Gender-Specific Application Behavior, Matching, and the Residual Gender Earnings Gap*

Benjamin Lochner[1] Christian Merkl[2]

[1]*Friedrich-Alexander-Universität Erlangen-Nürnberg (FAU) and Institute for Employment Research (IAB)*
[2]*FAU and Institute of Labor Economics (IZA)*

July 18, 2022

Abstract

This paper opens up the black box of gender-specific application and hiring behavior and its implications for the residual gender earnings gap. To understand the underlying mechanisms, we propose a two-stage matching model with testable implications. Using the German IAB Job Vacancy Survey, we show that the patterns in the data are in line with linear and nonlinear production functions at different jobs, where the latter imply compensating differentials of pay for higher employer-sided flexibility requirements. We show that women's application probability at high-wage firms is much lower than at low-wage firms. By contrast, women have the same probability of being hired as men at these firms once they applied. These patterns can be rationalized by high-wage firms asking for more employer-sided flexibility and thereby generating sorting of applicants into different types of jobs. We show that the share of male applicants increases in various dimensions of employer-sided flexibility requirements and therefore consider it as a proxy for flexibility. Adding this proxy to Mincer wage regressions reduces the residual earnings gap by around 50 to 60 percent. Women who match at jobs with a large share of male applicants earn 10 to 20 percent more than women at comparable jobs with more female-dominated application pools. In line with the nonlinear production function hypothesis, when women with children match at these jobs, they face large wage discounts relative to men.

Keywords: Job Search, Application Behavior, Gender Earnings Gap;
JEL Classifications: E24, J16, J31

*We thank Jesper Bagger, Steven Davis, Jason Faberman, Simon Jäger, Lisa Kahn, Philipp Kircher, Andreas Müller, Johannes Schmieder, Bastian Schulz, Linas Tarasonis, Rune Vejlin, and seminar participants from the Chair of Macroeconomics, the IAB, the Gender Gaps 2021 conference, the NUEdialog, the 12th ifo Conference on Macroeconomics and Survey Data, the DTMC Workshop on Interactions between Labor and Marriage Markets, and the APB/ IOS Conference on Gender Inequality in Central and Eastern Europe for valuable feedback. We thank Ipek Yükselen and Olga Zajkowska for discussing our paper.
Corresponding author: Benjamin Lochner, University of Erlangen-Nuremberg (FAU), Lange Gasse 20, 90403 Nuremberg, Germany. Email: Benjamin.Lochner@fau.de.

1 Introduction

After several decades of gender convergence, substantial differences in earnings between men and women remain. Part of this gap can be explained by men and women working in different occupations and industry sectors (Blau and Kahn, 2017), or in firms with different wage premiums (Card et al., 2016). However, even within narrowly defined sectors and occupations, a substantial gender earning gap remains. A recent strand of the literature analyzes the role of gender-specific search behavior of workers for the gender earnings gaps, combining search theory and newly available microeconomic datasets (Cortés et al., 2021; Faberman et al., 2017; Fluchtmann et al., 2020).

Our paper contributes to this stream of the literature, by using data from the German IAB Job Vacancy Survey which we link to administrative employment records. This unique combination allows us to observe important dimensions of the search and matching process such as characteristics of the hiring firm (e.g., wage premium), the hired worker (e.g., whether a woman is a mother), and the recruitment process itself (e.g., the gender distribution in the application pool).[1] Guided by our two-stage search and matching model, we show that men and women tend to apply at different firms[2] and for different jobs. These differences can explain a large part of the residual gender earnings gap. Specifically, we show that women in Germany have a lower probability to apply for jobs at firms with high wage premiums from a two-way fixed effect regression (Abowd et al., 1999; Card et al., 2013). At the same time, the probability of being hired at these high-wage firms conditional on having applied is similar for males and females. We argue through the lens of our theoretical model that these patterns are not reconcilable with taste-based discrimination at the hiring stage. By contrast, they can be explained by different job characteristics (Goldin, 2014), namely more employer-sided flexibility at high-wage firms. We show that the share of male applicants[3] increases with various employer-sided flexibility requirements (such as working irregular hours or at various locations). Adding the share of male applicants as a proxy for multidimensional flexibility requirements to standard Mincer wage regressions leads to a drop in the residual gender-wage gap of around 50 to 60 percent. Women who match in a pool with a large share of male applicants earn up to 20 percentage points higher wages than women that match in jobs with a low share of male applicants (controlling for a large set of worker observables). These patterns are in line with Goldin (2014)'s idea of nonlinear jobs that pay a disproportional premium for providing flexibility. We show that the wage discount is particularly strong for mothers with children in jobs with high employer-sided flexibility requirements. In line with our

