially change the types of workers who apply for these vacancies.

Since we are using predicted rather than actual gender preferences, equation (1) can be interpreted as the reduced-form from a two-stage-least-squares procedure in which we estimate the first stage models using only data from the pre-campaign period, and allow the effects of the endogenous variables in the outcome equation to vary between the pre-campaign and post-campaign periods. To formalize this reasoning, assume that the true model generating outcome y is:

$$y_j = \alpha_0 + \alpha_1 S_j^f + \alpha_2 S_j^m + \theta_1 S_j^f Post_j + \theta_2 S_j^m Post_j + X_j \gamma + \varepsilon_j \tag{2}$$

where S_j^f and S_j^m indicate actual SGP's in the pre-campaign period and *desired* SGP's in the post-campaign period (i.e., the preferences that employers would have stated if there was no campaign). In this model, θ_1 and θ_2 measure the changes in outcome y between vacancies in the pre-campaign period that were posted with given SGP and vacancies in the post-campaign period that would have been posted with that gender preference if such preferences were allowed. Notice that these effects are conceptually the same as the effects measured in the design of Kuhn and Shen (2021), which compares applications for the same vacancy after SGP's are eliminated.

Assume that the actual or desired SGP's are related to PGP's by a pair of simple models with constant coefficients between the pre- and post-campaign periods:

$$S_j^f = \pi_0 + \pi_1 D_j^f + \pi_2 D_j^m + X_j \pi_x + \xi_j^f \tag{3}$$

$$S_j^m = \psi_0 + \psi_1 D_j^f + \psi_2 D_j^m + X_j \psi_x + \xi_j^m \tag{4}$$

where ξ_j^f, ξ_j^m are prediction errors. Here π_1 represents the increment in the probability of an actual female SGP if the vacancy has a predicted female SGP relative to the omitted category of no predicted SGP, and π_2 represents the increment in the probability of an actual female SGP if the vacancy has a predicted male SGP, again relative to the case where it is predicted to have no SGP. Thus we expect π_1 to be positive and π_2 to be negative (though small in magnitude). Similar reasoning suggests that ψ_2 will be positive and ψ_1 will be negative.

Combining equation (2) with (3) and (4) shows that the difference-of-differences coefficients in (1) are:

$$\lambda_1 = \theta_1 \pi_1 + \theta_2 \psi_1 \tag{5}$$

$$\lambda_2 = \theta_1 \pi_2 + \theta_2 \psi_2. \tag{6}$$

Notice that if we ignore π_2 and ψ_1, then λ_1 is an attenuated version of θ_1 and λ_2 is an attenuated version of ψ_2, where the attenuation factors reflect the fractions of predicted vacancies with female and male preferences that actually have these preferences (conditional on the $X's$). More generally we expect π_2 and ψ_1 to be small negative numbers, so this intuition remains roughly correct.

Columns 1 and 3 of Table 4 present estimates of equations (3) and (4) using the observed SGP's in the 1999-2004 period and predictions from our first (leave-out-mean based) classification model. We see that $\pi_1 = 0.173$ and $\pi_2 = -0.035$, while $\psi_1 = 0.192$ and $\psi_2 = -0.027$. Thus, controlling for industry, occupation, and the firm's lagged gender composition, having a PGP of one gender raises the probability of an SGP of that gender by 17-19 percentage points relative to a vacancy with neither PGP, while having a PGP of the opposite gender lowers the probability by about 3 percentage points relative to the comparison group of neither PGP. (Columns 2 and 4 present estimates for vacancies by workplace context. We come back to these models when we discuss results that compare stereotypical and non-stereotypical preferences in Section 4.2.2.)

These estimates imply that the coefficients λ_1 and λ_2 are substantially attenuated – by a factor of roughly 80% – relative to the effects we would estimate if we could see the desired stated preferences of job recruiters in the post-campaign period. In other words, the true effects of eliminating stated gender preferences may be 5 times larger than the estimates we obtain from equation (1).

Before proceeding we note two other methodological points. First, we are assuming that the intervention of interest is the elimination of firms' abilities to declare their gender preferences in job advertisements. One concern with this assumption, particularly for our analysis of wage outcomes, is that the 2004 law change or the OET campaign may have also affected firms' pay decisions. While the OET's campaign did not directly address pay disparities, it is possible that it heightened awareness of gender-related pay gaps. Any differential changes in pay policies between firms that were more or less likely to use gender preference will be confounded in our difference in differences analysis.

Second, the framework of equations 2-4 makes clear that we are identifying the effect of stated gender preferences using interactions between occupation, industry, and the gender composition of the firm's workforce as instruments (since main effects for these 3 variables are included in all our models). Given the stark patterns in Figures 3 and 4 we believe this is a plausible strategy. As

Table 4: Relationship between Predicted and Actual Gender Preferences in Pre-campaign Period

	OLS Estimation			
Dependent Variable:	*Male SGP*		*Female SGP*	
	(1)	(2)	(3)	(4)
No PGP	omitted	omitted	omitted	omitted
Male PGP	0.192***		-0.035***	
	(0.004)		(0.001)	
Female PGP	-0.027***		0.173***	
	(0.002)		(0.004)	
M PGP in F workplace		0.266***		-0.077***
[Non-stereotypical Male PGP]		(0.009)		(0.005)
M PGP in M workplace		0.181***		-0.029***
[Stereotypical Male PGP]		(0.004)		(0.001)
F PGP in F workplace		-0.012***		0.164***
[Stereotypical Female PGP]		(0.002)		(0.004)
F PGP in M workplace		-0.068***		0.197***
[Non-stereotypical Female PGP]		(0.004)		(0.007)
Observations	217,568	217,568	217,568	217,568

Notes: OLS estimation. Regression of the actual SGP's on the predicted gender preferences (PGP's) using data from 2000-2004 only. Controls include occupation, industry and firm gender composition fixed effects as well as year FE. Beta coefficients reported and robust standard errors in parentheses. *** significant at 1 percent, ** significant at 5 percent, * significant at 10 percent.

Source: AMS-ASSD data, own calculations.

a check, we have re-estimated our main models for the gender of the newly hired worker adding industry×year effects, occupation×year effects, or workplace gender composition×year effects, and find that the implied impacts are quite similar (Figure D.4 in Appendix).

4.2.1 EFFECTS ON THE OF PROBABILITY OF HIRING A FEMALE CANDIDATE

We begin by using this simple framework to estimate the effects of eliminating the use of stated gender preferences on the probability of hiring a female worker. Before reporting difference-in-difference results, we present event study graphs in Figure 6. These are based on a slightly more general specification:

$$y_j = \beta_0 + \beta_{1t} D_j^f + \beta_{2t} D_j^m + X_j \delta + \varepsilon_j \tag{7}$$

with separate coefficients on D_j^f and D_j^m in each year. We estimate the pair of coefficients (β_{1t}, β_{2t}) for each year of our sample, then normalized the estimates relative to 2004 by forming $\widehat{\beta}_{1t} - \widehat{\beta}_{1,2004}$ and $\widehat{\beta}_{2t} - \widehat{\beta}_{2,2004}$. The left panel of the figure plots the values of the normalized series using our LOM classifier to predict SGP's, while the right panel shows estimates based on the Random Forest classifier.

The results using either classifier suggest that in the period up to 2004 the fraction of females hired for vacancies with a predicted female or male gender preference (PGP) were relatively constant. Starting in 2005, however, there was a clear shift, with a rise in the fraction of women hired to fill vacancies that were predicted to have a male PGP (blue lines in each panel) and a fall in the fraction of women hired to fill vacancies that were predicted to have a female PGP (red lines in each panel). The magnitudes of the shifts differ somewhat between the two prediction models, with the LOM classifier suggesting a roughly 2-3 percentage point rise in the fraction of women hired for openings with a predicted male PGP, and a 1-2 points fall in the fraction hired for openings with a predicted female PGP, and the RF classifier suggesting a smaller rise in the former group (on order of 1-2 percentage points) but a larger fall for the latter group (2-3 percentage points). Note that the difference in the probability of hiring a female worker between vacancies with predicted male and female preferences (i.e., the difference between the blue line and the red line) increased by 3-5 percentage points after the OET campaign using either classifier.

Figure 6: Stated Gender Preference Affects Female Hiring: Event History Results

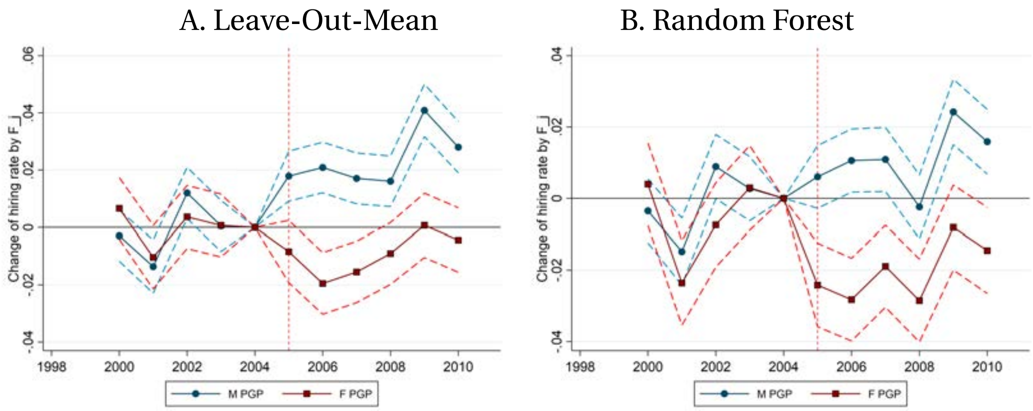

Note: This figure reports the regression coefficients capturing the effects of eliminating stated gender preferences on the hiring of females. Coefficients of the interaction term between year and indicators for job ads classified as advertising for women ("Female PGP"), or for men ("Male PGP") are reported. Dotted lines show the 95% confidence intervals. Controls include industry, occupation, lagged female share in the workplace and year fixed effects.A: classification uses regression, B: classification uses random forests.

Source: AMS-ASSD data, own calculations.

Next we turn to difference of difference estimates derived from equation (1). As noted we use the 2000-2004 data as the "pre" period and 2006-2010 as the "post" period, removing the transition year 2005 from the analysis. Figure 7 graphically summarizes the coefficients from equation (1). We report the main effects of the PGP variables ($\beta_1$and β_2) in the upper part of the figure. As a point of reference we also show a 0 coefficient representing the comparison group of vacancies with neither PGP. We report the interactions between the PGP dummies and the post-period dummy ($\lambda_1$and λ_2) in the lower part of the figure, again showing a 0 coefficient for the comparison group of vacancies with neither PGP. The left-hand figure (panel A) shows results based on our LOM prediction model, while the right-hand figure (panel B) shows results based on our RF model. (The main coefficients from these models (and all subsequent models) are presented in Appendix E: tables E.1, E.2, and E.3.)

Using the LOM classifier, in the pre period vacancies with a male PGP have a 6.5 percentage point (ppt) lower probability of hiring a female candidate relative to vacancies with neither PGP, while

Figure 7: Female Hiring (Difference in Differences Results)

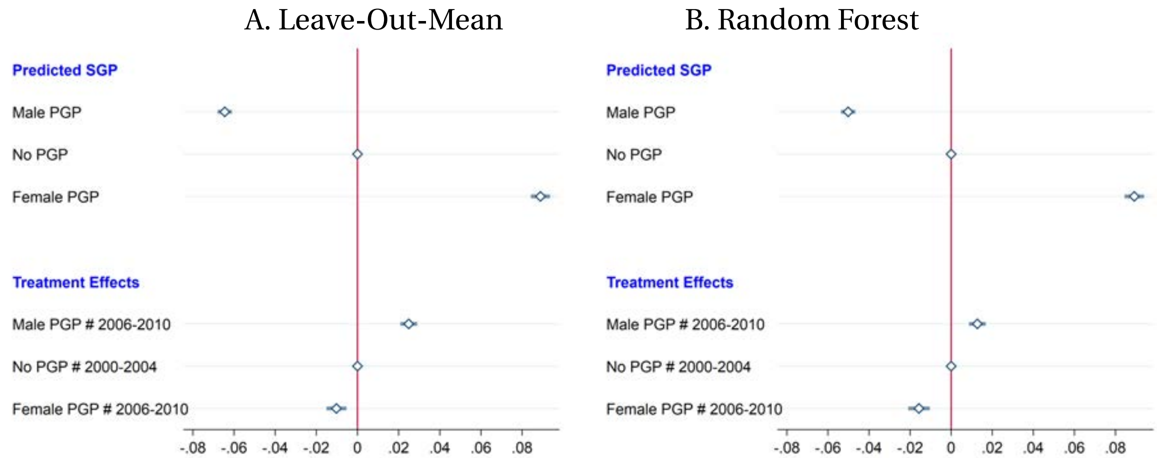

Notes: This figure reports the estimation results capturing the effect of eliminating stated gender preferences on the hiring of females. The regressions follow the model in equation 1. Diamond refers to the point estimates, horizontal lines show the 95% confidence intervals. Pre-campaign period refers to vacancies posted in 2000-2004, treatment effects refer to vacancies posted in 2006-2010. Vacancies posted in 2005 are dropped from the regression. Controls include industry, occupation, lagged female share in the workplace and year fixed effects. A: classification uses regression, B: classification uses random forests.

Source: AMS-ASSD data, own calculations.

vacancies with a female PGP have an 8.9 ppt higher probability of hiring a female. The estimates from the RF prediction model are similar, but show a slightly smaller impact of a male PGP (-5.5 ppts) and a slightly larger impact of a female PGP (10 ppt).

