The Gender Application Gap: Do men and women apply for the same jobs?*

Jonas Fluchtmann† Anita M. Glenny ‡ Nikolaj Harmon § Jonas Maibom ¶

January 10, 2023

Abstract

Men and women tend to hold different jobs. Are these differences present already in the types of jobs men and women apply for? Using administrative data on job applications made by the universe of Danish UI recipients, we provide evidence on gender differences in applied-for jobs for the broader labor market. Across a range of job characteristics, we find large gender gaps in the share of applications going to different types of jobs even among observationally similar men and women. In a standard decomposition, gender differences in applications can explain more than 70 percent of the residual gender wage gap.

Keywords: job search, wage decomposition, firm wage premium, gender earnings gap

JEL: E24, J29, J31, J71

*This paper has benefited from numerous questions and comments from participants at several seminars, conferences and workshops. We further thank Bas Van Der Klaauw, Ioana Marinescu, Morten Bennedsen, Michael Rosholm, Alexander Koch, Alexander Sebald, Steffen Altmann, Michèle Belot, Phillipp Kircher, Peter Kuhn, Roland Rathelot, Daphne Skandalis and Willem Adema for their input in this process. The paper was previously circulated under the title "Gender Gaps in Job Applications and Hiring Outcomes".

†Aarhus University; present: OECD; e-mail: jonas.fluchtmann@oecd.org; The views expressed here are those of the authors and do not necessarily reflect the offical views of the OECD or its member countries.

‡Aarhus University; e-mail: glenny@econ.au.dk

§University of Copenhagen; e-mail: nikolaj.harmon@econ.ku.dk

¶Aarhus University, IZA and ECONAU; e-mail: maibom@econ.au.dk; Jonas Maibom acknowledges support from the Carlsberg Foundation (CF15-0647) and the ECONAU project database (Department of Economics and Business Economics, Aarhus University).

1 Introduction

In most labor markets, there are large and persistent differences in the types of jobs men and women tend to hold. This is true both in terms of job characteristics, such as industries or occupations, as well as in terms of the typical wage level of the job. In most cases, women tend to hold jobs that pay systematically less.[1] These gender differences in job outcomes may arise at different possible stages of the labor market: They may arise through gender differences in wage bargaining or promotion rates in ongoing employment relationships. They may arise earlier at the hiring stage, through gender differences in how firms select among job applicants. Alternatively, gender differences in job outcomes may arise already at the job application stage when male and female job seekers decide which jobs to apply for.

In this paper, we track gender gaps in job outcomes back to the job application stage and ask a simple question: To what extent are gender gaps present already in the types of jobs men and women apply for? The answer to this question has implications for ongoing debates about the origin of gender gaps and for policies that aim to promote gender equality. Outside of some very specific settings, however, empirical evidence on gender gaps in applied-for jobs is virtually nonexistent. This is primarily due to data constraints: Detailed data on job applications is typically only available for a small and/or selected subset of individuals and jobs. The ability to applications with rich background characteristics and with actual job outcomes is even more limited.

We overcome these data constraints by exploiting a novel administrative data source that contains job application data for the universe of Danish unemployment insurance (UI) recipients. Since 2015, all UI recipients in Denmark have been required to systematically register applied-for jobs on an online web portal run by the Danish employment agency. Based on a unique person identifier, we are able to link the resulting data set of applied-for jobs with worker characteristics and hiring outcomes. Based on firm and job information in the job application data, we are further able to link each job application to administrative data on firms and jobs. The resulting data set allows us to examine the characteristics of the firms and jobs that male and female job seekers apply to and also compare them to the characteristics of the firms and jobs they eventually end up in.

The primary job characteristics we focus on are a set of standard wage determinants, including

[1] For a broad overview of the recent literature on gender differences in earnings and wages, see e.g. Blau and Kahn (2017) or Olivetti and Petrongolo (2016).

the industry and occupation of the job. In addition, a particular advantage of our linked administrative data is that we can construct measures of firm type to examine whether an applied-for job is at a high- or low-wage firm. We do this based on the firm fixed effects from a two-way fixed effects regression for wages (Abowd et al., 1999; Card et al., 2013). Finally, we use administrative data on actual wage payments to compute a predicted wage for each type of job given its characteristics. We refer to this as the *typical wage* for each type of job and use it to measure gender differences in the typical wage levels of applied-for jobs.

The first part of our analysis documents gender differences in job applications. We split jobs according to the different job characteristics and then compute differences in the share of applications that men and women send to each type of job. We refer to these differences as *gender application gaps*. Across all the job characteristics we focus on, gender application gaps are substantial and remain also when we condition on a very rich set of individual labor market observables. Gender application gaps tend to closely mirror gender gaps in job outcomes. Across most characteristics, women tend to target jobs that pay lower wages. These gender differences in targeted job types therefore add up to a stark gender gap in the typical wages of the jobs men and women apply for: After conditioning on individual observables, women apply to jobs with a typical wage that is 4.5 percent lower than men.

Next, we use the linked nature of our data to quantify how much of the observed gender gaps in job outcomes may be explained by the gender application gaps that we find. Because our data contains information on both job applications and job outcomes for the same individuals, we can apply standard decomposition methods from the wage gap and inequality literature. We use a sequential implementation of the semi-parametric decomposition of DiNardo et al. (1996) which allows us to focus on the part of the gender gap that is not explained by individual labor market observables (as in Butcher and DiNardo, 2002 and Altonji et al., 2012). Specifically, we first condition out individuals' observable characteristics and then examine how much of the residual gender gaps in job outcomes can be explained by gender differences in where men and women apply.

The decomposition shows that differences in applied-for jobs are able to explain 79 percent of the residual gender gap in typical wages and 72 percent of the residual gender gap in realized starting wages. Results for the gender gap in the employing firm's type are similar. For gender segregation across industries and occupations, however, the results are more varied: For some industries and occupations, gender differences in applied-for jobs can explain virtually all of the observed gender segregation. For others, the observed gender segregation appears largely unrelated to application behavior. Drawing on previous evidence regarding industry and occupation-specific discrimination, we find that this may reflect the direct effects of different forms of hiring discrimination; after accounting for application behavior, our decomposition shows women to be underrepresented in industries and occupations where previous work have found discrimination against women, while the reverse is true in industries and occupations where previous work have found discrimination in favor of women.

Overall, our results suggest that gender gaps in job outcomes are for the most part present already at the application stage. We therefore finish the paper by discussing possible explanation for *why* men and women are applying to such systematically different jobs. We also present some suggestive evidence on this question, although we note that our observational data is less well suited to tease out mechanisms. Possible explanations include the possibility that women rationally choose to apply less for job types where they correctly believe they have a lower likelihood of being hired (so-called "self-fulfilling discrimination", Lundberg and Startz, 1983; Glover et al., 2017; Coate and Loury, 1993); the possibility that men and women differ in their degree of (over)confidence, their information, risk preferences or personality (Cortes et al. 2020; Flinn et al. 2020); and the possibility that men and women differ in their valuation of non-wage job characteristics (Le Barbanchon et al., 2021; Wiswall and Zafar, 2018; Maestas et al., 2019; Hotz et al., 2018). Within the limitations of our data, we are only able to find evidence for the latter of these explanations: Women indeed target systematically different non-wage job characteristics by sending more of their applications to jobs that are part-time, that involve a shorter commute and that are at more family-friendly firms. These jobs also tend to pay lower wages. We also find some support for the idea that these differences in job application behavior may be partly related to motherhood (Kleven et al., 2019; Hotz et al., 2018; Lundborg et al., 2017). At the same time, however, we can not rule out that other explanations contribute to our main results.

Our paper is directly motivated by the large literature on gender gaps in the labor market (see Blau and Kahn (2017) or Olivetti and Petrongolo (2016) for an overview). A pervasive finding in this literature is that the types of jobs that men and women hold are important for the overall gender gap in earnings (see Gallen et al. (2019) for Danish evidence). Moreover, recent work has also emphasized differences in the types of firms that men and women work at (Card et al., 2016). This motivates our focus on tracking gender differences in job outcomes back through the job search and job application process. It also motivates our particular focus on the firm type.

Relative to existing work using micro data on job search, a major contribution of our paper lies in the coverage of the data we use. This is true both in terms of of the covered individuals and the covered job search activities. In terms of individuals, much of the previous work study gender differences in job search study only among users of particular online platforms (e.g. Kuhn et al. (2018), Gee (2019), Banfi et al. (2019) and Rousille (2021)), applicants to a particular job or career (e.g. Flory et al. (2014) and Samek (2019)), or survey respondents from a single student body (e.g. Barbulescu and Bidwell (2013) and Cortes et al. (2020)). In contrast, our data allow us to study the universe of UI recipients. This alleviates many concerns about differential selection by gender into surveys or particular search platforms. It also allows us to provide evidence on gender gaps for the broader labor market: UI recipients constitute a large and well-defined subgroup of workers and are quantitatively very important for the overall labor market.[2]

Our data also stands out by containing actual applied-for jobs of all types. This is in contrast to previous work using self-reported information about which jobs workers would be willing to accept, or using information from a systematically selected subset of actual applied-for jobs. In earlier work on gender, Eriksson and Lagerström (2012) use data from the job website of the Swedish employment agency to show that female UI recipients report being interested in jobs involving shorter commutes than men, while Caliendo et al. (2017) use survey data from Germany to show that female UI recipients report lower reservation wages. In work concurrent with our own, Le Barbanchon et al. (2021) replicate both of these findings using data on UI recipients' self-reported job preferences from the French employment agency and apply a novel structural framework to show that gender differences in commuting preferences can explain about 10 percent of the gender wage gap.

Most closely related to our application data, both Marinescu and Skandalis (2021) and supplementary results from Le Barbanchon et al. (2021) use data on job applications sent within the specific search platform of the French employment agency. By being limited to applications and

[2]Hires out of unemployment cover about half of all new hires in the Danish labor market and roughly half of a given cohort receive UI at some point during their labor market career.

jobs on this specific platform however, these data only cover a relatively small and selected subset of applied-for jobs.[3] This is substantially different from our data, in which UI recipients are asked to register applied-for jobs from all sources. Indeed, based on a range of validity checks we estimate that our application data cover between 69 and 80 percent of all applied-for jobs and that the covered subset is highly representative. Combined with the fact that our linked data contain information on a wide range of job characteristics, this makes our data uniquely suited to examine differences in the types of jobs workers apply for.

Our results have implications for the design and evaluation of policies to promote gender equality. Chiefly, our results imply that to close the gender gap in job outcomes, policy initiatives must create substantial changes in application behavior. This has particular implications for commonly-discussed policies - such as gender quotas or bias awareness training - whose first order effect is not to change application behavior but to change how women are treated in the hiring process. Our results suggest that the effectiveness of such policies will depend crucially on whether men and women internalize their effects on hiring and adapt their job application behavior in response. One implication of this is that the full effect of such policies may materialize only after some time has passed and beliefs have adapted. Another implication is that the effect of such policies may depend on how salient their implementation is for job seekers.

Finally, we note that our paper is related to a large empirical literature that estimates gender discrimination in the hiring process using either natural (e.g. Goldin and Rouse, 2000) or controlled experiments (see e.g. Neumark, 2004; Neumark, 2018; Rich, 2014; Riach and Rich, 2002). The central focus in this literature has been to estimate how the probability of being hired (or interviewed) differs across similar men and women when they apply for the same job. This is directly complementary to the present paper which documents to what extent men and women in fact tend to apply to the same types of jobs.

[3] Marinescu and Skandalis (2021) report that about one-fifth of all unemployed workers send any applications through the French PES platform and that 5 percent of new hires from unemployment appear to have stemmed from an application in the PES data. In the analysis sample of Le Barbanchon et al. (2021), less than 25 percent of unemployed workers find a job through the PES platform and 44 percent of vacancies report the minimum wage as their posted wage. Importantly, the PES data offers other advantages however. First, the fact that the PES data contain direct information about reservation wages and reservation commutes is key for the structural framework developed in Le Barbanchon et al. (2021). Second, because of the significant variation in usage, the number of applications sent on the PES platform turn out to be a useful for proxy for search effort (see Marinescu and Skandalis (2021)). Inferring search effort from the Danish Joblog data is more difficult as the number of registered application appears mainly driven by the reporting requirements (see Section 3.3).

2 Aim and roadmap of the empirical analysis

To frame our empirical analysis and its implications, consider two unemployed workers who are searching for a job. One is male and the other is female. Otherwise the workers have the same labor market observables including previous employment trajectories. Given the well-documented gender gap, we would expect these two workers to have very different job outcomes, with the male worker ending up in a better paid job.[4] Our empirical analysis asks to what extent there are differences already in the jobs these workers apply for during the job search process. The answer to this question is informative about why and where gender gaps exist.

At one extreme, it is possible that gender differences in applied-for jobs are negligible and that gender gaps in job outcomes occur only as a direct consequence of employers treating male and female applicants differently. This would point to discrimination or possibly differences in unobservable skills as the sole drivers for the overall gender gap. Alternatively, men and women could in fact be applying to very different jobs already at the application stage. This would imply that gender differences in preferences or beliefs may also contribute to the overall gender gap, either in addition to employer decisions, or partly as a result of them. In particular, if women or men internalize later discrimination, they may correctly believe their hiring chances to be lower for certain jobs and may avoid these jobs already at the application stage (i.e. a form of self-fulfilling discrimination).

To help distinguish between these different possibilities, the rest of the paper examines gender differences in application behavior. In Section 3 we describe our data and institutional framework and discuss the validity of the data for measuring individual job application behavior and for analyzing gender gaps. In Section 4, we present descriptive evidence and document gender differences in applications across a range of job characteristics. In Section 5 we use a decomposition method to show that the observed differences in application behavior are likely to be quantitatively important for overall gender gaps. Finally, Section 6 discusses possible reasons why men and women may be applying for different jobs, including showing some limited but suggestive evidence based on our data.

[4] In Section 3.5, we return to the question of how gender gaps in the overall labor market translate to the specific sample of unemployed workers transitioning into new jobs.

3 Data and institutional setting

In Denmark, UI is available for up to two years at a replacement rate of 90 percent of previous income and a cap of 18.500 DKK (2.500 Euro in 2017). The cap is binding for the majority of workers. UI eligibility requires membership and quarterly membership fees to one of the 24 different UI funds sufficiently well in advance of becoming unemployed. Although such membership is voluntary, a large majority of Danish employees are members of a UI funds and UI recipients make up the vast majority of unemployed job seekers.[5]

To remain eligible for UI while unemployed, UI recipients have to document that they are actively searching for jobs.[6] Since 2015 this documentation has been centralized through an online system called *Joblog*. The Joblog system works as follows: To register an application in the system, unemployed workers need to log in to the central online platform of the Danish public employment service *(Jobnet)*. This platform serves as the main means of communication between UI recipients and public authorities and also functions as a job board, where job seekers can find most posted vacancies in Denmark. After entering the Joblog system, unemployed workers fill in a form describing their job application. It is mandatory to provide information on the applied-for job, including the job-title and hours (part-/full-time), and about the potential employer, including firm name and address. This information serves as the basis for our analysis. We note that although UI recipients are required to log in to the online platform of the Danish public employment service, UI recipients are asked to register *any* type of applied-for job in the system. They are not limited to register jobs that are posted on this platform.

Administration and payout of UI in Denmark is carried out by the UI funds. This includes administration of the job search documentation requirements in Joblog. During a UI recipient's first weeks of unemployment, the UI fund is legally required to instruct the UI recipient in the use of the Joblog system. Over the subsequent unemployment spell, the fund is required to assess whether the UI recipient is complying with the documentation requirements necessary to maintain eligibility.

[5] In 2015, 76 percent of Danish employees were members of a UI fund. Among the gross unemployed 70 percent were currently UI recipients. Of the remaining 30 percent, however, more than two-thirds received mean-tested social assistance, which typically means that they were former UI recipients whose UI benefits had expired (see e.g. Danish Economic Council (2014)).

[6] Additional requirements for maintaining eligibility are that the UI recipient accepts appropriate job offers and participates in activities (such as meetings and activation programs) at the municipal job centers and at the UI funds, see e.g. Kreiner and Svarer (2022); Maibom (2022).

Formally, this is to be done on case-by-case basis, however, as a general rule of thumb, UI recipients are instructed that they need to register somewhere between 1.5 and 2 applications per week in the Joblog system to maintain eligibility.[7] Failure to comply with documentation requirements results in sanctions in the form of lost or reduced UI payments.[8] UI recipients thus face a clear economic incentive to comply with the requirements and register submitted job applications in Joblog.

3.1 Selecting the analysis sample

Our base sample is constructed from administrative data on UI payments and consists of all UI recipients of Danish nationality entering new UI spells from September 2015 to September 2017. September 2015 is the time where the Joblog system was fully operational and September 2017 is the last month where we have labor market information. For each UI recipient, we use a unique person identifier to identify all applied-for jobs that have been registered in the Joblog system during the unemployment spell. We further use this person identifier to merge in data from a wide range of other administrative data sets maintained by Statistics Denmark (DST). These data sets include demographic information, education and the full history of public benefit payments and employment, including information on occupation, hours, wages and firm identifiers for the employing firms (see additional details on the data sources and data construction in Appendix A.1). In selecting the final sample, we make four sample restrictions:

1. We only consider UI spells lasting at least 8 weeks.
2. We only consider individuals who registered at least 4 job applications during their spell.
3. We only consider UI spells that end with the individual finding a job within one year.
4. We exclude job applications made in the last four weeks before a transition to employment.

[7] The law always requires the UI fund to specify a minimum amount of weekly or monthly applications that each individual needs to register, however, this amount should in principle be based on a specific assessment of the workers' education, work experience and competencies, as well as the demand for labor in the area that the worker needs to be available for. Despite the lack of a formal universal threshold of registration requirements, the vast majority of UI funds post general guidelines and it is generally well-known that registering between 1.5 and 2 applications per week should be sufficient for recipients to fulfill eligibility requirements.

[8] In the case of non-compliance with the job search requirements, UI recipients will typically be given a short time period to prove eligibility and register previously unregistered (or ongoing) job search after which the UI fund will make its final assessment. The size of the sanctions ranges from a loss of benefits for a couple of days to a permanent loss of benefits depending on the severity of the non-compliance. In cases where registered job applications are not considered adequate (due to e.g. an assessed risk of proforma search, fake applications), similar requirements apply.

Restriction 1 ensures that we are not looking at individuals who already have a new job lined up when leaving their old employer but receive UI while they wait for this job to start. Restriction 2 removes a small number UI recipients who never start using the Joblog registration system before exiting from UI.[9] Restriction 3 reflects that we are interested in gender gaps in earnings or wages conditional on being employed. Our analysis thus aims to understand how gender differences in job search relate to differences in the jobs men and women are hired for, rather than whether their search results in a hire. Finally, Restriction 4 gets rid of applications that UI recipients are making after successfully landing a job but before this job has actually started.[10]

The top part of Table A.1 in Appendix A.1 show how each of these restrictions affect our sample. After imposing all restrictions we are left with a sample of 105,879 individuals, covering 114,375 UI spells with a total of 2,911,585 job applications. Each of the UI spells in this data ends with the individual transitioning into a job. In the rest of the paper we refer to these jobs as the UI recipients' *new jobs.* In Sections 3.4 and 5 we impose some additional trimming on our sample (more details in the respective sections). As evident at the bottom of Table A.1 in Appendix A.1, this further reduces our sample for these exercises.

As noted in the introduction, a key contribution of our paper is the coverage and representativeness of the data sources we use. To ensure that this representativeness is not foiled by the sample restrictions we impose, Appendix B.1 conducts an extensive set of robustness checks to examine whether any of the sample restrictions substantially affect our results. None of the conclusions presented later are sensitive to the sample restrictions.

3.2 Measuring job characteristics and wages

Our analysis uses data on a range of characteristics of the jobs that men and women apply for and the new jobs they are hired into. For each applied-for job in the data, we use string matching on

[9]Although the UI funds are required to instruct UI recipients in the use of the Joblog system at the beginning of their UI spells, the data shows that unemployed are usually subject to a "phasing-in" period in which they slowly get introduced to Joblog and other components of the UI system, and realize that they have to register applications regularly. For some individuals who leave UI very quickly this can imply that they only register very few applications before exiting.

[10]Many jobs do not start right away which implies that UI recipients typically continue receiving UI for some weeks after they have accepted a new job. In the data, we see a clear drop in the number of applications that people register in Joblog about one month before they enter employment, likely reflecting that the individuals have already accepted their new job at this point in time and are simply waiting for it to start. Applications made while waiting for the new job to start may not represent an individual's general application behavior.

the job title, firm name and firm address to determine the occupation of the job and to merge in firm information from the administrative data. We successfully match 86 percent of applications to a firm and 82 percent to an occupation. Additional details of the matching process are outlined in Appendix A.2.[11]

After matching, we immediately observe the industry and occupation for each UI recipient both for the applied-for jobs and the new job they are actually hired into. To further construct a measure for whether jobs are at a high- or low-paying firm, we use the matched administrative data and estimate an (AKM) log wage regression with worker and firm fixed effects on the universe of Danish workers and firms (Abowd et al. (1999), details of the procedure are given in Appendix A.3). We use the estimated firm fixed effects from this regression as our measure of whether applied-for and actual new jobs are at high- or low-paying firms. Since we conduct separate analyses focused on industry, we standardize the estimated firm fixed effects within industry so that they reflect within-industry differences in firms' wage levels and use these as a main input in our analysis. However, for completeness we also report gender gaps based on the non-standardized firm effects further below.

We also use information on wages in our analysis. For each new job, we observe monthly earnings in the administrative registers and construct a measure of the hourly wage that is paid to the individuals one month after entering a new job. For job applications, we have no direct measure of the wage an unsuccessful applicant would have received if hired. Instead, we use the actual wages and characteristics of the new jobs that UI recipients are hired into to estimate a model that predicts the wage in a given job from the full set of job characteristics. We use this model to compute a predicted wage for each of the jobs in our job application data as well as for each of the actual new jobs. We refer to this prediction as the *typical wage* for the job given its type. Appendix A.4 discusses the details of the prediction procedure. Given that it is often necessary to infer wages from other job characteristics in micro data on applications and vacancies, Appendix A.4 also present some results that speak to the applicability of our prediction method for other settings. Appendix B.6 also presents results from an alternative approach that computes typical wages separately for men and women.

[11] In our analyses of the various job characteristics, we exclude applications with missing information on the relevant job characteristic. In Table A.2 in Appendix A.1 we show the share of missing values for different job characteristics. We have also tried versions where we included a separate category for missing values. This does not change any of our conclusions.

Since we use the predicted *typical wages* when estimating the gender gap in the wage of applied-for jobs, we note that our estimates here will not include potential wage differences across jobs that are unrelated to all of the job characteristics we observe. When we perform our wage decomposition exercise, however, we show results for both typical and actual wages.[12]

3.3 Coverage of the Joblog application data

Relative to many other data sets with information on job search, a key advantage of the data sources we use is their coverage, both in terms of individuals and in terms of job search activities. In terms of individuals, the administrative data sets we build on include the universe of UI recipients in Denmark by construction. This differs from many existing micro data sets on job search which only cover a very selected subset of individuals. In terms of job search activities, a particular advantage of the Joblog application data is that they include all types of applied-for jobs. This differs from most existing data sets which are limited to job applications submitted via a certain channel or platform and/or applications made to a specific subset of potential jobs.

At the same time, since registering jobs in the Joblog data is done entirely by UI recipients themselves, the coverage and validity of these data warrant further discussion and analysis. A priori, a reassuring feature of the Joblog data is that UI recipients face very clear incentives to register job applications in the data and to do so truthfully. As discussed previously, UI recipients face sanctions if they fail to register the required number of applications or if they are caught registering fictitious applications. These incentives are borne out in a very high level of registration activity. In the raw data 96 percent of new UI recipients register at least one applied-for job during the UI spell.

In terms of the number of jobs that each UI recipient applies to, the rule of thumb requirements to register between 1.5 and 2 jobs per week is also strongly borne out in the data. In our final sample, the average number of applied-for jobs per week is strongly centered slightly above 1.5 applications per week (see Figure A.2 for a histogram of weekly applications). This indicates that

[12]When estimating the average applied-for wage, we need wage information for all the jobs that workers apply for, including jobs that the worker did not get. This is why we can only measure the applied-for wage gap based on predicted typical wages. The decomposition method uses propensity score reweighting to construct the counterfactual reemployment wages that women would receive if they applied to jobs with the same characteristics as men. This only requires us to observe wages in the actual job that each person is hired into and thus we can perform the decomposition exercise either using the typical wage measure or actual observed wage payments.

many UI recipients respond to the registration incentives by registering just enough jobs to satisfy requirements. To the extent that job seekers sometimes apply to more jobs than the required number, however, this means the data may not cover all applied-for jobs. In Appendix A.9.1 we use auxiliary survey data on Danish UI recipients to look closer at the degree of coverage. We find a high level of coverage: Survey results suggest that between 69 and 80 percent of all applied-for jobs are registered in the Joblog data. Coverage also appears quite similar across gender: Depending on the method we use, we find that the Joblog data contains between 68 and 82 percent of all applications made by women and between 72 and 76 percent of all applications made by men. Given that differential sample selection by gender is a pervasive issue in the literature, these similar coverage rates are reassuring.[13]

Since the focus of our analysis is on *where* individuals send their applications rather than the total number of applications sent, any potential lack of full coverage is less problematic as long as the subset of applications that are being registered is representative of overall application behavior. A particular concern is the possibility that men and women differ systematically in which types of applications they log. Again the incentive structure around Joblog is somewhat reassuring however. Danish UI recipients face no formal incentives to selectively register some applications over others. Moreover, UI recipients who register fictitious or erroneous applications would be subject to economic sanctions if discovered. Since we cannot rule out that some selective logging occurs, however, we have subjected the data to a range of validity checks. These checks exploit the fact that - independently of the application data - we also observe actual job outcomes. The checks are summarized below but are presented at length in Appendix A.9.2.

First, we show that our data on applied-for jobs is highly predictive of later job outcomes; data on applied-for jobs predicts the characteristics of a UI recipient's new job about as well as the characteristics of their previous job. Moreover, the data on applied-for jobs continue to be predictive even after conditioning on the characteristics of the previous job. We also find that the predictive power of applied-for jobs is very similar among both male and female UI recipients.

Second, we examine how often we are able to trace a new hire back to a job application that is contained in our data. Specifically, for each UI recipient who finds a job at some firm, we check

[13] The concerns are particularly salient in the literature on gender differences in job search. Gender biases in reporting is a major concern for previous work using survey data. Moreover, both surveys and data limited to individual job search platforms face concerns about selection into the sample that could vary by gender.

whether we see that the UI recipients has previously applied for a job at this firm according to our data. To see the usefulness of this exercise, consider the following: Since our raw data is estimated to cover between 69 and 80 percent of all applications and since we successfully match 86 percent of these applications to the corresponding firm, our firm-matched sample of applications should cover between 59 and 69 percent of all applications. Survey data from Denmark suggests that around 73 percent of new hires out of unemployment involve the job seeker actively applying for the job.[14] This suggest a simple test of the representativenes of the application data: If the subset of applied-for jobs contained in our data are a representative subset of all applied-for jobs, the likelihood that we are able to match a given new hire to an application in our data should be between $0.59 \cdot 0.73 = 0.43$ and $0.69 \cdot 0.73 = 0.50$. In contrast, if the application data is non-representative, the data is likely to either over- or under-represent applications that end up turning into a new hire, which would imply a higher or lower match rate. Looking at the new hires in our data, however, the share of new hires that we are able to link to an application is in fact 0.47, consistent with the data being representative.[15] Appendix A.9.2 provides more details and presents additional validity checks.

While we can never rule out that some selective reporting occurs, our overall takeaway is that the Joblog data both has high coverage and appears highly representative of all applied-for job. This in turn makes the data uniquely well-suited to study differences in the types of jobs unemployed workers apply for.

3.4 Conditioning on observables

In our analysis, we primarily want to focus on differences in job applications and hiring outcomes among men and women with the same labor market observables. Throughout the main text, we therefore analyze gender differences in application behavior after conditioning out other labor

[14] The remaining 27 percent reflect jobs or internships assigned to job seekers by their educational institution or through a temp agency. It also reflects instances where the job may be offered to the worker without the worker making an active application (i.e., when a firm directly recruit workers via headhunting, recalls past workers or recruits workers actively through social networks). See Appendix A.9.2 for additional details.

[15] We can also compute this match rate separately for male and female job seekers. Doing so we find a corresponding match rate of 53 percent for women and and 41 percent for men. Qualitatively, the difference here is consistent with previous evidence that among all workers, men are somewhat more likely to find a job in ways that do not involve a formal job application (Engman (2019)), which should mechanically lower the match rate. Since we do not have reliable data on how often this occurs for male and female UI recipients specifically, however, we cannot perform the quantitative benchmarking exercise for men and women separately in our sample of unemployed workers.

market observables.

To condition out observables we employ a standard propensity score reweighting procedure and reweigh the women in our sample to have the same distribution of observables as the men.[16] We opt for propensity score reweighting because this ties in naturally with the decomposition methodology we use later in the paper (see Section 5).

In terms of which exact observables to condition on, the richness of our administrative data means that we have access to an extremely large number of potentially relevant variables. Rather than making ad hoc decisions about which subset of variables to include, we instead use a Machine Learning procedure to discipline variable selection, following recent suggestions in the literature (see e.g. Athey and Imbens, 2019; Angrist and Frandsen, 2022; Mullainathan and Spiess, 2017). Specifically, we construct a very large baseline set of potential variables and then use the double-LASSO of Belloni et al. (2014) to select the subset of these variables that is most important for explaining wage differences between men and women. We then condition on this set of variables throughout our analysis.

Our baseline set of potential variables consists of 4,196 variables containing detailed information on age, work experience, level and field of education, prior industries and occupations, and dependence on public transfers.[17] Importantly, these variables both include detailed information about time-invariant characteristics such as education, as well as detailed information about recent labor market history, which helps deal with the possibility that men and women may self-select into UI at particular point in their labor market career.

From the baseline set of 4,196 variables, the double-LASSO selects 332 variables as being important. In our main analysis we reweight the women in our sample based on estimated propensity scores using these 332 variables. To avoid the usual issues of non-overlapping support when

[16]Our propensity score reweighting works as usual: For some individual in our data, let m be an indicator for being male, and let x be the vector of labor market observables we wish to condition on. The propensity score is now defined as the probability of being male given these other characteristics, $p = P(m = 1|x)$. For each person in the data, we compute an estimated propensity score, $\hat{p}$, and then reweight each woman in the data by $\frac{\hat{p}}{1-\hat{p}}$.

[17]To capture educational differences, the baseline set contains years of education, as well as dummies for the field of study. To capture additional differences in general human capital, the baseline set includes age, total work experience and work experience over the last five years. To capture additional differences in specific human capital, the baseline set includes dummies for the sector, industry and occupation of the previous job as well as continuous measures for the total work experience over the last five years in each of the different industries and occupations. To capture differences in dependence on public transfers, the baseline set includes the total time spent receiving unemployment insurance, social assistance and other public transfers over the last five years. Finally, all variables are also interacted with both age, years of education, total work experience and work experience over the last five years. Appendix A.5 provides additional details.

propensity score-reweighting, we exclude individuals with a propensity score above 0.99 or below 0.01 throughout our descriptive analysis. This reduces our sample of UI spells by 5.4 percent (see Table A.1). Appendix A.5 and A.6 provides additional details regarding the reweighting procedure and the choice of variables to include. Table 1 provide descriptive statistics for the men and women in our analysis sample both before and after reweighting.