[1] To the best of our knowledge, we are the first to use data that has information on the pool of applicants for a particular job in a particular firm.

[2] Although we refer to firms, the IAB data rather identifies plants/establishments, i.e., individual production units. We use these terms interchangeably throughout the paper.

[3] We residualize the share of male applicants by controlling for occupation, sector, and firm size.

model, if mothers match at these nonlinear jobs, they are more likely unable to provide the desired flexibility and thereby earn a lower wage.

We motivate and structure our empirical exercise with a simple two-stage search and matching model. In the first stage, searching workers have to decide whether they want to apply for a particular job profile. Facing heterogeneous application costs, they will apply whenever the expected returns from the application are larger than application costs. In the second stage, only those worker-firm pairs with a positive surplus will form a match. Worker-firm pairs draw an idiosyncratic match-specific training cost shock. Only a certain fraction of workers will be selected in the model (see Chugh and Merkl (2016), or Carrillo-Tudela et al. (2020) for selection models). In our model, the male and female application behavior is a function of the expected match surplus. Thus, a high share of male applicants shows that men (on average) perceive a higher surplus for certain job types. We analyze two scenarios. In the first scenario, we assume taste-based discrimination at the hiring stage. Employers only recruit women if they are compensated by higher profits for their distaste. This scenario leads to lower female application rates at discriminating firms, lower selection rates and lower wages for women at discriminating employers. In the second scenario, we assume nonlinear and linear jobs as proposed by Goldin (2014). In nonlinear jobs, higher input (e.g., in terms of providing more working hours or more employer-sided flexibility) leads to a more than proportional increase in output. We assume that the desired input level among men and women is heterogeneous. If there is a larger fraction of women who is willing/able to provide lower input, this will generate a sorting equilibrium with more women applying for linear jobs and more men applying for nonlinear jobs. Under strong sorting (i.e., workers who are unable to provide a large input apply predominantly for linear jobs), firms with nonlinear production functions would predominantly receive applications from workers that are willing and able to provide a high input. Thus, men and women who actually apply at these nonlinear firms would have similar selection rates and wages.

In the first step of the empirical analysis, we sort different hiring firms along AKM firm wage effect deciles, which we obtained from a two-way fixed effects regressions (Abowd et al., 1999; Card et al., 2013). We find that the probability for women to apply for a job decreases almost monotonically in the firm wage premium. After taking into account differences in sectors, occupations, and firm size, women have a 10 percentage points higher probability to apply in the lowest AKM firm decile and a 6 percentage points lower probability to apply in the highest AKM decile.[4] Interestingly, within these deciles, employers seem not to discriminate against women, as we observe indistinguishable male vs. female selection rates in the second stage of the application process (after controlling for sectors, occupations, and firm size).

[4]Moreover, this pattern holds qualitatively within well-defined occupational task complexities. See Appendix B.2.

In the second step, we show that the (residualized) share of male applicants increases in various indicators for employer-sided flexibility (e.g., mobility, longer working hours, changes in working hours). Therefore, we argue that the share of male applicants is a suitable encompassing proxy for employer-sided flexibility, which we use in subsequent regressions.