Note that if we assume an 80% attenuation in the effects of PGP status relative to true SGP status, then these estimates imply that stating a male SGP leads to 35 ppt lower probability of hiring a female relative to no SGP, while stating a female SGP leads to a 45 ppt increase. The implied gap in the probability of female hiring between female and male SGP's is therefore around 80 ppt, consistent with the patterns observed in Figures 3 and 4.

Looking next at the treatment effects, the estimates in Panel A (based on the leave-out-mean classifier) show that eliminating gender preferences increased the hiring of women by 2.5 percentage points for vacancies with a male PGP (standard error = 0.002), which is about 40% of the pre-campaign gap in the female hiring rate relative to vacancies with neither PGP. The estimated effect on the probability of hiring a female for vacancies with a predicted female PGP is -1.0 ppt (s.e. = 0.002), which is about 12% of the pre-campaign gap for these vacancies relative to the comparison group. RF estimates in Panel B suggest a slightly smaller gain (+1.6 ppts) in the hiring of females for vacancies with a male PGP, and a slightly larger reduction in the hiring of females (-2.7 ppts) for vacancies with a female PGP, consistent with the patterns in Figure 7.

Accounting for the attenuation arising from the slippage in our prediction models, these implied treatment effects are large in magnitude. The LOM-based estimates suggest that the elimination of gender preferences in job ads led to a roughly 15 ppt increase in the probability of hiring women to fill jobs that would have been advertised with a male gender preference in the absence of the new law, and 5 ppt increase in the hiring of men to fill jobs that would have been advertised with a male gender preference. The combination of the two effects therefore closed about one quarter (=(15+5)/80) of the previous gap in the rate of hiring women between jobs that specified a gender preference for men versus women.

As noted in the discussion of Table 1, vacancies filled by job searchers who had obtained services from the AMS differ somewhat from the overall population of vacancies. To assess the possible impact of these differences on our main models, we developed a set of weights reflecting the inverse probability that a given vacancy was filled by an AMS client, and re-estimated the models in Figure 7. Reassuringly, the results, reported in Appendix Figure D.1, are qualitatively very similar to those in Figure 7.

HETEROGENEITY BY FIRM SIZE Gender based hiring might differ between small and large firms. To probe the sensitivity of our results to firm size, we exclude firms that had less than five employees during the year before the vacancy was posted. For simplicity, in this analysis and all subsequent analysis we limit attention to models based on the LOM classifier. Figure 8 summarizes estimation results for our difference in differences model applied to larger firms. The pattern of results is very similar to the baseline pattern (Figure 6), suggesting that our results are not much affected by including very small firms.

Figure 8: Results for Large Firms

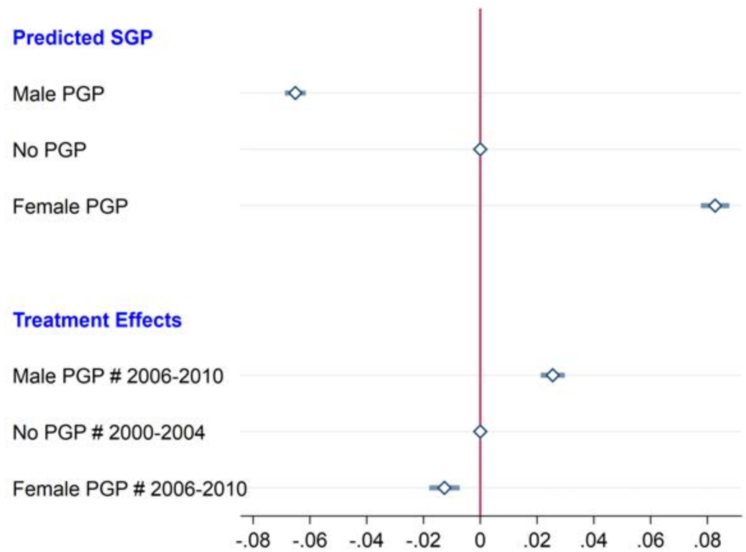

Note: This figure reports the estimation results capturing the effect of eliminating stated gender preferences on the hiring of females for firms with more than 5 employees. The regression follows the model in equation 1. Diamond refers to the point estimates, horizontal lines show the 95% confidence intervals. Pre-campaign period refers to vacancies posted in 2000-2004, treatment effects refer to vacancies posted in 2006-2010. Vacancies posted in 2005 are dropped from the regression. Controls include industry, occupation, lagged female share in the workplace and year fixed effects.

Source: AMS-ASSD data, own calculations.

VIENNA VS. THE REST OF AUSTRIA: Vienna is by far the largest labor market in Austria, with a relatively high share of employment in the services sector and a lower share in manufacturing and other more traditional sectors. We therefore estimated our baseline model separately for Vienna and the rest of Austria, with results in Figure 9.

Looking at the pre-campaign differences in the upper part of the figure, we see that vacancies with a male PGP had only about a 2.5 ppt lower probability of hiring a female in Vienna relative to vacancies in the comparison group, versus a roughly 7 ppt gap in the rest of Austria. In contrast, for vacancies with a female PGP the pre-campaign gap relative to the comparison group was about 10 ppt in Vienna versus 9 ppt in the rest of Austria.

After the campaign the probability of hiring a female for male PGP vacancies increased by about 1 ppt in Vienna – closing about 1/3 of the previous gap – while in the rest of the country it rose by about 2.5 ppt – again closing about 1/3 of the gap. For female PGP's the probability of hiring a female fell by about 3 ppt in Vienna (nearly one quarter of the pre-campaign gap) while in the rest of the country the decline was smaller (about a 1 ppt reduction, closing only 1/10 of the pre-campaign gap relative to the comparison group).

Figure 9: Vienna vs the rest

(a) Vienna (b) the rest of Austria

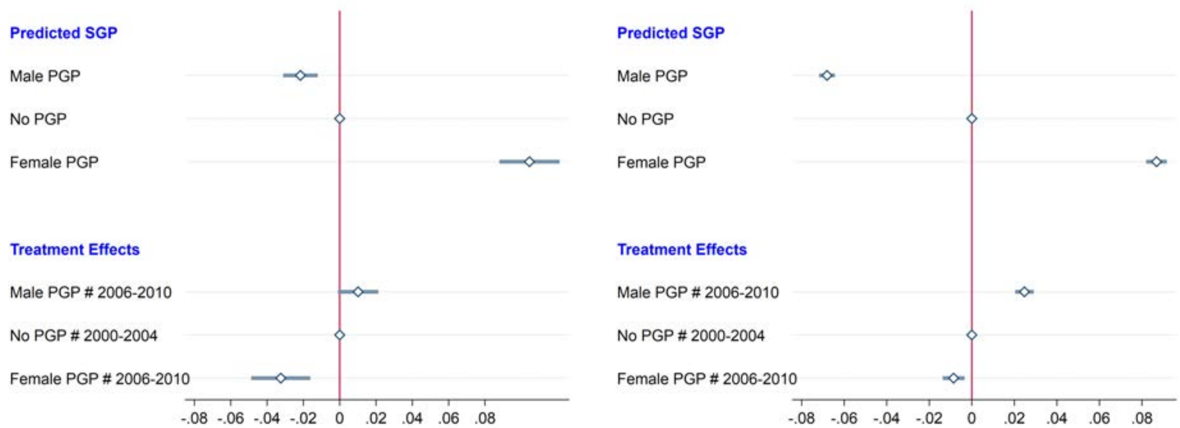

Note: This figure reports the estimation results capturing the effect of eliminating stated gender preferences on the hiring of females for firms in Vienna (panel A) and in the rest of Austria (panel B). The regression follows the model in equation 1. Diamond refers to the point estimates, horizontal lines show the 95% confidence intervals. Pre-campaign period refers to vacancies posted in 2000-2004, treatment effects refer to vacancies posted in 2006-2010. Vacancies posted in 2005 are dropped from the regression. Controls include industry, occupation, lagged female share in the workplace and year fixed effects.

Source: AMS-ASSD data, own calculations.

These results suggest that in the pre-campaign period the impacts of SGP's varied somewhat across labor markets, but that in all cases the elimination of posted gender preferences led to some gender diversification in hiring. The effects were particularly large for jobs with male PGP's outside of Vienna, and for jobs with female PGP's in Vienna. We speculate that more heavily male blue collar jobs were particularly segregated in the smaller towns and rural areas of Austria, and opened up more after the campaign, while more heavily female pink collar jobs in services and government were more segregated in Vienna, and also opened up more after the campaign.[20]

[20]Another concern with our baseline analysis is the presence of a fair amount of seasonal work in Austria. In a sensitivity analysis, we remove vacancies that offer seasonal jobs (12.1% in 2000-2004, 11.8% in 2006-2010), and find the same pattern of results as in our baseline analysis. Results are available upon request.

4.2.2 Stereotypical vs Non-stereotypical Vacancies

Our baseline difference of differences specification makes no distinction between SGP's that specify a preference for a gender that matches the majority of the existing workforce of the firm (stereotypical preferences) or for the opposite gender (non-stereotypical preferences). Our simple conceptual model of SGP's, however, suggests that vacancies with a predicted non-stereotypical gender preference could be most affected by the elimination of gender preferences. Moreover, the contrast between stereotypical and non-stereotypical preferences is interesting because the former help maintain gender segregation, while the latter promote gender diversity. In this section we expand our framework and explore heterogeneity along this dimension.

To proceed we label workplaces with more than 50% females as mainly female, indicated by a dummy variable $C_j^f = 1[F_j >= 0.5]$, and those with less than 50% females as mainly male, indicated by a dummy variable $C_j^m = 1[F_j < 0.5]$.[21] A vacancy with a given predicted SGP is stereotypical if the predicted gender preference is concordant with the firm's existing gender composition ($D_j^f C_j^f = 1$, stereotypical female; or $D_j^m C_j^m = 1$, stereotypical male) and non-stereotypical if the predicted gender preference is discordant with the existing workforce composition ($D_j^f C_j^m = 1$, non-stereotypical female; or $D_j^m C_j^f = 1$, non-stereotypical male).

Table 5 presents some descriptive information on stereotypical and non-stereotypical vacancies in the pre-campaign period (where we can see actual gender preferences as well as our predicted preferences). Columns 1 and 2 show characteristics of vacancies with a male ***predicted*** gender preference (PGP), classified further by whether the employer had a mainly female workplace (column 1) or a mainly male workplace (column 2). Note that only 7% of male PGP vacancies (in column 1) are non-stereotypical. Similarly, columns 4 and 5 show characteristics of vacancies with a female PGP, classified by whether the firm has a majority female workplace (column 4 - the stereotypical case) or a mainly male workplace (column 5 – the non-stereotypical case). Again, only a small share (12.8%) of female PGP vacancies are non-stereotypical.

Panel A of the table shows the shares of each PGP that had actual stated preferences for males,

[21] This classification is the same as we used on Table 3 to examine vacancy filling times in the pre-campaign years.

females, or neither. Interestingly, we see that the fraction of concordant actual and predicted gender preferences is about the same whether the PGP is stereotypical or not. For example, 58.81% of stereotypical male PGP's have an actual male SGP (column 1), versus 60.73% of non-stereotypical male PGP's. The gap is only slightly wider between stereotypical and non-stereotypical female PGP's (54.45% versus 48.18%).

Panel B shows the shares of female workers at the posting firms and in the relevant occupations for each PGP group. By construction stereotypical and non-stereotypical vacancies come from firms with much different shares of female workers. The differences in gender shares in the recruited occupations, however, are much smaller. For example, a typical occupation for a non-stereotypical female PGP vacancy (column 5) was 67.03% female, whereas a typical occupation for a stereotypical female PGP vacancy (column 4) was 70.75% female. The gap is somewhat larger for male PGP vacancies (~18 ppt) but is still only about 1/3 as large as the gap in the share of women at the firm's workplace (50 ppts).

Table 5: Descriptive Statistics for Stereotypical and Non-Stereotypical Predicted Gender Preferences

	Predicted Gender Preference				
	M PGP in F WP [Non-stereo Male PGP]	**M PGP in M WP** [Stereo Male PGP]	**Neutral**	**F PGP in F WP** [Stereo Female PGP]	**F PGP in M WP** [Non-stereo Female PGP]
	(1)	(2)	(3)	(4)	(5)
Panel A: SGP					
Male SGP	0.5881	0.6073	0.1820	0.0168	0.0617
Neutral	0.3719	0.3802	0.6314	0.4387	0.4565
Female SGP	0.0400	0.0126	0.1867	0.5445	0.4818
Panel B: Context					
Share of women in the firm	0.6349	0.1330	0.4696	0.8810	0.3068
Share of women in occupation	0.3728	0.1952	0.5203	0.7075	0.6703
Observations	3,676	51,841	114,273	41,745	6,033
Shares	0.0169	0.2383	0.5252	0.1919	0.0277

This table shows summary statistics for stereotypical (Male PGP in Male Workplace in column (2) and Female PGP in Female Workplace column (3)) and non-stereotypical (Male PGP in Female Workplace in column (1) and Female PGP in Male Workplace column (4)) PGP's. Period 2000-2004.