3.5 Gender gaps in the analysis sample and overall

Our data allows us to examine gender gaps among UI recipients transitioning into new jobs. These make up a substantial part of the overall labor market; hires out of unemployment cover about half of all new hires in the Danish labor market and in historical data roughly half of a given cohort receives UI at some point during their labor market career. At the same time, however, new hires from UI may differ from the overall stock of employees and we do also impose some additional sample restrictions on the base UI data. In this section, we therefore compare gender gaps in our analysis sample to the broader labor market.

In Panel A of Figure 1, the bar furthest to the left shows the overall gender wage gap in Denmark, estimated on the stock of all employees in August 2015. Overall, men in Denmark are paid around 12 percent (log-points) more than women.[18] This overall gap reflects both wage differences in longer-running employment relationships as well as differences in starting wages among men and women who just joined their employer. In the second bar, we restrict attention to the latter by including only newly hired employees. We see a gender gap in starting wages of 7.7 percent, suggesting that about a third of the overall gender gap reflect gender differences in wage growth within jobs. Since our analysis deals with starting wages in new jobs our analysis of job applications will have little to say about this source of gender differences in wages (although we do provide some direct evidence that women apply for jobs offering lower wage growth in Appendix C.6).

In addition to focusing on new hires, our data only allow us to examine new hires coming from UI. The third bar in Panel A of Figure 1, shows the gender wage gap in starting wages among all new hires in our baseline UI data. With a gender wage gap of 8.2 percent, new hires from UI turn

[18] A 12 percent wage gap is in the range of the numbers reported by Larsen and Larsen (2018) who report the wage gap for different measures of wages. Larsen and Larsen (2018) also compare the gender wage gap in Denmark to other European and Scandinavian countries. The gender wage gap in Denmark is close to the EU average, similar to the gap in e.g. Norway, France and smaller than the gap in Germany. Gallen et al. (2019) show the evolution of the gender wage gap in Denmark 1980-2010.

out to be quite similar to all new hires, at least in terms of the wage gap. Finally, the fourth bar shows the gender gap in wages in our final analysis sample after imposing the sample restrictions described earlier. Imposing these restrictions changes the gender wage gap little. Our final analysis sample shows a raw gender gap of 6.9 percent.[19]

Besides the sample restrictions, two additional aspects of our analysis imply that we focus on a particular part of the gender wage gap: i) For many results, we focus on the typical wages for a given type of job instead of the realized wage, ii) For all our main results, we examine gender differences after conditioning on labor market observables. Panel B of Figure 1 compares gender gaps in realized and typical wages in our analysis sample both before and after conditioning on observables. Comparing the black and grey bars, we see that the gender gap in typical wages is about half of the gap in actual wages. This suggests that about half of the gender wage gap is related to differences in the types of jobs men and women hold, along the dimensions we consider. Comparing the two bars on the left to the ones on the right, we see that conditioning on observables reduces gender gaps somewhat but substantial gaps remain. In our final analysis sample of new hires out of UI, men face a starting wage that is 5.6 percent higher than observationally similar women in their new jobs. This in part reflects that these men are hired into types of jobs whose typical wage is 2.7 percent higher. In the rest of the paper, we examine to what extent these gender gaps are present also in the types of jobs men and women apply for.

4 Descriptive results: Do men and women apply to the same jobs?

We begin our analysis by providing descriptive evidence on differences in the type of jobs that men and women apply for. To do this, we first categorize jobs according to one of the job characteristics we consider (occupation, industry, firm wage-level or typical wage). For each category of jobs, we then compute what share of their applications women and men are on average sending to jobs in this category. We refer to the difference between these averages as the gender gap in applications for this type of job. Throughout the main text, we focus on gender gaps computed *after* reweighting the sample on observables (see Section 3.4).

[19]As we show in Appendix B.1, our results are also highly robust to changes in sample restrictions.

4.1 Occupation and industry

In Panels (a) and (b) of Figure 2 we examine gender gaps in applications across occupations and industries. Specifically, the gray bars in the two panels show the female-male gap in the average share of applications going to each one-digit occupation and industry (the black bars in the figure show corresponding gender gaps in hiring shares; we return to these further below).

There are substantial gender gaps in applications across both occupations and industries. In terms of occupations, for example, the average woman sends almost 12 percentage points more of her applications to service occupations than does the average man. In contrast, she sends 4.7 and 6.1 percentage points less of her applications to machine and craft occupations respectively. For industries, gender gaps in applications are also substantial. For example, the average woman sends about 5.6 percentage point fewer of her applications to jobs in construction than does the average man. Importantly, because these results are conditional on observables (including several measures of past education and occupation), these differences are not simply explained by men or women being more likely to have education or experience from a particular occupation or industry.[20]

To assess the magnitudes, it is instructive to compare the observed gender gaps in applications to the observed gender gaps in the likelihood of ending up in a particular occupation or industry. For this purpose, the black bars in Panels (a) and (b) of Figure 2 show corresponding gender gaps in the share of women vs. men that ends up being hired into each occupation and industry. Comparing the gray and black bars visually, two broad patterns stand out. First, in terms of their sign and relative magnitude, gender gaps in hiring outcomes closely mirror gender gaps in applications; the industries and occupations that women are much less likely to end up in are also the ones that women are applying much less to (and vice versa). Second, in terms of their absolute magnitudes, gender gaps in applications tend to be systematically *larger* than gaps in hiring outcomes.

Table 2 and 3 quantifies these two patterns.[21] In Table 2, we apply a standard measure of

[20]We note that our reweighting procedure need not impose exact balance on all covariates in finite samples. As shown in Table 1, however, gender differences in the previous occupation are much smaller than the observed gender gaps in occupations in the reweighted sample. Changing our reweighting procedure to ensure exact balance on previous occupation or industry also does not eliminate the gender gaps in applications to different occupations and industries (see Appendix B.3).

[21]In the corresponding tables standard errors are obtained via bootstrapping. Specifically, we create 2000 bootstrap samples by sampling unemployment spells with replacement. For each bootstrap sample we restimate propensity scores using the selected variables from our initial variable selection procedure (see Section 3.4). We then calculate reweighted application and hiring gaps and continue our bootstrapping procedure. In each sample we exclude individuals with a propensity score above 0.99 or below 0.01.

industry or occupational segregation - the Duncan index - to measure the extent of segregation in both applications and hiring outcomes (Duncan and Duncan, 1955).[22] Across both industries and occupations, the Duncan index suggests that gender segregation in applications is larger than gender segregation in actual hiring outcomes.

In Table 3, we compute the correlation between gender gaps in applications and in hiring outcomes across occupations and industries. As noted, we see a strong positive correlation for all the different measures.

Finally, it is natural to ask whether the gender application gaps across occupations and industries reflect that women are applying less to high-paying industries or occupations. With this purpose in mind, the occupations and industries on the x-axis in Panels (a) and (b) of Figure 2 have been arranged left-to-right in terms of their average wage. Visual inspection of Panel (a) thus shows some tendency that women apply more to lower paying occupations. This pattern is less clear for industries in Panel (b) .

4.2 Firm wage levels

Next we examine differences in the wage levels of the firms that men and women are applying to. The gray bars in Panel (c) of Figure 2 shows gender gaps in applications to jobs categorized by the deciles of the employing firms' wage level, as measured by their (standardized) AKM fixed effect (see Section 3.2). As before, the black bars show corresponding gender gaps in the share of workers who end up working at firms in each decile.

Gender gaps in applications to firms with different wage levels follow a striking pattern. Starting at the bottom two deciles, we see a gender gap in applications of about 0.7-1.4 percentage points. Thus, the average woman is sending slightly more of her applications to jobs at these low-paying firms. Moving up the distribution of firms' wage levels to higher deciles, however, the gender gap in applications decreases until it reaches about -2.2 percentage points at the two top deciles. This adds up to a very systematic difference in the types of firms that men and women are applying

[22] Applied to applications, the Duncan index measures which fraction of the average woman's (or man's) applications must be changed in order to eliminate the gender application gap. Applied to hiring outcomes, the index measures which fraction of men and women need to be moved to a different job in order to achieve gender balance across occupations or industries. Formally, let y index job types, let a_w^y and a_m^y be the share of applications sent to job y by the average woman and man respectively and let let s_w^y and s_m^y be the share of women and men that are hired into job y. The Duncan index for applications is then defined as $\frac{1}{2}\sum_y |a_w^y - a_m^y|$, while the Duncan index for hiring outcomes is defined as $\frac{1}{2}\sum_y |s_w^y - s_m^y|$.

for. Overall, women are sending 3.2 percentage points more of their applications to firms that are in the bottom three deciles and 4.6 percentage points less of their applications to firms that are in the top three deciles.

Contrasting gender application gaps with the corresponding gender gaps in hiring outcomes in Panel (c) (black bars), we again see that the two follow each other closely. In Table 3, we compute the correlation between gender gaps in applications and in hiring shares across the deciles, which turns out to be 0.945. In terms of the relative magnitudes of the two gaps, however, the pattern is less systematic: For some deciles, application gaps are larger than gaps in hiring shares. For other deciles the reverse is true. To get an overall sense of these magnitudes, in Table 2 we compute gender gaps in the wage level of both the average firm applied-for and in the wage level of the actual firms where men and women end up. On average, men are applying to firms with a wage level that is 1.6 percent higher (or 0.075 standard deviations, see Section 3.2) higher than women. The gap in the wage level of the firms they end up at is slightly lower at at 1.5 percent (0.073 standard deviations).

4.3 Jobs' typical wage levels

The results in the previous sections showed that when it comes to occupation and firm type, women tend to apply more to types of jobs that pay less. We now examine how such gender differences in the characteristics of applied-for jobs add up to create differences in the overall wage levels of the jobs men and women apply for. In Panel (d) of Figure 2 we show gender gaps in applications to jobs that fall in different deciles of the distribution of typical wages. Recall that a job's typical wage refers to the wage level that is typical for a job given its characteristics (see Section 3.2). The gender gaps in Panel (d) thereby summarize how women are targeting jobs with high-paying or low-paying characteristics relative to men.

There are striking gender gaps in the typical wage level of applied-for jobs. Starting at the bottom, women are sending 8.2 percentage points more of their applications to jobs whose typical wage are in the bottom decile. Moving up the wage distribution, the gender gap in applications decreases almost monotonically and eventually flips. For jobs that are in the top decile, women are sending about 4.8 percentage points less of their applications to these jobs.

Contrasting the gender gap in applications in Panel (d) (gray bars) with the corresponding gender gaps in hiring outcomes (black bars) shows a familiar pattern: Gender gaps in applications closely mirror hiring outcomes; across the deciles of jobs' typical wages, the correlation between gender gaps in applications and in hiring is 0.897 (Table 3). Moreover, gender gaps in applications are on average slightly higher than gender gaps in hiring outcomes; while women are on average applying for jobs with a typical wage that is 4.5 percent lower than men, the gender gap in typical wages for the jobs that men and women end up in is 2.7 percent (second panel of Table 2).

4.4 Additional results and robustness

We finish this section by discussing some additional results and robustness checks regarding the documented gender gaps in applications that are presented at length in the appendix.

First, we show that our overall conclusions are robust to a range of alternative empirical approaches. This includes different approaches to addressing covariates (Appendix B.3 and B.4), different sample selection criteria (Appendix B.1), focusing on the number of applications going to different jobs as opposed to shares (Appendix B.5), focusing only on workers who become unemployed due to a mass-layoff (Appendix B.2), and using alternative measures of job's and firms typical wage levels that allow for wages in a given job to differ by gender (Appendix B.6).

Second, we consider potential gender differences in the dynamics of job search. As we show in the appendix, however, men and women change behavior in almost exactly the same way over time (Appendix C.1). Examining gender gaps in applications at different times throughout the unemployment spell therefore yields very similar results.[23]

Third, our measure of typical wages captures differences in the typical *starting* wages for different jobs. In addition to differences in starting wages, however, different jobs may also differ in their wage trajectories over time. Using a simple firm-based measure of on-the-job wage growth, we find that women are in fact both applying to and getting hired into lower wage-growth jobs, consistent with the gender wage gap further growing over the course of an employment spell (Appendix C.6).

Fourth, rather than grouping workplaces by industry, Appendix C.3 presents application and hiring gaps between public and private sector workplaces. Consistent with previous work on gender

[23]In unreported results, we have conducted our main analysis separately for subsamples of job applications that are made after individuals have been unemployed for a specific amount of time. Such changes in time horizons yielded no qualitative differences in the observed patterns. See also Section B.1 Figure B.2 where we show that results are similar if we instead focus on individuals who find employment within 26 weeks.

and public sector jobs, we find that women are substantially more likely to both apply-for and get hired into public sector jobs.

Finally, in Appendix C.2 we examine the characteristics of applied-for jobs relative to those of the previous job. We see that the gender gap in applied-for wages exists also when examining wages relative to the previous job. Men are significantly more likely than women to apply to jobs whose typical wage is higher than that of their previous job. The gender gap in applied-for wages thus does not simply reflect that men had a higher wage in their previous job and are using this higher wage as a reference point in their job search. We also see a similar pattern for occupations: men are more likely to apply to a higher-ranked occupation than that of their previous job.

5 Decomposition: Can applications explain gender gaps in hiring?

As documented in the previous section, there are substantial differences in the jobs men and women apply to. Moreover these gender gaps in applications closely mirror gender gaps in hiring outcomes. A natural question to ask is therefore to what extent the observed differences in application behavior are capable of explaining the observed gender gaps in hiring outcomes. To answer this question, we decompose the observed gender gaps in hiring outcomes (including wages) into a part that can be explained by gender differences in applications and a part explained by other factors.

Our decomposition exercise leverages the fact that our data contains joint information on both job applications and actual hiring outcomes for the same individuals. This allows us to apply the standard semi-parametric decomposition method introduced by DiNardo et al. (1996). Since this methodology has not previously been used to decompose job application behavior, we first lay out and discuss the relevant methodology for our setting over the next two (sub)sections.[24] We then present results.

5.1 Decomposition and counterfactuals, methodology

Consider some individual in our (unweighted) analysis sample. Let m be an indicator for being male, let x be the vector of other labor market observables that we have been conditioning on

[24] Banfi et al. (2022) also apply a weighting scheme based on DiNardo et al. (1996) to job application data from a online job board in Chile. The implementation details are different however because their aim is not to decompose application gaps.

throughout (see Section 3.4), and let a be some vector capturing which types of jobs the person has applied for. In our main results we let a consist of the share of applications sent to each two-digit occupation, each two-digit industry, each decile of the (standardized) firm wage distribution and each decile of the typical wage distribution. Let y be a measure of job type (capturing occupation, industry, wage-level etc.). For expositional convenience we will assume that y is discrete, while x and a are absolutely continuous.

Now let $P^M(y)$ and $P^W(y)$ denote the probability of being hired into a job of type y, conditional on being a man or a woman respectively. In large samples, these gender-specific hiring probabilities will equal the share of men and women hired into job type y, so we will refer to these probabilities interchangeably also as *hiring shares.* Next, let $P^M(y|a,x)$ and $P^W(y|a,x)$ be the corresponding conditional hiring probabilities when conditioning on observable characteristics x and applications a. To decompose gender gaps in hiring, we are interested in estimating counterfactuals of $P^W(y)$ that show how women's hiring outcomes would differ if they had the same observables and job application behavior as men. With this objective in mind, note that by definition $P^M(y)$ and $P^W(y)$ can be written as

$$P^M(y) = \iint P^M(y|a,x) f^M_{a|x}(a|x) f^M_x(x) \,\mathrm{d}a \,\mathrm{d}x$$

$$P^W(y) = \iint P^W(y|a,x) f^W_{a|x}(a|x) f^W_x(x) \,\mathrm{d}a \,\mathrm{d}x$$

Here f^M_x and f^W_x are the distributions of labor market observables among men and women respectively, while $f^M_{a|x}$ and $f^W_{a|x}$ are the conditional distributions of applications among men and women after conditioning on labor market observables.

Throughout the descriptive analysis in Section 4, we focused on gender gaps in applications and hiring among men and women with similar labor market observables. We adopt the same focus and approach in our decomposition. To do this we follow Butcher and DiNardo (2002) and Altonji et al. (2012) and apply a two-step version of the semi-parametric decomposition method of DiNardo et al. (1996) (see also Fortin et al. (2011) for a more recent treatment). In an auxiliary first step, we use the method to condition out labor market observables. In a second step, we then decompose how much of the remaining gender gap can be explained by differences in applications among men and women with similar labor market observables. Additional details regarding the implementation of the decomposition are given in Appendix Section A.7. In addition Appendix B.7 provides a decomposition of the raw gaps, without conditioning on observables.

Step 1: Measuring hiring gaps among men and women with similar characteristics

To set up the first step, let $P_X^{\tilde{W}}(y)$ be the hiring probability that women would have faced if they had the same distribution of labor market observables as men:

$$P_X^{\tilde{W}}(y) = \iint P^W(y|a,x) f_{a|x}^W(a|x) f_x^M(x)\, \mathrm{d}a\, \mathrm{d}x$$

Based on this, we can define the gender gap in the hiring probability for job type y after conditioning on labor market observables. Since the focus of our analysis will be on decomposing this gap, we will refer to this simply as our *baseline hiring gap:*

$$BaselineHiringGap(y) \equiv P_X^{\tilde{W}}(y) - P^M(y) \tag{1}$$

To compute this gap in the data, we need to compute (or estimate) $P^M(y)$ and $P_X^{\tilde{W}}(y)$. Computing $P^M(y)$ is straightforward: We simply compute the share of men in our sample who get hired into job type y. As for $P_X^{\tilde{W}}(y)$, the key insight from DiNardo et al. (1996) is that it can be reliably estimated by propensity score reweighting the sample of women to have the same labor market observables, x, as men and then computing the share of women who get hired into job type y in the reweighted sample.

Reweighting and computing these baseline hiring gaps constitutes the auxiliary first step in our decomposition. In the implementation, we follow the standard practice of using a logit model for the propensity score. The set of labor market observables we include is again the same set as we have used up until this point (see Section 3.4). This implies that the reweighting procedure used in this step of the decomposition is *exactly the same* as the reweighting procedure that was used for the descriptive results in Section 4. There is thus a direct link between the decomposition and the descriptive results presented previously: the descriptive gender gaps in hiring shares shown in Figure 2 are formally equivalent to the *baseline hiring gaps* that we focus on in our decomposition.[25] As in

[25] As we return to in the next section, implementation of the decomposition requires additional trimming of our

Section 4, we continue to trim the sample by dropping observations with an estimated propensity score above 0.01 and 0.99.

In addition to decomposing gender gaps in hiring into particular job types, we are also interested in decomposing the average gender gap in wages and firm wage levels, as well as the overall extent of gender segregation across industries and occupations (as measured by the Duncan index). This proceeds completely analogously in two steps. In the first step, we measure e.g. the baseline wage gap that would prevail if women had the same distribution of labor market observables as men. Letting $w(y)$ be the wage in a job of type y, the baseline wage gap is:

$$BaselineWageGap \equiv \sum_y w(y) \cdot P_X^{\tilde{W}}(y) - \sum_y w(y) \cdot P^M(y) \tag{2}$$

Similarly, letting y index industries and occupation, the baseline level of industry or occupation segregation is:

$$BaselineSegregation \equiv \frac{1}{2} \sum_y \left| P_X^{\tilde{W}}(y) - P^M(y) \right| \tag{3}$$

Again the equivalence of the reweighting procedure implies that these baseline wage gaps and the level of baseline segregation will be formally equivalent to the descriptive results presented previously in Table 2 if computed on the same sample.

Step 2: Decomposing the hiring gaps

$BaselineGap(y)$, $BaselineWageGap$ and $BaselineSegregation$ show gender gaps in hiring outcomes for men and women with similar observables. The key question that we answer in our second step of the decomposition is how much of these gaps can be explained by differences in application behavior. To do this, we define $P_{A,X}^{\tilde{W}}(y)$ to be the hiring probability which women would have faced if they had the same application behavior as men (as well as the same distribution of observables):

$$P_{A,X}^{\tilde{W}}(y) = \iint P^W(y|a,x) f_{a|x}^M(a|x) f_x^M(x) \, \mathrm{d}a \, \mathrm{d}x$$

analysis sample so in the implementation, the baseline hiring gaps we decompose will differ slightly from the gender gaps presented earlier in Figure 2 and Table 2. See Appendix A.1 and Table A.1 for additional details.

We can now define the gender gap in hiring (for job type y) that remains after conditioning on both applications and labor market observables. We refer to this as the *residual hiring gap* after application behavior (and observables) have been accounted for:

$$ResidualHiringGap(y) = P^{\tilde{W}}_{A,X}(y) - P^{M}(y) \tag{4}$$

The baseline gender gaps in hiring into job y can now be decomposed into a part that is explained by differences in applications and a residual gap as follows:

$$BaselineHiringGap(y) = \underbrace{\left(P^{\tilde{W}}_{X}(y) - P^{\tilde{W}}_{A,X}(y)\right)}_{\text{Explained by applications}} + \underbrace{ResidualHiringGap(y)}_{\text{Explained by other factors}} \tag{5}$$

This is the key equation that defines our decomposition of interest. Focusing on men and women with similar labor market observables, this equation decomposes gender differences in hiring shares into two parts. The first part is the part that can be explained by differences in the types of jobs men and women apply for as measured by the Joblog application data. The second part is a residual that contains all other factors that can explain gender gaps in hiring shares. This residual will include discrimination and any other differences in how employers treat observationally similar men and women.[26] The residual may also include other potential explanations, however, such as measurement error or gender differences in the likelihood of accepting a job offer.[27]

To implement Equation 5, we need an estimate of $P^{\tilde{W}}_{A,X}(y)$. As shown by Fortin et al. (2011), however, such an estimate can be obtained simply by propensity score reweighting the sample of women to have both the same observables, x, *and* the same application behavior, a, as the men, and then computing the share of women who get hired into job type y on the reweighted sample.

In implementing this second reweighting step, we again follow standard practice and estimate the propensity score using a logit model. We also continue to trim observations where the estimated propensity score is above 0.99 or below 0.01. This reduces our sample of UI spells by 7.1 percent

[26] Note that "observationally similar" refers to the rich set of observables that we include in the vector x (see Section 3.4). As usual, however, men and women may also differ in unobserved ways that may influence employers hiring decisions.

[27] As discussed in Section, 3.3, the Joblog application appears to have very high coverage and reliability so we expect measurement error in these data to be limited. In addition, UI recipients in Denmark have to accept all job offers that may be considered reasonable if they wish to maintain UI eligibility. This suggest that gender differences in the rejection of job offers may also play a limited role.

(see Table A.1 in the Appendix).[28]

Finally, we can extend this second step of the decomposition to decompose wage gaps and gender segregation by industry and occupation. Again letting $w(y)$ be some measure of wages in job type y, we can define the average residual wage gap as:

$$ResidualWageGap \equiv \sum_y w(y) \cdot P_{A,X}^{\tilde{W}}(y) - \sum_y w(y) \cdot P^M(y) \tag{6}$$

This immediately gives rise to the following decomposition for wages and firms' wage levels:

$$BaselineWageGap = \underbrace{\sum_y w(y) \cdot \left(P_X^{\tilde{W}}(y) - P_{A,X}^{\tilde{W}}(y)\right)}_{\text{Explained by applications}} + \underbrace{ResidualWageGap}_{\text{Explained by other factors}} \tag{7}$$

Letting y index occupations and industries, we further extend this approach to decompose gender segregation across industry and occupation:

$$ResidualSegregation \equiv \frac{1}{2} \sum_y \left| P_{A,X}^{\tilde{W}}(y) - P^M(y) \right| \tag{8}$$

$$BaselineSegregation = \underbrace{\frac{1}{2} \sum_y \left(\left| P_X^{\tilde{W}}(y) - P^M(y) \right| - \left| P_{A,X}^{\tilde{W}}(y) - P^M(y) \right| \right)}_{\text{Explained by applications}} + \underbrace{ResidualSegregation}_{\text{Explained by other factors}} \tag{9}$$

5.2 Identification of counterfactuals in the decomposition

The decomposition method relies on estimating the counterfactual hiring outcomes that women would have faced if they had the same application behavior as men (and the same labor market observables). As a result, it is worth discussing the underlying assumptions that identify these

[28] In total, we thus trim our sample twice before performing the decomposition. The first trimming removes observations where the estimated propensity score lies below 0.01 or above 0.99 when the propensity score is estimated only based on labor market observables (x). The second trimming removes observations where the estimated propensity score lies below 0.01 or above 0.99 when the propensity score is estimated based on application behaviour (a) or both labor market observables and application behavior (x and a). Since the second trimming step is not needed for the descriptive analysis in Section 4, the sample of UI spells used for the decomposition is 7.1 percent smaller than the sample used in Section 4. Comparing the (baseline) hiring gaps in Figure 2 and 3, however, we see that the additional trimming affects gender gaps little.

counterfactuals.

As discussed and formalized in Fortin et al. (2011), the key identifying assumption can be viewed as an ignorability or conditional independence assumption. In the derivations above, this is most clearly seen from the fact that the conditional hiring probability for women only depends on application behavior and labor market observables, $P^W(y|a,x)$. This in particular implies that any unmodeled or unobserved factors that affect hiring outcomes must be independent of application behavior (conditional on labor market observables). Put differently, if a woman who is currently not applying to "male jobs" were to start applying to such jobs, we are assuming that she will face the same hiring outcomes as those women in the data who are already applying to these jobs and have similar labor market observables.

While this identifying assumption is standard in decomposition exercises, it raises several concerns. Perhaps the most salient concern is the possibility that women who are currently applying to more male jobs differ from other women along dimensions not captured by our vector of labor market observables and that these differences affect hiring outcomes. Although our vector of observables aims to include very detailed measures of the relevant labor market observables (see Section 3.4), unobservable characteristics may still play a role.[29]

Regardless, we view the decomposition as providing a useful benchmark, especially given the lack of previous evidence. In particular, the decomposition shows how much of the observed gender gaps in outcomes can be explained by differences in applications under a standard assumption used in the literature.

5.3 Decomposing wage gaps and segregation measures

Table 4 presents decompositions for the overall level of gender segregation by industry or occupation (based on Equation 9) and for average gender gaps in wages or firms' wage levels (based on Equation 7). Here and throughout the decomposition exercise, standard errors and confidence intervals are obtained via bootstrapping.[30] The first four rows decomposes the level of gender segregation across

[29] Given the existing literature on gender differences in competitiveness (see e.g. Niederle and Vesterlund, 2007), we may for example expect that women who apply to more "male" jobs have a more competitive personality. This may impact their hiring probabilities as well, and may not be captured well by the set of variables we condition on in our analysis.

[30] Specifically, we create 2000 bootstrap samples by resampling unemployment spells. Within each bootstrap sample we then reestimate both the propensity score weights based on our set of variables from the initial variable selection procedure (see Section 3.4) and the propensity score weights that includes information about applied for

industries and occupations. Depending on the level of aggregation considered, gender differences in application are capable of explaining 28 to 40 percent of observed gender segregation across occupations among men and women with similar labor market characteristics. For segregation across industries, gender differences in applications are able to explain 20 to 25 percent of the observed industry segregation.

The next row shows the decomposition of the gender gap in the employing firms' wage levels. Results suggest that gender differences in the types of firms that men and women apply to can explain a large part of the observed gap. Among men and women with similar labor market observables, women on average work at firms whose wage level is 0.068 standard deviation lower than men. After accounting for differences in application behavior, however, this difference drops to 0.017 standard deviations, implying that differences in applications can explain 75 percent of the baseline gap.

Finally, the last two rows of the table focus on gender gaps in wages using two different measures. The first is our measure of the *typical wage* paid in jobs with different characteristics: Among men and women with similar labor market observables, women are on average hired into types of jobs that pay 2.7 percent less than men. The second to last row shows a decomposition of this gap. The second wage measure is the *actual wage* of the new job. Among newly employed men and women with similar labor market observables, women in our data are on average paid 5.6 percent less in their new job than are men. The last row decomposes this gap. In both decompositions, gender differences in applications can explain a very substantial part of the gender wage gap. After accounting for differences in where men and women are applying, the gender gap in typical wages drops to 0.6 percent, while the gender gap in actual wage drops to 1.6 percent. Overall, the gender application gap is capable of explaining about 79 percent of the gap in typical wages and 72 percent of the gap in actual wages.

5.4 Decomposition results for individual industries and occupations

While results in Table 4 suggest that gender differences in application behavior are capable of explaining most of the gender gap in wages, applications seem to play a more limited role in shaping observed gender segregation across industry and occupations.

---

jobs.

In Figure 3 we unpack the decomposition results for occupations and industries by considering hiring into different occupations and industries separately. This is based on applying the decomposition in Equation 5. The black bars in the figure show the baseline hiring gaps after conditioning on labor market observables. The solid-white and dashed-white bars then decompose these hiring gaps into a part that can be explained by gender differences in applications (solid bars) and a residual gap (dashed bars). Error bars show bootstrapped 95 percent confidence intervals.

Panel A shows results for occupations. The overall results presented earlier (Table 4) turn out to mask very heterogeneous patterns across occupations. For several of the occupations with significant baseline gaps, applications are in fact able to explain most or all of this gap. Among Service and Sales Workers, for example, the residual hiring gap after accounting for applications is only about a third of the baseline gap, while the residual gap is essentially zero for Process Control Technicians. For other occupations, results are starkly different. For the substantial baseline gap in Craft and Related Trades, for example, essentially none of it can be explained by gender differences in applications. Finally, the Professionals occupation also deserves particular mention. Here gender differences in applications "over-explain" the baseline gap: The baseline hiring gap among Professionals is negative, but the residual gap is positive. The decomposition results thus suggest that if women applied to the same jobs as men, women would go from being underrepresented to overrepresented in this occupational group (all else equal). Panel B shows results for industries and reveals a similar picture. For some industries gender differences in applications are able to explain most of the baseline gap, while for others applications matters little and large residual gaps remain. We also see an example of applications over-explaining the baseline gap for the Information and Communication industry.

At first glance, it might seem puzzling that gender differences in applications are capable of explaining a large majority of the gender wage gap when they only have limited explanatory power for gender differences in industries and occupations. Recalling that Figure 3 arranges industries and occupations left-to-right in terms of average pay, clarifies that this occurs because the unexplained residual gender gaps in industries and occupations are largely unrelated to wage levels: Looking left-to-right in Figure 3, it is not the case that residual gaps are systematically more negative for high-paying occupations and industries. Rather, we see positive and negative residual gaps among both high and low-wage job types.

5.5 Comparison to previous estimates of gender discrimination

Why do substantial residual gaps exist for some industries and occupations and not for others? By their nature, residual gaps may capture a wide range of factors that shape hiring gaps differently across occupations and industries. One particularly salient explanation however is that gender discrimination in the hiring process may vary across different types of jobs, with some jobs discriminating more against women and others discriminating more against men. This has been documented in previous work that uses audit studies or other quasi-experimental methods to measure gender discrimination in hiring (see Rich, 2014; Riach and Rich, 2002).

To examine the role of discrimination more directly, we compare our residual hiring gaps to estimates of gender discrimination from the audit study literature. We do this in two ways. First, we use data from Ahmed et al. (2021), to construct direct measures of gender discrimination at the occupation level. Ahmed et al. (2021)'s data cover three audit studies conducted in Sweden between 2016-2019. Importantly. the Swedish setting closely mimics Denmark in terms of both gender norms and female labor market outcomes so both the time period and setting should be comparable to our data. For five of the nine occupations we consider, the data from Ahmed et al. (2021) allow us to estimate gender discrimination in the first part of the hiring process via the gender gap in the call back rate: the difference in the likelihood of getting a positive initial reply for a female applicant relative to an otherwise identical male applicant (Appendix C.5 provides additional details).