In the third step of the empirical analysis, we estimate standard Mincer wage regression controlling for detailed worker, firm, and job characteristics. Next, we add the share of male applicants as proxy for employer-sided flexibility. We find that this proxy has significant explanatory power beyond standard observables. The residual gender earnings gap drops significantly in all specifications. It falls from around 14-15 percentage points to around 6-7 percentage points, i.e. by around 50-60 percent.

Importantly, the share of male applicants is also relevant for the level of earnings when we look at female matches only. Women who match in a pool with a low share of male applicants earn about 10-20 percent less than women who match in a pool with a high share of male applicants (again netting out worker, firm, and job characteristics). Furthermore, workers that match in a pool with a larger share of male applicants have on average larger AKM worker fixed effects. This is in line with what we expect based on the model. If a certain group of workers matches at firms with nonlinear production functions (proxied by a larger share of male applicants in the data), they will produce more on average and part of this larger production will be handed on to workers as higher wages. Thus, these patterns in the data provide further support for the sorting hypothesis from the theoretical model.

Finally, we show that the residual gender earning gaps is significantly larger for mothers than for women without children and that there is a strong interaction with flexibility requirements. If mothers match at high-flexibility jobs, they face much larger discounts both relative to men and relative to women without children. Again, this is in line with our hypothesis of nonlinear production functions. Women with children tend to be less flexible. If they match at nonlinear jobs, they produce significantly less and thereby face particularly large wage discounts.

Our findings are complementary to a recent strand of the literature that analyzes gender-wage gaps for specific industries or firms (Azmat and Ferrer, 2017; Bolotnyy and Emanuel, 2019; Cook et al., 2021). These authors find that once they control for the detailed working behavior (e.g. working longer hours or working night shifts), the gender wage gap shrinks considerably. While these studies have very detailed information on gender-specific behavior of workers within certain industries, we have a dataset that represents the entire economy, but at the same time contains information on application behavior and flexibility requirements that are both typically absent in standard datasets.

Our work is most closely related to another recent strand of literature that analyzes gender issues combining insights from search and matching theory with rich microeco-

nomic data. Faberman et al. (2017) document men and women's job search behavior and the implications for the gender wage gap using US survey data for workers. Cortés et al. (2021) show a substantial difference between men and women in terms of the timing of their job acceptances based on a sample of (former) undergraduate students. Xiao (2021) analyzes the gender wage gap from a life cycle perspective and finds that both statistical discrimination based on fertility concerns and different labor force attachments play an important role in explaining the gender wage gap in Finland. While these studies are similar in spirit to our paper, the unique combination between the tractable model and the IAB Job Vacancy Survey with its linkages with administrative data allows us to shed light on the intertwining of the gender-specific application of workers and the selection behavior of firms. Specifically, the data allows us to explore the role of job characteristics such as employer-sided flexibility requirements while simultaneously controlling for important worker and firm characteristics. Due to the cross-sectional nature of our data, we have less to say on the life cycle component. However, in Appendix B.1, we show that the residual gender earnings gap is particularly large for women who match in their 30s and 40s (when childcare considerations may matter most). In addition, we directly show that women with children face the largest earnings discount in male-dominated jobs. This observation is in line with Illing et al. (2021) who show that having children sharply increases the gender gap in earnings losses after displacement. Fluchtmann et al. (2020) are probably closest to our paper. They use Danish unemployment insurance (UI) recipients data to empirically show that the differences in the application behavior between males and females can explain large parts of the traditional gender wage gap. The data are very similar, however, ours are not limited to UI recipients as they contain all hires. In addition, we have specific information about the exact gender distribution of the pool of applicants for each specific recruitment process, which allows us, as we show below, to calculate important measures derived from our model which help explaining the gender wage gap.