We use an extended version of equation (1) to measure the effects of the 2004 campaign on 4 distinct types of vacancies based on PGP and workplace gender composition:

$$
\begin{aligned}
y_j &= \beta_0 + \beta_1 D_j^f C_j^f + \beta_2 D_j^f C_j^m + \beta_3 D_j^m C_j^m + \beta_4 D_j^m C_j^f \\
&+ \lambda_1 D_j^f C_j^f Post_j + \lambda_2 D_j^f C_j^m Post_j + \lambda_3 D_j^m C_j^m Post_j + \lambda_4 D_j^m C_j^f Post_j + X_j \delta + \varepsilon_j
\end{aligned} \tag{8}
$$

In this specification the β_k coefficients measure differences across vacancies in the outcome y before the campaign is implemented, while the λ_k coefficients measure changes in the the outcome after the introduction of the campaign (λ_1 and λ_3 for stereotypical vacancies, and λ_2 and λ_4 for non-stereotypical vacancies). As in the simpler model (1), the λ_k coefficients are attenuated by misclassification errors relative to the coefficients in a model that relates the outcome to the stated gender or desired gender preference. In the Appendix we present an analysis similar to that in equations (3)-(6) discussing the attenuation effects. As shown in columns 2 and 4 of Table 4, when we relate actual SGP's to predicted gender preferences, allowing different effects for stereotypical and non-stereotypical PGP's, we find that non-stereotypical PGP's of a given gender are stronger predictors of actual SGP's of that gender, but are also more negatively related to the SGP of the opposite gender. Allowing for the multiple channels we conclude that the expected attenuation of the effects of stereotypical and non-stereotypical vacancies in a model like equation (7) are likely to be similar, and in both cases close to a factor of 80%.

Estimation results for equation (7) are summarized in Figure 10.[22] The estimates show that in the pre-campaign period non-stereotypical PGP's had relatively large effects on the probability of hiring a female candidate (as would be expected from the "Z"-shaped pattern in Figures 3-4), whereas stereotypical PGP's had more modest effects, comparable to the overall effects from specification (1) reported in Figure 10. Specifically, females were 4.7 ppts less likely to be hired for vacancies with a stereotypical male PGP (relative to the omitted group of vacancies with no preference), but 18.7 ppts less likely to be hired for vacancies with a stereotypical male PGP. Conversely, females were were 5.7 ppts more likely to be hired for vacancies with a stereotypical female PGP and 19.3 ppts

[22]As in our previous models we include controls for time, industry (86 dummies), occupation (100 dummies), and the proportion of females at the workplace in the year prior to the posted vacancy (51 dummies).

more likely to be hired for vacancies with a stereotypical female PGP.

Figure 10: Effects on Female Hiring (Stereotypical vs Non-Stereotypical Vacancies)

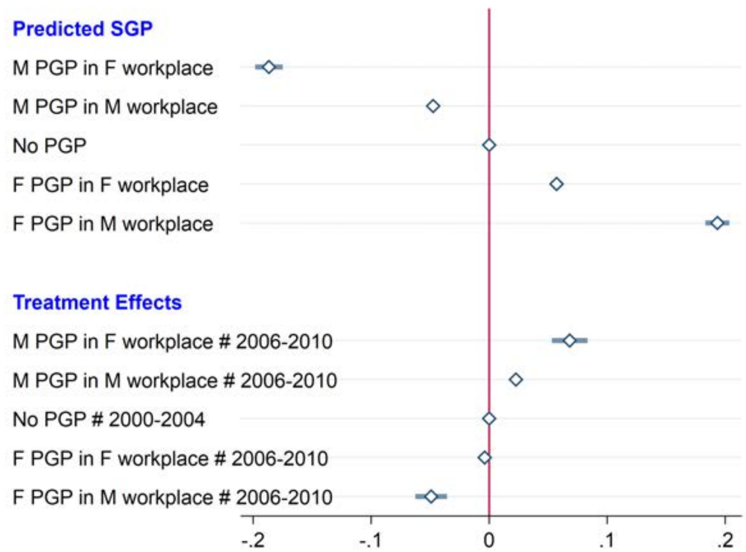

Notes: This figure reports the estimation results capturing the effect of eliminating stated gender preferences on the hiring of females for Stereotypical (Predicted M SGP in M Workplace and F SGP in F Workplace) and Non-Stereotypical (Predicted M SGP in F Workplace and F SGP in M Workplace) vacancies. The regression follows the model in equation 8. Diamond refers to the point estimates, horizontal lines show the 95% confidence intervals. Pre-campaign period refers to vacancies posted in 2000-2004, treatment effects refer to vacancies posted in 2006-2010. Vacancies posted in 2005 are dropped from the regression. Controls include industry, occupation, lagged female share in the workplace and year fixed effects.

*Source:*AMS-ASSD data, own calculations.

Likewise, the difference in differences estimates show bigger effects of the elimination of SGP's on the hiring outcomes of vacancies with non-stereotypical PGP's: a 6.8 ppt increase (s.e.=0.008) in the probability of hiring a female for a non-stereotypical male PGP versus a 2.3 ppt increase (s.e.=0.002) for a stereotypical male PGP; and a 4.9 ppt reduction (s.e.=0.007) in the probability of hiring a female for vacancies with a non-stereotypical female PGP's, versus an insignificant 0.004 ppt reduction for vacancies with a stereotypical female PGP.

These estimates imply that most of the overall increase in the hiring of females for vacancies with a male PGP was attributable to gains at majority-male workplaces (which tended to use stereotypical male gender preferences in the pre-campaign period). Interestingly, there was no corresponding increase in the hiring of males for vacancies with a female PGP at majority-female workplaces. Instead, the gains in hiring for men were driven entirely by a rise in the fraction of vacancies at mainly male workplaces with a predicted female preference that were filled by men once SGP's were eliminated.

4.2.3 Effects on Vacancy Durations and Other Job Outcomes

The elimination of gender preferences in recruiting could affect both the cost of forming new job matches and the quality of these matches. To assess these effects we examine three outcomes: the number of days required to fill a vacancy ("fill time"); the starting wage at the newly created job, and the duration of the newly created job. We interpret fill time as a measure of the cost of forming a match and the duration of the newly created job as a measure of the quality of the match. The interpretation of wage outcomes is less clear, though as we will see impacts of predicted gender preferences and the elimination of such preferences on wages are small. We conduct our analysis in the framework of equations (1) and (7), simply replacing the dependent variable with the match outcome measures. All estimates condition on the three sets of fixed effects included in our previous models (industry, occupation, and lagged female share) as well as year fixed effects.

Looking first at vacancy fill times, the results in Panel A of Figure 11, show that in the pre-campaign period vacancies with a female PGP had shorter fill times. The elimination of SGP's had little or no effect on filling times for vacancies with male PGP, but significantly increased the time to fill those with a female PGP. The estimates in Panel B, based on specification (7) provide some additional insights. In particular, we see that average fill times for both stereotypical and non-stereotypical male PGP's were about the same in the post-campaign period as before. For vacancies with a stereotypical female PGP we also see no significant change in fill times. All of the rise in average fill times for predicted female PGP's is attributable to an increase in vacancy durations for non-stereotypical female PGP's.

The lack of effects on vacancy filling times for stereotypical male PGP's is interesting. Contrary to what might have been expected given the assumptions of Kuhn and Shen (2013), we do not see a slow down in the vacancy filling rate caused by "congestion" in the candidate assessment process for such vacancies. Taken together with the fact that more female candidates were hired to fill these vacancies in the post-campaign period, the evidence suggests that some employers had been using invalid or out-of-date priors to limit their applicant pools. Once SGP's were eliminated, they learned that there were in fact acceptable female candidates and actually hired some of them, taking about the same time to reach a decision with a presumably larger applicant pool.

Figure 11: Effects on Vacancy Filling

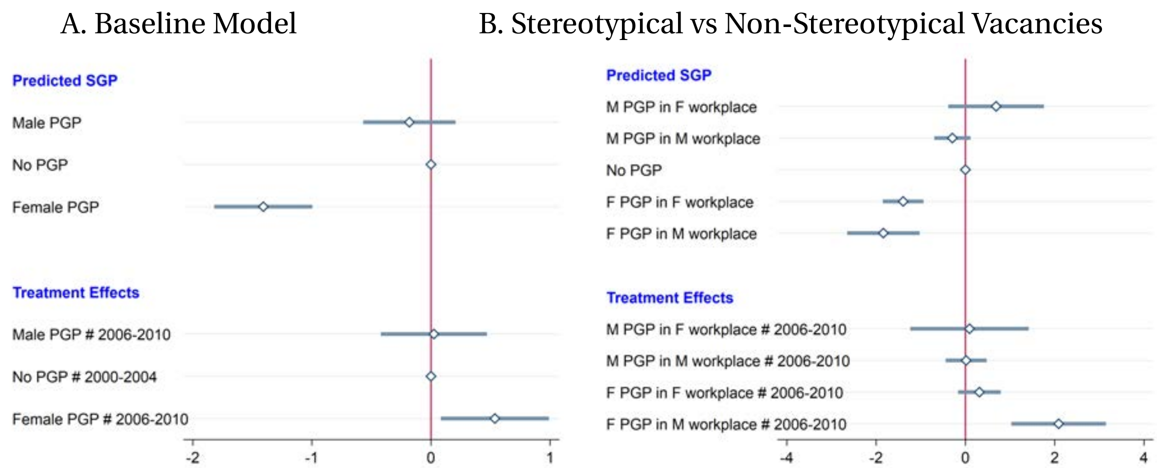

Notes: This figure reports the estimation results capturing the effect of eliminating stated gender preferences on the hiring of females for Stereotypical (Predicted M SGP in M Workplace and F SGP in F Workplace) and Non-Stereotypical (Predicted M SGP in F Workplace and F SGP in M Workplace) vacancies. The regression follows the model in equation 8. Diamond refers to the point estimates, horizontal lines show the 95% confidence intervals. Diamond refers to the point estimates, horizontal lines show the 95% confidence intervals. Pre-campaign period refers to vacancies posted in 2000-2004, treatment effects refer to vacancies posted in 2006-2010. Vacancies posted in 2005 are dropped from the regression. Controls include industry, occupation, lagged female share in the workplace and year fixed effects.

*Source:*AMS-ASSD data, own calculations.

The most notable post-campaign effect in Panel B of Figure 12 is the approximately 2 day rise in filling times for vacancies with non-stereotypical female PGP's (a roughly 7% increase). We saw in Figure 10 they these openings were more likely to be filled by men, suggesting that some employers with a majority male workforce that wanted to recruit a female worker experienced a reduction in applications from females in the post-2005 period and ultimately had to compromise on a male candidate, with some overall slow down in the decision process.

Next we turn to job durations and wages. In our main models for these outcomes we add controls for the gender, age and previous wage of the new recruit. While the characteristics of the newly hired worker are endogenous, we believe it is most straightforward to interpret the impacts of the elimination of SGP's on wages and job durations with these controls in place, since female workers have lower wages and longer job durations across all jobs, and a change in the gender of the new recruit would be expected to mechanically affect these outcomes. We present models without the extra controls in Appendix Figure D.3, and discuss the differences relative to models with these controls below.

Figure 12 summarizes the results for job durations. An issue in constructing the dependent variable for this analysis is right censoring of jobs that were still in progress at the time we extracted our sample. We address this by censoring all jobs at four years (1440 days). Since only about 10 percent of jobs last more than four years, we believe this censoring has little impact on our findings.

The estimates in panel A of Figure 12 show that in the pre-2005 period durations were insignificantly longer for jobs created by vacancies with a male PGP (relative to those with no predicted preference) but significantly shorter (by about 19 days) for jobs that filled vacancies with a female PGP. We also see (at the bottom of the figure) that jobs filled by women tend to last considerably longer (by about 45 days) conditional on industry, occupation and the gender composition of the employer's workforce – consistent with the findings in Table 2. In the post-campaign period there was no change in the average duration of jobs created by vacancies with a male PGP but a further $\sim$ 14 day reduction in the duration of jobs created by vacancies with a female PGP, equivalent to a roughly 7 percent effect relative to the median.

The results in Panel B, based on equation (7), provide some further insights into these treatment effects. In particular, we we can see that the relatively constant average duration of jobs arising from vacancies with a male PGP reflects a combination of a small (insignificant) increase in the duration of jobs associated with stereotypical male PGP's and a large reduction in the duration of jobs associated with non-stereotypical male PGP's. The decline in the average duration of jobs from vacancies with a female PGP, on the other hand, is entirely attributable to a fall in the duration of jobs associated with stereotypical female PGP's (i.e., jobs at mainly female workplaces).

We suspect that the shortened duration of jobs associated with non-stereotypical male PGP's is reflective of a reduction in match quality for these jobs (which are more likely to be filled by female candidates in the post-campaign period), coupled with a general reduction in job durations at mainly female workplaces that also affects the jobs associated with stereotypical female PGP vacancies. To examine this further, we conducted an event study analysis of job durations associated with male and female PGP's (see Appendix Figure D.5). A visual examination suggests that durations of jobs associated with female PGP's (the vast majority of which are at mainly female workplaces) were trending downward by 2-3 days per year throughout our sample period. This trend can account for most of the apparent treatment effect for stereotypical female PGP vacancies.