Panel A of Figure 4 plots our occupation-specific residual gender gaps against the female-male gap in the call back rate based on Ahmed et al. (2021). We see a strong positive relationship. Occupations with larger call back rate gaps have systematically larger residual hiring gaps in our decomposition. Notably, the pattern fits very well also for the Professionals occupation in which the residual gap has a different sign than the baseline hiring gap. The correlation in the figure thus does not simply reproduce a possible correlation between callback rates and baseline hiring gaps.

In Panel B we conduct a more indirect test. A robust finding in the audit study literature is that gender discrimination goes in favor of women in female-dominated jobs, while the opposite tends to be true in male-dominated jobs (Riach and Rich (2002); Rich (2014)). If residual hiring gaps differ across industries and occupations due to differences in discrimination, this implies that residual hiring gaps should be positive in industries and occupations with many women and negative in industries and occupations with few women. In Panel B, we therefore plot the residual hiring gap for each industry or occupation against the share of females employees in the corresponding industry or occupation. In line with the prediction from the audit study literature, we see a clear upward slope both across industries and occupations.[31]

Overall, the results are consistent with the idea that gender segregation by industry and occupation is in large part caused by direct gender discrimination rather than gender differences in application behavior.

6 Discussion: Why are men and women applying to such different jobs?

The results in the previous sections show that gender differences in application behavior are large and are capable of explaining a very large part of observed gender gaps in hiring outcomes. This raises the natural question of *why* men and women are applying to such different jobs. Unfortunately, while our data offer a unique opportunity to document gender differences in application behavior for the broader labor market, they are less well suited to analyze why these differences exists. Doing so typically requires convincing sources of (quasi-)experimental variation that is not readily available in our data. Nevertheless, in this section we discuss possible explanations for the observed gender differences, as well as some suggestive empirical checks on these explanations.

6.1 Gender differences in the valuation of non-wage job characteristics

Several recent papers have suggested that women have a higher valuation of a range of non-wage job characteristics (see e.g. Wiswall and Zafar, 2018, Maestas et al., 2019, Le Barbanchon et al., 2021, Hotz et al., 2018; Goldin, 2014). If these characteristics correlate negatively with jobs' wage levels, this can explain why women are applying to lower paying jobs more than men. To examine this explanation in the data, we construct measures of three non-wage job characteristics that have

[31] Note that this upward sloping relationship is not mechanical. The share of female employees in an occupation/industry depends on the hiring gap into the occupation/industry overall, as well as the relative tendency of women to transition out of it. In contrast, the residual hiring gaps from our decomposition are based on hiring gaps after conditioning out observables *and* accounting for application behavior. There is thus nothing that forces the residual gaps to line up with the overall share of female employees.

been emphasized in previous work and compute gender gaps in the share of applications going to jobs with these characteristics. As in Section 4, we compute these gender gaps after reweighting on labor market observables (see Section 3.4).[32]

The non-wage job dimensions we consider are part-time vs. full-time jobs, jobs with different implied commute lengths and jobs with different levels of family-friendliness. The first measure is readily available in our data: One of the mandatory fields that UI recipients fill in when registering a job application is whether the applied-for job is full-time (37 hours per week in Denmark) or part-time (less than full-time work). Further, using information on the address or zip code of applied-for jobs as well as information on the zip code of the UI recipient, we can compute the implied commuting for each applied-for job.[33] Finally, based on the linked administrative data we compute a simple measure of a firm's family-friendliness based on how much parental leave the average employee takes when they or their partner give birth. We correct the measure for the gender of the employees that take leave and treat firms that never experience an employee or partner giving birth as a separate category.[34] Appendix A.8.1 provides additional details.

Figure 5 shows gender gaps in applications and hiring to part-time vs. full-time jobs, to jobs with different implied commute lengths and to firms in different deciles in terms of family-friendliness. There is clear support for the idea that women are targeting jobs with different non-wage characteristics than men. Women are sending a markedly smaller share of their applications to full-time jobs, to jobs that are more than 60 minutes away and to firms that score low on our family-friendliness measure.[35] Further, in Appendix A.8.2 we verify that part-time jobs, jobs involving shorter commutes and jobs at family-friendly firms all tend to offer lower typical wages in our data. Overall, these results indicate that gender differences in the valuation of non-wage job characteristics may be one of the explanations for the gender application gap that we document in this paper.

[32]In addition to the non-wage job characteristics we focus on here, one could also imagine that men and women have different preferences over the gender composition of their coworkers. In Appendix C.10, we show that women indeed send systematically more applications to firms with a larger share of women in the workforce.

[33]Commuting distances are calculated with Google Maps API using the individuals' municipality of residence and the jobs' postal code, see Harmon (2015) for further details on this data.

[34]We adjust for the gender of the employee because - as is well-known - men take systematically much shorter leaves than women in Denmark. Average leave length is of course undefined for firms that never experience an employee or partner giving birth. Having no employees give birth is arguable an indicator that these firms have low levels of family-friendliness; if working at extremely family-unfriendly firms, employees might respond by not having children at all. We therefore opt to leave these firms in the analysis as a separate category.

[35]We see that women are particularly unlikely to apply to firms that do not experience a birth over the time period we consider. Among firms that do experience a birth, however, there is also a clear pattern of women applying more to firms where long parental leaves are common.

Additionally, as we show in Appendix C.7, we also find some support for the idea that gender differences in the valuation of these job characteristics are related to motherhood (see e.g. Kleven et al. (2019); Hotz et al. (2018)). Focusing on individuals in an age window around childbearing (age 25-40), gender gaps in job applications tend to be larger between men and women with young kids than between men and women without kids. At the same time, however, we find substantial gender gaps in job application behavior also among individuals without kids.

6.2 Other explanations

A number of other possible explanations exists for the observed gender differences in applications, including so-called self-fulfilling discrimination, gender differences in beliefs and information, (over)confidence and risk-aversion, or broad gender discrimination against women. Since we are not able to find any evidence for these explanations with our data, however, we only briefly discuss these below and relegate corresponding empirical results to Appendix C.

One possible explanation for the observed gender application gaps is a version of the "self-fulfilling discrimination" mechanism that has been proposed and documented in other settings (Lundberg and Startz, 1983; Coate and Loury, 1993; Glover et al., 2017). In the context of the job application process, the basic idea behind this is as follows: Gender discrimination in hiring implies that women face a lower likelihood of being hired when applying for some jobs. As a result women have an incentive to apply less to these jobs and more to jobs where the chance of being hired is higher. In this way, gender differences in the likelihood that an application turns into a hire may explain why there are gender differences in applications.

To look for evidence that this mechanism is at play in our data, we compute simple measures of how much more/less likely a woman is to be hired into a particular type of job when she applies for it, relative to when a man applies for it. We then correlate this measure with the observed gender gaps in applications to different jobs. To find evidence of self-fulfilling discrimination, we should see a positive correlation here; women should be applying less than men exactly to those jobs where women also face systematically lower returns to applying. This is not what we see, however. If anything the correlation is slightly negative in our data. While we thus find no evidence of self-fulfilling discrimination in our data, we note that these findings also do not rule out self-fulfilling discrimination. Ruling out this mechanism would require us to also examine gender differences in the counterfactual success probabilities of potential job applications that men and women *could have sent* but opted not to. This is not possible with our data. Appendix C.8 provides additional details and discussion.

Another possible explanation for the observed gender application gaps are gender differences in beliefs and risk tolerance (see e.g. Cortes et al., 2020), or possibly in personality more broadly (Flinn et al. (2020)). If men are more (over)confident and systematically judge their labor market prospect to be better than women do - or if men simply are less risk averse - this could lead men to systematically target more high-paying but harder-to-get jobs than women. As discussed in Section 4.4, we indeed find that men are more likely to apply upwards on the job ladder relative to their previous job.

To look for more direct empirical evidence for this mechanism, we examine the speed of job finding for men and women; if gender differences in beliefs or risk preferences are causing men to systematically target harder-to-get jobs, this implies that men should be finding jobs at a slower rate than women. This is not borne out in the data. If anything, men tend to find jobs faster than women both in the raw data and after conditioning on observables.

To test specifically for gender differences in beliefs, note that if unemployed women and men start out with different beliefs but learn over time, we should expect see different dynamics in application behavior. As noted in Section 4.4, this is not what we see; men and women in fact change their application-behavior in the same way over the course of an unemployment spell. While these findings are inconsistent with gender differences in beliefs or risk preferences being the only source of differences in application behavior, we again note that they do not rule out that gender differences in beliefs or risk preferences are at play. If other differences across men and women also affect search dynamics and job finding rates, these other differences may obscure some of the direct effects of gender differences in beliefs and risk aversion. Appendix C.1 and C.9 provide additional details and discussion.

Finally, we note that both gender differences in information and general discrimination against women in the hiring process could also contribute to the patterns we see in our data. If women are poorly informed about the available jobs or if women face a lower probability of being hired throughout the labor market, standard models would predict that women could respond by applying for less attractive jobs that are easier to get. This would explain why women tend to apply for lower-paying jobs. At the same time, if the information frictions or hiring discrimination faced by women is not fully offset by women targeting easier-to-get jobs, women might still end up with lower job finding rates than men.

7 Conclusion

In this paper we provide evidence on gender differences in the types of jobs men and women apply for. We do this by exploiting new administrative data on job applications made by the universe of Danish UI recipients which we link to additional administrative data on hires, firms and UI recipients characteristics.

We document substantial gender differences in applied-for jobs, even among men and women with similar labor market observables. These gender gaps in applications closely mirror observed gender gaps in actual hiring outcomes. In particular, women apply for systematically lower-paying jobs.

Applying a standard decomposition method, we show that gender differences in applications are able to explain a substantial part of observed gender gaps in hiring outcomes among men and women with similar labor market characteristics. For wages, the gender application gap can explain more than 70 percent of wage differences among observationally observationally similar men and women.

Finally, we discuss possible explanations for why men and women are applying to such different jobs and also provide some suggestive evidence on this question. In our data, we find support for the idea that women and men apply for different jobs because they have different valuations of non-wage job characteristics, such as hours, commuting time and family-friendliness. We note however that many other mechanisms may also be at play, including self-fulfilling discrimination, where women shy away from jobs where they expect to be discriminated against. Pinning down the exact reasons that men and women apply differently is a key objective for future work.

Our results show that in to obtain gender equality in the labor labor market, it will be necessary to instigate substantial changes in the application behavior of men and women. This has particular implications for policy initiatives whose primary aim is to change how women are treated in the the hiring process, such as gender quotas or bias awareness training. In particular, our results suggest that the ability of such policies to close gender gaps will depend crucially on whether they change not only hiring behavior but also manage to change application patterns by changing men and women's perceptions about their hiring chances. One implication of this is that it may be important to design and announce such policies to be as salient as possible among job seekers. Another implication is that if job seekers' perceptions change slowly, the full effect of policies like gender quotas will only materialize after some time.

References

Abowd, J. M., Kramarz, F., and Margolis, D. N. (1999). High wage workers and high wage firms. *Econometrica*, 67(2):251–333.

Ahmed, A., Granberg, M., and Khanna, S. (2021). Gender discrimination in hiring: An experimental reexamination of the swedish case. *PLOS ONE*, 16(1):1–15.

Ahrens, A., Hansen, C. B., and Schaffer, M. E. (2018). PDSLASSO: Stata module for post-selection and post-regularization OLS or IV estimation and inference. *Statistical Software Components*.

Ahrens, A., Hansen, C. B., and Schaffer, M. E. (2019a). Lassopack: Model Selection and Prediction with Regularized Regression in Stata. *IZA Discussion Papers*.

Ahrens, A., Hansen, C. B., and Schaffer, M. E. (2019b). LASSOPACK: Stata module for lasso, square-root lasso, elastic net, ridge, adaptive lasso estimation and cross-validation. *Statistical Software Components*.

Altonji, J. G., Bharadwaj, P., and Lange, F. (2012). Changes in the characteristics of American youth: Implications for adult outcomes. *Journal of Labor Economics*, 30(4):783–828.

Andrews, M. J., Gill, L., Schank, T., and Upward, R. (2008). High wage workers and low wage firms: Negative assortative matching or limited mobility bias? *Journal of the Royal Statistical Society. Series A: Statistics in Society*, 171(3):673–697.

Angrist, J. D. and Frandsen, B. (2022). Machine Labor. *https://doi.org/10.1086/717933*, 40(S1):S97–S140.

Athey, S. and Imbens, G. W. (2019). Machine Learning Methods That Economists Should Know About. *Annual Review of Economics*.

Banfi, S., Choi, S., and Villena-Roldan, B. (2019). Deconstructing Job Search Behavior. *Working Paper*.

Banfi, S., Choi, S., and Villena-Roldán, B. (2022). Sorting on-line and on-time. *European Economic Review*, 146:104128.

Barbulescu, R. and Bidwell, M. (2013). Do women choose different jobs from men? Mechanisms of application segregation in the market for managerial workers. *Organization Science*, 24(3):737–756.

Belloni, A., Chernozhukov, V., and Hansen, C. (2014). Inference on treatment effects after selection among high-dimensional controls. *Review of Economic Studies*, 81(2):608–650.

Belloni, A., Chernozhukov, V., Hansen, C., and Kozbur, D. (2016). Inference in High-Dimensional Panel Models With an Application to Gun Control. *Journal of Business and Economic Statistics*, 34(4):590–605.

Bertheau, A., Acabbi, E. M., Barceló, C., Gulyas, A., Lombardi, S., and Saggio, R. (2022). The Unequal Consequences of Job Loss across Countries. *American Economic Review: Insights (forthcomming).*

Blau, F. D. and Kahn, L. M. (2017). The gender wage gap: Extent, trends, & explanations. *Journal of Economic Literature*, 55(3):789–865.

Butcher, K. F. and DiNardo, J. (2002). The Immigrant and Native-Born Wage Distributions: Evidence from United States Censuses. *Industrial and Labor Relations Review*, 56(1):97.

Caliendo, M., Lee, W.-S., and Mahlstedt, R. (2017). The gender wage gap and the role of reservation wages: New evidence for unemployed workers. *Journal of Economic Behavior & Organization*, 136:161–173.

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. *Quarterly Journal of Economics*, 131(2):633–686.

Card, D., Heining, J., and Kline, P. (2013). Workplace Heterogeneity and the rise of Est German Wage Inequality. *Quarterly Journal of Economics*, 128(3):967–1015.

Coate, S. and Loury, G. (1993). Will affirmative-action policies eliminate negative stereotypes? *American Economic Review*, 83(5):1220–1240.

Cortes, P., Pan, J., Pilossoph, L., and Zafar, B. (2020). Gender differences in job search and the earnings gap: Evidence from business majors. *Working Paper*.

Dahl, M. and Krog, N. (2018). Experimental evidence of discrimination in the labour market: Intersections between ethnicity, gender, and socio-economic status. *European Sociological Review*, 34(4):402–417.

Danish Economic Council (2014). Danish Economy, Autumn Report.

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

Duncan, O. D. and Duncan, B. (1955). A Methodological Analysis of Segregation Indexes. *American Sociological Review*, 20(2):210.

Engman, T. S. (2019). Mænd får oftere job gennem netværk og bliver lidt oftere headhuntet end kvinder. *Statistics Denmark*.

Eriksson, S. and Lagerström, J. (2012). The Labor Market Consequences of Gender Differences in Job Search. *Journal of Labor Research*, 33(3):303–327.

Flinn, C., Todd, P., and Zhang, W. (2020). Personality Traits, Job Search and the Gender Wage Gap. *Working Paper*.

Flory, J. A., Leibbrandt, A., and List, J. A. (2014). Do Competitive Workplaces Deter Female Workers? A Large-Scale Natural Field Experiment on Job Entry Decisions. *The Review of Economic Studies*, 82(1):122–155.

Fortin, N., Lemieux, T., and Firpo, S. (2011). Decomposition Methods in Economics. *Handbook of Labor Economics*.

Gallen, Y., Lesner, R. V., and Vejlin, R. (2019). The labor market gender gap in Denmark: Sorting out the past 30 years. *Labour Economics*, 56(C):58–67.

Gee, L. K. (2019). The more you know: Information effects on job application rates in a large field experiment. *Management Science*, 65(5):2077–2094.

Glover, D., Pallais, A., and Pariente, W. (2017). Discrimination as a self-fulfilling prophecy: Evidence from french grocery store. *Quarterly Journal of Economics*, 132(3):1219–1260.

Goldin, C. (2014). A grand gender convergence: Its last chapter. *American Economic Review*, 104(4):1091–1119.

Goldin, C. and Rouse, C. (2000). Orchestrating impartiality: The impact of "blind" auditions on female musicians. *American Economic Review*, 90(4):715–741.

Harmon, N. a. (2015). Are Workers Better Matched in Large Labor Markets ? *Working Paper*.

Hotz, V. J., Johansson, P., and Karimi, A. (2018). Parenthood, Family Friendly Workplaces, and the Gender Gaps in Early Work Careers. *Working Paper*.

Kleven, H., Landais, C., and Søgaard, J. E. (2019). Children and gender inequality: Evidence from Denmark. *American Economic Journal: Applied Economics*, 11(4):181–209.

Kreiner, C. T. and Svarer, M. (2022). Danish Flexicurity: Rights and Duties. *Journal of Economic Perspectives*, 36(4):81–102.

Kuhn, P., Shen, K., and Zhang, S. (2018). Gender-targeted job ads in the recruitment process: Evidence from china. *IZA Discussion Papers*.

Lachowska, M., Mas, A., and Woodbury, S. A. (2020). Sources of displaced workers' long- term earnings losses. *American Economic Review*, 110(10):3231–3236.

Larsen, M. and Larsen, M. R. (2018). Forskelle mellem kvinders og mænds timeløn. Technical report, VIVE - The Danish Center for Social Science Research.

Le Barbanchon, T., Rathelot, R., and Roulet, A. (2021). Gender Differences in Job Search: Trading off Commute Against Wage. *The Quarterly Journal of Economics*, 136(1):381–426.

Lundberg, S. and Startz, R. (1983). Private Discrimination and Social Intervention in Competitive Labor Markets. *American Economic Review*, 73(3):340–47.

Lundborg, P., Plug, E., and Rasmussen, A. W. (2017). Can women have children and a career? IV evidence from IVF treatments. *American Economic Review*, 107(6):1611–1637.

Maestas, N., Mullen, K., Powell, D., von Wachter, T., and Wenger, J. (2019). The Value of Working Conditions in the United States and Implications for the Structure of Wages. *NBER Working Paper Series*.

Mahlstedt, R., Sebald, A., Cairo, S., and Altmann, S. (2019). Complexity and the Effectiveness of Public Policy. *Manuscript*.

Maibom, J. (2022). The Welfare Effects of Mandatory Reemployment Programs: Combining a Structural Model and Experimental Data. *International Economic Review*.

Marinescu, I. and Skandalis, D. (2021). Unemployment insurance and job search behavior. *Quarterly Journal of Economics*, 136(2):887–931.

Marinescu, I. E. and Wolthoff, R. (2019). Opening the Black Box of the Matching Function: the Power of Words. *Journal of Labor Economics*.

Mullainathan, S. and Spiess, J. (2017). Machine learning: An applied econometric approach. In *Journal of Economic Perspectives*, volume 31, pages 87–106.

Neumark, D. (2004). Sex differences in labor markets. *Routledge*, 111(3):1–416.

Neumark, D. (2018). Experimental research on labor market discrimination. *Journal of Economic Literature*, 56(3):799–866.

Niederle, M. and Vesterlund, L. (2007). Do women shy away from competition? Do men compete too much? *Quarterly Journal of Economics*, 122(3):1067–1101.

Olivetti, C. and Petrongolo, B. (2016). The Evolution of Gender Gaps in Industrialized Countries. *Annual Review of Economics*, 8(1):405–434.

Riach, P. A. and Rich, J. (2002). Field experiments of discrimination in the market place. *Economic Journal*, 112(483):F480–F518.

Rich, J. (2014). What Do Field Experiments of Discrimination in Markets Tell Us? A Meta Analysis of Studies Conducted Since 2000. *IZA Discussion Paper*, (8584):1–71.

Rousille, N. (2021). The central role of the ask gap in gender pay inequality. *Working Paper*.

Samek, A. (2019). Gender differences in job entry decisions: A university-wide field experiment. *Management Science*, 65(7):3272–3281.

Urminsky, O., Hansen, C., and Chernozhukov, V. (2016). Using Double-Lasso Regression for Principled Variable Selection. *SSRN Electronic Journal*.

Wiswall, M. and Zafar, B. (2018). Preference for the workplace, investment in human capital, and gender. *Quarterly Journal of Economics*, 133(1):457–507.

Table 1: Descriptives of the analysis sample

	Male Raw	**Women** Weighted	**Women** Raw		**Male** Raw	**Women** Weighted	**Women** Raw
Age	38.56	38.87	37.57	**Occupation**			
				Manager	0.03	0.03	0.02
Education: Level				Professional	0.16	0.16	0.24
Lower secondary	0.32	0.28	0.21	Associate professional	0.08	0.08	0.08
Upper secondary	0.48	0.51	0.50	Clerical support	0.08	0.08	0.13
Tertiary	0.18	0.20	0.28	Service/sales	0.14	0.15	0.28
Education length (years)	13.21	13.41	13.76	Agricultural	0.02	0.02	0.01
				Elementary	0.10	0.10	0.02
Education: Field				Plant/machine	0.08	0.10	0.03
General	0.50	0.46	0.42	Craft	0.17	0.14	0.10
Education	0.01	0.01	0.02				
Humanities	0.04	0.03	0.06	**Industry**			
Social sciences	0.02	0.02	0.03	Agriculture	0.03	0.02	0.01
Business	0.12	0.13	0.18	Manufacturing	0.14	0.15	0.07
Natural sciences	0.01	0.01	0.01	Construction	0.10	0.08	0.02
Information	0.01	0.01	0.00	Trade	0.23	0.24	0.20
Engineering	0.19	0.20	0.05	Communication	0.04	0.05	0.03
Agriculture	0.03	0.04	0.02	Finance	0.02	0.02	0.01
Health	0.03	0.03	0.15	Real estate	0.01	0.02	0.01
Service	0.04	0.04	0.04	Services	0.14	0.14	0.11
				Public	0.20	0.21	0.44
Region				Culture	0.05	0.05	0.06
Hovedstaden	0.30	0.28	0.32				
Midtjylland	0.22	0.22	0.22	**Sector**			
Nordsjælland	0.13	0.12	0.12	Private sector	0.71	0.69	0.50
Sjælland	0.14	0.15	0.13	Public sector	0.12	0.13	0.30
Syddanmark	0.22	0.22	0.21	Other sector	0.17	0.17	0.19
Previous year				**UI spell**			
Employment, weeks	35.87	35.38	33.64	UI length (weeks)	19.85	19.95	20.10
Public transfers, weeks	18.84	18.94	21.40	Applications per week	1.46	1.55	1.60

Notes: The table reports summary statistics for our analysis sample. Means or shares for the different variables are reported separately for men, for women before applying propensity score reweighting and for women after applying propensity score reweigthing. All characteristics, except for the UI spell details, are measured prior to the unemployment spell.

Table 2: Summary table for descriptive gender gap

	Applications	**Hiring**
Occupational segregation	0.187	0.082
(Duncan index, 1-digit)	*(0.004)*	*(0.006)*
Occupational segregation	0.253	0.123
(Duncan index, 2-digit)	*(0.005)*	*(0.007)*
Industry segregation	0.114	0.096
(Duncan index, 1-digit)	*(0.004)*	*(0.006)*
Industry segregation	0.132	0.100
(Duncan index, 2-digit)	*(0.004)*	*(0.007)*
Gap in mean firm wage level	0.075	0.073
(Male-female gap, std. AKM fixed effect)	*(0.017)*	*(0.010)*
Gap in mean firm wage level	0.016	0.015
(Male-female gap, AKM fixed effect)	*(0.001)*	*(0.001)*
Gap in mean typical wage	0.045	0.027
(Male-female gap, log typical wage)	*(0.001)*	*(0.002)*

Notes: In the top panel of this table we report the Duncan index measure of industry or occupational segregation at the 1- and 2-digit level, calculated based on application and hiring shares respectively (see footnote 22 for a definition of the Duncan index). The lower panel of the table reports the gender gaps in the mean firm wage level (standardized and non-standardized AKM fixed effect) and (log) typical wage across applied-for jobs and actual new jobs. The results are computed for the main analysis sample after applying propensity score reweighting to condition out differences in observables. Standard errors in parenthesis are calculated by bootstrapping unemployment spells and reestimating propensity scores and gaps (see footnote 21)

Table 3: Correlation between gender gaps in application and hiring shares

	Correlation
Occupations, 1-digit	0.953
	(0.03)
Occupations, 2-digit	0.863
	(0.03)
Industries, 1-digit	0.960
	(0.02)
Industries, 2-digit	0.938
	(0.03)
Firm wage level, deciles	0.945
	(0.05)
Typical wage level, deciles	0.893
	(0.04)

Notes: In this table we report the correlation between the gender gap in application shares and the gender gap in hiring shares across 1- and 2-digit occupations, across 1- and 2-digit industries, across deciles of the (standardized) firm wage level distribution and across deciles of the typical wage distribution. The results are computed for the main analysis sample after applying propensity score reweighting to condition out differences in observables. Standard errors in parenthesis are calculated by bootstrapping unemployment spells and reestimating propensity scores and gaps (see footnote 21)

Table 4: Decomposing gender gaps

	Baseline	**Explained by application**	**Residual**
Occupational segregation	0.068	0.028	0.040
(Duncan index, 1-digit)	*(0.006)*	*(0.008)*	*(0.007)*
		[0.40]	[0.60]
Occupational segregation	0.105	0.030	0.075
(Duncan index, 2-digit)	*(0.006)*	*(0.008)*	*(0.007)*
		[0.28]	[0.72]
Industry segregation	0.075	0.019	0.056
(Duncan index, 1-digit)	*(0.005)*	*(0.007)*	*(0.006)*
		[0.25]	[0.75]
Industry segregation	0.081	0.016	0.065
(Duncan index, 2-digit)	*(0.006)*	*(0.007)*	*(0.007)*
		[0.20]	[0.80]
Firm wage level	0.068	0.051	0.017
(Male-female gap, std. AKM fixed effect)	*(0.008)*	*(0.013)*	*(0.015)*
		[0.75]	[0.25]
Firm wage level	0.011	0.009	0.002
(Male-female gap, AKM fixed effect)	*(0.001)*	*(0.002)*	*(0.002)*
		[0.81]	[0.19]
Typical wage for job	0.027	0.021	0.006
(Male-female gap, log typical wage)	*(0.002)*	*(0.002)*	*(0.002)*
		[0.79]	[0.21]
Actual Wages	0.056	0.041	0.016
(Male-female gap, log wage)	*(0.005)*	*(0.005)*	*(0.004)*
		[0.72]	[0.28]

Notes: The table decomposes the baseline gaps in hiring outcomes after conditioning on observables. The gaps are decomposed into a part explained by applications and a residual gap (see equation 5). Brackets report the share of the baseline gap explained by each component. Standard errors are reported in parenthesis. They are calculated by bootstrapping unemployment spells and reestimating propensity scores and the decomposition within each sample (see footnote 30). Note that the baseline hiring gap differs slightly from Table 2 due to extra trimming in the decomposition, see footnote 25.

Figure 1: Gender wage gaps

(a) Gender gap in wages across samples

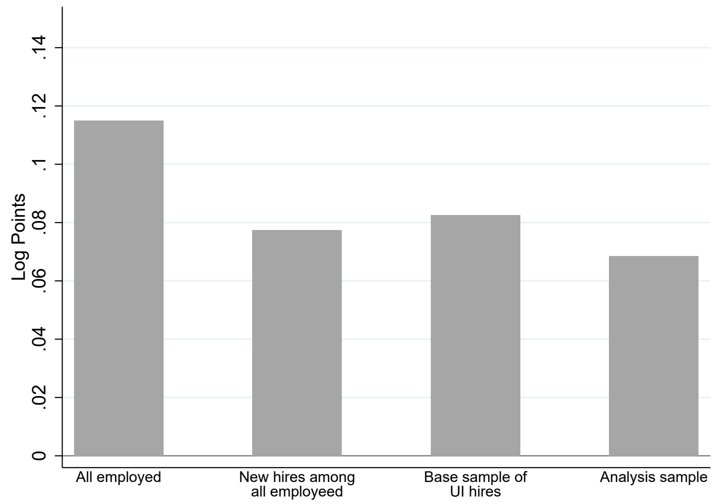

(b) Gender gaps in analysis sample

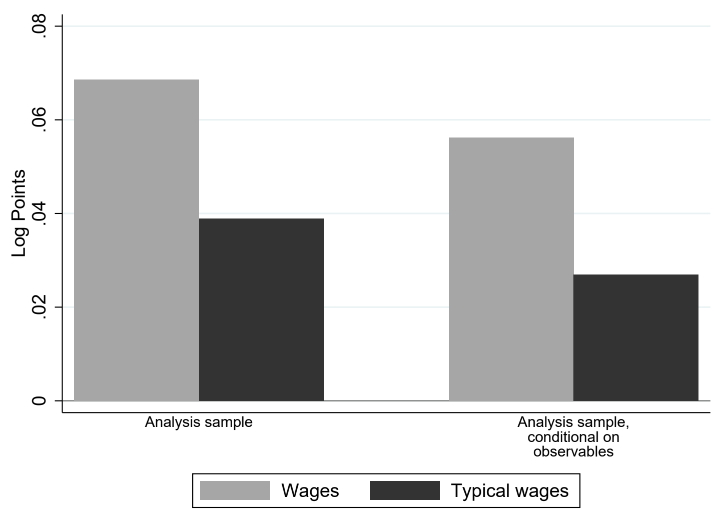

Note: Panel A plots raw (unweighted) gender gaps in wages for employed workers (”All employed”) and for all new employment spells (”New hires”) in august 2015. ”Base sample of UI hires” and ”Analysis sample” corresponds to our sample of unemployed before and after further sample restrictions. Wages are winzorised at the 5th and 95th percentile. Panel B plots raw (unweighted) gender gaps in wages and typical wages for our sample of unemployed (”Analysis sample”). The bars to the right (”Analysis sample, conditional on observables”) control for observable differences between men and women by applying propensity score reweighting.