Our paper also contributes to the recent literature on compensating differentials. Sorkin (2018) shows for the United States that compensating differentials can explain around two-thirds of the variance of firm-level earnings. Taber and Vejlin (2020) show for Denmark that preferences for non-pecuniary aspects are very important for job choices. Our empirical findings are in line with these findings. Women have a higher probability to apply for low-wage jobs and to get compensated in terms of less employer-sided flexibility requirements. Consistently, Budig and Hodges (2010) show that mothers are more willing than women without children to trade wages for family friendly employment and that this effect is largest at the lower part of the wage distribution.

Based on experimental data, Wiswall and Zafar (2017) show that women have a higher willingness to pay for non-wage job characteristics. In the same vein, Le Barbanchon et al. (2020) analyze gender differences in willingness to commute. They show for France that

women commute much shorter than men. Based on their search model, they find that 14 percent of the residualized gender wage gap can be explained by this mechanism. While the IAB Job Vacancy Survey does not contain any commuting times, we believe that this mechanism is included in our regressions when we use the share of male applicants in our Mincer type regressions. Matches that require longer commuting times can be expected to be disliked by women (in particular those with care responsibilities). In our view, this is another dimension of employer-sided flexibility requirements that is not directly measurable in our data. Thus, our proxy for employer-sided flexibility is more encompassing than pure commuting times. Against this background, it makes sense that the residual gender-earnings gap is reduced by a lot more in our regressions (by 50 to 60 percent) than in Le Barbanchon et al. (2020).

Our paper is also highly relevant from an economic policy perspective. In particular, the Covid-19 episode with work-from-home arrangement provided a test-laboratory whether more flexibility from the employee side is possible. Barrero et al. (2020) argue that these working from home arrangements have boosted productivity. To the extent that these arrangements have changed the production process and that they will stick permanently, the results from our paper imply that this leads to a decline in the residual earnings gap, as this would make certain jobs more accessible and attractive for women.

The rest of the paper proceeds as follows. Section 2 describes the model framework and derives theoretical implications for taste-based discrimination at the hiring stage and for different production functions. Section 3 provides details on the used datasets. Section 4 contains the empirical analysis on gender-specific application behavior, the estimated gender earnings gap, differences between male and female dominated jobs, and how flexibility requirements and being a woman with children interact. Section 5 briefly concludes.

2 Theory

We derive a theoretical model that allows us to interpret the patterns in the IAB Job Vacancy Survey through a gender-specific labor market flow perspective.[5] In the data, we observe the application behavior of males and females for particular jobs (both in terms of pay and flexibility requirements) and the hiring behavior of firms for particular jobs. Accordingly, our model assumes a two-stage decision problem (i.e., application and hiring/selection). In the first stage, we assume that applicants face ex-ante heterogeneous application costs. A worker only applies for a certain job if the expected return is larger than the application costs. Those workers who apply for a particular job draw an idiosyncratic match-specific cost shock in the second stage. Only those worker-firm pairs with

[5] In line with the cross-sectional dataset, the model is completely silent on some potentially important dimensions (e.g. the intertemporal life cycle perspective).

a positive match surplus will create a match, i.e. only a certain fraction of workers will be selected by firms.[6] We analyze the implications of two specific scenarios and compare them to patterns in the data. First, some firms may do taste-based discrimination at the hiring stage, i.e. they may dislike hiring women. Discrimination by firms leads to lower selection rates, lower female application rates and lower earnings. Second, following Goldin (2014), we assume that there are jobs with nonlinear and others with linear production functions. At nonlinear jobs, the output increases more than proportionally with input. Working hours are certainly an important dimension of input. However, we define input in a more multidimensional sense than Goldin (2014) and consider employer-sided flexibility requirements as particularly important dimension (e.g., ability to do business travels or to be available on short notice). If women are on average less able or willing to provide higher input levels (e.g. high flexibility is difficult if they bear a larger part of childcare), this leads to sorting into different jobs and also generates a gender earnings gap. On average female application rates would also be lower, but those women who apply at nonlinear firms would not necessarily have lower selection rates (as those are willing to provide large input) and earnings than men.