Figure 12: Effects on Job Duration

C. Baseline Model

D. Stereotypical vs Non-Stereotypical Vacancies

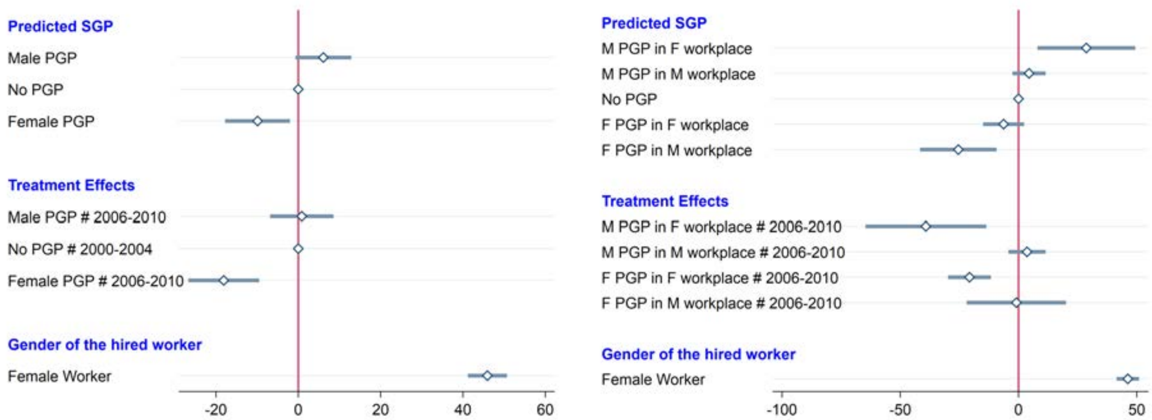

Notes: This figure reports the estimation results capturing the effect of eliminating stated gender preferences on completed job durations. In panel A the the model follows the baseline specification in equation 1. In panel B the model follows the specification in equation 8, distinguishing between Stereotypical (M PGP in M Workplace and F PGP in F Workplace) and Non-Stereotypical (M PGP in F Workplace and F PGP in M Workplace) vacancies. Diamond refers to the point estimates, horizontal lines show the 95% confidence intervals. Pre-campaign period refers to vacancies posted in 2000-2004, treatment effects refer to vacancies posted in 2006-2010. Vacancies posted in 2005 are dropped from the regression. Controls include previous wage, gender, age, as well as industry, occupation, lagged female share in the workplace and year fixed effects.

Source: AMS-ASSD data, own calculations. calculations

The elimination of stated gender preferences might also affect the wages offered to new recruits (conditional on gender, age, and past wage). As shown in panel A of Figure 13 we find that, prior to the campaign, daily earnings were about 3 percentage points lower for jobs associated with vacancies with a female PGP, and around 1 percentage point higher than in the comparison group for jobs arising out of vacancies with a male PGP. Neither set of jobs experience significant changes in pay after the campaign. Looking in panel B of the figure, we see that in the pre-2005 period there was a significant wage disadvantage for non-stereotypical male PGP jobs (mostly jobs held by men at mainly female workplaces) as well as for jobs associated with both stereotypical and non-stereotypical female PGP's. In the post-campaign period there was a modest, marginally significant gain in wages for non-stereotypical male PGP jobs but no large change for any other type of new job. Indeed, apart from the gain in wages for non-stereotypical male PGP jobs, we can rule out wage changes that are bigger than ±2 percentage points at conventional levels of statistical significance.[23]

Figure 13: Effects on Wages

A. Baseline Model

B. Stereotypical vs Non-Stereotypical Vacancies

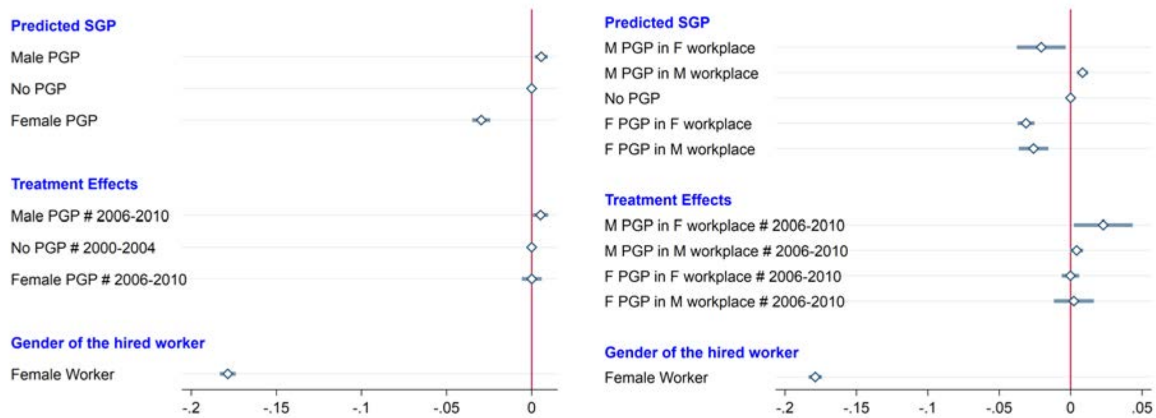

Notes: This figure reports estimation results capturing the effect of eliminating stated gender preferences on wages for newly filled jobs. In panel A the the model follows the baseline specification in equation 1. In panel B the model follows the specification in equation 8, distinguishing between Stereotypical (M PGP in M Workplace and F PGP in F Workplace) and Non-Stereotypical (M PGP in F Workplace and F PGP in M Workplace) vacancies. Diamond refers to the point estimates, horizontal lines show the 95% confidence intervals. Pre-campaign period refers to vacancies posted in 2000-2004, treatment effects refer to vacancies posted in 2006-2010. Vacancies posted in 2005 are dropped from the regression. Controls include previous wage, gender, age, as well as industry, occupation, lagged female share in the workplace and year fixed effects.

Source: AMS-ASSD data, own calculations. calculations

It is worth emphasizing that the models in Figure 13 control for the gender age and previous wage of the new recruits. When we exclude these controls, we find that mean wages for jobs associated with non-stereotypical male PGP's fell in the post-campaign period while wages for jobs associated with non-stereotypical female PGP's rose. We interpret both of these impacts as a result

[23]There are some differences in the part-time versus full time status of male versus female PGP's: vacancies with a female PGP tend to have a larger share of part time work, while those with a male PGP have a smaller share of part time work.

of compositional effects: in the former case because of a rise in the share of women hired to fill vacancies with non-stereotypical male PGP's; in the latter because of a rise in the share of men hired to fill vacancies with non-stereotypical female PGP's.

4.3 Effects on Workplace Diversity

As a final exercise, we return to the issue of the gender of the newly hired work and ask explicitly how the new recruit affected the gender diversity of the employer's workforce. To quantify this effect we define a simple measure $d_j \equiv |H_j - F_j|$, where H_j is a dummy equal to 1 if the new hire for vacancy j is female, and F_j represents the share of female workers at the workplace posting the vacancy (measured in the previous year). This indicator, which ranges from 0 (a completely segregated workplace hires a new worker of the same gender as all existing workers) to 1 (a completely segregated workplace hires a new worker of the opposite gender), summarizes the deviation of the gender of the new hire from the composition of the workforce. The index will be higher if new hires are more likely to work with co-workers of the opposite gender.[24] The mean value of the index during the pre-2005 period was about 0.22.

Figure 14 summarizes the estimates from our simpler specification (1) (in Panel A) and from our richer specification (7) that distinguishes between stereotypical and non-stereotypical PGP's (in

[24] Specifically, rearranging the indicator shows $E(d_j) = E(M_j|H_j = 1)P(H_j = 1) + E(F_j|H_j = 0)P(H_j = 0)$, where $M_j \equiv 1 - F_j$ is the share of males at the workplace. The first term in this expression is the share of men at workplaces hiring women times the share of new hires that are women; the second term is the share of women at workplaces hiring men times the share of new hires that are men. Holding constant the gender composition of all new hires, this indicator will rise if the coworkers of new hires are more likely to be of the opposite gender. A second way to measure diversity is, $e_j \equiv H_j(1 - C_j^f) + (1 - H_j)C_j^f$ where C_j^f identifies female firms. This second index identifies female hires in male firms, $H_j(1 - C_j^f)$, and male hires in female firms, $(1 - H_j)C_j^f$. This index bears a relationship with the Duncan index of gender segregation index (Duncan and Duncan, 1955), $D_j = \frac{1}{2}\sum_j |\frac{m_j}{M} - \frac{f_j}{F}|$, where m_j is the number of men, f_j is the number of women employed at workplace j, and M and F are the total number of men and women. Suppose the total numbers of men and women are identical, e.g. $M = F$, then the Duncan index then is $D_j = \frac{1}{2F}\sum_j [(m_j - f_j)(1 - C_j^f) + (f_j - m_j)C_j^f]$, capturing by how much men outnumber women in male firms, and women outnumber men in female firms, on average. Suppose we compare segregation over two time periods in a situation where firms hire only one person, and no person leaves the firm. The change in segregation across these two time periods is (up to normalization) $\Delta D_j \propto \sum_j [(\frac{1}{2} - H_j)(1 - C_j^f) + (H_j - \frac{1}{2})C_j^f] = \sum_j \frac{1}{2}[(1 - C_j^f) - C_j^f - e_j]$. In the appendix (Figure D.6), we report results for $\frac{1}{N}\sum_j e_j$, and these results can be interpreted as the reduction in segregation, as measured by the Duncan index, induced by the campaign. Results are consistent across both types of indices.

Panel B). Looking first at Panel A, we see that in the pre-campaign period, use of SGP's was associated with reductions in diversity, relative to the comparison group of vacancies with neither PGP. This reflects the fact that, e.g. vacancies with male PGP recruit men into male workplaces, and similarly for women. The effect of eliminating SGP's was to increase the probability of a diversity-enhancing hire for both male and female PGP vacancies, with an effect of +0.007 (s.e.=0.001) for male PGP's and +0.006 (s.e.=0.002) for female PGP's. Recalling that these effects are attenuated by a factor of roughly 80% relative to a specification using stated/desired SGP's, the estimates suggest relatively large increases in the diversity of new hires (+.035 for male male PGP's and +0.030 for female SGP's) relative to a pre-campaign mean of 0.22. Stated differently, we estimate that the campaign to eliminate stated gender preferences reduced the tendency to hire non-diversifying candidates at workplaces that would use these preferences by a factor of 1/6. Of course many other factors that contribute to non-diverse hiring (such as location of the workplace, schedule of work, and attitudes of managers and coworkers) were unaffected by the campaign.

Figure 14: Effects on Workforce diversity

A. Baseline Model B. Stereotypical vs Non-Stereotypical Vacancies

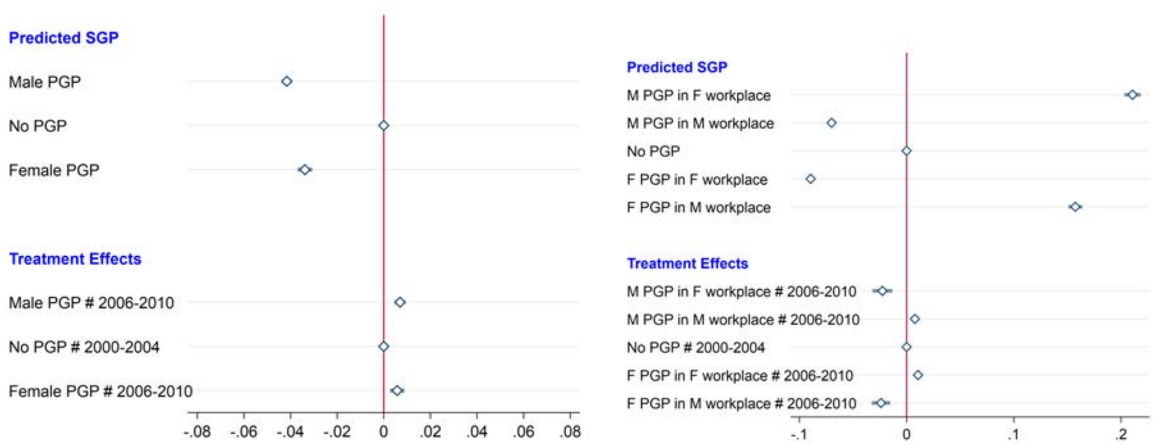

Notes: This figure reports the estimation results capturing the effect of eliminating stated gender preferences on workplace composition. In panel A the model follows the baseline specification in equation 1. In panel B the model regression follows the specification in equation 8, distinguishing between Stereotypical (Predicted M SGP in M Workplace and F SGP in F Workplace) and Non-Stereotypical (Predicted M SGP in F Workplace and F SGP in M Workplace) vacancies. Diamond refers to the point estimates, horizontal lines show the 95% confidence intervals. Pre-campaign period refers to vacancies posted in 2000-2004, treatment effects refer to vacancies posted in 2006-2010. Vacancies posted in 2005 are dropped from the regression. Controls include industry, occupation, lagged female share in the workplace and year fixed effects.

Source: AMS-ASSD data, own calculations.

Looking at the estimates from our richer specification (Panel B), we see a sharp contrast between stereotypical and non-stereotypical PGP's in both the pre-period and the post-campaign period. When stated SGP's were allowed, our estimates suggest that non-stereotypical SGP's substantially increased the probability of a diversity-enhancing hire, while stereotypical SGP's had the opposite effect. The effect of eliminating stated gender preferences was to increase the probability of a diversity-enhancing hire for vacancies that are predicted to have had a stereotypical gender preference, but to reduce the probability of a diversity-enhancing hire for those predicted to have had a non-stereotypical gender preference.[25] These patterns are consistent with the effects of the campaign on the gender of the candidates hired to fill stereotypical and non-stereotypical PGP's noted in Figure 10. But since stereotypical vacancies account for 95% of all vacancies, the increase in diversity on stereotypical hires dominates the loss in diversity on non-stereotypical vacancies. Interestingly, we also see that non-stereotypical vacancies continue to reduce segregation to an important extent, even after the campaign.

5 CONCLUSION

In early 2005 the *Ombud for Equal Treatment* (OET) agency in Austria began a publicity campaign to inform employers and newspapers that gender preferences in job advertising were illegal. Over the next year the use of stated gender preferences on the job board of the Austrian Employment Service fell from around 40% of all vacancies to less than 5%. We use data on filled vacancies from this board to study how the elimination of stated gender preferences affected hiring outcomes and job quality measures. To focus on affected firms, we use pre-2005 data on the occupation of the job opening, the firm's industry, and its lagged female employment share to predict the gender preferences (female, male, or neither) in a given vacancy. We further classify predicted preferences into stereotypical or non-stereotypical, based on the concordance between the predicted preference and the

[25] The result on stereotypical female vacancies is noteworthy. Firms that posted stereotypical female vacancies did not hire significantly more men (Figure 10B), but the male and female hires entered more workplaces with many co-workers of the opposite gender. Abolishing stereotypical job ads can increase diversity even without affecting hires directly.

majority gender at the firm's workplace. Then we estimate simple difference in differences models, comparing pre- and post-campaign hiring outcomes by predicted gender preference group. These models can be interpreted as reduced form estimates from a system that expresses hiring outcomes in terms of latent gender preferences that are only observed in the pre-2005 period.