Figure 2: Gender gaps in applications and hiring outcomes

(a) Occupations

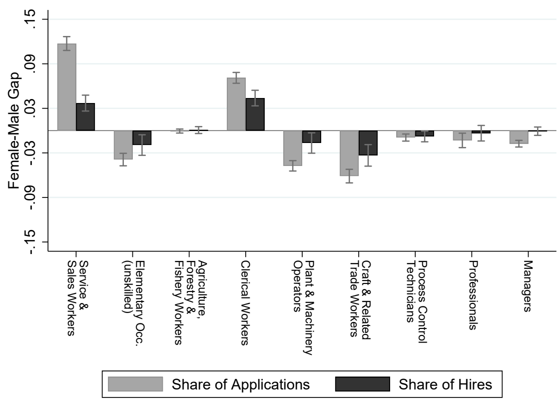

(b) Industries

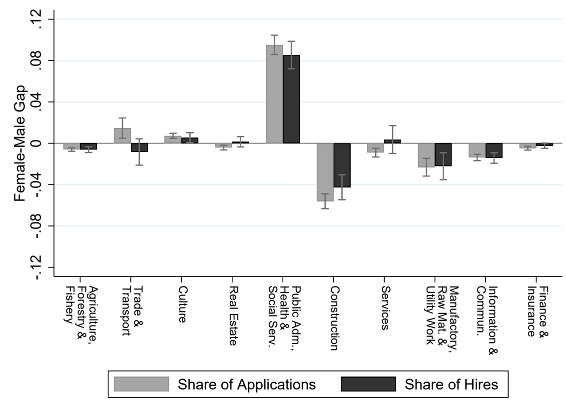

(c) Firm wage level

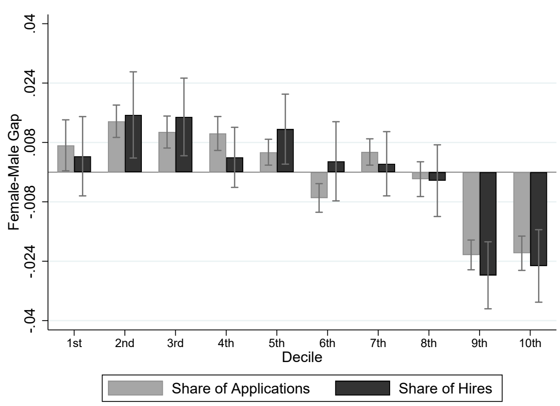

(d) Typical wages

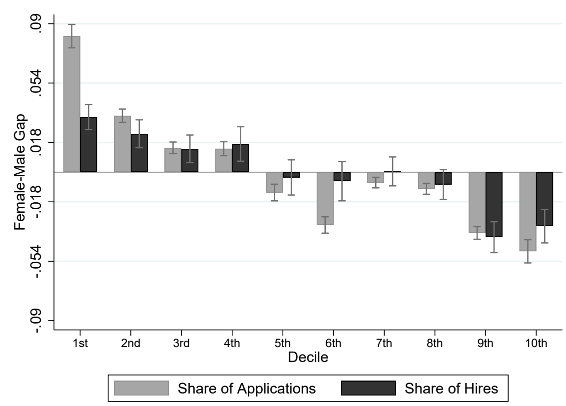

Notes: The figures plot gender gaps in the share of applications going to specific types of jobs and corresponding gender gaps in hiring shares. All gaps are based on the reweighted sample and are therefore conditional on observables. Occupations and industries on the x-axis are ordered from lowest paying on the left to highest paying on the right. The 95 percent confidence bars are based on standard errors clustered on the individual level.

Figure 3: Decomposing hiring share gaps for individual industries and occupations

(a) Occupations

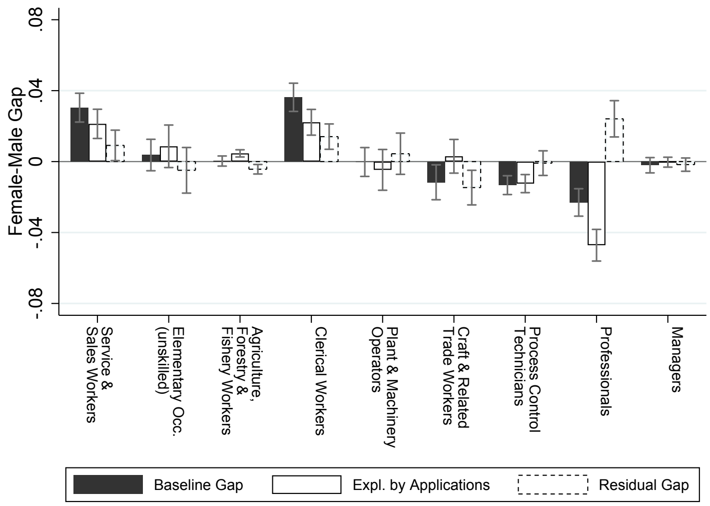

(b) Industries

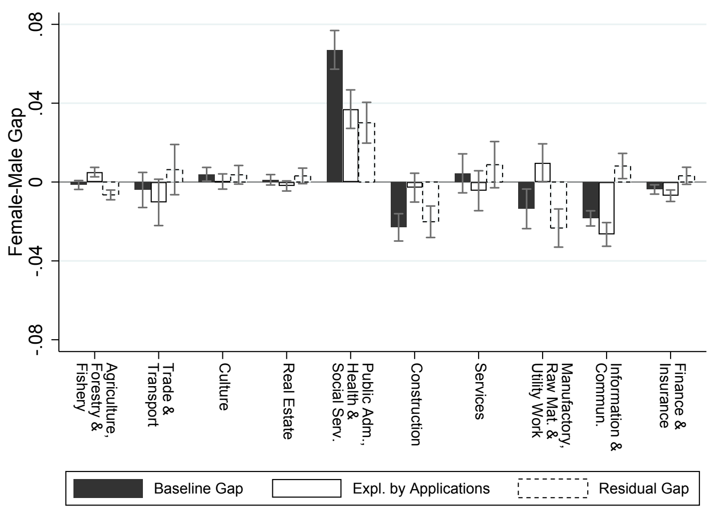

Note: Figure decomposes baseline gaps in the share of men and women hired into different types of jobs after conditioning on labor market observables. The gaps are decomposed into a part which is explained by differences in applications and a residual gap (see Equation 5). Standard errors are reported in parenthesis and are calculated by bootstrapping individuals and reestimating propensity scores and the decomposition within each sample (see footnote 30). Note that the baseline hiring gap differs slightly from Figure 2 due to extra trimming in the decomposition (see footnote 25).

Figure 4: Comparing residual hiring gaps to audit study findings

(a) Residual gaps and estimated call back rate gaps

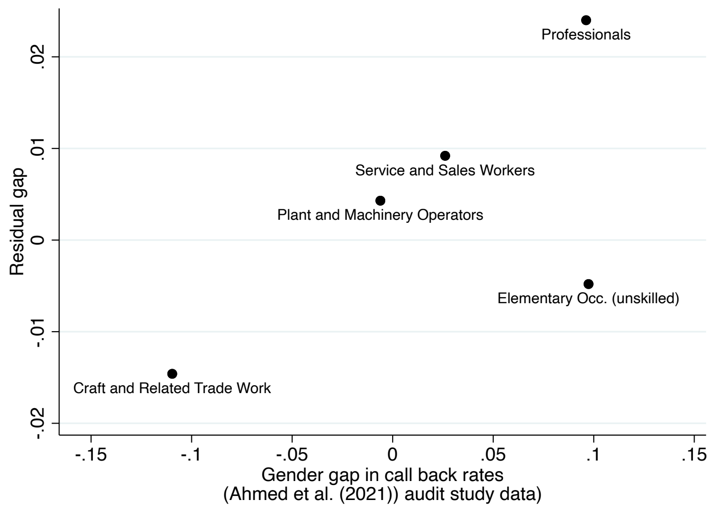

(b) Residual gaps and share female employees

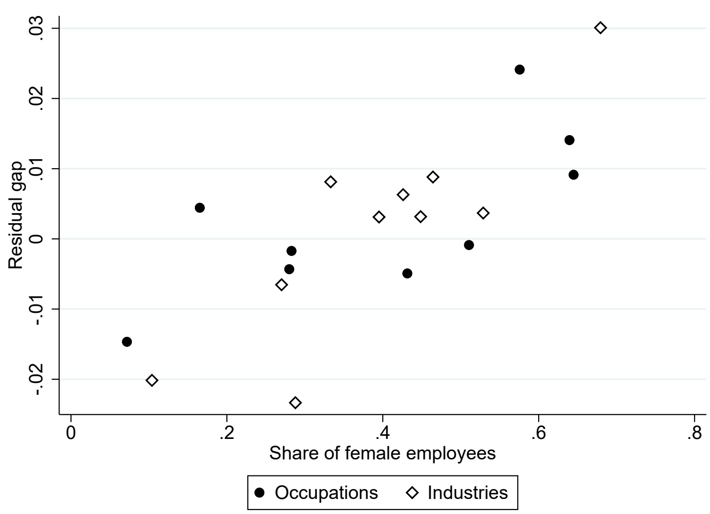

Note: Panel A plots estimates call back rates from Ahmed et al. (2021) against the residual hiring gaps (as defined in Equation 6) across different industries and occupations. Panel B plots the residual hiring gaps for industries and occupations against the share of female employees in the the corresponding industry or occupation.

Figure 5: Gender gaps in applications and hiring across non-wage job characteristics

(a) Part-time vs. full-time

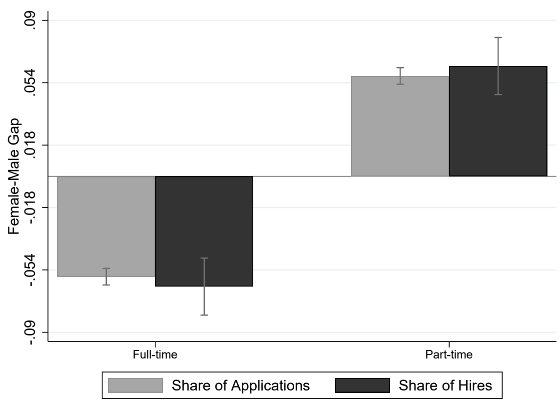

(b) Commuting

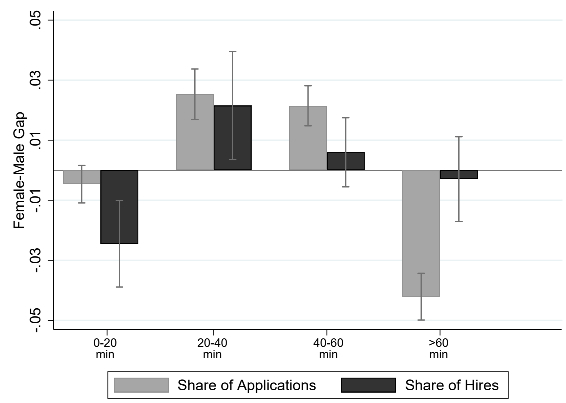

(c) Family-friendliness

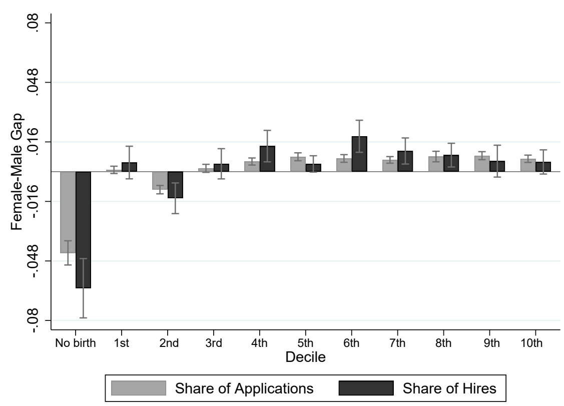

Note: The figures plot gender gaps in shares of applications going to specific types of jobs and corresponding gender gaps in hiring outcomes. Family-friendliness is measured based on the average parental leave taken when an employee or their partner gives birth. The measure is corrected for the gender of the employee. Firms where no employee or partner gives birth is included as a separate category as indicated on the x-axis (see Appendix A.8.1 for details). All gaps are based on the reweighted sample so are conditional on observables. The 95 percent confidence bars are based on standard errors clustered on the individual level.

A Online Appendix: Data and measurements

A.1 Data sets and sample selection steps

Our sample consists of UI recipients of Danish nationality entering new UI spells from September 2015 to September 2017. A new UI spell is defined when an individual who has not received UI benefits in the previous 4 weeks is observed with at least 4 consecutive weeks of UI payouts. For each UI recipient we identify and add all submitted job applications that have been registered in the Joblog system during the unemployment spell. The raw Joblog data is organized in several different databases with information on each edit (and save) of a given application entry. To generate the Joblog data used in this paper, we therefore pre-process these data sources and only include applications that were actually sent, and we further only select the first version of a given Joblog entry.[36]

To construct our final data set we make some additional sample restrictions, the effect of which we show in Table A.1 and also discussed in the main text. First, from the sample of new UI spells, we only consider the UI spells lasting at least 8 weeks. Second, we restrict our sample to individuals who register at least 4 applications in Joblog during their respective unemployment spell. Note that this condition is effective after dropping the last 4 weeks of applications, the individuals therefore need to have at least 4 registered applications during their UI spell. Third, we restrict our sample to individuals who leave UI for employment within the first year of their unemployment spell. Fourth, we drop all applications made in the last four weeks before entering employment. The data shows a drop in the number of applications that people register in Joblog about one month before they enter employment, reflecting that individuals have already accepted their new job at this point and are just waiting for it to start. We therefore drop applications from the last four weeks before the new jobs start based on the median transition time between a successful application and starting a job, as applications made while waiting for the new job to start may not represent an individual's general application behavior.[37]

[36]Besides documenting search activity to qualify for UI, the Joblog section of the Jobnet website was developed with the goal of helping job seekers keep track of their job search. In addition to submitting information on jobs that the worker has applied for, workers can also use the Joblog form to register and keep track of vacancies that the worker is considering applying for in the future and to register other job search events such as being called for an interview or being rejected. However, the coverage of these other events is much lower, and we only use data on the formal job application that UI recipients report that they have applied for.

[37]The median transition time is 6 weeks, but we assume it takes 2 weeks from the application to the eventual job

To the job application data we add data from three administrative data bases: These data sources are IDA, BFL and DREAM. IDA, the Integrated Database for Labor Market Research, is a matched employer-employee panel containing socioeconomic information on the entire Danish population. BFL, the Employment Statistics for Employees, contains monthly data on jobs, paid hours of work and earnings.[38] DREAM, is an event-history data set created by the ministry of employment tracing the participation of individuals in public income support programs at a weekly level. All data sets are available through servers at Statistics Denmark (see `https://www.dst.dk/en/TilSalg/Forskningsservice`). We link applied-for firms to firms in the BFL registers using a sting matching procedure which we explain in Section A.2.

Our final data set allows us to examine the characteristics of the firms and jobs that all male and female UI recipients apply to and compare them to the characteristics of the firms and jobs they eventually end up in.[39] For some of our measures of job characteristics we have missing values for some of the applications, in Tables A.2 we document the prevalence.[40] In the analysis we simply leave out job applications when the job characteristics in question are missing.

In Figure A.1 we plot the survivor function for a version of our main analysis sample where we do not require individuals to find employment within 52 weeks. We also plot the average number of registered applications for each week in unemployment for the main analysis sample. As we discuss in Section 3.3 the average weekly number of applications during the unemployment spell is around 1.5 applications. Finally, in Figure A.2 we report the distribution of the number of submitted applications per week in our initial and final sample. Note that we discuss the dynamics in job

offer. The law requires active job search until the actual UI benefits stop (and the job starts), but naturally the type of jobs applied to are different, and the UI funds explicitly write that permanent positions should not be applied to when future employment is already secured. From 2019 the requirement was changed such that individuals with less than 6 weeks until the beginning of a new job will no longer have to register applications.

[38] The hourly wage measure we use is based on recorded monthly earnings divided by the recorded monthly hours in the job in the first full month of employment. This measure has the highest coverage in our sample as it is available already after one month of employment. The gender wage gap based on e.g. an average over 4 months of earnings is very similar, and the gap is likewise similar when we exclude observations where hours worked have been imputed by Statistics Denmark, see `https://www.dst.dk/da/Statistik/dokumentation/Times/beskaeftigelse-for-loenmodtagere/ajo-loentimer`.

[39] The occupation and industry classifications are available to several degrees of detail, grouped in major, sub-major and minor groups. Occupations are based on the Danish version of the ISCO classifcation (DISCO) and are grouped with 9, 55 or 153 respective occupations (referred to below as 1-, 2- or 3-digit groupings). The industries are based on NACE Rev. 2 and are grouped in 10, 21 or 38 respective industries.

[40] The typical reason for missing job characteristics is that we were unable to link the application to either the firm id or the specific occupation, see Section A.2. The higher shares of missing applications for the industry and firm wage levels (AKM firm fixed effects) reflects that either the firm match was unsuccessful or the firm is so new (small) that e.g. the industry affiliation is not recorded in the employment register (BFL) or it is not a part of the connected set. See also Section A.3.

search further in Section C.1.

A.2 Data matching algorithm

Before matching reported job titles and firms to official classifications and registers, we perform an extensive cleaning of these entries. In this step, we streamline the notation between source and target files and correct obvious spelling mistakes.

As a first step in the actual matching, we use the self-reported job titles and link these to the official Danish occupational codes (DISCO). We exploit that many of the self-reported job titles have the actual occupation as a part of the self-reported title. Thus, as a first step we identify occurrences of the DISCO occupations in the reported job titles. We only consider as 1:1 matches in this step (43.4 percent), i.e. if a certain job title links to several occupations we do not treat it as a match. For remaining unmatched entries, we manually match some job titles to occupations as many job titles use acronyms that do not match to the standard classification.[41] This adds about 27.2 percent to the matches. Finally, we also use some fuzzy matching techniques on the remaining unmatched observations to circumvent misspellings in the job titles, adding the manual titles from the step before. We rank the potential matches along several scoring functions and only pick consistently high-ranked matches. For this we use compget, speedist and soundex routines from SAS as well as sub-string occurrences which adds 10.9 percent. Overall, we can thus map 81.5 percent of the applications to a DISCO group.

As the second matching step, we link the reported firm information to firm identifiers. With the mandatory reporting of firm name, zip code and city, we develop a matching procedure which matches this information to the official firm registers recording all Danish firms (CVR-register).[42] We can then use these links to identify firms in the registers at Statistics Denmark (BFL). Our matching procedure on firms also starts with perfect matches, using both firm name and zip codes. Here we have a 1:1 match for 66.3 percent of the applications in Joblog. We further add the sub-string matches which are spatially the closest to the reported firm address (13.9 percent). To link

[41] For example 'social og sundhedshjælper', Danish for social and health care workers, are most often reported as 'sosu-hjælper'.

[42] The Danish central firm register (CVR-register in short) contains information on companies officially registered in Denmark. The register covers all firms, with the exception of privately held companies with an annual turnover below 50,000 DKK (about 7,500 USD). Each firm is registered with a uniquely identifiable CVR number that's linkable to Danish administrative data sets.

applications which we cannot match exactly on firm names, we employ a fuzzy matching procedure using the matchit command in STATA to identify the 50 closest matches. We then test these 50 potential matches using several scoring functions besides the one obtained from matchit. For each of the scores (5 in total), we calculate the ranking of the 50 potential matches (rank 1 is the best) and identify the "correct" match as the match which receives the best average rank (the scores we use are Bi-gram Similscore, Token, TokenSound from matchit and the compget and speedist functions in SAS). This adds further 6.2 percent to the matches, so we end up with an overall firm match of 86.4 percent.

A.3 Measuring firm wage levels using AKM firm fixed effects

We use our matched employer-employee data to estimate an AKM model (Abowd et al., 1999) and use the estimated firm fixed effects as a measure of the firms wage premium. The AKM model captures implied firm fixed effects on wages, i.e. the firms' wage premium, by identifying moves of workers from one firm to another while simultaneously absorbing individual wage components in worker fixed effects. The separate identification of firm fixed effects relies on the connection between firms in terms of worker movements. Thus, the firm fixed effects can only be recovered for the set of connected firms. We take advantage of the rich administrative data on the whole Danish working population, in particular the BFL data set (see Section A.1) covering monthly salaries, to construct a matched employer-employee panel based on earnings in March and September within a given year. We focus on the years 2008 to 2017 and arrive at 265,425 (connected) firms for whom we get estimates of firm effects (this amounts to around 95 percent of all firms present in the data in the sampling period).

Our estimation uses hourly wages and in addition to worker and firm effects we also include calendar month fixed effects to absorb any aggregate time trends. The hourly wage measure we use is based on recorded monthly earnings divided by the recorded monthly hours. After the estimation, we subtract industry specific averages of firm effects from the estimated firm effects, and divide through by industry specific standard deviations, to ensure that the rankings we obtain account for industry differences in the distribution of firm effects. Further, to guarantee equal size of the decile bins, we employment weight these rankings with the number of employees in each firm as of August 2015, the month before we observe applications in Joblog.[43] For completeness we also report gender gaps based on the non-standardised firm effects when relevant.[44]

A.4 Constructing typical wages

To construct our measure of the typical wage paid in a job with certain characteristics, we use data on the new jobs in our analysis sample to estimate a model that predicts wages based on the characteristics of the job. Since this is a pure prediction problem, we use a LASSO-based machine learning approach. Specifically, we consider a linear regression with log wages as the outcome variable and a very large number of potential explanatory variables based on the available job characteristics in our data. We then use LASSO estimation to select the subset of these variables that most efficiently trades off predictive power in-sample against the risk of overfitting.

As the baseline set of explanatory variables we include dummies for the industry and occupation of the job at both the 1-, 2- and 3- digit levels and the within-industry-standardized AKM firm fixed effect.[45] To handle firm's where no AKM fixed effect can be estimated, we normalize their AKM fixed effect to zero and include a dummy for the fixed effect being missing. Finally, we include all pairwise interactions between the variables for a total of 10,407 baseline explanatory variables. We rely on the Rigorous-LASSO approach of Belloni et al. (2016) to choose the regularization parameters for the LASSO estimation. Because some individuals show up with several UI spells in our data, we allow for clustered disturbances at the individual level in estimation. The estimation was conducted using the LASSOPACK implementation of Ahrens et al. (2019a,b).

Out of the 10,407 baseline explanatory variables, the Rigorous-LASSO selects 233 variables. As the final step, we run a standard OLS regression with log wage as the outcome variable and these 233 variables as explanatory variables (so-called Post-LASSO OLS) to arrive at our final prediction model.[46] Table A.3 summarizes the estimation and final model.

[43] This additionally implies that we ensure that our results are not driven by estimated firm fixed effects from smaller firms which are known to be imprecisely estimated, see e.g. Andrews et al. (2008).

[44] As a robustness check we have also repeated the above steps separately by gender to obtain gender specific firms effects. For these results and further discussion, see Section B.6.2.

[45] Note that the resulting set of dummies thus exhibit perfect multicollinearity by definition. Because of the penalization term in the LASSO objective function, however, perfect does not create a problem for estimation. Including all the dummies, however, is advantageous because the LASSO aims to select a parsimonious model with high predictive power. An obvious candidate for such a model is one that includes dummies for most of the aggregate 1-digit categories but in addition includes dummies for some 2- or 3-digit categories. The final prediction model indeed has this flavor (see Table A.3).

[46] As is common, some of the 233 variables selected by the Rigorous-LASSO turn out to be collinear and thus drop

In Table A.4, we examine the performance of our prediction model. In addition, to validating the use of the prediction model for our analysis, this allows us to draw lessons for future analyses on related data; the need to infer wages from job characteristics arises often when analyzing data on vacancies and/or applied-for jobs (see e.g. Marinescu and Skandalis (2021). To validate model performance we compute an out-of-sample R^2 based on a standard sample-splitting procedure: We randomly mark 80 percent of the new hires data as a training subsample. We then use only this subsample to estimate (train) the model. Finally we use the model to predict for the remaining 20 percent of the data and compute the R^2 for these predictions.

The first row of Table A.4 show that our main LASSO prediction model has an out-of-sample R^2 of 0.201. As a natural benchmark, the second row shows results for a standard OLS approach that simply regresses log wages on the standardized AKM fixed effects as well as industry and occupation dummies at the 3 digit level. Unsurprisingly, this OLS performs reasonably well with an out-of-sample R^2 of 0.192, however, as expected our main LASSO-based approach performs noticeably better. This illustrates the advantage of using a modern Machine Learning approach to prediction. Finally, the last two columns show how much predictive power is lost if the AKM fixed effect is not used or if all firm information is excluded and only occupation information is used. This is informative for settings where linked employer-employee information is not available or where no firm information is available at all. As expected the AKM fixed effect is very important for the models predictive power; dropping just this variable cuts the out-of-sample R^2 down to 0.164. We note, however, that a LASSO model based only on occupations still retains an R^2 of 0.138. In line with previous results relating wages to occupations and job titles, meaningful information about wages can be learned simply from the occupation of a vacant job (Marinescu and Wolthoff (2019)).

After estimating our main LASSO prediction model, we use it to compute our measures of the *typical* wage for new jobs as well as for applied-for job in our application data by simply predicting the log wage from the job's characteristics. Since some applied-for jobs in our application data cannot be linked with firm/or occupation information (see Section A.2), we estimate two alternative prediction models by applying exactly the same procedure as above but excluding either firm or occupation information from the baseline set of explanatory variables. We use these alternative models to fill in typical wages also for the applications that cannot be linked with firm/or occupation

out in the OLS regression. As a result, the total number of variables in the final OLS regression is 117.

information.

We note that the prediction models we use to compute typical wages only ever include job and firm characteristics but never include any worker characteristics. Our measure of the typical wage in a given job thus makes no attempt to capture that workers with different characteristics might face different wages in the same job.[47]Throughout our analysis of gender differences in applications and hiring outcomes we condition on observables so that we are in fact comparing men and women with similar labor market observables. A particular issue arises, however, if the typical wage offered in a particular type of job depends directly on gender (see Section 4.4). As a robustness check we therefore also present results where our measure of typical wages in a given job is based only on the wage paid to either men or women (See Appendix B.6). We do this simply by redoing the final Post-LASSO OLS estimation on either the male or female half of the sample.

In unreported results, we have also experimented with using alternative prediction approaches to the construction of our typical wage measure, including other machine learning approaches or simple linear regressions with a smaller set of variables. Our results are not sensitive to using these other approaches.

A.5 Selecting our set of conditioning variables

To discipline which labor market observables we condition on we follow recent suggestions in the literature (e.g. Angrist and Frandsen, 2022; Athey and Imbens, 2019; Mullainathan and Spiess, 2017) and rely on a Machine Learning procedure. Specifically we use the double-LASSO procedure of Belloni et al. (2014) to select the most important variables for explaining the gender wage gap. We start by specifying a very large baseline set of variables that ex ante could be important to condition on when analyzing the gender wage gap. Using data on all the individuals in our analysis sample, along with the wages in their new jobs, the double-LASSO procedure then involves two separate LASSO regressions: First the LASSO is applied to a regression that has log wage in the new job as the outcome variable and includes the full set of baseline variables as regressors. Intuitively, this step selects out any variables in the baseline set that are relevant predictors of wages. Second, the LASSO is applied to a regression that has a female dummy as the outcome

[47] Formally, our measure of typical wage in a given job reflects the wage paid to the average individual that is hired into this type of job in the data.

variable but again includes the full set of baseline variables as regressors. This step selects out variables that are significant correlates of gender. Combining all the variables selected in each of these two steps then gives the set of most important variables for explaining the gender gap in wages.[48] In each of the steps in the double-LASSO, a data-driven procedure is used to determine the penalty parameter for the LASSO estimation. See Belloni et al. (2014) and Urminsky et al. (2016) for additional discussion and formal results.

The baseline set of potential variables that we include in the two regressions consists of a set of 4,196 variables: To capture educational differences, the set contains years of education, as well as dummies for the field of study. To capture additional differences in general human capital, the set includes age, total work experience and total work experience over the last five years. To capture additional differences in specific human capital, the set includes dummies for the sector, industry and occupation of the previous job as well as continuous measures for the total work experience over the last five years in each of the different industries and occupations. To capture differences in dependence on public transfers, the set includes the total time spent on unemployment insurance, social assistance and other public transfers over the last five years. When including dummies for industry, occupation or field of education and when including continuous measures of industry or occupation-specific experience, we always include all possible measures at both the 1-, 2- and 3-digit level.[49] Finally, all variables are also interacted with both age, years of education, total work experience and work experience over the last five years.[50]

The double-LASSO selects 332 of these variables which we use as our observable characteristics to condition on throughout the main analysis.[51] The estimation was carried out using the

[48]The intuition here is that in order to play a significant role in explaining gender differences in wages a variable has to be strongly correlated with either wages or gender. Variables that are weakly correlated with both however should not play an important role in explaining gender differences in wages. These are exactly the variables excluded by the double-LASSO.

[49]Our coding of occupations, industries and fields of education are based on the official definitions by Statistics Denmark, see also footnote 39. We note that the resulting set of included dummies exhibit perfect multicollinearity by definition. Because of the penalization term involved in the LASSO objective function, however, perfect multicollinearity does not create a problem for estimation. Including all the dummies, however, is advantageous because the LASSO aims to select a parsimonious set of variables with high predictive power. An obvious candidate for such a a set is one that includes dummies for most of the aggregate 1-digit categories but in addition includes dummies for some 2- or 3-digit categories. The final set of selected variables indeed has this feature (see Table A.5).

[50]In implementing this, we allow for variables to be interacted with themselves so that our baseline set includes squared terms in age, years of education, total work experience and work experience over the last five years.

[51]As is common, some of the 332 variables selected by the Rigorous-LASSO turn out to be collinear and thus in practice drop out from our conditioning set when we implement our propensity weighting procedure (see Section A.6). As a result, the total number of variables in the model that we use to estimate propensity scores is 302.

PDSLASSO implementation of Ahrens et al. (2018). Table A.5 summarizes the final selected set of variables that we condition on.

A.6 Propensity score reweighting for descriptive results

As discussed in Section 3.4 and Section A.5, we use propensity score reweighting in all of our analysis. Using the notation introduced in Section 3.4, the reweighting scheme involves reweighting woman i by a weight equal to $\frac{\hat{p}_i}{1-\hat{p}_i}$, where $\hat{p}_i$ is an estimate of the conditional probability of being male given observables, $P(m_i = 1|x_i)$.

After selecting the set of variables to include in our vector of observables, x_i, we follow the standard in the literature and estimate a logit model for the probability of being male, using the variables in x_i as our explanatory variables. We then obtain the $\hat{p}_i$s as the predicted probabilities from this model and use these to reweight the women in our sample.

In Figure A.3 we show the distribution of the estimated propensity scores in our sample. Note that we trim our sample to avoid very small or very large weights, see Table A.1. Specififcally we trim all observations with an estimated propensity score larger than 0.99 or smaller than 0.01.

A.7 Decomposition additional details

In this section we briefly translate the main insights and methodology of DiNardo et al. 1996 and Fortin et al. 2011, showing how propensity score reweighting can be used to construct estimates of the counterfactual hiring probabilities $P^{\tilde{W}}_{A,X}(y)$ and $P^{\tilde{W}}_{X}(y)$ underlying our decomposition exercise and introduced in Section 5.

We start by considering the counterfactual hiring probability for women if they had the same distribution of observables as men.:

$$P^{\tilde{W}}_{X}(y) = \iint P^{W}(y|a,x) f^{W}_{a|x}(a|x) f^{M}_{x}(x) \,\mathrm{d}a \,\mathrm{d}x$$

Multiplying and dividing by $f^{W}_{x}(x)$ inside the integral, we can rewrite this as follows:

$$P^{\tilde{W}}_{X}(y) = \iint P^{W}(y|a,x) f^{W}_{a|x}(a|x) f^{W}_{x}(x) \Psi_X(x) \,\mathrm{d}a \,\mathrm{d}x$$

Here we have defined $\Psi_X(x) = \frac{f^{M}_{x}(x)}{f^{W}_{x}(x)}$. The first insight is that $P^{\tilde{W}}_{X}(y)$ is simply a weighted expectation of $P^W(y|a,x)$ over the set of all women weighted by $\Psi_X(x)$:

$$P_X^{\tilde{W}}(y) = E\left[\Psi_X(x) P^W(y|a,x)|m=0\right]$$

It follows that if the weighting function $\Psi_X(x)$ was known, $P_X^W(y)$ could be estimated by applying the weighting function and then simply computing the share of women hired into job type j in the weighted sample.[52] Now, by an application of Bayes rule, $\Psi_X(x)$ is proportional to a simple function of the conditional probability for being male conditional on observable characteristics x (the propensity score):

$$\Psi_X(x) \propto \frac{P(m=1|x)}{1-P(m=1|x)}$$

It follows that$P_X^{\tilde{W}}(y)$ can be estimated via propensity score reweigthing. We estimate a logit model for the likelihood of being male as a function of our observable characteristics x and then use the predicted probabilities from this model to reweight the women in our sample before computing hiring probabilities. As noted in the main text, this is equivalent to the propensity score reweighting we use to condition out observable differences between men and women, as introduced in Section 3.4.