2.1 Model Environment

We assume that there is a continuum of different job profiles, where y_p denotes the output level for each of these jobs. For simplicity but without loss of generality, we derive a static model and we exclude the possibility of multiple vacant jobs for one worker, i.e. one random job is visible for each searching worker. In the first stage, workers have to decide whether to apply for particular jobs or not. They will do so if application costs e are smaller than the expected return from this application. The ex-ante application costs e are drawn from a stable density function, $g(e)$. The application costs are sunk at the time of application, i.e. they will not play any role for the surplus in the second stage.

In the second stage, workers j that decided to apply for a particular job profile p, draw a match-specific training cost shock upon contact with a firm. We denote this shock by $\varepsilon_{p,j}$. The ex-post shock is drawn from a stable density function, $f(\varepsilon)$. Only those worker-firm pairs with a positive joint surplus will create a match.

2.1.1 Application Decision

Worker j will apply for a particular job p whenever the expected returns from a match are larger than the application costs:

$$E\eta_{p,j}\bar{w}\left(\tilde{\varepsilon}_{p,j}\right) - \xi_j > e_{p,j}. \tag{1}$$

[6] For other selection models, see Brown et al. (2016), Chugh and Merkl (2016), or Carrillo-Tudela et al. (2020)

The left-hand side of the equation shows the expected returns from a match, where $\eta_{p,j}$ is the hiring rate in the second stage, $\bar{w}\left(\tilde{\varepsilon}_{p,j}\right)$ is the expected wage conditional on being hired that will be defined below (which is a function of the cutoff point in the second stage, $\tilde{\varepsilon}_{p,j}$), and E is the expectations operator.[7] ξ is the worker's value of unemployment (e.g., home production and benefits). e are application costs that are drawn from a stable density function.

Thus, there is a certain cutoff point level, $\tilde{e}_{p,j}$, up to which workers will apply for job type p:

$$\tilde{e}_{p,j} = E\eta_{p,j}\tilde{w}\left(\varepsilon_{p,j}\right) - \xi_j. \tag{2}$$

Above $\tilde{e}_{p,j}$, application costs are larger than returns. Below this threshold, workers will apply for job p. The application rate of group j for a particular job p is the integral from the lower support of the distribution up to the cutoff point:

$$\alpha_{p,j} = \int_{e_p^{\min}}^{\tilde{e}_{p,f}} g\left(e\right) de. \tag{3}$$

2.1.2 Hiring Decision

Upon contact, each worker-firm pair draws an idiosyncratic match-specific cost shock, $\varepsilon_{p,j}$, which we interpret as training costs. Some workers require little training, others require a lot of training to do the same job. Once a match is formed, each job profile produces a certain output level $y_{p,j}$, which may be dependent on the willingness of the worker to provide effort (to be discussed and specified below). In addition, there may be taste-based discrimination of employers at the hiring stage against certain worker groups. This means that the firm will only hire from this group if there is a compensation in the amount of $t_{p,f}$ for the distaste. The joint match surplus between workers and firms is defined as:

$$\Pi_{p,j} = y_{p,j} - \varepsilon_{p,j} - t_{p,j} - \xi_j > 0. \tag{4}$$

The wage is bilaterally efficient as long as both workers and firms have a positive match surplus (in case of a positive joint surplus)

$$w\left(\Pi_{p,j}\right) - \xi_j \geq 0, \tag{5}$$

$$y_{p,j} - w\left(\Pi_{p,j}\right) - \varepsilon_{p,j} - t_{p,j} \geq 0. \tag{6}$$

[7] In the first stage, workers do not know their shock realization in the second stage yet. However, they know the output level of the job, y_p, and the properties of the training costs distribution. Therefore, under rational expectations, they know the average expected hiring probability and wage.