Prior to the OET's campaign we find that most stated gender preferences (90% or more) were aligned with the gender composition of the recruiting firm's existing workforce. Thus, on average, stated gender preferences tended to reinforce existing patterns of gender segregation. Nevertheless, a small fraction of employers posted non-stereotypical preferences, requesting candidates that were the opposite gender of the majority of their workforce. Consistent with a series of papers looking at call-back rates to job applicants (Kuhn and Shen, 2013; Delgado Helleseter *et al.*, 2018; Kuhn *et al.*, 2020) we find that stated gender preferences were strongly predictive of actual hiring outcomes. We also find that vacancies with both stereotypical and non-stereotypical gender preferences had faster filling times, suggesting that gender preferences served as reliable and salient signals to job searchers, and that there were many workers available to fill even non-stereotypical job openings.

Our difference in differences models show that the elimination of stated gender preferences led to a significant 2.5 percentage point increase in the fraction of women hired to fill vacancies with a predicted male preference, and a smaller (but still significant) 1 point increase in the fraction of men hired to fill vacancies with a predicted female preference. Scaling these effects to reflect the attenuation between predicted and actual gender preferences we infer that the OET campaign had relatively large effects on job opportunities, particularly for women.

Looking further into stereotypical versus non-stereotypical preferences, we find that the gains for women were driven by a rise in hiring to fill job openings at majority-male workplaces with predicted male preferences (i.e., stereotypical male preferences). There was also an increase in the hiring of women at majority-female workplaces that were predicted to use male preferences, offset by a reduction in hiring at majority male workplaces that were predicted to use female preferences. The rise in female hiring for openings with stereotypical male preferences implies a reduction in gender segregation across workplaces, while the shifts in hiring for vacancies with non-stereotypical preferences work in the opposite direction. Given the relatively small shares of non-stereotypical preferences, however, the net effect was still toward desegregation. Examining segregation effects directly using a variant of the Duncan index, we calculate that eliminating gender preferences reduced gender segregation by about 16% at workplaces that would have used gender preference statements.

Looking at filling times, wages and durations of the newly created jobs, we find little evidence that the elimination of stated preferences led to a reduction in the speed of matching or the quality of job matches for the large group of vacancies with stereotypical gender preferences. For vacancies that were predicted to have non-stereotypical preferences we see some increase in recruiting times of openings with a predicted preference for men in female firms, and declines in the duration of jobs of openings with a predicted preference for men at a mainly-female workplace. Such effects are not necessarily surprising, given the relative rarity of such vacancies and the difficulty in signaling the firm's intended hiring target given other information about the job and the workplace.

This pattern of evidence suggest that both job seekers and firms adapted to the removal of gender preference statements after the OET campaign. On the worker side, our findings on hiring outcomes, and related findings from Kuhn and Shen (2021) on the responses of job applicants to the removal of SGP's, suggest that applicant pools became more gender-diverse after the campaign. In our setting this change may have been driven in part by changes in the referrals made by Employment Service case-workers. On the firm side, our most surprising finding (again, consistent with findings reported by Kuhn and Shen, 2021 based on employer call-backs) is that some firms that would have specified a gender preference prior to the campaign responded to the newly diverse applicant pool by actually hiring workers of the opposite gender, rather than simply ignoring their applications. In most cases the net effect of these responses was to increase the gender diversity of hiring, though in cases where an employer would have used a non-stereotypical gender preference the effect was to reduce gender diversity of hiring.

Overall, our findings suggest that any negative effects of eliminating early job market signals of gender preference were mainly confined to the set of firms with non-stereotypical hiring goals. For the large set of firms with stereotypical male preferences, eliminating the ability to advertise these preferences led to an increase in hiring of women. Taken together with the finding of no degradation in match quality for the associated jobs, we conclude that many of the stereotypical preferences observed before the campaign were likely based on outdated priors, rather than on true productivity gaps or on rigidly held discriminatory beliefs that would be immune to policy.

REFERENCES

Agan, A. and Starr, S. (2018). Ban the box, criminal records, and racial discrimination: A field experiment. *The Quarterly Journal of Economics*, **133**(1), 191–235.

Agan, A., Cowgill, B., and Gee, L. K. (2020). Do workers comply with salary history bans? a survey on voluntary disclosure, adverse selection, and unraveling. In *AEA Papers and Proceedings*, volume 110, pages 215–19.

Azmat, G. and Petrongolo, B. (2014). Gender and the labor market: What have we learned from field and lab experiments? *Labour Economics*, **30**(C), 32–40.

Baert, S. (2017). Hiring Discrimination: An Overview of (Almost) All Correspondence Experiments Since 2005. GLO Discussion Paper Series 61, Global Labor Organization (GLO).

Ballance, J., Clifford, R., and Shoag, D. (2020). "no more credit score": Employer credit check bans and signal substitution. *Labour Economics*, **63**, 101769.

Bayard, K., Hellerstein, J., Neumark, D., and Troske, K. (2003). New evidence on sex segregation and sex differences in wages from matched employee-employer data. *Journal of labor Economics*, **21**(4), 887–922.

Beller, A. H. (1983). The effects of title vii of the civil rights act of 1964 on women's entry into non-traditional occupations: An economic analysis. *Law & Ineq.*, **1**, 73.

Bergmann, B. R. (1974). Occupational segregation, wages and profits when employers discriminate by race or sex. *Eastern economic journal*, **1**(2), 103–110.

Blau, P. M. (1977). *Inequality and heterogeneity: A primitive theory of social structure*, volume 7. Free Press New York.

Booth, A. and Leigh, A. (2010). Do employers discriminate by gender? A field experiment in female-dominated occupations. *Economics Letters*, **107**(2), 236–238.

Bos, M., Breza, E., and Liberman, A. (2018). The labor market effects of credit market information. *The Review of Financial Studies*, **31**(6), 2005–2037.

Botosaru, I. and Gutierrez, F. H. (2018). Difference-in-differences when the treatment status is observed in only one period. *Journal of Applied Econometrics*, **33**(1), 73–90.

Brown, C. (1984). Black-white earnings ratios since the civil rights act of 1964: The importance of labor market dropouts. *The Quarterly Journal of Economics*, **99**(1), 31–44.

Cahuc, P., Carcillo, S., Minea, A., and Valfort, M.-A. (2019). When Correspondence Studies Fail to Detect Hiring Discrimination. CEPR Discussion Papers 14028, C.E.P.R. Discussion Papers.

Card, D. (1992). Using regional variation in wages to measure the effects of the federal minimum wage. *Ilr Review*, **46**(1), 22–37.

Card, D., Cardoso, A. R., and Kline, P. (2016). Bargaining, Sorting, and the Gender Wage Gap: Quantifying the Impact of Firms on the Relative Pay of Women. *The Quarterly Journal of Economics*, **131**(2), 633–686.

Cortes, P. and Pan, J. (2017). Occupation and Gender. IZA Discussion Papers 10672, Institute of Labor Economics (IZA).

Darity, W. A. and Mason, P. L. (1998). Evidence on Discrimination in Employment: Codes of Color, Codes of Gender. *Journal of Economic Perspectives*, **12**(2), 63–90.

Delgado Helleseter, M., Kuhn, P., and Shen, K. (2018). The Age Twist in Employers' Gender Requests: Evidence from Four Job Boards. *Journal of Human Resources*, pages 0416–7836R2.

D'Haultfœuille, X. and Rathelot, R. (2017). Measuring segregation on small units: A partial identification analysis. *Quantitative Economics*, **8**(1), 39–73.

Doleac, J. L. and Hansen, B. (2020). The unintended consequences of "ban the box": Statistical discrimination and employment outcomes when criminal histories are hidden. *Journal of Labor Economics*, **38**(2), 321–374.

Duncan, O. D. and Duncan, B. (1955). A methodological analysis of segregation indexes. *American Sociological Review*, **20**(2), 210–217.

Gerard, F., Lagos, L., Severnini, E., and Card, D. (2018). Assortative Matching or Exclusionary Hiring? The Impact of Firm Policies on Racial Wage Differences in Brazil. NBER Working Papers 25176, National Bureau of Economic Research, Inc.

Goldin, C. (2014). A Grand Gender Convergence: Its Last Chapter. *American Economic Review*, **104**(4), 1091–1119.

Groshen, E. L. (1991). The structure of the female/male wage differential: Is it who you are, what you do, or where you work? *Journal of Human Resources*, pages 457–472.

Hellerstein, J., Neumark, D., and McInerney, M. (2008). Changes in Workplace Segregation in the United States between 1990 and 2000: Evidence from Matched Employer-Employee Data. In *The Analysis of Firms and Employees: Quantitative and Qualitative Approaches*, NBER Chapters, pages 163–195. National Bureau of Economic Research, Inc.

Helleseter, M. D., Kuhn, P., and Shen, K. (2016). The Age Twist in Employers' Gender Requests: Evidence from Four Job Boards. NBER Working Papers 22187, National Bureau of Economic Research, Inc.

Hersch, J. and Shinall, J. B. (2015). Fifty years later: The legacy of the civil rights act of 1964. *Journal of Policy Analysis and Management*, **34**(2), 424–456.

Kettemann, A., Mueller, A. I., and Zweimüller, J. (2018). Vacancy durations and entry wages: evidence from linked vacancy-employer-employee data. Technical report, National Bureau of Economic Research.

Kline, P., Rose, E., and Walters, C. (2021). Systematic discrimination among large u.s. employers. NBER Working Paper 29053, National Bureau of Economic Research.

Kuhn, P. and Shen, K. (2013). Gender discrimination in job ads: Evidence from china. *The Quarterly Journal of Economics*, **128**(1), 287–336.

Kuhn, P., Shen, K., and Zhang, S. (2020). Gender-targeted job ads in the recruitment process: Facts from a chinese job board. *Journal of Development Economics*, **147**, 102531.

Kuhn, P. J. and Shen, K. (2021). What happens when employers can no longer discriminate in job ads? Technical report, National Bureau of Economic Research.

Lujanski-Lammer, E. (2006). Stellenausschreibungen – ein spiegelbild der gesellschaft [job advertisements – a mirror image of society]. Diplomarbeit, University of Graz, Institute for Pedagogy.

Mueller, A. I., Osterwalder, D., and Zweimueller, J. (2019). The Austrian AMS Vacancy Database. Mimeo, University of Zurich.

Mumford, K. and Smith, P. N. (2009). What determines the part-time and gender earnings gaps in britain: Evidence from the workplace. *Oxford Economic Papers*, **61**, i56–i75.

Petersen, T. and Morgan, L. A. (1995). Separate and unequal: Occupation-establishment sex segregation and the gender wage gap. *American Journal of Sociology*, **101**(2), 329–365.

Wozniak, A. (2015). Discrimination and the effects of drug testing on black employment. *Review of Economics and Statistics*, **97**(3), 548–566.

Ziegler, A. (2021). Skill demand and wages. evidence from linked vacancy data. IZA Discussion Papers 14511, Institute of Labor Economics (IZA).

Zweimüller, J., Winter-Ebmer, R., Lalive, R., Kuhn, A., Wuellrich, J.-P., Ruf, O., and Büchi, S. (2009). Austrian social security database.

APPENDIX

A Database Construction

The AMS vacancy database contains the stock of all the open vacancies with a monthly frequency. AMS also records the inflow (posting date) and the outflow (closing date), as well as the outcome of each vacancy. Each vacancy can result with (i) the hiring of a worker through direct mediation of the AMS system, (ii) the hiring of a worker through different channels or (iii) no hiring (or no information about the hiring)[26].

The full database contains about 13.9 millions observations for the period 1997-2013, among which 5.2 millions are recorded as outflows. We consider only vacancy outflows, and first step toward the construction of our database is providing an opening date for each vacancy in the outflow subsample. AMS record a vacancy identifier, however, the identifier is not always reliable because an previously used identifier can be re utilized for a new opening in the future. We then match vacancy in time (inflows and outflows) by using the vacancy identifier (*vdgnr*) and other 10 time invariant vacancy specific variables. In this process the sample size reduces by 118,136 units (2%).

Our empirical analysis focuses only on vacancies filled through AMS, since they contain a person identifier (*penr*) that allows us to get information on the hired person, the gender in particular. This subsample contains 1.2 millions vacancies, corresponding to about one quarter of all the hires[27].

The second step of our database construction is matching the AMS vacancy outflows with the employment spells in the ASSD database (*qualifikation*). This is the most challenging part of the process since the firm identifiers in the AMS (*btrnr*) and in the ASSD (*benr*) database differ. We proceed as follows:

- We first "treat" the employment spells database. We normalize the spells by merging together all the spells between the same firm-worker pair with a break lower than 70 days and we drop

[26]This usually happens when the firm withdraw the job ad from the AMS because it is not interested in finding a new worker anymore or when the firm can not be contacted by the AMS any longer.

[27]The portion of lapsed vacancies is about 13% of the total number of outflows.

all the spells related to social security institutions and to sick leave, unemployed or retired status, self employed and civil servants.

- We then match each individual hired in the vacancy database with the spell database through the worker id (*penr*). At this point for each vacancy we have a list of all the employment spells of the hired worker and we need to identify the spell that the vacancy refers to. To do that we consider only the spells starting around the closing date of the vacancy (-40, +90 days) and we drop all spells ending before the closing date of the vacancy.