Next consider the counterfactual hiring probability that women would have faced if they also had the same distribution of application behavior as men:

$$P_{A,X}^{\tilde{W}}(y) = \iint P^W(y|a,x) f_{a|x}^M(a|x) f_x^M(x)\, \mathrm{d}a\, \mathrm{d}x$$

Letting $f_{a,x}^M$ and $f_{a,x}^M$ denote the joint distribution of application behavior and observables for men and women respectively, we rewrite this in a similar way as before:

[52] To see this more clearly, let $I(y)$ be an indicator for ending one's UI spell by being hired into job type y and note that we have:

$$\begin{aligned} P_X^{\tilde{W}}(y) &= E\left[\Psi_X(x) P^W(y|a,x)|m=0\right] = E\left[\Psi_X(x) E\left[I(y)|a,x\right]|m=0\right] \\ &= E\left[E\left[\Psi_X(x) I(y)|a,x\right]|m=0\right] = E\left[\Psi_X(x) I(y)|m=0\right] \end{aligned}$$

The direct empirical counterpart of the last expectation is then the share of women hired into job type y after applying the $\Psi_X(x)$ weights: $\frac{1}{N_W}\sum_{m_i=0} I(y_i)\Psi_X(x_i)$ (here subscript i refers to individuals in the data, and N_W is the total number of women in the data).

$$P^{\tilde{W}}_{A,X}(y) = \iint P^W(y|a,x) f^W_{a,x}(a,x) \Psi_{A,X}(a,x)\, \mathrm{d}a\, \mathrm{d}x$$

Here we have defined $\Psi_{A,X}(a,x) = \frac{f^M_{a,x}(a,x)}{f^W_{a,x}(a,x)}$. Similar to before we see that this implies that the counterfactual hiring probability can be estimated by reweighting the women according to the weighting function $\Psi_{A,X}(a,x)$. Also as before, an application of Bayes rule shows that the weighting function is proportional to a simple function of the conditional probability for being male conditional on both observable characteristics x and application behavior a (a different propensity score):

$$\Psi_{A,X}(a,x) \propto \frac{P(m=1|a,x)}{1-P(m=1|a,x)}$$

It follows that$P^{\tilde{W}}_X(y)$ can also be estimated using propensity score reweigthing. We estimate a logit model for the likelihood of being male as a function of our observable characteristics x and application behavior and then use the predicted probabilities from this model to reweight the women in our sample before computing hiring probabilities.

A.8 Additional details on non-wage characteristics

A.8.1 Measuring family-friendliness

To construct a simple measure of how family-friendly a firm is, we use data on how much parental leave employees at the firm tend to take when they become parents. The basic idea here is that employees will tend to take longer leave if their firms offers more generous parental leave terms and/or are very tolerant towards employees going on leave.[53] Family-friendly firms that offer generous leave packages and are supportive of employees leave-taking should thus see longer parental leave periods among their employees.

We extract information on the duration of leave through the Danish register on sickness ben-

[53] During our sample window, government-mandated parental leave rules are as follows: Mothers are entitled to the following weeks on leave with compensation by the government: 4 weeks of leave just before birth, 14 weeks of maternity leave post birth, and subsequently 32 weeks of parental leave which can also be used by the father. Fathers are further entitled to 2 weeks of paternity leave immediately after birth. The government compensation during leave is at the UI benefit level, however most employment contracts in Denmark offer periods of leave where the worker receives additional compensation by the firm. The duration of these employer coverage periods differ by firms, occupations and collective agreements. During the employer coverage periods the government subsidy is instead paid to the firm, and the firm simply continues to pay the worker his/her wage. When the employer coverage period ends the government payment is redirected and paid directly to the worker.

efit claims (SGDP) which contains all information about benefits paid out by the government in connection with sickness and childbirths. Because this data covers both reimbursements made to firms with workers on leave and payments made directly to the workers on leave, it can be used to infer how much parental leave an employee takes. We focus on benefit claims after 2011 for both men and women and select all payouts which are related to childbirth and where the worker eventually returns to the same pre-birth employer. We then accumulate days with payments within individuals. As our data does not come with readily available information on birthdays for children, we infer these from starting a new parental leave spell (or having more than half a year between payments).[54] Finally, we calculate the average days on leave per birth at the firm level.

In Figure A.4 we plot the distribution of average leave lengths at the firm level. Before computing the the average leave taken at each firm, we correct for the gender of the employee to account for the fact that women take much longer leaves than men. We do this by simply demeaning the leave periods by gender.

Obviously, the average parental leave length will be undefined for firms that do not experience an employee or their partner giving birth at any point during the time period we consider. Since such firms may be likely to be at the bottom of the distribution in terms of family-friendliness,[55] however, we do not exclude them from the analysis but include "no birth" as a separate category when analyzing family-friendliness.

A.8.2 Non-wage characteristics and wage correlations

In Table A.6 we show the result from simple linear regressions where the dependent variable is the typical wage associated with a given application (see Section 3.2) and the independent variables are our selected non-wage characteristics.[56] The sample contains all the the submitted applications, and we cluster standard errors at the level of the unemployment spell. The estimates show that part-time jobs, jobs involving shorter commutes and jobs at family-friendly firms all tend to offer lower typical wages in our data. As a result, if women send more applications to jobs characterized by e.g. shorter commute compared to males, this should translate into gender application gaps

[54] Focusing on mothers this corresponds well with the official birth statistics. Note that some women may transition from one leave to another, thus we censor the length of leave at 14 months and regard any subsequent leave as a new birth/spell. This is however very limited in our data.

[55] This type of firm likely also include rather new firms or firms which are so small that births are rather infrequent.

[56] We remove applications where the relevant job characteristics are missing from the analysis when necessary.

suggesting that women to a larger extent apply to jobs with lower typical wages.

A.9 Further documentation of application behavior and data

A.9.1 Survey about Joblog usage

Table A.7 and A.8 present results from a survey conducted among Danish UI recipients by Mahlstedt et al. (2019) in March 2018.[57] Table A.7 reports survey answers about how individuals log applications in Joblog. Looking at the first column, 41 percent of respondents report that they always log all the jobs they have applied for in Joblog regardless of whether they have fulfilled the logging requirements. An additional 21 percent report that they only log applications up to the point where they have satisfied their logging requirements but that they rarely apply for more jobs than what is required. Putting these together suggest that Joblog has close to full coverage for 63 percent of respondents. For the remaining 37 percent, however, the survey responses suggest that the Joblog data often misses some job applications that they have made beyond the required number. The second and third columns report corresponding numbers by gender. These show a similar pattern overall, although men are somewhat more likely to say that they often apply to jobs that they do not register. 42 percent of men say that they often apply to more jobs than they register, while only 32 percent of women say so.

To get a sense of how many applications may be missed by the Joblog data, Table A.8 presents survey responses about the total number of job applications sent the past month and the number of job applications sent that were not registered in Joblog. In addition, the Table also shows the actual number of registered Joblog applications made by the survey respondents in the month before the survey. This was computed by linking survey responses with the actual Joblog data. On average, survey respondents report applying for 11.5 jobs in total over the past month. Of those jobs, survey respondents on average say they failed to register 2.4 jobs in Joblog. This suggest that Joblog covers 80 percent of actual applications. The bottom of the table instead shows that average number of jobs respondents actually registered in Joblog was 8.0. Relative to the total number of reported applications, this suggest that respondents on average failed to register 3.5 applications, implying that Joblog on average covers 69 percent of all applications.[58] The table also reports

[57] We thank the authors for making this data available.

[58] This difference could reflect imperfect recall among survey respondents or could relate to measurement error from

separate numbers for men and women. Women self-report failing to register 2.1 jobs on average, implying that joblog covers 82 percent of applied-for jobs for women. Alternatively, comparing total reported applications to actual joblog registrations, suggest that women on average fail to register 3.8 applications, corresponding to a coverage of 68 percent. Corresponding calculations for men suggest that Joblog covers between 72 and 76 percent of applied-for jobs for men.

A.9.2 Coverage and representativeness of the Joblog data

In this section we present additional evidence and validity checks regarding the quality of the Joblog application data. In doing so we exploit the fact that the data is linked to actual job outcomes.

First, we verify how the application data relates to actual hiring outcomes. If the Joblog application data accurately capture actual application behavior, we would expect the application data to be highly predictive of the type of job that each UI recipient ends up being hired into. To assess whether this is the case, we benchmark the predictive value of the application data against a known strong predictor of job outcomes: the characteristics of UI recipients previous job.

Table A.9 compares how the Joblog application data and prior job characteristics predict respectively the industry, the occupation, the firm wage level or the typical wage level of a UI recipient's new job. Each column of the table corresponds to a different prediction model estimated on our analysis sample. When predicting the industry of the new job, we use a simple multinomial logit model that includes either dummies for the industry of the previous job, the share of applications going to each industry or both sets of variables. Similarly, when predicting the occupation of the new job we use a multinomial logit model that includes either dummies for the occupation of the previous job or the share of applications going to each occupation. For both industries and occupations, we exclude one small industry/occupation for which the model obtains near perfect predictions for a few observations.[59] When predicting the firm wage level or the typical wage of

the timing of registered jobs and/or the precise interpretation of the survey question. Registering applied-for jobs in Joblog can be done retroactively so the interpretation of the survey question could either refer to the date at which applications were sent or to the date at which the application was entered into the Joblog system. Additionally, the fact that UI recipients are able to register other activities besides formal job applications introduces some ambiguity about the interpretation of the survey question (if for example UI recipients have registered that they reached out to a friend about a specific job).

[59]Specifically, we exclude individuals from the sample who find a job in the smallest industry or occupation, respectively, as well as individuals whose prior job was in this industry or occupation. Results are almost identical if these observations are included, however, we see indications that the likelihood function becomes ill-behaved in some specificaions in this case, reflecting that some observations are predicted nearly perfectly.

the new job, we use a simple linear regression that includes either the firm wage level or the typical wage of the previous job, or includes the mean of the firm wage level or typical wage across the applied-for jobs. For the linear regression models, we measure the predictive power simply using the regression R^2. For the multinomial logit models, we use McFadden's *pseudo*-R^2.

Looking across Table A.9, we see that models that predict job outcomes only using application data perform quite similarly to models that instead use prior job characteristics. Application data does do markedly worse than prior job characteristics when predicting the occupation of the new job (column (4) vs (5)) but only slightly worse for firm wage level and the typical wage (columns (7) and (10) vs. (8) and (11)). At the same time, application data actually does better than prior job characteristics when predicting the industry of the new job (column (1) vs. (2)).

For models that include both prior job characteristics and application data (columns (3), (6), (9) and (12)), we see that application data remains highly predictive even after prior job characteristics have been conditioned on. Adding the application variables alongside prior job characteristics always leads to sizeable increases in the (*pseudo*-)R^2 relative to models that only use prior job characteristics. Moreover the application variables are always highly statistically significant in the combined models. Overall, we conclude that the Joblog application data is highly predictive of later job outcomes.

In Tables, A.10 and A.11, we repeat the prediction exercise separately for the men and women in the data. Results are very similar and we see little indication of systematic differences in the predictiveness of job applications for later job outcomes. While job application information appears to predict occupational outcomes slightly better for men, they instead appear to predict wage outcomes slightly better for women. Moreover, the differences are small throughout.

As a further check on the representativeness and quality of the Joblog application data, we examine how often we are able to trace a new hire back to a job application that is contained in our data. For 47 percent of the new hires, we are able to identify a previous application that the UI recipient sent to the firm in question. This is informative about the representativeness of the data. To see why, assume that the Joblog data covers a share r of all applications and that the share of applications that we successfully match to firms in our data matching procedure is s. In this case, our data will contain firm information for a share $s \cdot r$ of all applied-for jobs. Next assume that the fraction of jobs that stem from a job application is j. If the applied-for jobs in our data are a representative subset of all applied-for jobs, the share of new hires that we should be able to trace back to an application, t, should then be:

$$t = j \cdot r \cdot s$$

Based on independent survey data from Table A.8 we estimated that the raw Joblog data contain between 69 and 80 percent of all applied-for jobs, that is r is between 0.69 and 0.80. Furthermore, as described in Section 3.1, $s = 0.86$ in our data matching procedure. Finally, to gauge the share of hires that stem from a job application, j, we rely on Statistics Denmark's official survey *Arbejdskraftsundersøgelsen* on how unemployed Danes report landing their first job out of unemployment (Engman (2019)). In these data, 11 percent of respondents report landing their job in a way that is very unlikely to have involved the worker applying for the job (the job resulted from work at a temp agency, they got the job via their educational institution as an internship or the job seekers themselves advertised publicly), while 58 percent of respondents report landing their job in a way that almost surely involved making a formal job application (they themselves applied to a posted position, they applied to a firm with no posted positions or they were directed to the job by the employment agency or other authorities).[60] To arrive at an estimate for the fraction of hires that stem from workers applying for the job, we simply assume that half of the remaining jobs involved a job application. This implies that about 73 percent of new hires out of unemployment involve the worker applying for the job at some point so that $j = 0.73$.[61]

Plugging in these values, we see that if the applied-for jobs in our data is representative, the share of new hires that we should be able to match, t, should be between $0.73 \cdot 0.68 \cdot 0.86 = 0.43$ and $0.73 \cdot 0.80 \cdot 0.86 = 0.50$. As noted, we in fact have $t = 0.47$ in our data, consistent with the data containing a representative subset of all applied-for jobs.

Finally, we can also compute how often we are able to trace a new hire back to a job application that is contained in our data for men and women respectively. For 41 percent of all new male hires, we are able to identify a previous application that the UI recipient sent to the firm in question.

[60] The remaining respondents report landing their jobs through channels that may have involved applying for the job application but may also have involved receiving a job offer more directly. This includes finding the job through an acquaintance or finding a job after having been contacted by the firm.

[61] Alternatively, we could use 0.58 as a lower bound on r and use 0.89 as an upper bound. Plugging into the formulate above, we then see that if the applied-for jobs in our data is representative, the share of new hires that we should be able to match, t, should be between $0.58 \cdot 0.69 \cdot 0.86 = 0.34$ and $0.89 \cdot 0.80 \cdot 0.86 = 0.61$.

For women, the corresponding number is 53 percent. A likely reason for this difference is that Danish men are more likely to find jobs in ways that do not involve a formal job application (see e.g. Engman (2019) for evidence of this for the overall labor market). This mechanically implies that we should be able to match a smaller share of new male hires to a job application in our data. Because we do not have reliable data on how often this occurs for unemployed men and women, however, we cannot assess quantitatively how the gender-specifc match rates match up with our other data.

Table A.1: Sample selection and trimming

	Individuals	Spells	Applications
Inflow	**227.515**	**261.529**	**7.422.148**
Minimum 8 weeks spell length	177.145	194.660	7.019.513
Spells with > 4 applications	170.304	185.959	7.008.284
Employment within 52 weeks	105.879	114.375	3.439.690
Censoring last 4 weeks of applications	105.879	114.375	2.911.585
Analysis sample	**105.879**	**114.375**	**2.911.585**
After trimming for descriptive analysis†	100.268	108.173	2.790.250
After trimming for decomposition ††	93.550	100.504	2.631.684

Notes: The top of the table shows the number of individuals, unemployment spells as well as number of applications in the base UI inflow data and after applying each of our four sample restrictions. The bottom of the table shows how the analysis sample changes when we apply trimming to remove observations with extreme propensity scores. † refers to the trimming we use in our descriptive analysis where propensity scores are estimated based only on labor market observables (see Section 3.4). †† refers to the additional trimming we use in the decomposition exercise where propensity scores are estimated based on both labor market observables and application behavior (see Section 5).

Table A.2: Share of missing job characteristic

	Firm ID	**Occupation**	**Industry**	**Firm Wage Level**	**Typical Wage**	**Wage**
Application	0,185	0,193	0,329	0,330	0,068	1,000
Hires	0,001	0,083	0,059	0,060	0,000	0,102

Notes: The table shows the share of missing job characteristics both for applications and for the jobs that UI recipients are hired into in the analysis sample.

Table A.3: Prediction model summary

Model and estimation summary:	
Explanatory variables in baseline set:	10,407
Variables selected in Rigorous-LASSO:	233
Parameters in final Post-LASSO OLS model:	117
R^2 for Post-LASSO OLS model (in sample):	0.202

Summary of selected variables:
Standardized firm fixed effect
7 dummies for 1-digit occupations
7 dummies for 2-digit occupations
17 dummies for 3-digit occupations
7 dummies for 1-digit industries
7 dummies for 2-digit industries
7 dummies for 3-digit occupations
161 occupation-industry interactions
9 occupation-firm fixed effect interactions
10 industry-firm fixed effect interactions

Notes: The table summarizes the main prediction model used to construct the measures of typical wages. The difference between the number of selected variables in the Rigorous LASSO and the number of parameters in the final Post-LASSO OLS model reflect that some of the selected variables are perfectly multicollinear (see footnote 46).

Table A.4: Prediction model performance

	Out-of-sample R^2	Estimator	Variables
Main LASSO model	0.201	Post-LASSO OLS (Rigorous penalty)	All occupation dummies (1, 2 and 3 digits), All industry dummies (1, 2 and 3 digits), AKM fixed effects + pairwise interactions
OLS benchmark	0.192	OLS	Occupation dummies, 3-digits Industry dummies, 3-digits AKM fixed effects
LASSO, no AKM fixed effects	0.164	Post-LASSO OLS (Rigorous penalty)	All occupation dummies (1, 2 and 3 digits), All industry dummies (1, 2 and 3 digits), + pairwise interactions
LASSO, occupation only	0.138	Post-LASSO OLS (Rigorous penalty)	All occupation dummies (1, 2 and 3 digits)

Notes: The table shows model performance for four different prediction models. The models are either based on OLS or on Post-LASSO OLS using the rigorous penalty approach of Belloni et al. (2016). The models differ also in terms of the set of baseline variables included, as shown in the last column. The out-of-sample R^2 of each model was computed based on an 80-20 train-test split of the data.

Table A.5: Summary of observables selected in double-LASSO

Non-interacted continous variables:	
	Age, years of education
	2 continous measures of experience from 1-digit industries
	9 continous measures of experience from 2-digit industries
	12 continous measures of experience from 3-digit industries
	4 continous measures of experience from 1-digit occupations
	1 continous measures of experience from 2-digit occupations
	32 continous measures of experience from 3-digit occupations
Non-interacted dummy variables:	
	4 dummies for the 1-digit occupation of most recent job
	13 dummy for the 2-digit occupation of most recent job
	29 dummies for the 3-digit occupation of most recent job
	3 dummies for the 1-digit industry of most recent job
	6 dummies for the 2-digit industry of most recent job
	10 dummies for the 3-digit industry of most recent job
	2 dummies for the sector of the most recent job
	3 dummies for education field at the 1-digit level
	8 dummies for education field at the 2-digit level
	17 dummies for education field at the 3-digit level
Interactions involving only continous variables:	
	5 interactions involving only combinations of age, years of education or experience
	2 interactions involving continous measures of past receipt of public transfers
	47 interactions involving continous measures of experiences from specific occupations
	10 interactions involving continous measures of experiences from specific industries
Interactions involving discrete variables:	
	42 interactions involving dummies for the occupation of most recent job
	16 interactions involving dummies for the industry of most recent job
	3 interactions involving dummies for the sector of the most recent job
	50 interactions involving dummies for field of education

Notes: The table summarizes the set of 332 variables selected in the double-LASSO. Interaction terms always involve either age, years of education, total work experiene or work experience over the last five years as one of the two interacted variables. We also allow for variables to be interacted with themselves, corresponding to a squared term (see footnote 50).

Table A.6: Correlation between job characteristics and typical wages

	(1) Log typical wage	(2) Log typical wage	(3) Log typical wage	(4) Log typical wage
Commute	0.000240***			0.000227***
(minutes)	(4.78e-06)			(4.58e-06)
Part-time		-0.0712***		-0.0704***
(indicator)		(0.000584)		(0.000576)
Family-friendliness			-0.000143***	-0.000143***
(days of leave)			(3.70e-06)	(3.60e-06)
Constant	5.152***	5.179***	5.180***	5.173***
	(0.000474)	(0.000412)	(0.000445)	(0.000522)
Observations	2,638,100	2,638,073	2,638,100	2,638,073
R-squared	0.016	0.024	0.007	0.047

Notes: ***p<0.01, **p<0.05, *p<0.1. The table shows the result of regressing the log typical wages of an applied-for job in the main anlysis sample on different non-wage job characteristics. Jobs with missing data on any of the involved variables have been dropped in all specificatons. Standard errors are clustered at the spell level.

Table A.7: Survey question "Which of these statements best describes your use of Joblog?"

Answer:	**Share Overall:**	**Share Men:**	**Share Women:**
Fulfill requirements, often applied to more jobs	0.36	0.42	0.32
Fulfill requirements, rarely applied to more jobs	0.21	0.19	0.23
Always register all applied-for jobs	0.41	0.38	0.44
Never register applications	0.1	0.2	0.1
Number of respondents	1236	515	721

Notes: The table shows answers to the question "Which of these statements best describes your use of Joblog?" based on the survey of UI recipients conducted in Mahlstedt et al. (2019).

Table A.8: Self-reported and registered applications in the previous month

	Mean, overall	**Mean, men**	**Mean, women**
Survey answers			
# of applied-for jobs	11.5	11.2	11.7
# of applied-for jobs not registered	2.4	2.7	2.1
Joblog data			
# of applied-for jobs	8.0	8.1	7.9

Notes: The to part of the table shows the reported number of job applications sent over the last month and the reported number of these jobs applications that were not registered in Joblog based on the based on the survey of UI recipients conducted in Mahlstedt et al. (2019). The bottom part of the table shows the actual number of jobs registered in Joblog by the survey respondents in the month prior to the survey.

Table A.9: Predicting job outcomes from application data vs. prior job characteristics

Job outcome: Model:	Industry (1 digit) Multinomial logit			Occupation (1 digit) Multinomial logit			Firm wage level Linear regression			Typical wage Linear regression		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
Explanatory variables												
Characteristics of previous job	No	Yes	Yes	No	Yes	Yes	No	Yes	Yes	No	Yes	Yes
Characteristics of applied-for jobs	Yes	No	Yes	Yes	No	Yes	Yes	No	Yes	Yes	No	Yes
# of parameters	99	99	189	63	63	119	2	2	2	2	2	2
(pseudo-)R-squared	0,292	0,245	0,368	0,331	0,395	0,497	0,056	0,077	0,112	0,268	0,285	0,392
p-value, test of excluding applied-for job variables	$< 0{,}01$		$< 0{,}01$	$< 0{,}01$		$< 0{,}01$	$< 0{,}01$		$< 0{,}01$	$< 0{,}01$		$< 0{,}01$

Notes: The table examines the predictiveness of job applications and past job characteristics for the sample overall. Columns (1)-(3) correspond to multinomial logit models for the 1-digit industry of the UI recipients new job. Explanatory variables in these models are dummies for the 1-digit industry of the previous job or the share of job applications sent to jobs in each 1-digit industry. Columns (4)-(6) correspond to multinomial logit models for the 1-digit occupation of the UI recipients new job. Explanatory variables in these models are dummies for the 1-digit occupation of the previous job or the share of job applications sent to jobs in each 1-digit occupation. Specifications in Columns (1)-(6) exclude individuals who found a job in the smallest industry or occupation, or who had their previous job in this industry or occupation (see footnote \ref{fn:PredValExcl}. Columns (7)-(9) correspond to linear regressions where the outcome variable is the (standardized) firm fixed effect for the UI recipients new job. Explanatory variables in these models are the firm fixed effect for the UI recipients previous job or the average firm fixed effect across all the applied-for jobs. Columns (10)-(12) correspond to linear regressions where the outcome variable is the typical wage of the UI recipients new job. Explanatory variables in these models are the typical wage of the UI recipients previous job or the average typical wage across all the applied-for jobs. The table reports the R^2 for the linear regression models. For the multinomial logit models, the table reports the McFadden's pseudo-R^2. The last row of the table show the p-value for testing the exclusion of all explanatory variables pertaining to applied-for jobs.

Table A.10: Predicting job outcomes from application data vs. prior job characteristics, men only

Job outcome:	Industry (1 digit)			Occupation (1 digit)			Firm wage level			Typical wage		
Model:	Multinomial logit			Multinomial logit			Linear regression			Linear regression		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
Explanatory variables												
Characteristics of previous job	No	Yes	Yes	No	Yes	Yes	No	Yes	Yes	No	Yes	Yes
Characteristics of applied-for jobs	Yes	No	Yes	Yes	No	Yes	Yes	No	Yes	Yes	No	Yes
# of parameters	99	99	189	63	63	119	2	2	2	2	2	2
(pseudo-)R-squared	0,276	0,249	0,362	0,300	0,402	0,486	0,049	0,094	0,121	0,196	0,267	0,346
p-value, test of excluding applied-for job variables	$< 0,01$		$< 0,01$	$< 0,01$		$< 0,01$	$< 0,01$		$< 0,01$	$< 0,01$		$< 0,01$

Notes: The table examines the predictiveness of job applications and past job characteristics for the male half of the sample. Columns (1)-(3) correspond to multinomial logit models for the 1-digit industry of the UI recipients new job. Explanatory variables in these models are dummies for the 1-digit industry of the previous job or the share of job applications sent to jobs in each 1-digit industry. Columns (4)-(6) correspond to multinomial logit models for the 1-digit occupation of the UI recipients new job. Explanatory variables in these models are dummies for the 1-digit occupation of the previous job or the share of job applications sent to jobs in each 1-digit occupation. Specifications in Columns (1)-(6) exclude individuals who found a job in the smallest industry or occupation, or who had their previous job in this industry or occupation (see footnote \ref{fn:PredValExcl}. Columns (7)-(9) correspond to linear regressions where the outcome variable is the (standardized) firm fixed effect for the UI recipients new job. Explanatory variables in these models are the firm fixed effect for the UI recipients previous job or the average firm fixed effect across all the applied-for jobs. Columns (10)-(12) correspond to linear regressions where the outcome variable is the typical wage of the UI recipients new job. Explanatory variables in these models are the typical wage of the UI recipients previous job or the average typical wage across all the applied-for jobs. The table reports the R^2 for the linear regression models. For the multinomial logit models, the table reports the McFadden's pseudo-R^2. The last row of the table show the p-value for testing the exclusion of all explanatory variables pertaining to applied-for jobs.

Table A.11: Predicting job outcomes from application data vs. prior job characteristics, women only

Job outcome:	Industry (1 digit)			Occupation (1 digit)			Firm wage level			Typical wage		
Model:	Multinomial logit			Multinomial logit			Linear regression			Linear regression		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
Explanatory variables												
Characteristics of previous job	No	Yes	Yes	No	Yes	Yes	No	Yes	Yes	No	Yes	Yes
Characteristics of applied-for jobs	Yes	No	Yes	Yes	No	Yes	Yes	No	Yes	Yes	No	Yes
# of parameters	99	99	189	63	63	119	2	2	2	2	2	2
(pseudo-)R-squared	0,271	0,213	0,339	0,328	0,367	0,483	0,063	0,063	0,107	0,307	0,281	0,409
p-value, test of excluding applied-for job variables	$< 0{,}01$		$< 0{,}01$	$< 0{,}01$		$< 0{,}01$	$< 0{,}01$		$< 0{,}01$	$< 0{,}01$		$< 0{,}01$

Notes: The table examines the predictiveness of job applications and past job characteristics for the female half of the sample. Columns (1)-(3) correspond to multinomial logit models for the 1-digit industry of the UI recipients new job. Explanatory variables in these models are dummies for the 1-digit industry of the previous job or the share of job applications sent to jobs in each 1-digit industry. Columns (4)-(6) correspond to multinomial logit models for the 1-digit occupation of the UI recipients new job. Explanatory variables in these models are dummies for the 1-digit occupation of the previous job or the share of job applications sent to jobs in each 1-digit occupation. Specifications in Columns (1)-(6) exclude individuals who found a job in the smallest industry or occupation, or who had their previous job in this industry or occupation (see footnote \ref{fn:PredValExcl}. Columns (7)-(9) correspond to linear regressions where the outcome variable is the (standardized) firm fixed effect for the UI recipients new job. Explanatory variables in these models are the firm fixed effect for the UI recipients previous job or the average firm fixed effect across all the applied-for jobs. Columns (10)-(12) correspond to linear regressions where the outcome variable is the typical wage of the UI recipients new job. Explanatory variables in these models are the typical wage of the UI recipients previous job or the average typical wage across all the applied-for jobs. The table reports the R^2 for the linear regression models. For the multinomial logit models, the table reports the McFadden's pseudo-R^2. The last row of the table show the p-value for testing the exclusion of all explanatory variables pertaining to applied-for jobs.

Figure A.1: Survival rates and registered applications

(a) Kaplan-Meier survivor functions in non-employment (unweighted)

(b) Average number of logged applications

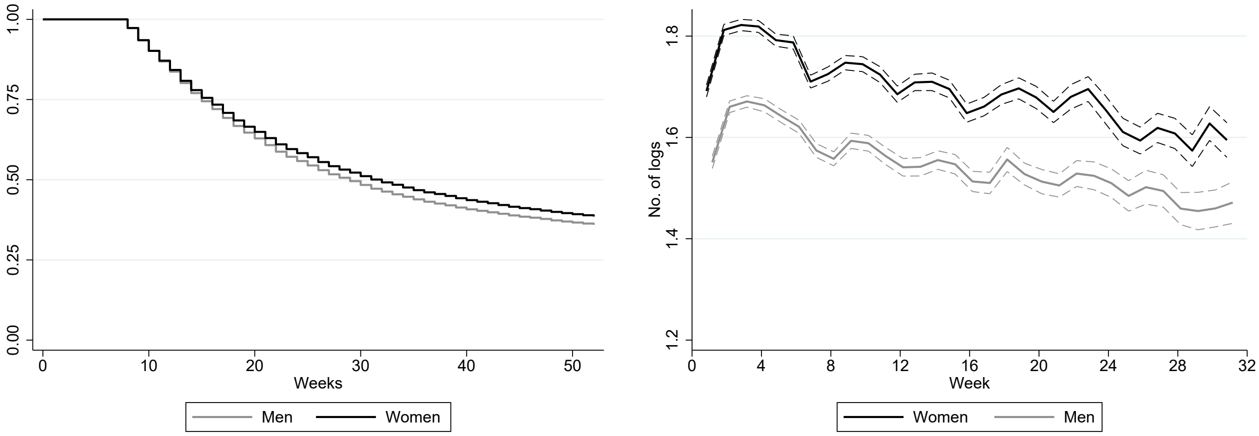

Note: The figures plot gender-specific Kaplan-Meier survival rate in nonemployment estimates (left) and the average number of registered applications (right). The X-axis measure weeks since the start of the UI spell. The Kaplan-Meier estimates of the survivor function in nonemployment are estimated on a version of the main analysis sample where the requirement of finding a job within a year has not been imposed. Average number of registered applications is shown for the main analysis sample.