Under a bilaterally efficient wage formation process, there will be production whenever $\Pi_{p,j} > 0$. Thus, we can calculate the cutoff point for idiosyncratic match-specific costs up to which workers and firms are willing to produce:

$$\tilde{\varepsilon}_{p,j} = y_{p,j} - t_{p,j} - \xi_j. \tag{7}$$

The selection rate of a worker from group j at job p is the integral from the lower support of the idiosyncratic cost function up to the cutoff point:

$$\eta_{p,j} = \int_{\varepsilon_p^{\min}}^{\tilde{\varepsilon}_{p,j}} f(\varepsilon)\, d\varepsilon. \tag{8}$$

2.1.3 Wage Formation

In order to be able to define the wage and the application rate, we need to take a stance on wage formation. Without loss of generality, we assume Nash bargaining between workers and firms. This leads to the plausible outcome that wages are a function of firm-specific output, the idiosyncratic training costs realization and workers' fallback options.

Under Nash bargaining, workers and firms maximize their joint Nash product, Λ, with respect to the wage:

$$\Lambda = \left(w\left(y_{p,j}, \varepsilon_{p,j}, \xi_j\right) - \xi_j\right)^{\alpha} \left(y_{p,j} - w\left(y_{p,j}, \varepsilon_{p,j}, \xi_j\right) - \varepsilon_{p,j} - t_{p,j}\right)^{1-\alpha}, \tag{9}$$

where α is workers' bargaining power.

This yields the following wage:

$$w\left(y_{p,j}, \varepsilon_{p,j}, \xi_j\right) = \alpha\left(y_{p,j} - \varepsilon_{p,j} - t_{p,j}\right) + (1-\alpha)\,\xi_j. \tag{10}$$

Equations (5) and (6) establish conditions under which wage formation is bilaterally efficient. They hold under Nash bargaining.

Based on the wage formation mechanism, we calculate the expected wage conditional on being hired for a particular job that we require for the first stage of the decision process:

$$\bar{w}\left(\tilde{\varepsilon}_{p,j}\right) = \frac{\int_{\varepsilon_{p,j}^{\min}}^{\tilde{\varepsilon}_{p,j}} w(\varepsilon)\, f(\varepsilon)\, d\varepsilon}{\eta_{p,j}}, \tag{11}$$

where $\tilde{\varepsilon}_p$ are the cutoff points in the second stage and $\varepsilon_p^{\min}$ is the lower support of the density function.