- At this point we have a unique vacancy-spell match for most of the vacancies. For the remaining ones, we compare the pair AMS firm identifier (*btrnr*) and ASSD firm identifier (*benr*) with the "*cross-walk*" document provided by the *Bundesministerium fuer Arbeit, Soziales und Konsumentenschutz*[28]. For vacancies with at least a match with the cross-walk document we get the first one in chronological order among the matched ones, for the others we get the first in chronological order among all the identified spells.

At the end of this procedure we are able to have a pair vacancy-employment spell for 88% of the sub-sample of vacancies for which we observe a worker identifier.

The last step is getting the firm information, the worker information and the earnings. We do this by matching our database with employers database and workers database through the firm id (*benr*) and the worker id (*penr*) contained in the employment spell data. The final *vacancy-employer-employee* sample contains 987,271 observations for the period 1997-2013.

Table A.1 shows how the number of observations decreases at each of the steps above, while tables A.2 and A.3 compare the distribution of industries and occupations, respectively, across the sub-samples.

[28]This document matches the two identifiers, even though we tested the validity of the crossing with poor results, we believe that it contains some information that we can exploit to refine our matching algorithm.

Table A.1: Observation by subsamples

	Observations
1. Vacancies from September 1997 to December 2013 (after restrictions)	13,906,275
2. Outflows of Vacancies from September 1997 to December 2013	5,214,539
3. Outflows of Vacancies from Sept. 1997 to Dec. 2013 with Inflow date	5,096,403
(a) With non missing key variables (4,998,146)	
4. Hired by AMS from September 1997 to December 2013	1,169,203
5. Drop 414 duplicates	1,168,789
6. Merge with Employment Spells Database	1,042,794
7. Merge with Firms Database	1,041,094
8. Generate key variables, including occupations codes	1,016,843
9. Get worker personal information (gender, birthdate)	1,014,701
10. Get wages	987,271

Notes: This table shows the number of observations in the database at each stage of the sample construction for the period 1997-2013.

Source: AMS-ASSD data, own calculations.

Table A.2: Job Ads by Industry

Industry	Subsample			
	All Outflows (1)	Hired (2)	Hired AMS (3)	Matched (4)
Agriculture, hunting and forestry	0.24%	0.25%	0.17%	0.17%
Mining and quarrying	0.10%	0.10%	0.14%	0.15%
Manufacturing	10.47%	10.81%	14.16%	14.62%
Electricity, gas and water supply	0.09%	0.10%	0.09%	0.09%
Construction	7.23%	7.05%	8.73%	8.81%
Wholesale and retail trade	15.22%	15.82%	16.90%	17.37%
Hotels and restaurants	24.20%	25.39%	19.77%	19.26%
Transport, storage and communication	3.74%	3.87%	4.18%	4.26%
Financial intermediation	0.95%	0.83%	0.64%	0.65%
Real estate, renting and business activities	27.07%	25.44%	22.68%	21.92%
Public administration and defence; compulsory social security	1.99%	1.85%	2.38%	2.44%
Education	0.95%	0.86%	0.76%	0.77%
Health and social work	3.47%	3.47%	4.78%	4.88%
Other community, social and personal service activities	4.27%	4.14%	4.59%	4.60%

Notes: This table shows the share of job postings by broad categories of industries for four different subsamples. Period 1997-2013.

Source: AMS-ASSD data, own calculations.

Table A.3: Job Ads by Occupation

Occupation	Subsample			
	All Outflows (1)	Hired (2)	Hired AMS (3)	Matched (4)
Legislators, senior officials and managers	1.00%	0.99%	0.63%	0.65%
Professionals	1.96%	1.73%	0.87%	0.83%
Technicians and associate professionals	9.95%	9.31%	6.64%	6.82%
Clerks	6.79%	7.01%	8.20%	8.50%
Service workers and shop and market sales workers	33.89%	35.25%	30.52%	30.45%
Skilled agricultural and fishery workers	0.42%	0.42%	0.41%	0.41%
Craft and related trades workers	21.01%	19.88%	19.80%	20.06%
Plant and machine operators and assemblers	5.97%	6.06%	7.54%	7.67%
Elementary occupations	19.02%	19.34%	25.40%	24.61%

Notes: This table shows the share of job postings by broad categories of occupations for four different subsamples. Period 1997-2013.

Source: AMS-ASSD data, own calculations.

B Attenuation with (Non-)stereotypical Vacancies

We are using predicted SGP's rather than actual gender preferences, so equation (2) can be interpreted as the reduced-form from a two-stage-least-squares procedure in which we estimate the first stage models using only data from the pre-campaign period, and allow the effects of the endogenous variables to vary between the pre-campaign and post-campaign periods. To formalize this, let S_j^f and S_j^m represent dummies for actual SGP's in the pre-campaign period or *desired* SGP's in the post-campaign period (i.e., the preferences that employers would have stated if there was no effort to eliminate SGP's). Let C_j^f represent workplaces with a majority of women (F workplace), and C_j^f represents workplaces with a majority of men (M workplace). Assume that the true model generating outcome y is:

$$\begin{aligned} y_j = \alpha_0 + \alpha_1 S_j^f C_j^f + \alpha_2 S_j^f C_j^m + \alpha_3 S_j^m C_j^m + \alpha_4 S_j^m C_j^f \\ + \theta_1 S_j^f C_j^f Post_j + \theta_2 S_j^f C_j^m Post_j + \theta_3 S_j^m C_j^m Post_j + \theta_4 S_j^m C_j^f Post_j + X_j \gamma + \varepsilon_j \end{aligned} \tag{B.0.1}$$

This has the same form as equation (2) but relates the outcome to true (or desired) SGP's. Assume that the actual/desired SGP's are related to predicted preferences by a pair of simple models with constant coefficients between the pre- and post-campaign periods:

$$S_j^f = \pi_0 + \pi_1 D_j^f C_j^f + \pi_2 D_j^f C_j^m + \pi_3 D_j^m C_j^m + \pi_4 D_j^m D_j^f + X_j \pi_x + \xi_j^f \tag{B.0.2}$$

$$S_j^m = \psi_0 + \psi_1 D_j^f C_j^f + \psi_2 D_j^f C_j^m + \psi_3 D_j^m C_j^m + \psi_4 D_j^m C_j^f + X_j \psi_x + \xi_j^m \tag{B.0.3}$$

where ξ_j^f, ξ_j^m are prediction errors. Here π_1 and $\pi_2$represent the increment in the probability of an actual female SGP if the vacancy (π_1 for the stereotypical, and π_2 for the non-stereotypical vacancies) relative to the omitted category of a prediction of no SGP. π_3 and π_4 represent the increment in the probability of an actual female SGP if the vacancy has a predicted male stereotypical SGP (π_3 for the stereotypical, and π_4 for the non-stereotypical vacancy), again relative to the case where it is predicted to have no SGP. Thus we expect π_1 and π_2 to be positive, whereas $\pi_3$and π_4 will be negative. Similar reasoning suggests that ψ_3 and ψ_4 will be positive and $\psi_1$and ψ_2 will be negative.

Combining equation (B.0.1) with (B.0.2) and (B.0.3) shows that the difference-of-differences coefficients in (1) are:

$$\lambda_1 = \theta_1\pi_1 + \theta_4\psi_1 \tag{B.0.4}$$

$$\lambda_2 = \theta_2\pi_2 + \theta_3\psi_2 \tag{B.0.5}$$

$$\lambda_3 = \theta_3\psi_3 + \theta_2\pi_3 \tag{B.0.6}$$

$$\lambda_4 = \theta_4\psi_4 + \theta_1\pi_4 \tag{B.0.7}$$

Notice that if we ignore ψ_1 and ψ_2,then λ_1 is an attenuated version of θ_1 and λ_2 is an attenuated version of ψ_2, where the attenuation factors reflect the fractions of predicted vacancies with female stereotypical (π_1) or non-stereotypical (π_2) preferences (conditional on the $X's$). Also, if we ignore π_3 and π_4,then λ_3 is an attenuated version of θ_3 and λ_4 is an attenuated version of θ_4, where the attenuation factors reflect the fractions of predicted vacancies with male stereotypical (ψ_3) or non-stereotypical (ψ_3) preferences that actually have these preferences (conditional on the $X's$). More generally we would expect ψ_1, ψ_2 π_3, and π_4 to be small in magnitude (though in each case negative), so the intuition of the benchmark case remains true.

Columns 2 and 4 of Table 4 present estimates of equations (B.0.3) and (B.0.3) using the observed SGP's in the 1999-2003 period and predictions from our first (leave-out-mean based) classification model. We see that $\pi_1 = 0.171$, $\pi_2 = 0.230$ and, while $\psi_1 = -0.012$ and $\psi_2 = -0.082$. Thus, controlling for industry, occupation, and the firm's lagged gender composition, having a stereotypical female predicted SGP raises the probability of an SGP of that gender by 17 percentage points relative to a vacancy that is predicted to have no SGP, while having a male non-stereotypical SGP of the opposite gender lowers the probability by 1 percentage point relative to the no-SGP base.

C Prediction Quality

Figure C.1: Prediction Quality

(a) Regression: With LOO

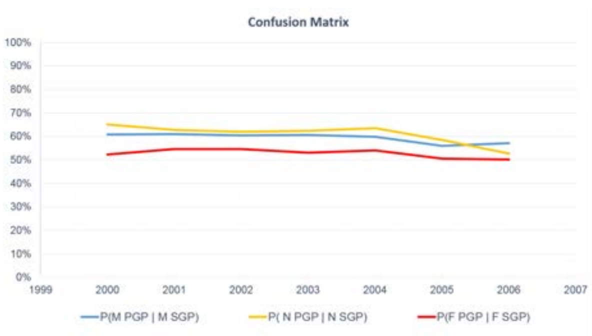

(b) Regression: Without LOO

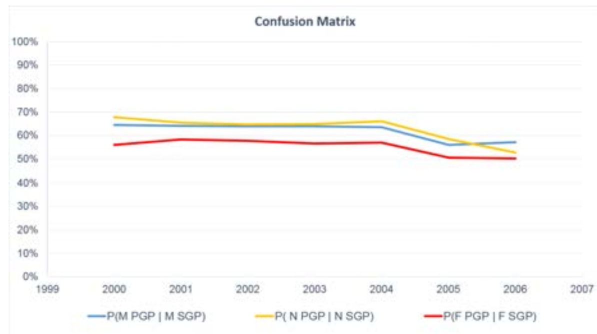

(c) Random Forest: With KFOLD

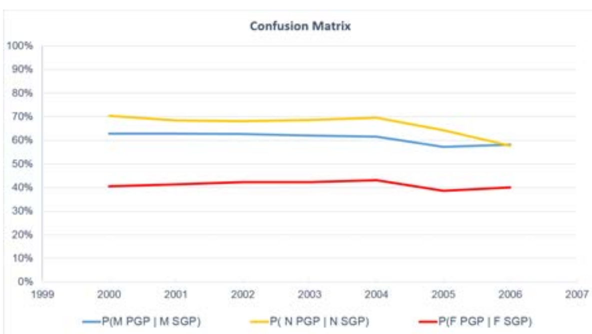

(d) Random Forest: Without KFOLD

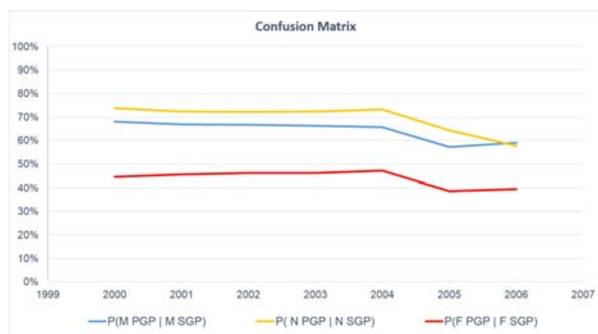

Notes: This figure shows the share of Stated Gender Preferences (SGP) that are correctly predicted by our Predicted Gender Preference (PGP), by year. The classification period is 2000-2004. Classification in panels (a) and (b) is performed through a fully saturated regression model, in panel (c) and (d) through the random forest model. In panel (b) we correct estimates through the leaveone-out mean mthod, in panel (c) we implement a rotating K-fold 1/4 vs 3/4 alghorithm.

Source: AMS-ASSD data, own calculations.