Figure A.2: Distribution of registered applications per week

(a) Baseline sample

(b) Analysis sample

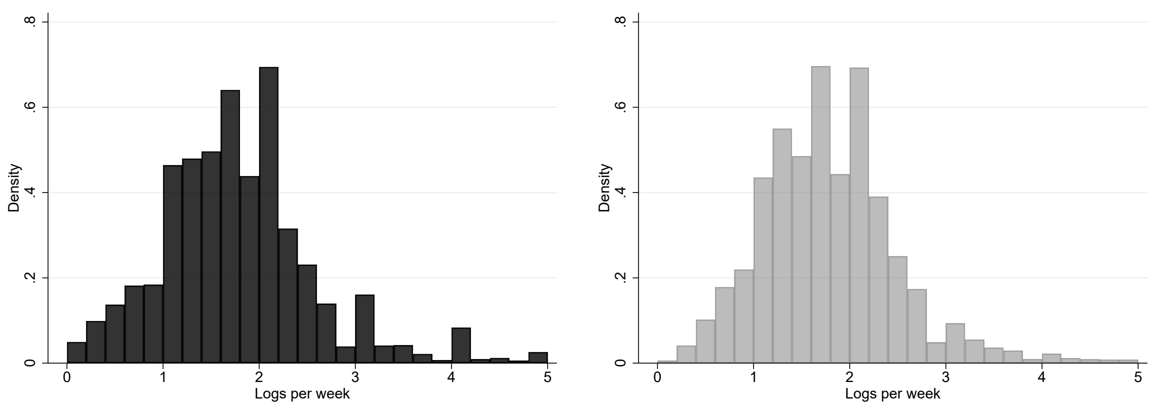

Note: Figure plots the distribution of average applications per week for the unrestricted sample (left) and the analysis sample (right).

Figure A.3: Distribution of estimated propensity scores for descriptive analysis

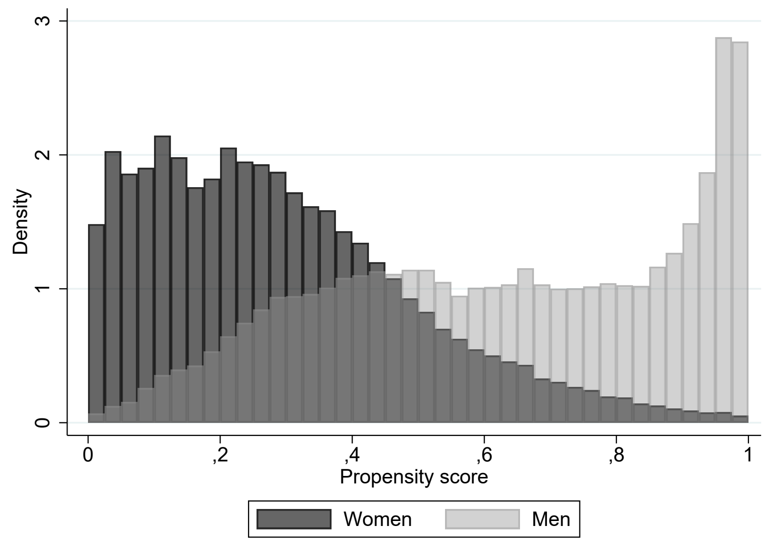

Note: Figure plots the distribution of male propensity score estimates for men and women. Propensity scores outside the range [0.01,0.99] have been trimmed to avoid extreme weights.

Figure A.4: Parental leave: Distribution of average length

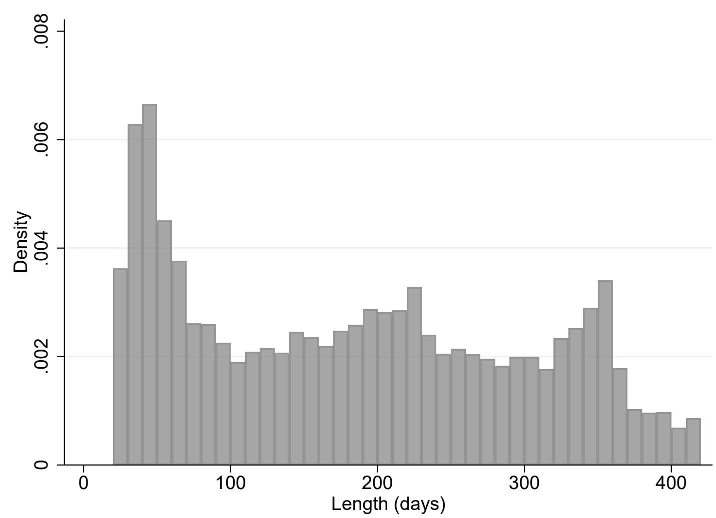

Notes: The figure plots the distribution of the average length of parental leave in days across firms with more than 28 days of leave on average.

B Online Appendix: Robustness checks

B.1 Alternative sample definition

We impose several restrictions on our main sample as laid out in Section 3.1. Below we show that our findings are not sensitive to these restrictions. Figures B.1 to B.5 replicate the female-male gaps in average application and hiring shares for samples based on different sample selection criteria than the ones used in our main analysis. Figure B.1 is based on a sample where we do not exclude the last 4 weeks of applications, whereas in Figure B.2 we focus on unemployment spells that find a job within 26 weeks (in contrast to the 52 week requirement used in the main text). In Figure B.3 we relax our sample restriction of only including unemployment spells with at least 4 registered applications by selecting all spells that have at least one application instead.[62] Figure B.4 removes the restriction of at least 8 weeks of unemployment to enter the sample, and instead include all available unemployment spells in the sample period.[63] Figure B.5 replicates the results for a sample that is not restricted to end in a new hire. Here we treat unemployment as a separate category to the hiring outcomes. Common to all of our robustness tests on the sample selection criteria is that the results do not change qualitatively. In fact, female-male gaps are remarkably stable across the different samples.

B.2 Displacement sample

In this section we repeat our main results for a subsample of individuals who are identified as being a part of a mass-layoff using standard definitions from the literature. The motivation for doing so is to asses the role of differential selection into UI by gender. Our main sample does not have any restrictions on the cause of entry into UI (besides of course UI eligibility) so in principle there may, for example, be slightly more males who enter UI more or less voluntary (through a quit) whereas women to larger extent enter involuntary (layoff). If the cause of entry into UI also affects application and hiring behavior this may contribute to the gender application gap. By focusing on individuals who enter unemployment through a "mass layoff" we try to focus on a group of individuals with a similar cause of entry.

[62]Note that registering at least one application is necessary to appear in the Joblog data we use.

[63]Nevertheless, these unemployment spells need to be at minimum 4 weeks long in order for us to properly identify them, see Section 3.1.

To identify spells in our estimation sample (see Section A.1) which can be linked to a mass layoff we use standard definitions from the literature (see e.g. Lachowska et al. (2020); Bertheau et al. (2022)). We only consider spells where the pre-displacement firm have at least 5 workers employed prior to the displacement event and where they experience a reduction in firm size of more than 30 percent between the displacement year (i.e. the year where the unemployment spell start) and the subsequent year.[64] We focus on all spells who do not return to the pre-displacement firm within our sample window, and we further require that the worker worked in the pre-displacement firm for at least 2 consecutive years prior to unemployment. In the end we identify 6217 spells (around 6 percent of our sample after trimming) satisfying the above criteria and for these spells we then proceed by creating application and hiring gaps.

In figure B.6 we report application and hiring gaps for the subsample of spells which can be linked to a mass layoff. Compared to Figure 1 in the main text we see broadly similar patterns , however precision is markedly lower in the displacement sample due to the reduction in sample size. This suggests that our main results are not explained by differential selection at entry into UI.

B.3 Conditioning on different observables

In the main text, our descriptive analysis conditions on observables by propensity score reweighting on a set of variables selected through a LASSO procedure. We have, however, also experimented with several other approaches to conditining on observables. None of these change our conclusions. In this section we present results from some of these alternative approaches.

In Figure B.7 we present results after propensity score reweighting only on the 3-digit industry of the previous job, thus imposing exact balance on previous industry across men and women. Similarly, in Figure B.8 we present results after propensity score reweighting only on the 3-digit occupation of the previous job, thus imposing exact balance on previous occupation. Finally, in Figure B.9 in which we use the same set of conditioning variables as the main analysis but include dummies for the quarter of entry into UI in the propensity score estimation, see Section 3.4. The purpose of this is to control for seasonality, i.e. whether entering the sample at different times is

[64] Since average firm size is low in Denmark, requiring firms to have more employees would reduce the sample size quite dramatically.

important for the differences in application behavior and hiring outcomes we observe.

Throughout the various alternative approaches, we see a similar pattern of gender gaps as in the main text.

B.4 Raw gender gaps in application and hiring outcomes

In Figure B.10, we show raw gender gaps in applications and hiring outcomes without conditioning on unobservables. We see that the overall patterns are similar to the conditional results presented in the main text but that, unsurprisingly, the raw gaps tend to be much larger in magnitude.

B.5 Gender gaps in number of applications instead of shares

In our main analysis we measure gender gaps in the share of applications going to different jobs. In Figure B.11 we instead show gender gaps in the absolute number of applications sent to different jobs. We see that the overall patterns of results from the main analysis remains.

B.6 Gender-specific measures of wages

B.6.1 Jobs' typical wage level

One drawback of the measure of a job's typical wage level that we use in the main analysis is that it does not allow for the possibility that men and women face different wages in the same job. A particular concern here is the possibility that some types of jobs tend to pay high wages to women but not men, while for other jobs it is the reverse. If this is the case we might expect women to apply much more for jobs in which they face higher wages and vice versa for men.

To address this possibility, this section presents results based on two alternative measures wages that capture the typical wage level faced by either women or men in a given job. We construct these measures simply by repeating the last step of the wage prediction exercise underlying our typical wage only for the male or female half of the sample (see Appendix A.4).

Figure B.12 shows gender gaps in applications and hiring splitting jobs according to the typical wage level they pay to either women or men. Although the exact size of the application and hiring gaps change slightly, the overall pattern of results is identical to what we see in the main analysis.

B.6.2 Gender-specific AKM firm effects

In similar spirit to Section B.6 this section contains results based on gender specific measures of firm effects thereby analyzing the importance of potential gender differences in firm effects and thus rankings across firms. We construct these measures simply by repeating our AKM estimation only for the male or female half of our AKM estimation sample (see Appendix A.3). After estimation we then rank all new jobs and applications according to the estimated firm effects and construct application and hiring gaps. The results are presented in Figure B.13. We see very similar overall patterns regardless of whether we rank based on firm effects estimated for men or women only. Although the exact size of the application and hiring gaps change slightly, the overall pattern of results is very similar to what we see in the main analysis. Note that precision is lower since the set of connected firms (i.e. firms for whom we can actually estimate firm effects) is substantially reduced when we restrict the sample to only males or females.[65]

B.7 Decomposing raw gender gaps

In this section, we present a decomposition of the raw gender gaps in hiring outcomes. We do this by simply forgoing the first step of the decomposition used in the main analysis: starting from the raw gender gaps in hiring outcomes in the data, we propensity score reweight the women in the sample to the same application behavior as men to see how much of this raw gap applications can explain.

Table B.1 shows the decomposition. Unsurprisingly, we see that raw gender gaps before conditioning on observables are generally larger than the baseline hiring gaps considered in the main analysis. We also see that applications are capable of explaining a larger share of these raw gaps. The only exception to this is firm-wage level, where the raw gap is smaller than the baseline gap and where applications are capable of explaining a smaller share of the raw gap.

[65]Coverage is lower for firm effects estimated for women only. Compared to the firm effects estimated based on the full sample we lose 33 percent of firm effects when estimation is based on women only. The same number for men is 14 percent.

Table B.1: Decomposing gender gaps without conditioning on observables

	Baseline	**Explained by application**	**Residual**
Occupational segregation	0.228	0.178	0.050
(Duncan index, 1-digit)	*(0.003)*	*(0.006)*	*(0.006)*
		[0.78]	[0.22]
Occupational segregation	0.351	0.225	0.126
(Duncan index, 2-digit)	*(0.003)*	*(0.007)*	*(0.006)*
		[0.64]	[0.36]
Industry segregation	0.253	0.218	0.036
(Duncan index, 1-digit)	*(0.003)*	*(0.006)*	*(0.006)*
		[0.86]	[0.14]
Industry segregation	0.263	0.205	0.058
(Duncan index, 2-digit)	*(0.003)*	*(0.006)*	*(0.006)*
		[0.78]	[0.22]
Firm wage level	0.026	0.012	0.014
(Male-female gap, std. AKM fixed effect)	*(0.004)*	*(0.017)*	*(0.017)*
		[0.46]	[0.54]
Firm wage level	0.015	0.012	0.003
(Male-female gap, AKM fixed effect)	*(0.001)*	*(0.002)*	*(0.001)*
		[0.81]	[0.19]
Typical wage for job	0.039	0.034	0.005
(Male-female gap, log typical wage)	*(0.001)*	*(0.002)*	*(0.002)*
		[0.87]	[0.13]
Actual Wages	0.069	0.055	0.014
(Male-female gap, log wage)	*(0.002)*	*(0.003)*	*(0.003)*
		[0.80]	[0.20]

Notes: The table decomposes the raw gaps in hiring outcomes. The gaps are decomposed into a part explained by applications and a residual gap. Standard errors based on bootstrapping individuals are shown in parenthesis. Brackets report the share of the raw gap explained by each component. To ease comparability, the sample is the same as the sample used for the decomposition in Table 4.

Figure B.1: Gender gaps in applications and hiring outcomes, no exclusion of last 4 weeks

(a) Occupations

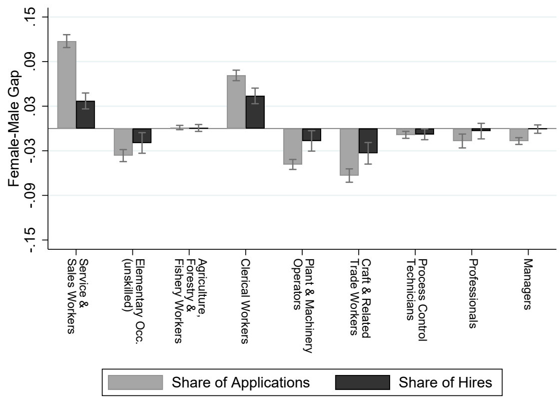

(b) Industries

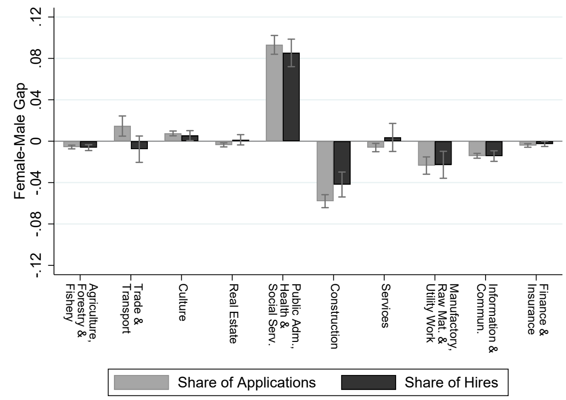

(c) Firm wage level

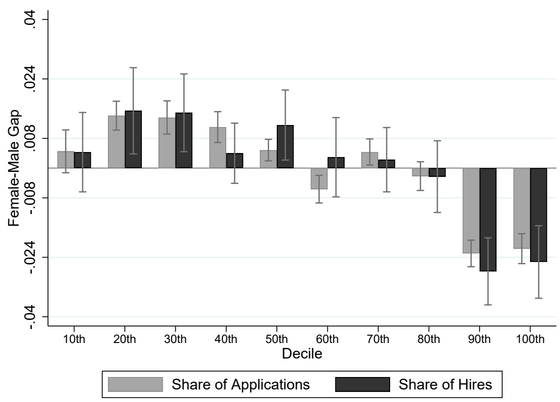

(d) Typical wage

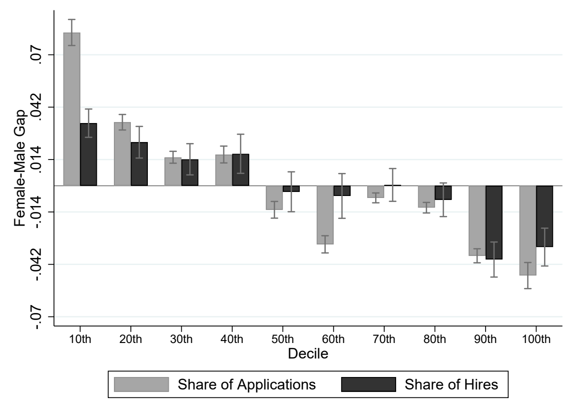

Notes: The figures plot gender gaps in the share of applications going to specific types of jobs and corresponding gender gaps in hiring shares. The figure relaxes the sample restriction on the last 4 weeks of unemployment and consider all applications. All gaps are based on a reweighted sample so are conditional on observables. Occupations and industries on the x-axis are ordered from lowest paying on the left to highest paying on the right. The 95 percent confidence bars are based on standard errors clustered on the individual level.

Figure B.2: Gender gaps in applications and hiring outcomes, 26 week sample

(a) Occupations

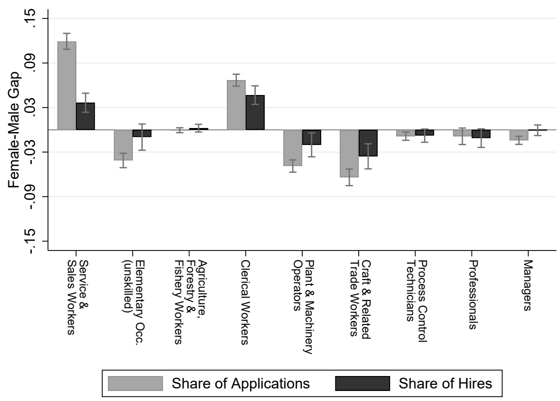

(b) Industries

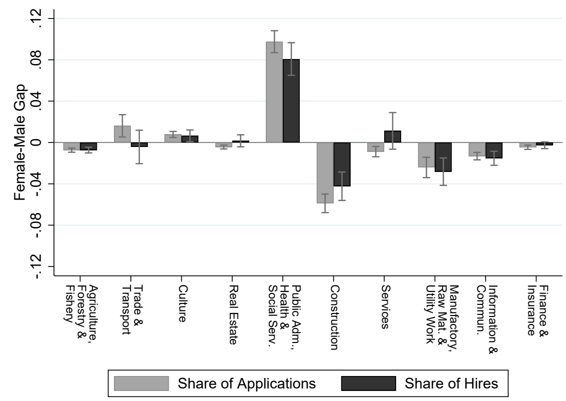

(c) Firm wage level

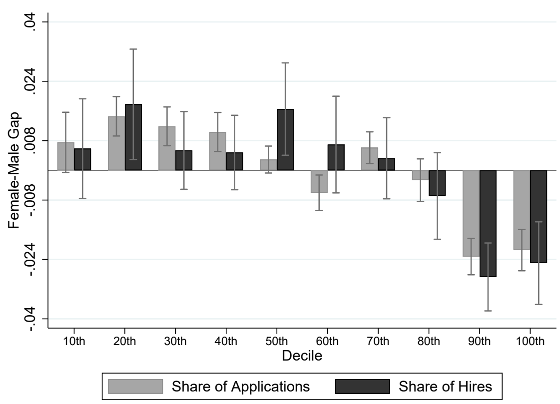

(d) Typical wage

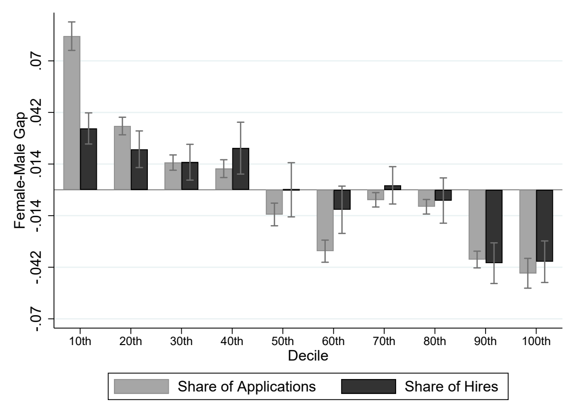

Notes: The figures plot gender gaps in the share of applications going to specific types of jobs and corresponding gender gaps in hiring shares. The figure changes the sample to only consider spells lasting at most 26 weeks. All gaps are based on a reweighted sample so are conditional on observables. Occupations and industries on the x-axis are ordered from lowest paying on the left to highest paying on the right. The 95 percent confidence bars are based on standard errors clustered on the individual level.

Figure B.3: Gender gaps in applications and hiring outcomes, only 1 application registration requirement

(a) Occupations

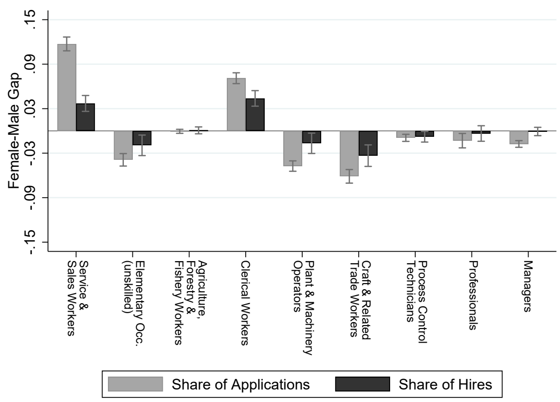

(b) Industries

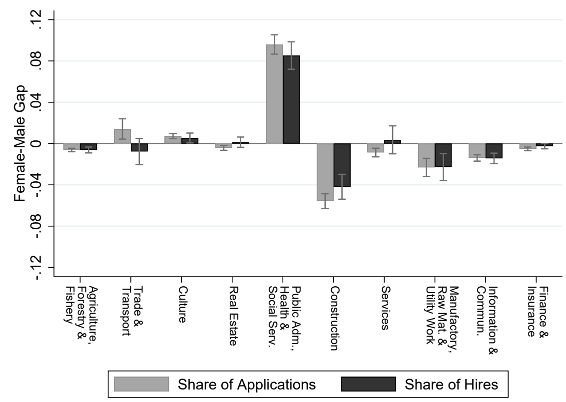

(c) Firm wage level

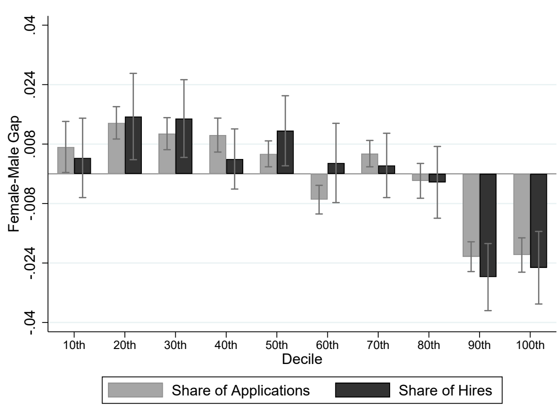

(d) Typical wage

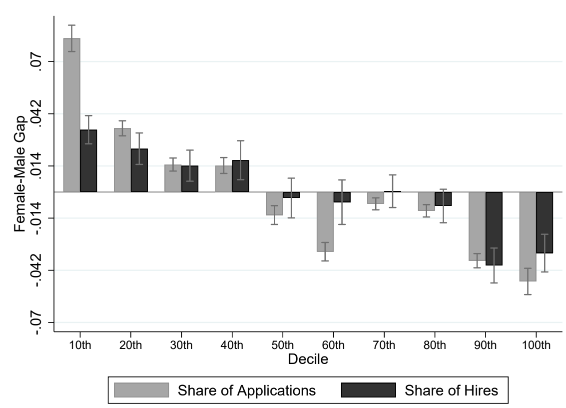

Notes: The figures plot gender gaps in the share of applications going to specific types of jobs and corresponding gender gaps in hiring shares. The figure relaxes the samples requirement of having min. 4 registered applications to at least 1 application. All gaps are based on a reweighted sample so are conditional on observables. Occupations and industries on the x-axis are ordered from lowest paying on the left to highest paying on the right. The 95 percent confidence bars are based on standard errors clustered on the individual level.

Figure B.4: Gender gaps in applications and hiring outcomes, no spell length requirement

(a) Occupations

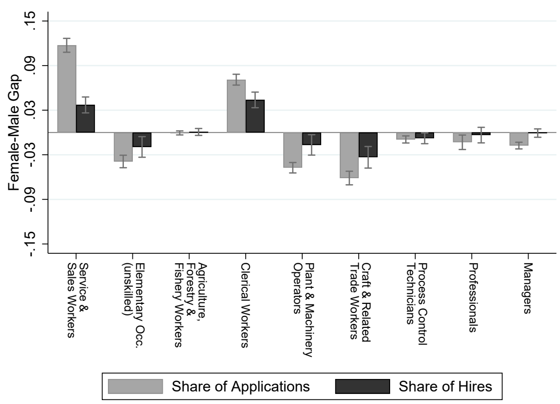

(b) Industries

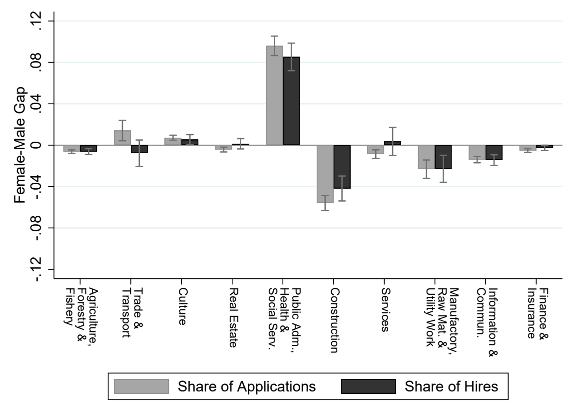

(c) Firm wage level

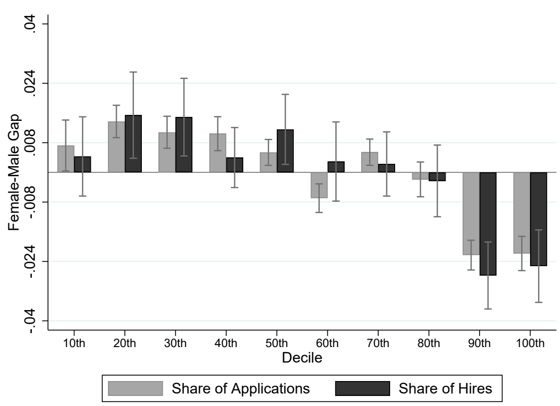

(d) Typical wage

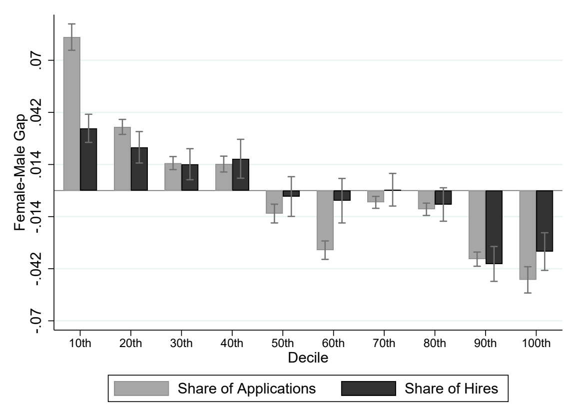

Notes: The figures plot gender gaps in the share of applications going to specific types of jobs and corresponding gender gaps in hiring shares. The figure relaxes the sample restriction on min. 8 weeks spell length and considers all spells. All gaps are based on a reweighted sample so are conditional on observables. Occupations and industries on the x-axis are ordered from lowest paying on the left to highest paying on the right. The 95 percent confidence bars are based on standard errors clustered on the individual level.

Figure B.5: Gender gaps in applications and hiring outcomes, no employment requirement

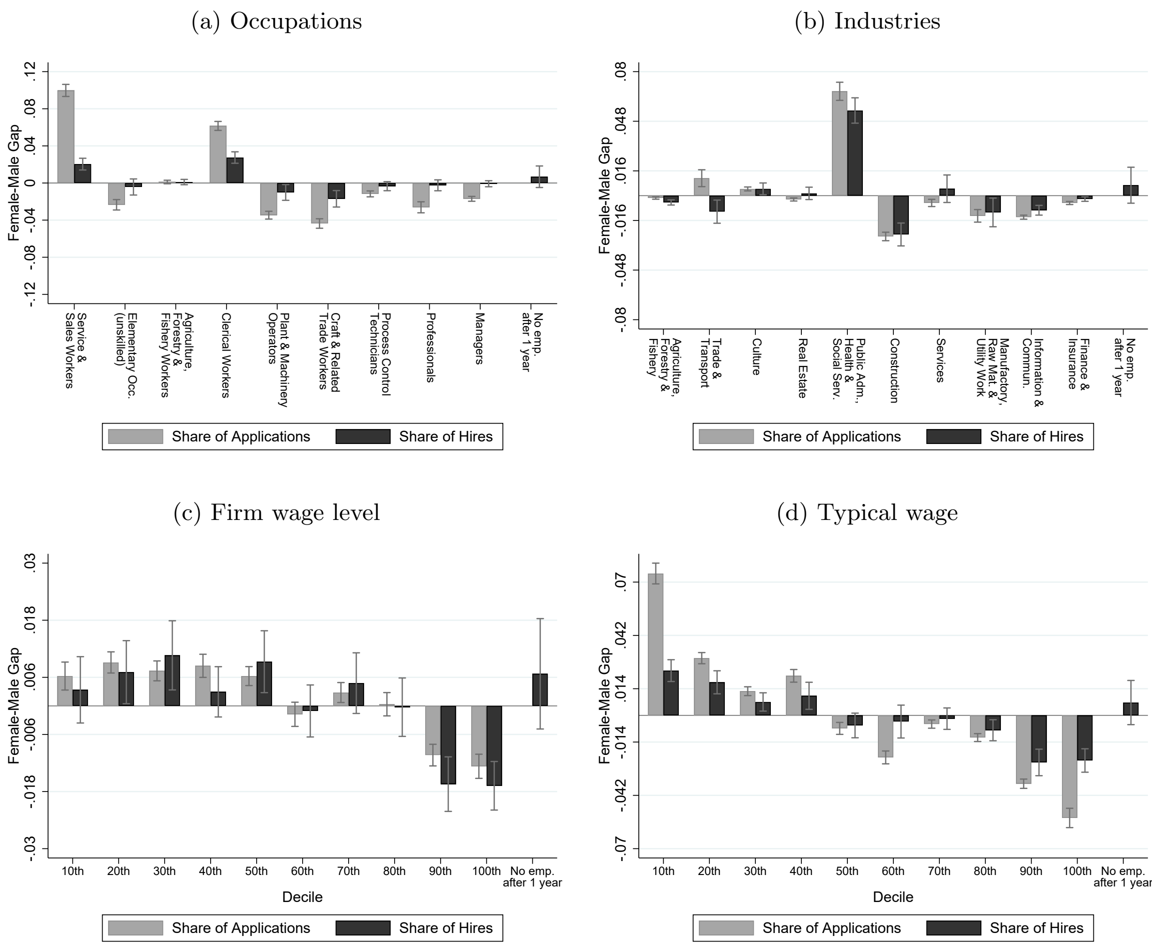

Notes: The figures plot gender gaps in the share of applications going to specific types of jobs and corresponding gender gaps in hiring shares. The sample also includes spells that do not end in employment. Staying unemployed is treated as a separate category throughout. All gaps are based on a reweighted sample so are conditional on observables. Occupations and industries on the x-axis are ordered from lowest paying on the left to highest paying on the right. The 95 percent confidence bars are based on standard errors clustered on the individual level.

Figure B.6: Gender gaps in applications and hiring outcomes, conditioning only on previous industry

(a) Occupations

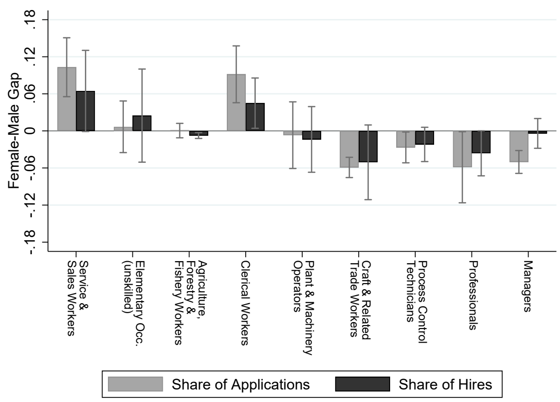

(b) Industries

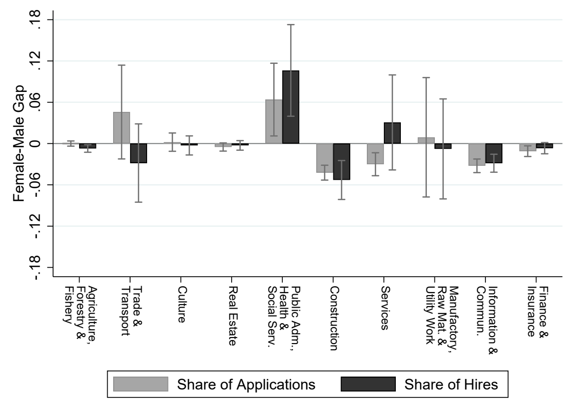

(c) Firm wage level

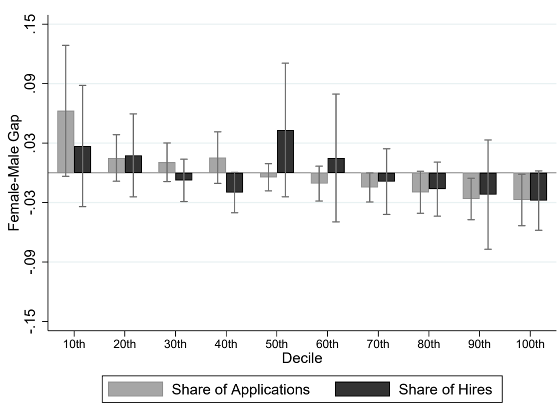

(d) Typical wage

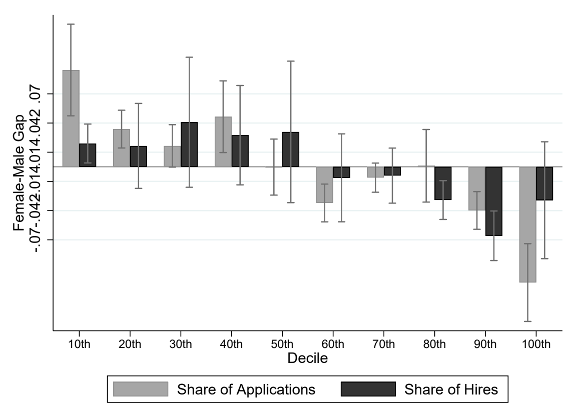

Notes: The figures plot gender gaps in the share of applications going to specific types of jobs and corresponding gender gaps in hiring shares. The sample consists of spells who can be linked to a mass-layoff, see Section B.2. All gaps are computed after propensity score reweigthing on only the 3-digit industry of the previous job. Occupations and industries on the x-axis are ordered from lowest paying on the left to highest paying on the right. The 95 percent confidence bars are based on standard errors clustered on the individual level.

Figure B.7: Gender gaps in applications and hiring outcomes, conditioning only on previous industry

(a) Occupations

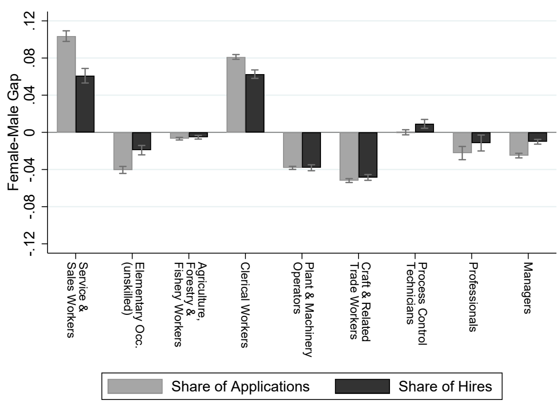

(b) Industries

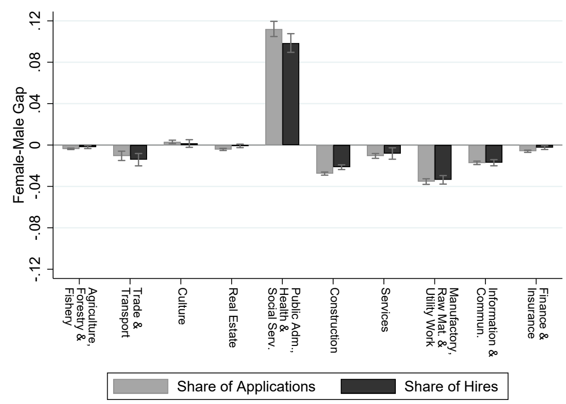

(c) Firm wage level

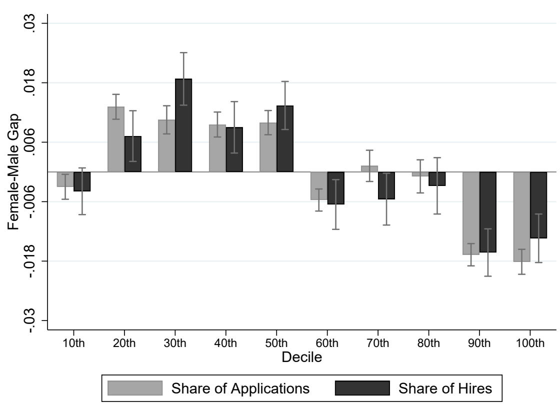

(d) Typical wage

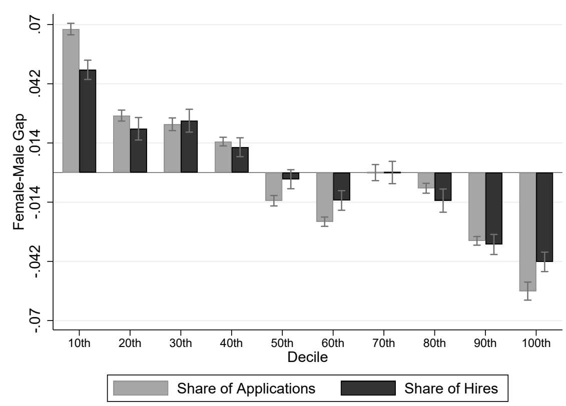

Notes: The figures plot gender gaps in the share of applications going to specific types of jobs and corresponding gender gaps in hiring shares. All gaps are computed after propensity score reweigthing on only the 3-digit industry of the previous job. Occupations and industries on the x-axis are ordered from lowest paying on the left to highest paying on the right. The 95 percent confidence bars are based on standard errors clustered on the individual level.

Figure B.8: Gender gaps in applications and hiring outcomes, conditioning only on previous occupation

(a) Occupations

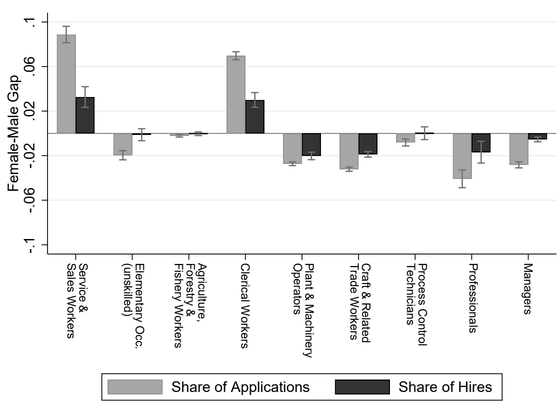

(b) Industries

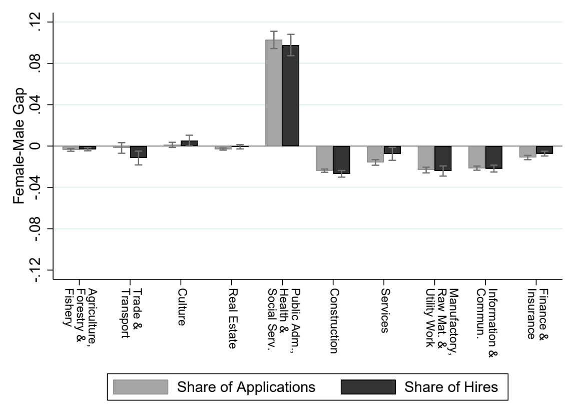

(c) Firm wage level

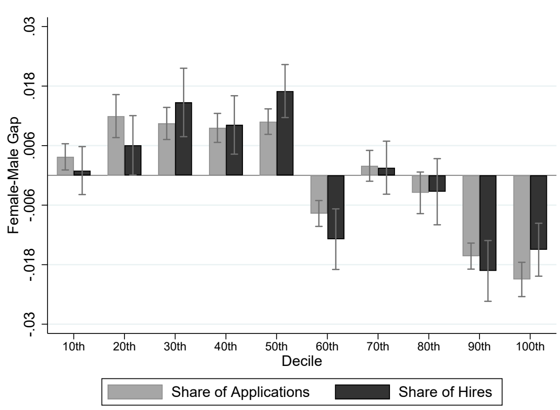

(d) Typical wage

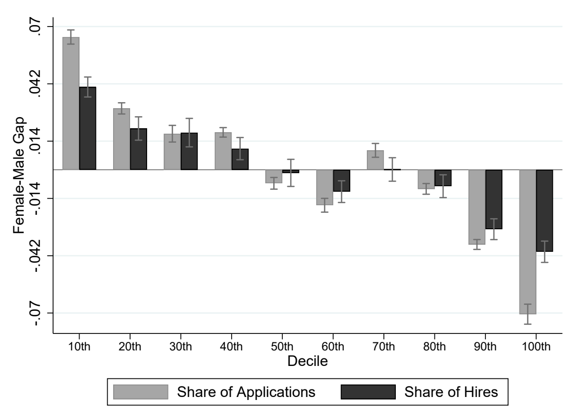

Notes: The figures plot gender gaps in the share of applications going to specific types of jobs and corresponding gender gaps in hiring shares. All gaps are computed after propensity score reweigthing on only the 3-digit occupation industry of the previous job. Occupations and industries on the x-axis are ordered from lowest paying on the left to highest paying on the right. The 95 percent confidence bars are based on standard errors clustered on the individual level.

Figure B.9: Gender gaps in applications and hiring outcomes, seasonality controls

(a) Occupations

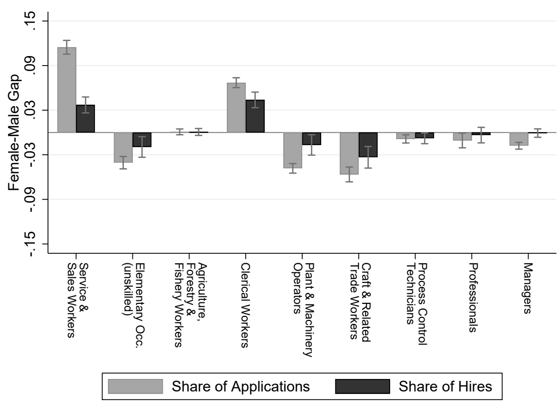

(b) Industries

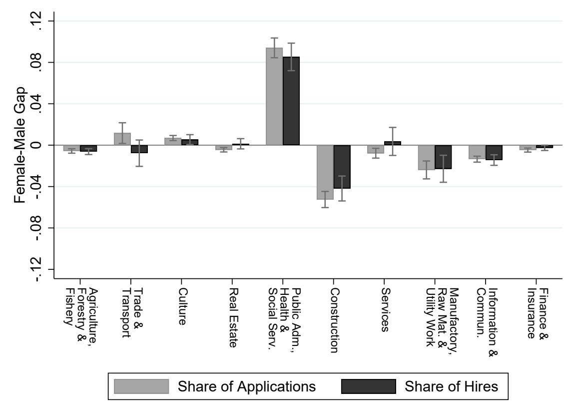

(c) Firm wage level

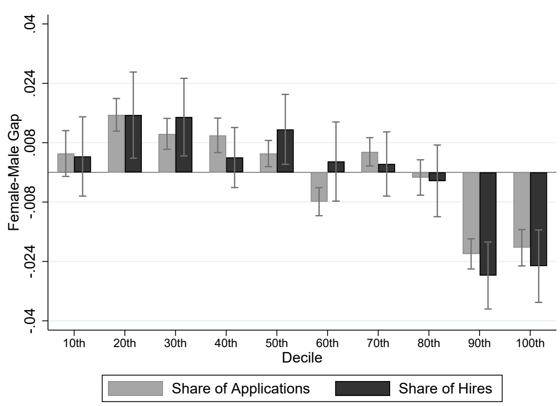

(d) Typical wage

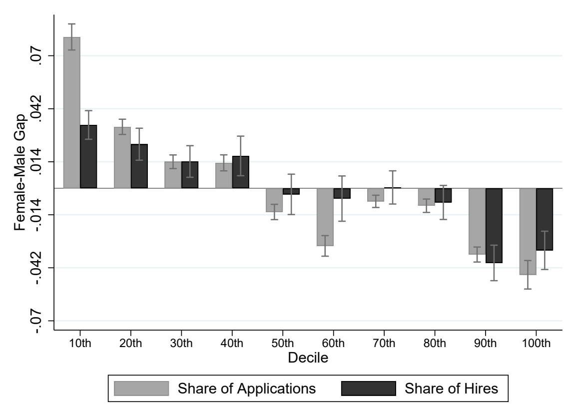

Notes: The figures plot gender gaps in the share of applications going to specific types of jobs and corresponding gender gaps in hiring shares. All gaps are computed after propensity score reweigthing on the quarter of inflow into unemployment along with all the conditioning variables used in the main analysis. Occupations and industries on the x-axis are ordered from lowest paying on the left to highest paying on the right. The percent confidence bars are based on standard errors clustered on the individual level.

Figure B.10: Gender gaps in applications and hiring outcomes, raw

(a) Occupations

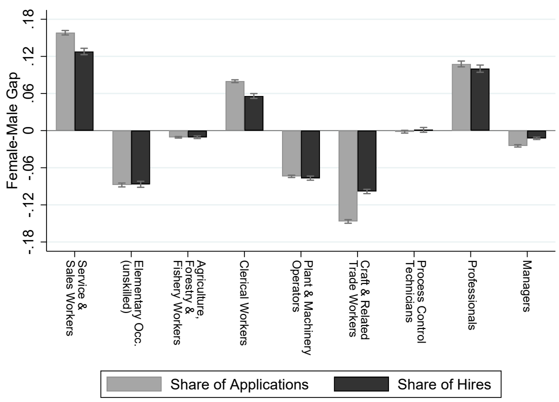

(b) Industries

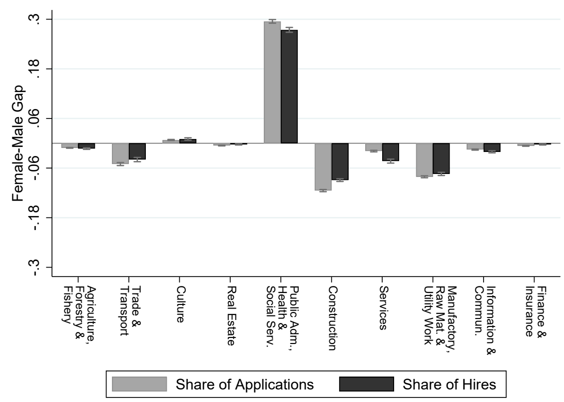

(c) Firm wage level

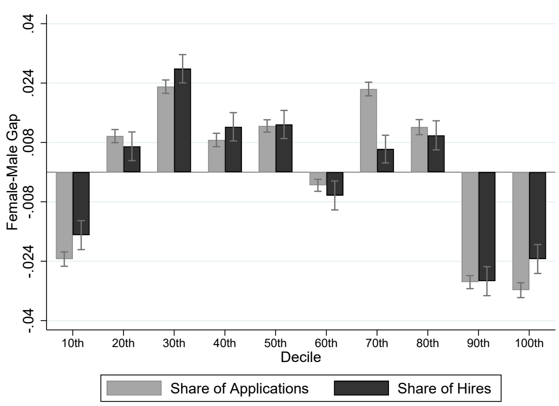

(d) Typical wage

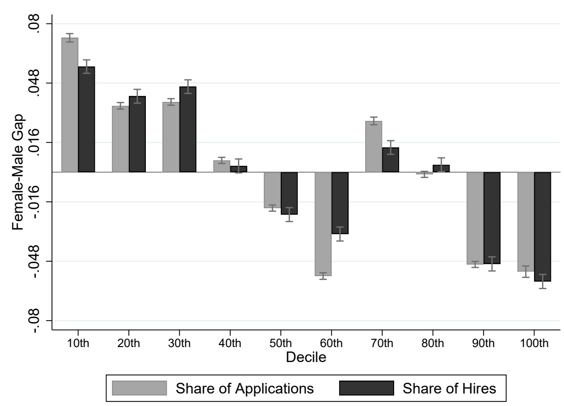

Notes: The figures plot gender gaps in the share of applications going to different types of jobs along with corresponding gender gaps in hiring shares. The figure shows raw gaps without conditioning on observables. Occupations and industries on the x-axis are ordered from lowest paying on the left to highest paying on the right. The 95 percent confidence bars are based on standard errors clustered on the individual level.

Figure B.11: Gender gaps in applications and hiring outcomes, absolute measure for applications

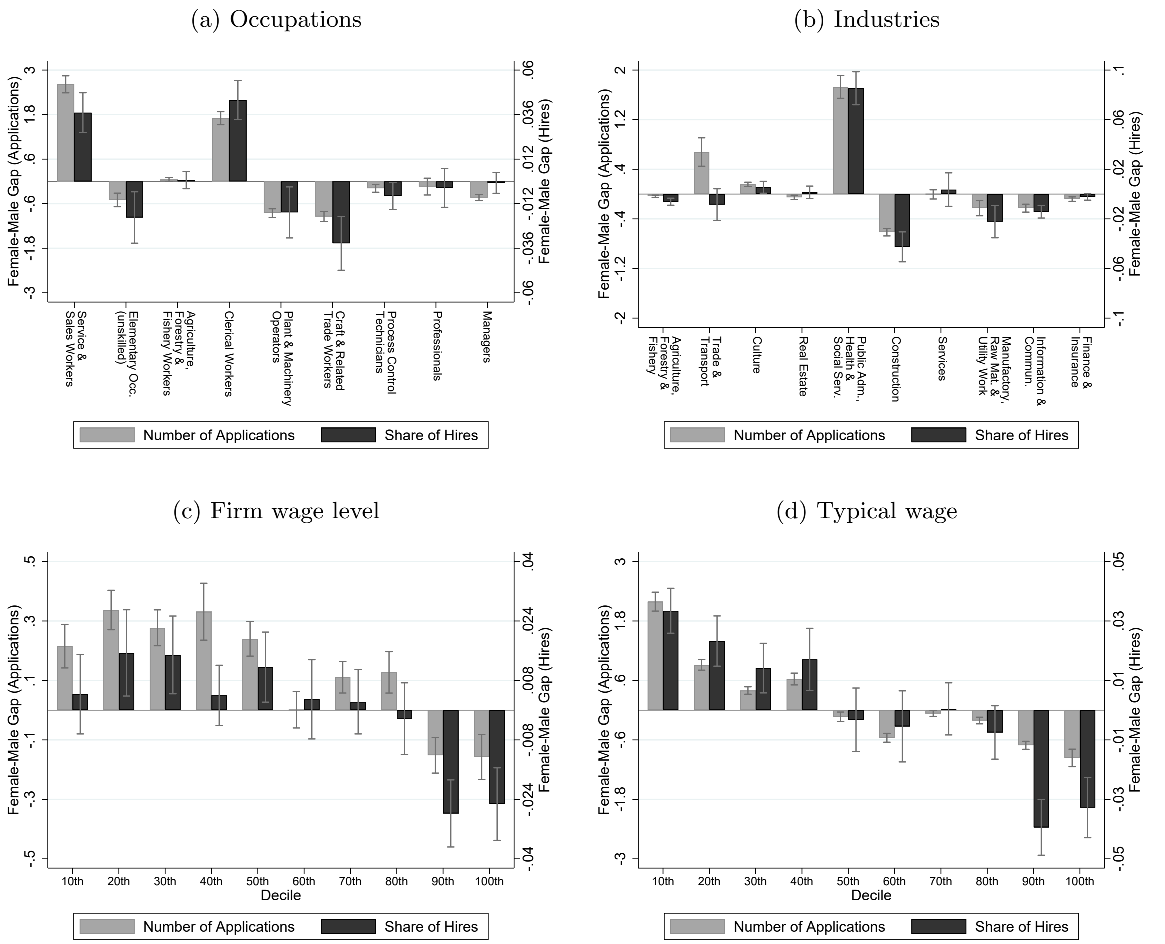

Notes: The figures plot gender gaps in the number of applications going to different types of jobs (left axis) along with corresponding gender gaps in hiring shares (right axis). All gaps are computed after propensity score reweigthing so are conditional on observables. Occupations and industries on the x-axis are ordered from lowest paying on the left to highest paying on the right. The 95 percent confidence bars are based on standard errors clustered on the individual level.

Figure B.12: Gender gaps in job applications and hiring, gender-specific typical wage measures

(a) Typical wage paid women

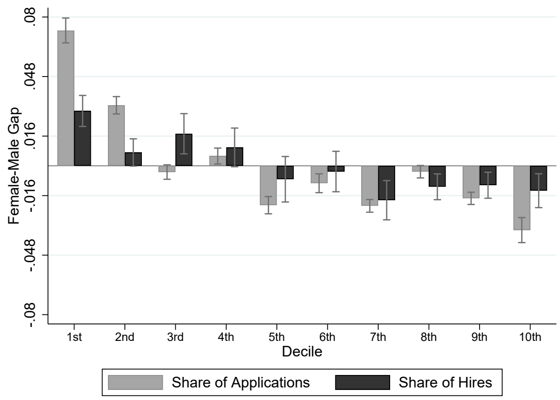

(b) Typical wage paid to men

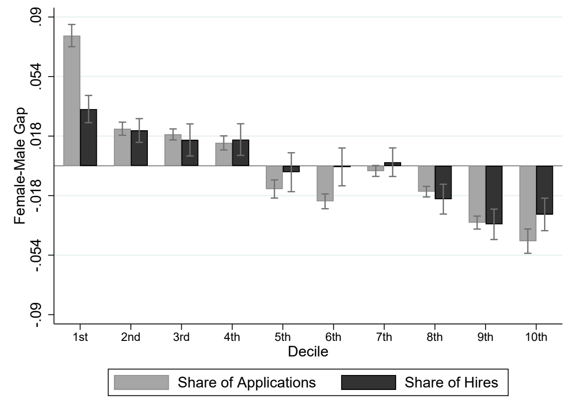

Notes: The figures plot gender gaps in the share of applications going different jobs with different level of typical wages paid to women (left) and men (right). All gaps are computed after propensity score reweigthing so are conditional on observables. Occupations and industries on the x-axis are ordered from lowest paying on the left to highest paying on the right. The 95 percent confidence bars are based on standard errors clustered on the individual level.

Figure B.13: Gender gaps in job applications and hiring, gender-specific firm effects

(a) Firm effects based on women only

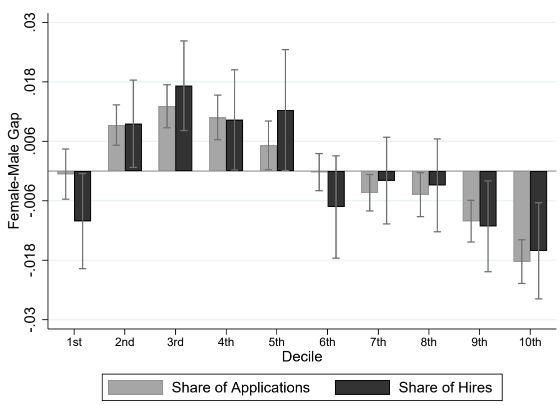

(b) Firm effects based on men only

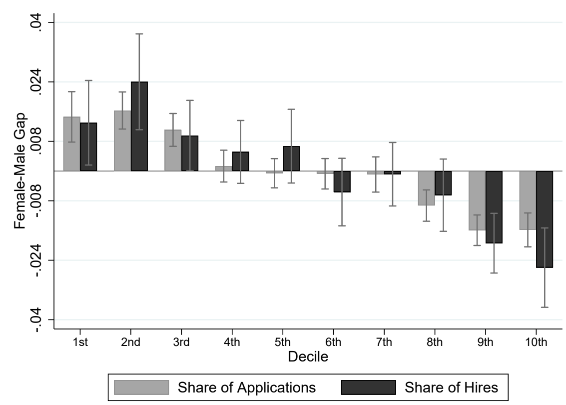

Notes: The figures plot gender gaps in the share of applications going different jobs with different level of firm effects estimated only on a sample of women (left) or men (right). Subsequently the obtained firm effects are then used to rank all applications and jobs. All gaps are computed after propensity score reweigthing so are conditional on observables. See Section A.3 regarding the estimation sample for the firm effects. The 95 percent confidence bars are based on standard errors clustered on the individual level.

C Online Appendix: Additional results

C.1 Changes in application behavior over time

In our analysis we pool all applications sent during the unemployment spell and analyze the composition of this pool across different individuals. In principle, however, there could be interesting gender differences in how application behavior changes over the course of an unemployment spell. To present some evidence on the importance of such dynamics we create a panel date that for each person and unemployment spell contains an observation that for each month of the spell. We then run an twoway fixed effect regression that uses the monthly average typical wage of applied-for jobs as the outcome variable and includes dummies for the number of months since entering unemployment alongside person-by-spell fixed effects.[66] We run this separately for men and women. The results of this regression shows how men and women change their application behavior over time within an unemployment spell.

Figure C.1 shows the estimates of gender specific time profile of job search from this regression. Changes in application behavior over time are modest and are very similar for men and women. Examining gender gaps in applications at different times throughout the unemployment spell therefore yields very similar results to the pooled results presented in the main text. This is true also for other job characteristcs besides wages.

C.2 Occupation and wages of applied-for jobs relative to previous job

In Figure C.2 we show gender application and hiring gaps when grouping applications and hires by whether the typical wage is higher or lower than the previous job. The figure reveal a stark difference between women and men, where men target higher typical wages compared to the typical wage of their previous job.

In Figure C.3 we repeat the same exercise for occupations. Specifically, we rank occupations according to their average wage and then compare applied-for occupations' rank to the rank of the previous occupation. We again see stark gender differences in the extent to which men and women apply "up" or "down" relative to their previous job. In particular women are more likely to submit

[66] We use the month of entry into unemployment as the baseline and omit a dummy for this month in the regression. We let the fixed effect be specific to each unemployment spell because some individuals in our data show up with more than one unemployment spell. We cluster standard errors at the level of the unemployment spell.

applications which would involve occupational downgrading and are also more likely to be hired in this type of jobs.

C.3 Applications to public vs. private sector jobs

In Figure C.4 we report the application and hiring gaps when we group firms by whether they are private or public.[67] Consistent with previous work on gender and public sector jobs, we see that women send about 9 percentage points more of their applications to public sector jobs and are also about 9 percentage points more likely to be hired into these jobs.

C.4 Decomposing hiring gaps across deciles of the firm and wage distribution

Mirroring the decomposition results for individual industries and occupations in Section 5.4, Figure C.5 decomposes hiring gaps into each decile of the distribution of firm wage-levels (Panel A) and typical wages (Panel B). For almost all deciles, gender differences in applications are able to explain the majority of the baseline hiring gap.

C.5 Comparison to audit study estimates of occupation-specific gender discrimination

For understanding our decomposition results across occupations, varying degrees of direct gender discrimination is a particularly salient possible explanation. A natural way of examining this explanation is to compare our residual hiring gaps to direct estimates of gender discrimination from the audit study literature (Riach and Rich (2002); Rich (2014); Neumark (2018)). Audit studies are controlled experiments that send fictitious job applications to actual vacancies to measure gender discrimination between otherwise identical applicants.

Unfortunately, no existing audit study (or set of studies) provide reliable estimates of discrimination across industries or occupations in Denmark.[68] The audit study data collected by Ahmed et al. (2021), however, provides a useful benchmark. Ahmed et al. (2021) combine data from three audit studies conducted in Sweden between 2016 and 2019 thus the setting and time period is

[67] We distinguish firms by whether they are private or public through their industry affiliation and classify public administration, education, health, culture and services as public firms.

[68] The closest is the Danish audit study of Dahl and Krog (2018). Since they only sent applications to 400 different jobs, however, their sample size is not well suited to construct occupation-specific measures of discrimination.

comparable to our data. The data contain 3,214 applications, covering 15 different types of jobs. Because the gender of the applicant is randomly assigned, gender discrimination in the initial hiring stage can be measured by looking at how often female and male applications received a positive reply. We refer to as the gender gap in callback rates (although note that employer responses were not restricted to come in the form of phone calls).

To construct occupation-level measures of the gender gap in callback rates, we need to map the 15 jobs considered in the audit studies to the occupation groups used in our analysis. Table C.1 shows the result of this mapping. For 12 of the 15 jobs considered, we are uniquely able to map them to 5 different of the occupation group in our data. For the remaining 3 jobs, it is not possible to uniquely map them to an occupation group.

Figure 4a in the main text plots the resulting occupation-specific gender gaps in callback rates against the occupation-level residual hiring gaps from our decomposition.

C.6 Do women apply for jobs with lower wage growth?

In Section 3.5, we saw that men and women apply to jobs with different typical wages. These differences largely correspond to differences in the eventual hiring outcome, leading to the observed gender wage gaps in typical wages. As described in Section, 3.2, however, the wage measures used in our main analysis reflect the wage level at the start of a workers new job. If there are also substantial differences in how fast wages increase over the initial years in a firm these gender gaps may further change over time during employment spells.

To get a sense on whether this is important, this section constructs a measure of the wage growth that a given job offers and examine whether there are gender gaps in applications and hiring also into high vs low wage growth jobs. Specifically, we calculate firm specific 1 and 5 year wage growth rates based on all jobs starting between 2008 and 2016 in a given firm. In order to separate general time trends from the growth rates as well as structural differences between industries, we control for year by industry fixed effects. As many individuals will have left their jobs by the one and, especially, the five year mark, we censor individuals that are no longer employed by the firm at these times.[69]

[69] Obviously it is not random who stays in a firm up to 5 years, and our numbers may therefore partly be driven by dynamic selection. Results should be interpreted with this in mind.

In Figure C.6 we split firms into deciles based on the typical wage growth they offer over either 1 or 5 years. We see that women are more likely to apply to firms with lower wage growth rates, with the exception of the lowest decile. Likewise, men apply substantially more to those firms in the very top deciles. In addition to the fact that women are applying and getting hired more into jobs with lower staring wages, women thus are thus also applying and getting hired into firms that offer lower rates of wage growth. These differences in applications and hiring outcomes contribute to a widening of the gender gap over time.

C.7 Are gender gaps in applications related to motherhood?

Several recent papers have emphasized that gender gaps in the labor market are particularly related to motherhood and its effects on the valuation of non-wage job characteristics. In Figures C.7 and C.8, we attempt shed some light on how motherhood relates to gender differences in job applications. In these figures, we limit our sample to UI recipients in an age window around the prime childbearing years, specifically 25-40 years. We then repeat our descriptive analysis of gender gaps in application and hiring separately for men and women with young children (0-5 years) and for men and women without children. For most of the non-wage job characteristics considered in Section 6.1 and for typical wages, we see that gender gaps in both applications and hiring outcomes tend to be larger when comparing men and women with young children. At the same time, however, we note that very substantial gender differences do exist in both groups.

C.8 Gender differences in the returns to applications (self-fulfilling discrimination)

One explanation for the observed gender application gap is a version of the "self-fulfilling discrimination" mechanism that has been proposed and documented in other settings (Lundberg and Startz, 1983; Coate and Loury, 1993; Glover et al., 2017). In the context of the job application process, the basic idea behind this is as follows: Gender discrimination in hiring implies that there are gender differences in the likelihood that an application turns into a hire for some jobs. As a result women have an incentive to apply less to these jobs and more to jobs where the chance of being hired is higher. In this way, gender differences in the likelihood than an application turns into a hire may explain why there are gender differences in applications.