Table C.1: Predicted vs Observed SGP

Year	Observed	Prediction			
		Male PGP	**No PGP**	**Female PGP**	**Total**
2000	**Male SGP**	8,989	5,504	266	14,759
	No SGP	3,592	13,636	3,734	20,962
	Female SGP	237	5,328	6,130	11,695
	Total	12,818	24,468	10,130	47,416
2001	**Male SGP**	6,439	3,933	187	10,559
	No SGP	3,568	12,639	3,905	20,112
	Female SGP	127	3,942	4,896	8,965
	Total	10,134	20,514	8,988	39,636
2002	**Male SGP**	6,089	3,781	189	10,059
	No SGP	4,518	14,672	4,487	23,677
	Female SGP	148	3,833	4,782	8,763
	Total	10,755	22,286	9,458	42,499
2003	**Male SGP**	6,200	3,847	186	10,233
	No SGP	4,701	15,066	4,379	24,146
	Female SGP	147	4,279	5,010	9,436
	Total	11,048	23,192	9,575	43,815
2004	**Male SGP**	5,927	3,728	246	9,901
	No SGP	4,696	16,134	4,563	25,393
	Female SGP	139	3,951	4,818	8,908
	Total	10,762	23,813	9,627	44,202
2005	**Male SGP**	4,753	3,563	177	8,493
	No SGP	7,184	18,780	6,167	32,131
	Female SGP	170	3,742	4,011	7,923
	Total	12,107	26,085	10,355	48,547
2006	**Male SGP**	1,181	830	54	2,065
	No SGP	12,881	24,707	9,283	46,871
	Female SGP	56	736	798	1,590
	Total	14,118	26,273	10,135	50,526
2007	**Male SGP**	313	250	9	572
	No SGP	11,877	26,094	10,394	48,365
	Female SGP	9	233	230	472
	Total	12,199	26,577	10,633	49,409
2008	**Male SGP**	22	33	1	56
	No SGP	11,382	26,894	10,345	48,621
	Female SGP	7	45	32	84
	Total	11,411	26,972	10,378	48,761
2009	**Male SGP**	34	14	1	49
	No SGP	9,874	22,679	9,345	41,898
	Female SGP		42	30	72
	Total	9,908	22,735	9,376	42,019
2010	**Male SGP**	31	37		68
	No SGP	10,755	23,632	9,247	43,634
	Female SGP	2	27	18	47
	Total	10,788	23,696	9,265	43,749

D Additional Results

Figure D.1: Female Hiring - Weighted Regression Results

A. Leave-Out-Mean

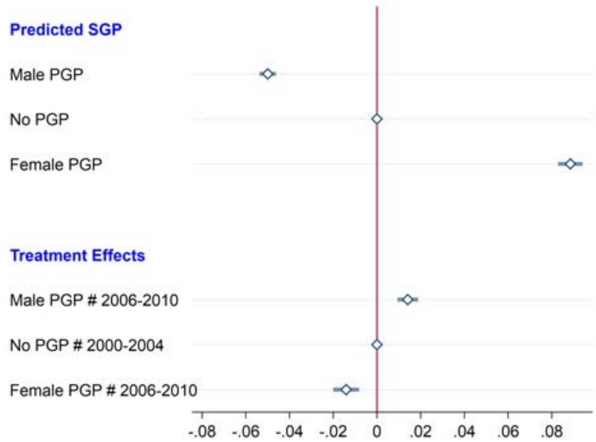

Notes: This figure reports the estimation results capturing the effect of eliminating stated gender preferences on the hiring of females. The regressions follow the model in equation 1. Diamond refers to the point estimates, horizontal lines show the 95% confidence intervals. Pre-campaign period refers to vacancies posted in 2000-2004, treatment effects refer to vacancies posted in 2006-2010. Vacancies posted in 2005 are dropped from the regression. Controls include industry, occupation, lagged female share in the workplace and year fixed effects. Observations are weighted using representative sample regression weights. Classification uses regression.

Source: AMS-ASSD data, own calculations.

Figure D.2: Effects Vacancy Filling (controlling for Hire)

A. Baseline Model

B. Stereotypical vs Non-Stereotypical Vacancies

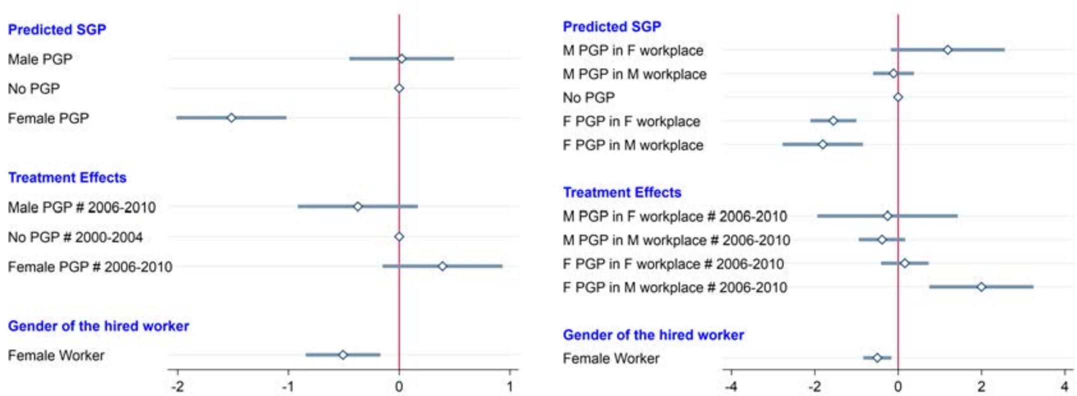

Notes: This figure reports the estimation results capturing the effect of eliminating stated gender preferences on vacancy filling times. In panel A the the model follows the baseline specification in equation 1. In panel B the model follows the specification in equation 8, distinguishing between Stereotypical (M PGP in M Workplace and F PGP in F Workplace) and Non-Stereotypical (M PGP in F Workplace and F PGP in M Workplace) vacancies. Diamond refers to the point estimates, horizontal lines show the 95% confidence intervals. Pre-campaign period refers to vacancies posted in 2000-2004, treatment effects refer to vacancies posted in 2006-2010. Vacancies posted in 2005 are dropped from the regression. Controls include previous wage, gender, age, as well as industry, occupation, lagged female share in the workplace and year fixed effects.

Source: AMS-ASSD data, own calculations.

Figure D.3: Effects on Wages and Job Duration: Composition Effects (NOT controlling for Hire)

A. Job Duration - Baseline Model

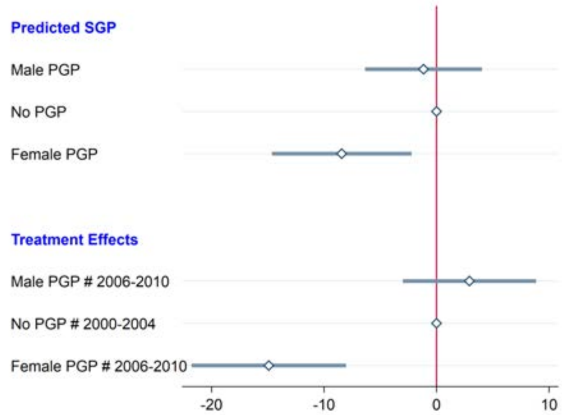

B. Job Duration - Stereotypical vs Non-Stereotypical Vacancies

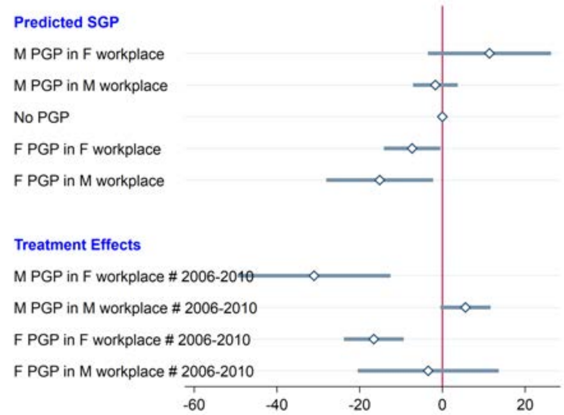

C. Wages - Baseline Model

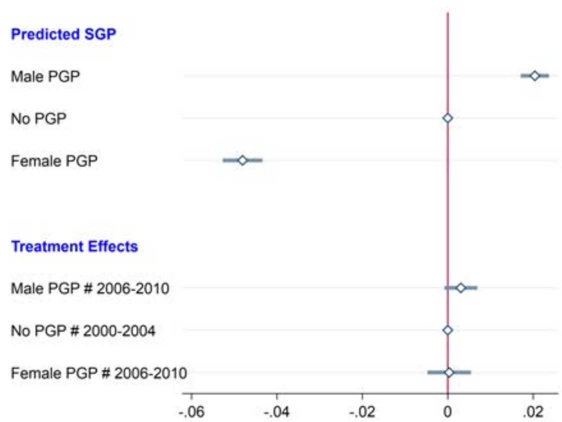

D. Wages - Stereotypical vs Non-Stereotypical Vacancies

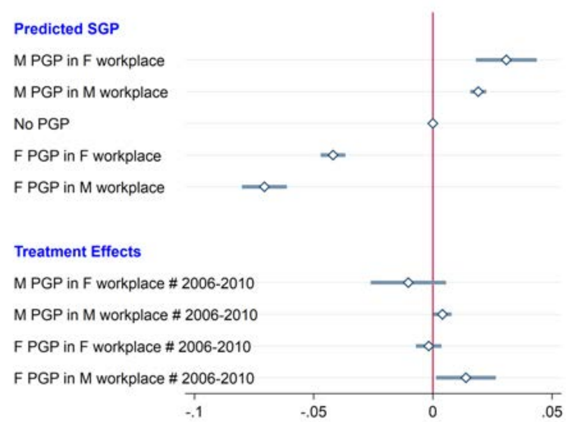

Notes: This figure reports the estimation results capturing the effect of eliminating stated gender preferences on completed job durations (Panels A, B) and wages on the newly filled job (Panels C, D). The models follows the specification in equation 8, distinguishing between Stereotypical (M PGP in M Workplace and F PGP in F Workplace) and Non-Stereotypical (M PGP in F Workplace and F PGP in M Workplace) vacancies. Diamond refers to the point estimates, horizontal lines show the 95% confidence intervals. Pre-campaign period refers to vacancies posted in 2000-2004, treatment effects refer to vacancies posted in 2006-2010. Vacancies posted in 2005 are dropped from the regression. Controls include industry, occupation, lagged female share in the workplace and year fixed effects.

Source: AMS-ASSD data, own calculations.

Figure D.4: Female Hiring (Difference in Differences Results)

Baseline + Industry-X-Year FE

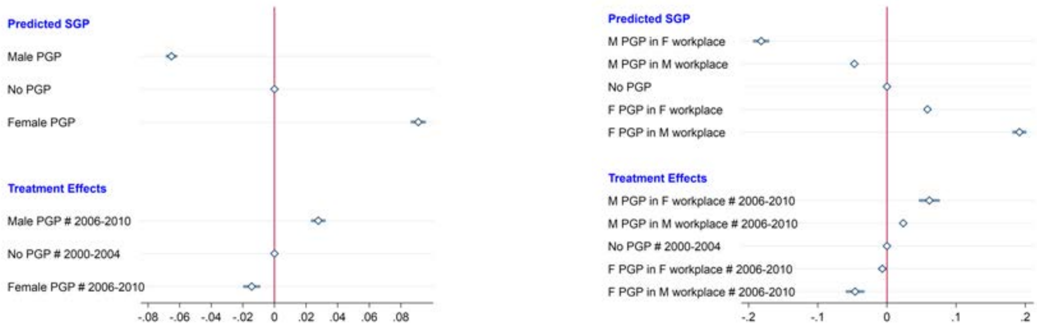

Baseline + Occupation-X-Year FE

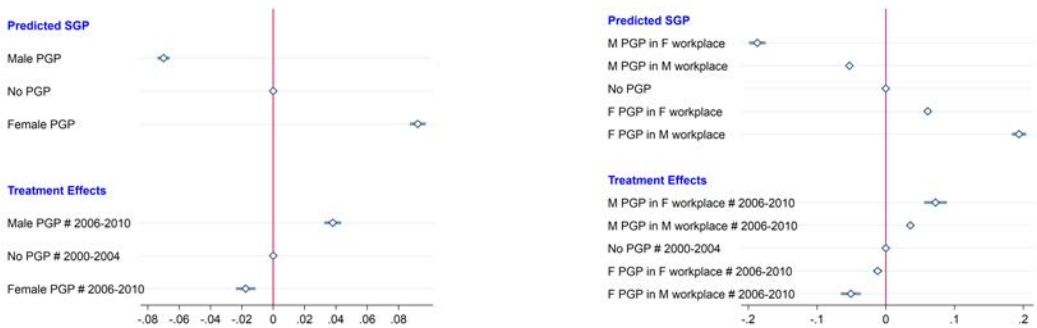

Baseline + Industry-X-Year FE + Occupation-X-Year FE

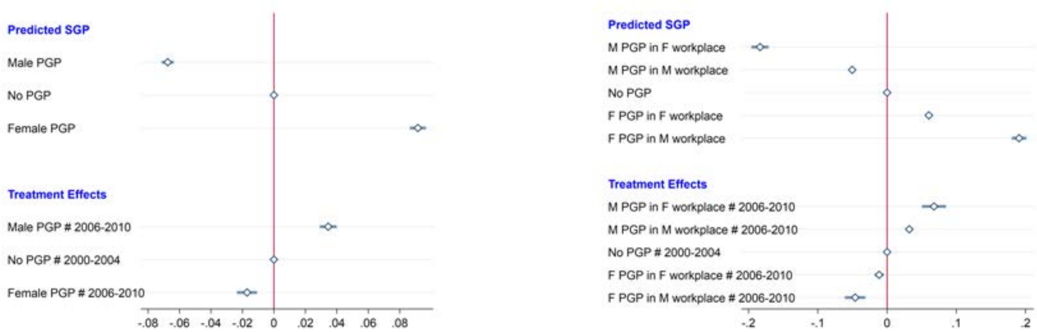

Baseline + Industry×Year Effects +
Occupation×Year Effects + FemaleShare×Year Effects

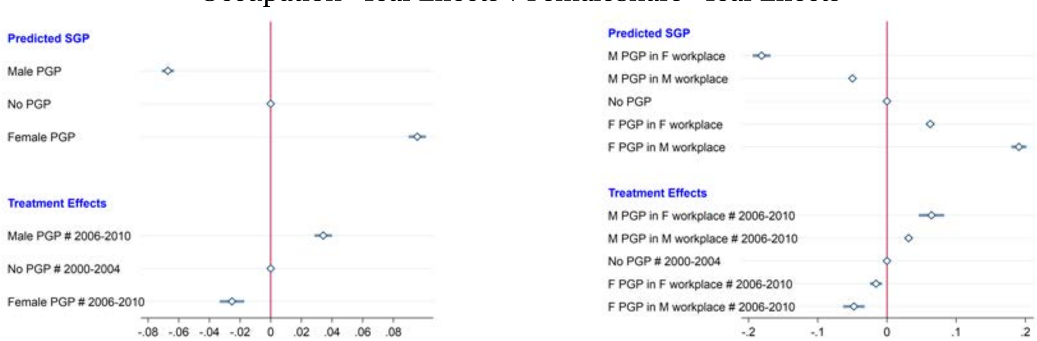

Notes: This figure reports the estimation results capturing the effect of eliminating stated gender preferences on female hiring controlling for time variant industry, occupation and female share FE. The models follows the specification in equation1 in the left column and 8 in the right column. Diamond refers to the point estimates, horizontal lines show the 95% confidence intervals. Pre-campaign period refers to vacancies posted in 2000-2004, treatment effects refer to vacancies posted in 2006-2010. Vacancies posted in 2005 are dropped.