To check whether we can detect this phenomenon in our data, we first construct measures of gender differences in the likelihood that an application turns into a hire. We do this in the context of a simple linear regression. For some individual in our analysis sample, let y denote a type of job (an occupation, industry or decile of a wage distribution), let a^y be the share of their applications that the individual sent to jobs of type y, and let d^y be an indicator for whether the individual was hired into a job of type y. We then consider estimating the following regression on our weighted analysis sample:

$$d^y = \beta_0^y + \beta_1^y a^y + \varepsilon \tag{10}$$

If a^y is measured in percentage points, the coefficient β_1^y in this regression captures how much the likelihood of being hired into job type y increases if one additional percentage point of applications is targeted to this type of job. If given a causal interpretation, this is a measure of the returns to applications for job type y. To obtain a simple measure of gender differences in the likelihood that an application turns into a hire for job type y, we therefore estimate equation 10 separately for men and women and compute the difference in the estimate of β_1^y across genders. We refer to this as the *gender gap in returns to applications* for job type y. We note that there are obvious concerns with treating the estimate of β_1^y as causal. In particular, if individuals tend to send more applications to jobs they are more likely to get due to e.g. differences in unobservables, we would expect β_1^y to overstate the returns to search.[70] Absent a source of exogenous variation in application behavior, we have no way of removing this potential bias. To the extent that the resulting bias is relatively constant across job types, however, the estimates from Equation 10 may still give a useful ranking of the types of jobs where men face relatively higher returns to search than women.

In Figure C.9 we examine whether differences in the returns to applications appear to explain the observed gender gaps in applications. Each data point in the figure corresponds to a job type, defined as either a two-digit occupation, a two-digit industry, a decile of the firm wage level distribution or a decile of the typical wage distribution. The y-axis shows the gender gap in applications to the different job types after conditioning on labor market observables. Finally, the

[70] Since we estimate Equation 10 on the reweighted sample, we are ensuring that labor market observables are balanced across men and women. There may of course still be unobserved differences across individuals that affect the likelihood of a successful application and likely correlates with application behavior as well.

x-axis contains our (standardized) measure of the gender gap in returns to applications for the different job types.[71] If the observed gender gaps in applications were driven by women applying more to jobs where their applications have a higher relative likelihood of turning into a hire, we would expect to see an upward sloping relationship in Figure C.9. This is not what we see. If anything the relationship seems to be slightly downward sloping.[72]

The fact that gender gaps in applications are not positively correlated with gender gaps in the likelihood that an application turns into a hire is also borne out in the relative magnitudes of the application and hiring gaps documented in Section 4. If the jobs that women apply less to than men are also the ones where they face a lower chance of being hired when they apply, this will compound to make overall gender gaps in hiring outcomes even larger than gender gaps in applications.[73] As shown in Section 4, however, gender gaps in applications actually tend to be larger than gender gaps in hiring.

In sum, with the data we have available, we cannot find evidence that "self-fulfilling discrimination" underlie the observed patterns of application behavior across genders in our data. As discussed, however, this does not rule out that "self-fulfilling discrimination" is at play and could be detected with different data.[74]

C.9 Gender differences in beliefs, overconfidence and risk preferences

Another possible explanation for our results is that gender differences in preferences and beliefs lead men to systematically apply for more high-paying and harder-to-get jobs than women. If men and women have different beliefs about their general labor market prospects such that men are

[71] Specifically, we standardize the measure within each category of job types (across industries, across occupations, across deciles of firm wage levels and across deciles of typical wages). Standardizing the measure jointly across all the categories or using an non-standardized version of the measure does not change the results.

[72] The correlation between the gender gap in applications and our measure of the gender gap in returns to applications for the different job types is -0.368 overall and range from -0.762 (typical wage deciles) to 0.257 (firm wage level deciles) if computed across one of the four job categories.

[73] Put differently, overall gender gaps in hiring outcomes is the product of two gaps: 1) the gap in how likely women are to apply for a particular type of job and 2) the gap in how likely an application from a woman is to result in a hire for a particular type of job. If the types of jobs where gap 1 is big are also the ones where gap 2 is big, the aggregate gap in hiring outcomes should be bigger than the application gap alone (gap 1).

[74] As noted, our simple regression-based measure of returns to applications does not account for unobservable characteristics that correlate with both application behavior and the likeihood of getting hired. In addition, the most extreme models of "self-fulfilling discrimination", can imply that women *never* apply to any jobs where they face a lower probability being hired than men do. In this case it would never be possible to measure any meaningful differences in the likelihood of being hired conditional on applying. In our data, women are of course sending many applications to each of the observed job types we consider, however, this does not rule out that men and women are differentially applying to jobs with different unobserved characteristics within these job types.

more (over)confident than women, this could lead men to systematically target more higher-paying but harder-to-get jobs.[75] Similar predictions also arise if men are less risk-averse than women or if some form of social norms lead women to hold themselves to a higher standard when deciding where to apply for jobs. The existence of these types of gender differences have received significant attention in previous work and have also found empirical support in some settings (see in particular Cortes et al. (2020)).

To provide some evidence on this possible mechanism, we can examine the speed with which men and women find jobs in our data. To this end, we construct a version of our main analysis sample where we drop the restriction that individuals must find employment within a year (see Section 3.1).[76] For this sample we then examine how many men and women still have not found employment after each week of their unemployment spell. Specifically, Figure C.10 shows Kaplan-Meier estimates of the survivor function in nonemployment, estimated separately for men and women. To account for differences in observable characteristics of men and women, the survivor cuves are estimated after reweighting on observables (see Section 3.4).[77]

If simple gender differences in beliefs or preferences are causing women to apply to less ambitious and easier-to-get jobs than men, we should see women find jobs faster than men. Looking at Figure C.10, however, this is not what we see in the data. Throughout the unemployment spell, the survivor curve for women is slightly above the survivor curve for men, implying that if anything women in fact find jobs slightly more slowly than men do.

In sum, we cannot find evidence that simple gender differences in beliefs or risk preferences underlie the observed patterns of application behavior across genders in our data. Of course, this not rule out that gender differences in beliefs or risk preferences exist and contribute to the observed application behavior - if men and women differ in other dimensions that influence job finding rates, for example, this may obscure the effect of beliefs or risk preferences on job finding rates.

[75]It is commonplace to refer to this explanation for gender differences as reflecting male overconfidence. Of course, it is in principle also possible that men have approximately correct or even downward biased beliefs about their labor market prospects, while women are (more) underconfident and have (more) downward biased beliefs about their labor market prospects. This could also explain the gender application gaps that we see in our data through exactly the same mechanisms.

[76]We find a similar pattern if we instead focus on the main analysis sample including the restriction that individuals must find employment within a year.

[77]Figure A.1 shows corresponding estimates without reweighting. This leads to a larger difference between men and women but still with women finding jobs slower than men.

C.10 Gender-shares at applied-for firms (preferences for diversity)

In Figure 5, we focus on a set of non-wage job amenities that have been emphasized in previous work on gender gaps. Of course, there could be many other non-wage job amenities that men and women value differently. One salient example here is the gender composition of the workforce. In particular, we might think that women prefer working in more gender diverse workplaces, or alternatively that they simply prefer working with many female coworkers.

In Figure C.11 we provide some evidence in this regard. For each applied-for job, we compute the share of female employees at the firm. We then split firms into deciles of this measure and compute gender application and hiring gaps (after conditioning on observables using reweighting). As the figure shows, women are systematically more likely to apply and get hired at firms with more female employees. Moreover, the relationship between hiring/application gaps and female share is completely monotonic. If we interpret this as reflecting preferences, it would thus suggest that female UI seekers simply prefer working at firms with more female employees, not at more diverse workplace (the median firm has a female share of 33 percent).

We note, however, that the figure is also consistent a situation in which women do not care about workforce composition but simply for some other reason prefer working at certain firms. Absent very particular turnover patterns, these firms should end up with more female employees in steady state, which would generate the pattern in Figure C.11.

Table C.1: Mapping job titles from Ahmed et al. (2021) to occupations

Occupation Group:	Job Title:	Observations:
Craft and Related Trades	*Vehicle mechanic*	214
Elementary Occupations	*Cleaner*	434
	Warehouse worker	141
Plant and Machinery Operators	*Truck driver*	337
Professionals	*Enrolled nurse*	206
	IT developer	153
	Pre-school teacher	270
Service and Sales	*Chef*	392
	Customer service	51
	Store clerk	127
	Telemarketing	43
	Waitstaff	497
Umatched	*Accounting clerk*	166
	B2B sales	152
	Childcare	71

Notes: The table shows our mapping of job titles in the audit study data of Ahmed et al. (2021) into 1-digit occupation groups in our data. The last column shows the number of applications sent to each job title in the audit study data.

Figure C.1: Change in job search over time, event study estimates

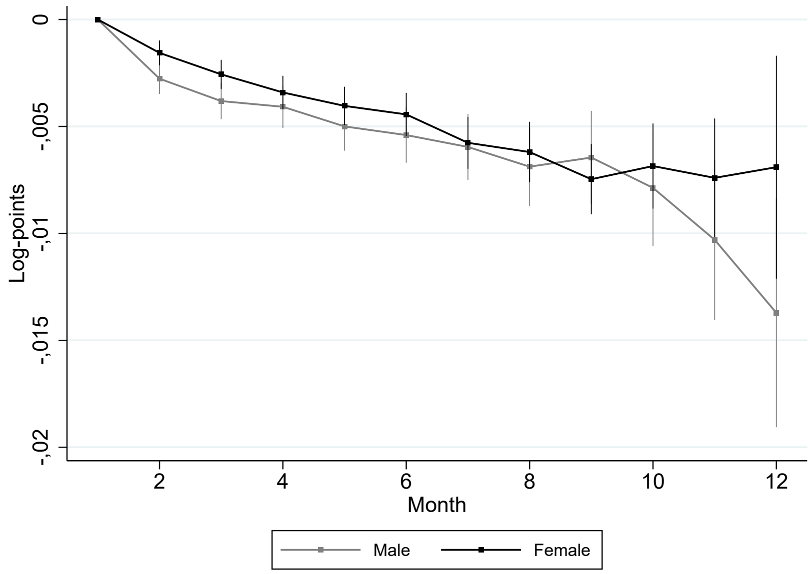

Notes: The figure plots estimates from an event study regression that regresses the average log typical wage of applied-for jobs on dummies for the number of months since entering unemployment. The figure plots the estimated coefficients on the month dummies. The regression includes a person-by-spell fixed effect and uses the month of entry (month 0) as the omitted baseline month. Results are shown separetely for men and women. Standard errors are clustered at the spell level.

Figure C.3: Gender differences in applications up and down the occupational job ladder

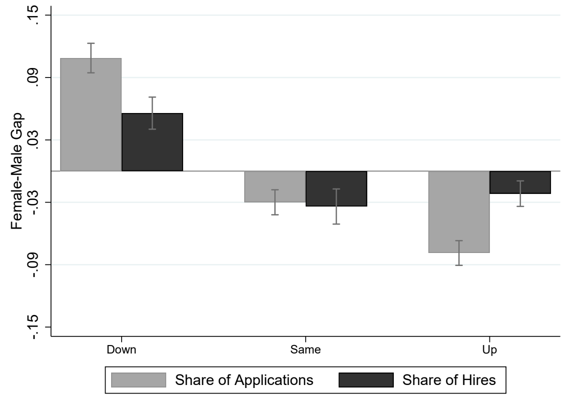

Figure C.2: Gender gaps in job applications and hiring, typical wage ranks relative to previous job

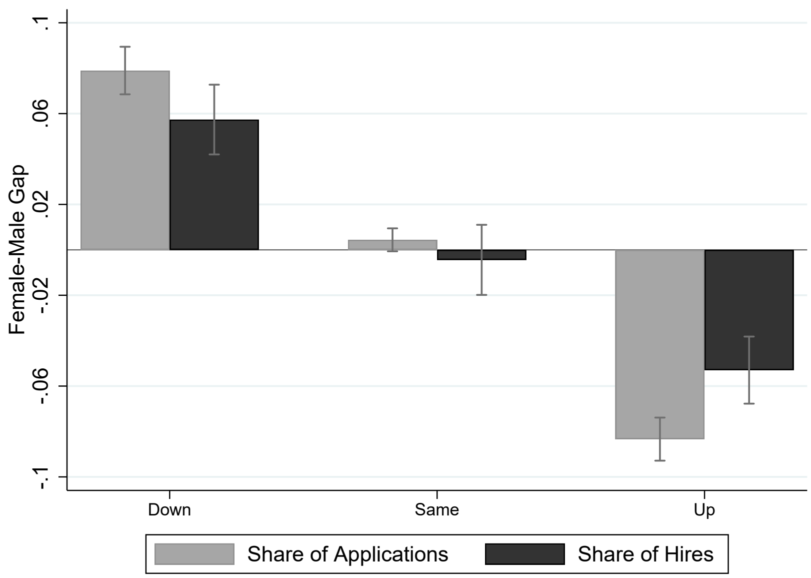

Notes: The figures plot gender gaps in the share of applications going to jobs that are in a higher, lower or the same decile as the previous job. The figure also plots corresponding gender gaps in hiring shares. All gaps are computed after propensity score reweigthing so are conditional on observables. The 95 percent confidence bars are based on standard errors clustered on the individual level.

Figure C.4: Gender gaps in job applications and hiring, public vs private sector

(a) Public Private

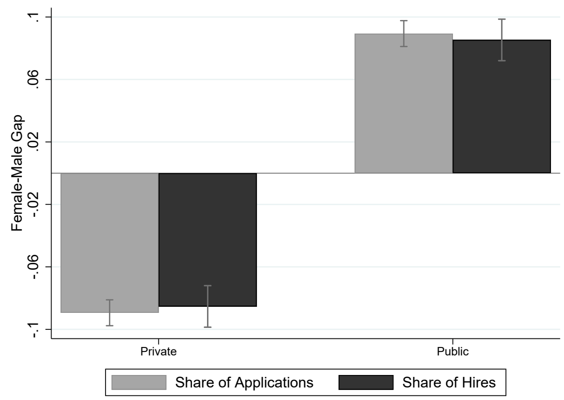

Note: Figure plots gender gaps in the share of applications going to public or private firms. We also plot corresponding hiring shares. The 95 percent confidence bars are based on standard errors clustered on the individual level.

Figure C.5: Decomposing hiring share gaps across deciles of firm types and wages

(a) Firm wage level

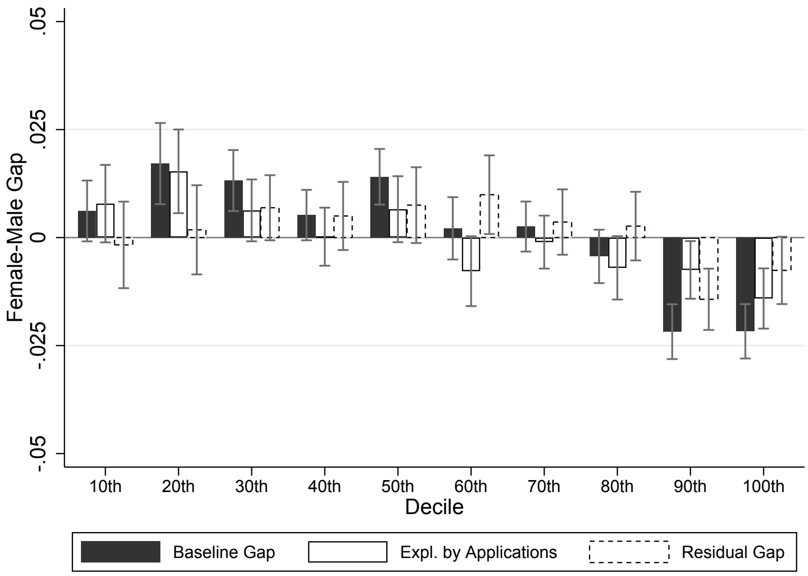

(b) Typical wages

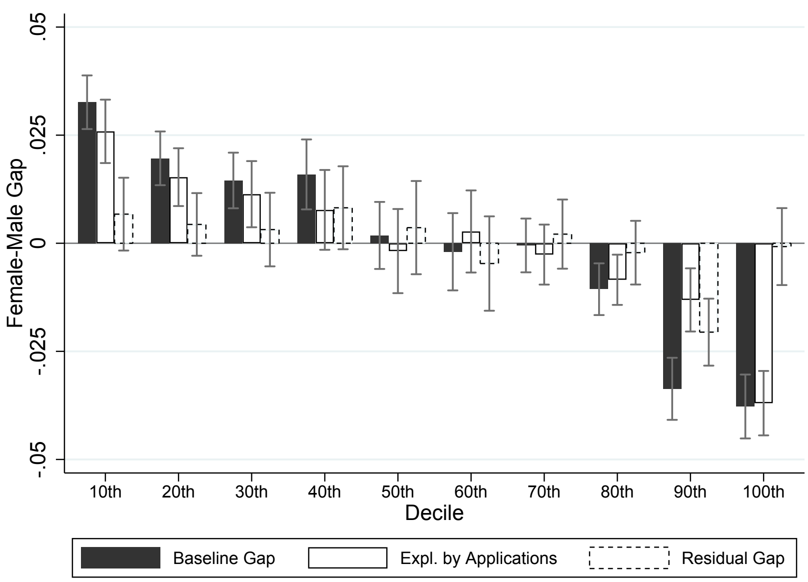

Note: Figure decomposes baseline gaps in the share of men and women hired into different types of jobs after conditioning on labor market observables. The gaps are decomposed into a part that explained by applications and a residual gap (see Equation 5).

Figure C.6: Gender gaps in job applications and hiring, wage growth

(a) Wage growth (1 year) (b) Wage growth (5 year)

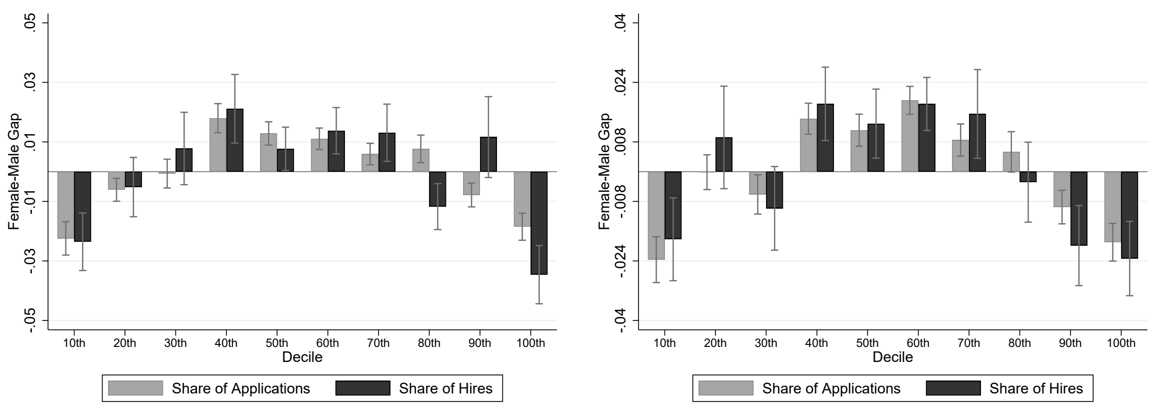

Note: Figure plots gender gaps in shares of applications going to specific wage growth deciles and gaps in which decile job-seekers are hired. Wage deciles are computed as the relative difference between the starting wage with the wage one year (left) or five years (right) after entering the respective job. The 95 percent confidence bars are based on standard errors clustered on the individual level.

Figure C.7: Gender gaps in applications and hiring outcomes, individuals with/without young children I

(a) Hours, individuals with young children

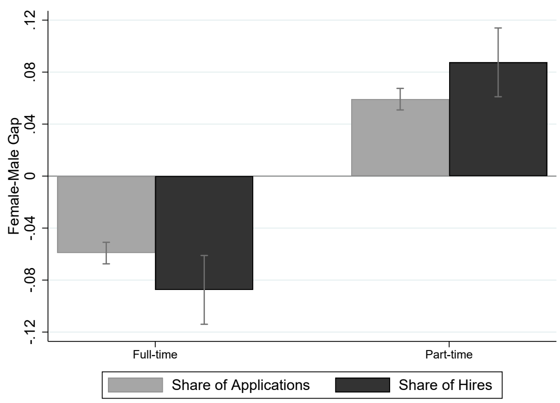

(b) Hours, individuals without young children

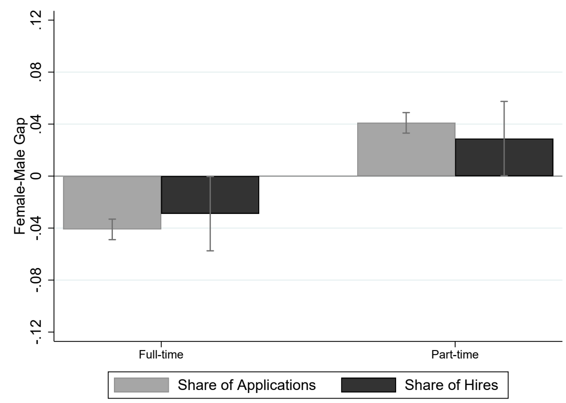

(c) Commuting, individuals with young children

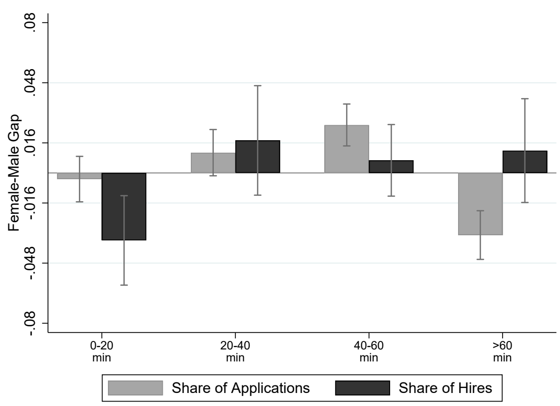

(d) Commuting, individuals without young children

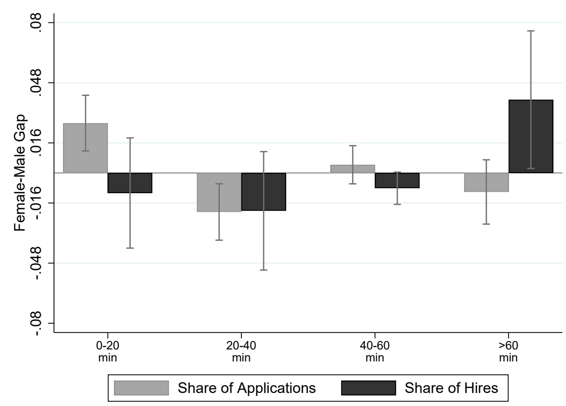

The figures plot gender gaps in shares of applications going to specific types of jobs and corresponding gender gaps in hiring outcomes separately for individuals (age 25-40) with young (0-5 years) children and without children. All gaps are based on the reweighted sample so are conditional on observables. The 95 percent confidence bars are based on standard errors clustered on the individual level.

Figure C.8: Gender gaps in applications and hiring outcomes, individuals with/without young children II

(a) Family friendly, individuals with young children

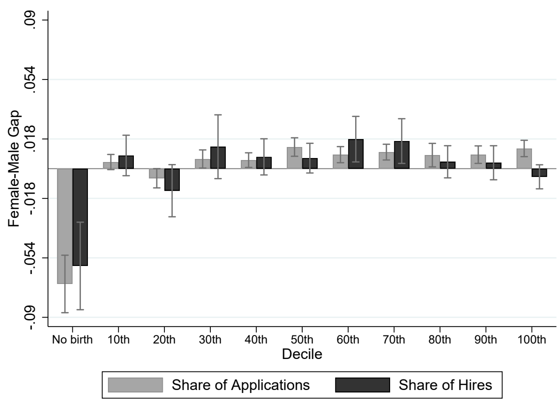

(b) Family friendly, individuals without young children

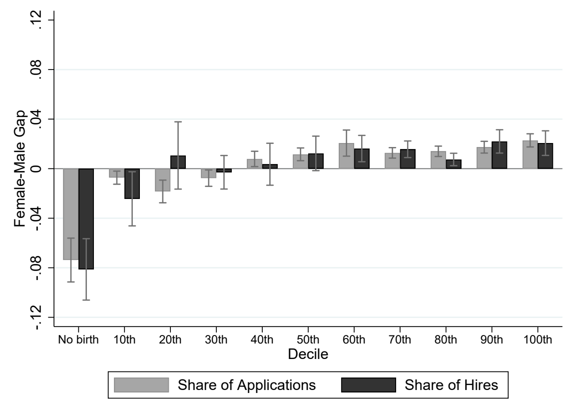

(c) Typical wages, individuals with young children

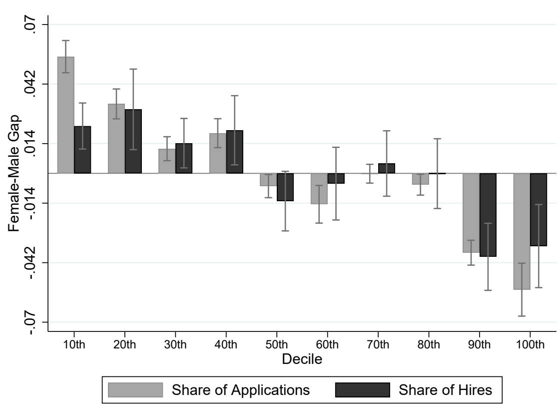

(d) Typical wages, individuals without young children

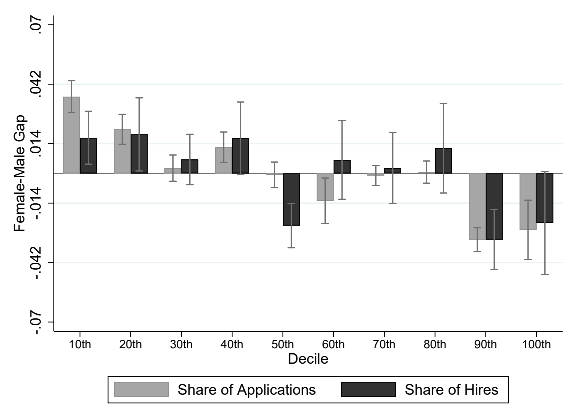

The figures plot gender gaps in shares of applications going to specific types of jobs and corresponding gender gaps in hiring outcomes separately for individuals (age 25-40) with young (0-5 years) children and without children. See Appendix A.8.1 for details on the measure of family friendliness. All gaps are based on the reweighted sample so are conditional on observables. The 95 percent confidence bars are based on standard errors clustered on the individual level.

Figure C.9: Gender application gaps and the return to applications

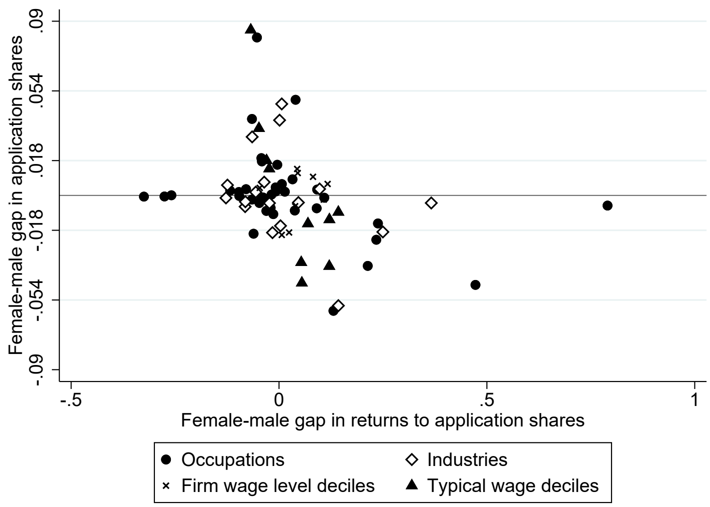

Note: The figure plots the corresponding gender application gap and the gender gap in the return to applications. Each data point is thus a job type defined as either a two-digit occupation, a two-digit industry, a decile of the firm wage level distribution or a decile of the typical wage distribution. Military occupations and job types, where fewer than 100 individuals finds employment are excluded. The y-axis shows the gender gap in applications to the different job types after conditioning on labor market observables. The x-axis contains our (standardized) measure of the gender gap in returns to applications for the different job types, see Equation 10.

Figure C.10: Kaplan-Meier survivor functions in nonemployment

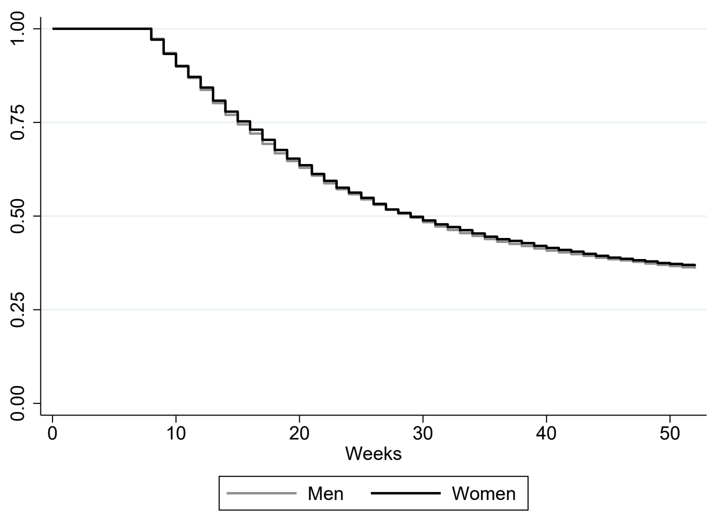

Note: The figure plots Kaplan-Meier estimates of the survivor function in nonemployment, estimated separately for men and women. The two curves thus show the share of men and women that still have not found a new job after each week since starting their UI spell. The curves are estimated on a version of the main analysis sample where the requirement of finding a job within a year has not been imposed. The curves are estimated after reweigthing on observables.

Figure C.11: Gender gaps in job applications and hiring, share of females at firm

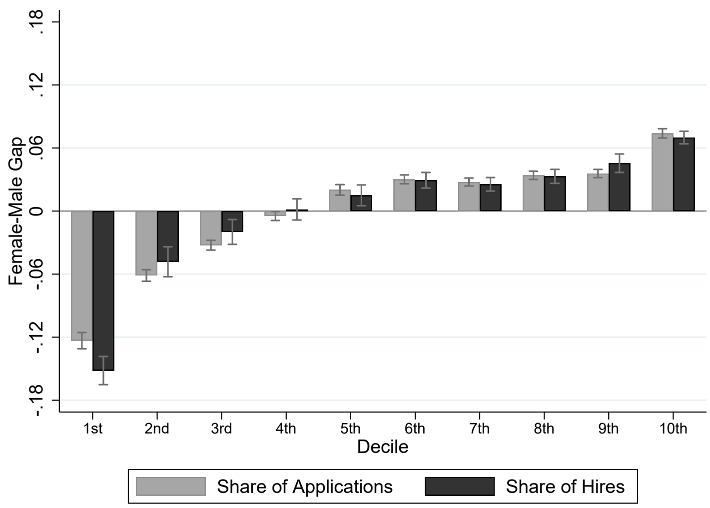

Notes: The figures groups firms into deciles in terms of the share of women in their workforce (i.e. the 10th decile contains the firms with the highest share of female employees) and then plots gender gaps in the share of applications going to each of these deciles . The figure also plots corresponding gender gaps in hiring shares. All gaps are computed after propensity score reweigthing so are conditional on observables. Occupations and industries on the x-axis are ordered from lowest paying on the left to highest paying on the right. The 95 percent confidence bars are based on standard errors clustered on the individual level.