Source: AMS-ASSD data, own calculations.

Figure D.5: Effects on Job Duration (Event History Results)

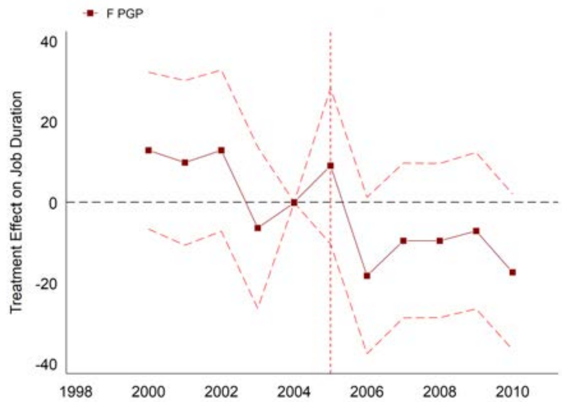

Note: This figure reports the regression coefficients capturing the effect of eliminating stated gender preferences on completed job durations. Coefficients of the interaction term between year and indicators for job ads classified as advertising for women ("Female PGP"), are reported. Dotted lines show the 95% confidence intervals. Controls include industry, occupation, lagged female share in the workplace and year fixed effects.

*Source:*AMS-ASSD data, own calculations. calculations

Figure D.6: Effects on Workforce Diversity ($e_j \equiv H_j(1-C_j^f)+(1-H_j)C_j^f$)

A. Baseline Model

B. Stereotypical vs Non-Stereotypical Vacancies

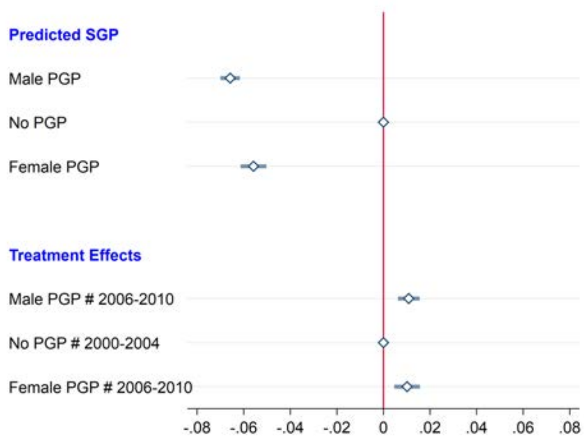

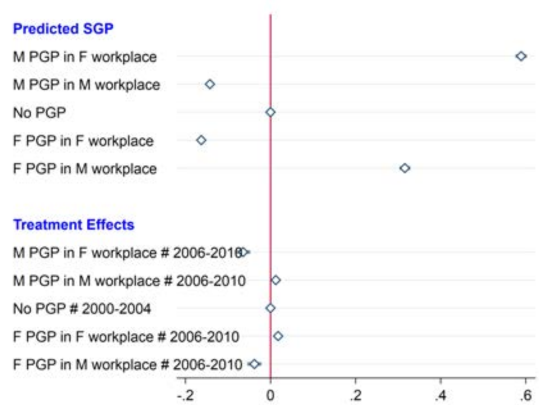

Notes: This figure reports the estimation results capturing the effect of eliminating stated gender preferences on workplace composition using the discrete indicator. In panel A the model follows the baseline specification in equation 1. In panel B the model regression follows the specification in equation 8, distinguishing between Stereotypical (Predicted M SGP in M Workplace and F SGP in F Workplace) and Non-Stereotypical (Predicted M SGP in F Workplace and F SGP in M Workplace) vacancies. Diamond refers to the point estimates, horizontal lines show the 95% confidence intervals. Pre-campaign period refers to vacancies posted in 2000-2004, treatment effects refer to vacancies posted in 2006-2010. Vacancies posted in 2005 are dropped from the regression. Controls include industry, occupation, lagged female share in the workplace and year fixed effects.

Source: AMS-ASSD data, own calculations.

E TABLES MAIN RESULTS

Table E.1: Effect of Eliminating Stated Gender Preferences on Female Hiring

	OLS Estimation				
Dependent Variable:	*Female Hiring*			*Workplace Diversity Index*	
	(1)	(2)	(3)	(4)	(5)
No PGP	omitted	omitted	omitted	omitted	omitted
Male PGP	-0.065***	-0.050***		-0.042***	
	(0.002)	(0.002)		(0.001)	
Female PGP	0.089***	0.089***		-0.034***	
	(0.002)	(0.002)		(0.002)	
Male PGP × 2006-2010	0.025***	0.013***		0.007***	
	(0.002)	(0.002)		(0.001)	
Female PGP × 2006-2010	-0.010***	-0.016***		0.006***	
	(0.002)	(0.003)		(0.002)	
M PGP in F workplace			-0.187***		0.211***
[Non-stereotypical Male PGP]			(0.006)		(0.004)
M PGP in M workplace			-0.047***		-0.070***
[Stereotypical Male PGP]			(0.002)		(0.001)
F PGP in F workplace			0.057***		-0.090***
[Stereotypical Female PGP]			(0.003)		(0.002)
F PGP in M workplace			0.193***		0.157***
[Non-stereotypical Female PGP]			(0.005)		(0.003)
M PGP in F workplace × 2006-2010			0.068***		-0.023***
[Non-stereotypical Male PGP × 2006-2010]			(0.008)		(0.005)
M PGP in M workplace × 2006-2010			0.023***		0.008***
[Stereotypical Male PGP × 2006-2010]			(0.002)		(0.001)
F PGP in F workplace × 2006-2010			-0.004		0.011***
[Stereotypical Female PGP × 2006-2010]			(0.003)		(0.002)
F PGP in M workplace × 2006-2010			-0.049***		-0.024***
[Non-stereotypical Female PGP × 2006-2010]			(0.007)		(0.004)
Observations	452029	452029	452029	452029	452029

Notes: OLS estimation results capturing the effect of eliminating stated gender preferences on the hiring of females (columns 1-3), and on the changes in workplace composition (columns 4-5). Regressions in columns 1,2 and 4 follow the specification in equation 1. Regressions in columns 3 and 5 follow the specification in equation 8. Classification in column 2 uses random forest. Industry, occupation, lagged female share in the workplace and year fixed effects are included as controls. Pre-campaign period refers to vacancies posted in 2000-2004, treatment effects refer to vacancies posted in 2006-2010. Vacancies posted in 2004 are dropped from the regression. Beta coefficients reported and robust standard errors in parentheses. *** significant at 1 percent, ** significant at 5 percent, * significant at 10 percent.

Source: AMS-ASSD data, own calculations.

Table E.2: Effect of Eliminating Stated Gender Preferences on Other Outcomes

	OLS Estimation					
Dependent Variable:	*Vacancy Filling*		*Daily Wage*		*Job Duration*	
	(1)	(2)	(3)	(4)	(5)	(6)
No PGP	omitted	omitted	omitted	omitted	omitted	omitted
Male PGP	-0.180		0.020^{***}		-1.153	
	(0.198)		(0.002)		(2.653)	
Female PGP	-1.409^{***}		-0.048^{***}		-8.429^{***}	
	(0.210)		(0.002)		(3.169)	
Male PGP × 2006-2010	0.024		0.003		2.921	
	(0.228)		(0.002)		(3.024)	
Female PGP × 2006-2010	0.538^{**}		0.000		-14.915^{***}	
	(0.232)		(0.003)		(3.505)	
M PGP in F workplace		0.686		0.031^{***}		11.378
[Non-stereotypical Male PGP]		(0.546)		(0.007)		(7.595)
M PGP in M workplace		-0.292		0.019^{***}		-1.693
[Stereotypical Male PGP]		(0.207)		(0.002)		(2.762)
F PGP in F workplace		-1.396^{***}		-0.042^{***}		-7.318^{**}
[Stereotypical Female PGP]		(0.233)		(0.003)		(3.487)
F PGP in M workplace		-1.839^{***}		-0.071^{***}		-15.152^{**}
[Non-stereotypical Female PGP]		(0.414)		(0.005)		(6.592)
M PGP in F workplace × 2006-2010		0.091		-0.010		-31.024^{***}
[Non-stereotypical Male PGP × 2006-2010]		(0.677)		(0.008)		(9.427)
M PGP in M workplace × 2006-2010		0.014		0.004^{**}		5.581^{*}
[Stereotypical Male PGP × 2006-2010]		(0.234)		(0.002)		(3.099)
F PGP in F workplace × 2006-2010		0.313		-0.002		-16.587^{***}
[Stereotypical Female PGP × 2006-2010]		(0.244)		(0.003)		(3.690)
F PGP in M workplace × 2006-2010		2.088^{***}		0.014^{**}		-3.429
[Non-stereotypical Female PGP × 2006-2010]		(0.541)		(0.006)		(8.701)
Observations	452029	452029	452027	452027	452029	452029

Notes: OLS estimation results capturing the effect of eliminating stated gender preferences on vacancy filling time (columns 1-2), on the wage of the newly filled job (columns 3-4) and on completed duration of the newly filled job (columns 5-6). Regressions in columns 1, 3 and 5 follow the specification of equation 1. Regressions in columns 2, 4 and 6 follow the specification in equation 8. Industry, occupation, lagged female share in the workplace and year fixed effects are included as controls. Pre-campaign period refers to vacancies posted in 2000-2004, treatment effects refer to vacancies posted in 2006-2010. Vacancies posted in 2004 are dropped from the regression. Previous wage available only for 68 percent of the sample. Beta coefficients reported and robust standard errors in parentheses. *** significant at 1 percent, ** significant at 5 percent, * significant at 10 percent.

Source: AMS-ASSD data, own calculations.

Table E.3: Effect of Eliminating Stated Gender Preferences on Other Outcomes, with controls for the Characteristics of the Hired Worker

	OLS Estimation					
Dependent Variable:	*Vacancy Filling*		*Daily Wage*		*Job Duration*	
	(1)	(2)	(3)	(4)	(5)	(6)
No PGP	omitted	omitted	omitted	omitted	omitted	omitted
Male PGP	0.023		0.006***		6.082*	
	(0.242)		(0.002)		(3.462)	
Female PGP	-1.515***		-0.030***		-9.899**	
	(0.254)		(0.003)		(4.024)	
Male PGP × 2006-2010	-0.373		0.005**		0.855	
	(0.277)		(0.002)		(3.931)	
Female PGP × 2006-2010	0.391		0.000		-18.109***	
	(0.277)		(0.003)		(4.396)	
M PGP in F workplace		1.190*		-0.021**		28.680***
[Non-stereotypical Male PGP]		(0.698)		(0.009)		(10.542)
M PGP in M workplace		-0.111		0.008***		4.456
[Stereotypical Male PGP]		(0.251)		(0.002)		(3.583)
F PGP in F workplace		-1.556***		-0.031***		-6.297
[Stereotypical Female PGP]		(0.281)		(0.003)		(4.433)
F PGP in M workplace		-1.809***		-0.026***		-25.454***
[Non-stereotypical Female PGP]		(0.492)		(0.005)		(8.261)
M PGP in F workplace × 2006-2010		-0.255		0.023**		-39.175***
[Non-stereotypical Male PGP × 2006-2010]		(0.860)		(0.011)		(13.037)
M PGP in M workplace × 2006-2010		-0.388		0.004*		3.598
[Stereotypical Male PGP × 2006-2010]		(0.284)		(0.002)		(4.016)
F PGP in F workplace × 2006-2010		0.160		-0.000		-20.731***
[Stereotypical Female PGP × 2006-2010]		(0.291)		(0.003)		(4.628)
F PGP in M workplace × 2006-2010		1.996***		0.002		-0.907
[Non-stereotypical Female PGP × 2006-2010]		(0.638)		(0.007)		(10.721)
Female Worker	-0.507***	-0.500***	-0.179***	-0.179***	45.937***	46.199***
	(0.172)	(0.173)	(0.002)	(0.002)	(2.426)	(2.428)
Observations	308075	308075	308075	308075	308075	308075

Notes: OLS estimation results capturing the effect of eliminating stated gender preferences on vacancy filling time (columns 1-2), on the wage for the newly created job (columns 3-4), and on the duration of the newly created job (columns 5-6). Regressions in columns 1, 3 and 5 follow the specification of equation 1. Regressions in columns 2, 4 and 6 follow the specification of equation 8. Previous gender, wage and age of the hire, as well as industry, occupation, lagged female share in the workplace and year fixed effects are included as controls. Pre-campaign period refers to vacancies posted in 2000-2004, treatment effects refer to vacancies posted in 2006-2010. Vacancies posted in 2004 are dropped from the regression. Previous wage available only for 68 percent of the sample. Beta coefficients reported and robust standard errors in parentheses. *** significant at 1 percent, ** significant at 5 percent, * significant at 10 percent.

Source: AMS-ASSD data, own calculations.