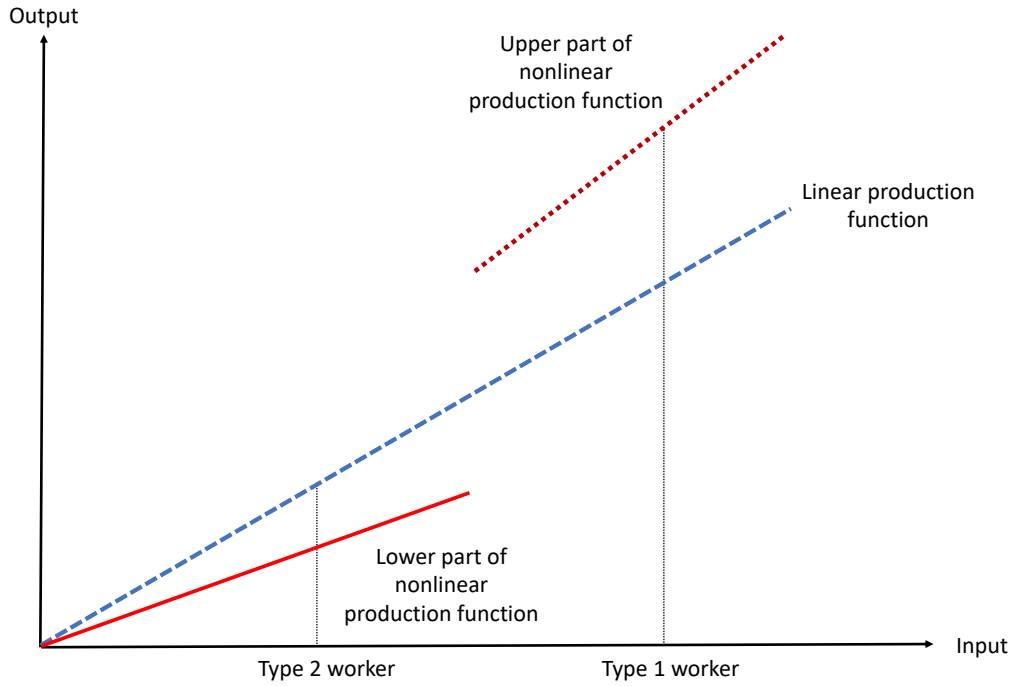
2.1.4 Production

We allow for two scenarios in terms of production. Either there is a fixed production level for each job profile, y_p, or there may be two types of production functions. The second

case will be derived below.

Following Goldin (2014), we assume that there may be firms with different production functions and that workers can choose some input margin, λ_j.[8] Note that we interpret input in a broader sense than Goldin (2014). Input may be working hours, but it may also be other employer-sided flexibility requirements such as working in different locations or such as being available on short notice.

Figure 1: Nonlinear and Linear Jobs



Note: The figure illustrates the output as a function of input for a linear and a nonlinear production function. It illustrates the input-output connection for a worker who is willing to provide a high input (type 1) and for a worker who is willing to provide a lower input (type 2.)

For jobs with a nonlinear production function, nl, output is defined as:

$$y_{nl,j} = \lambda_j a_{nl} \text{ if } \lambda_j > \lambda^* \tag{12}$$

$$y_{nl,j} = \lambda_j a_{nl} (1 - \delta) \text{ if } \lambda_j \leq \lambda^* \tag{13}$$

In addition, there are other jobs where the output is linear, l:

$$y_{l.j} = \lambda_j a_l \tag{14}$$

As in Goldin (2014), we assume that $\lambda_j a_{nl} > \lambda_j a_l$ for $\lambda_j > \lambda^*$ and $\lambda_j a_{nl} < \lambda_j a_l$ for $\lambda_j < \lambda^*$. Figure (1) illustrates the nature of the two production functions. If a worker is

[8] As we focus on workers' application behavior in a partial setting, we abstract from the question under which circumstances these nonlinear and linear firms coexist in a full general equilibrium setting.

willing to provide working hours/flexibility beyond the minimum threshold λ^*, this leads to more output at nonlinear firms than at linear firms. If not, there is more production at linear firms.

2.1.5 Equilibrium

The labor market equilibrium is described by the application cutoff point in equation (2), the application rate (3), the cutoff point for the idiosyncratic match-specific cost shock (7), the corresponding selection rate (8), and the wage expectations conditional on being hired (11). Output per job is either exogenous. Alternatively, production may be governed by different types of (non)linear production functions and the willingness of applicants to provide certain effort levels.

2.2 Model Implications

Our model allows us to analyze how different scenarios affect the application rate, the selection rate and the wage for different worker groups j. So far, we have not taken any stance yet on what the index j stands for. Therefore, we now look at two scenarios. First, we analyze what happens if there is taste-based discrimination against women in high-productivity jobs. Women earning systematically less than men (controlling for observables), may be driven by discrimination at firms that produce a large output level per worker. Second, we analyze the implications of our model with nonlinear and linear jobs. Goldin (2014) argues that some jobs show a nonlinear production function, where higher inputs increase the output more than proportionally.[9]

2.2.1 Taste-Based Discrimination

Let's start by assuming that workers are ex-ante homogeneous and production per job is exogenous, y_p. Applicants only differ in terms of their gender. For the sake of the argument, assume further that employers at certain firms/jobs discriminate against women at the hiring stage ($t_{p,f} > 0$, $t_{p,m} = 0$, where f stands for female and m for male).

Taste-based discrimination of females would reduce the joint surplus in case of a female match and thereby reduce the cutoff point for the idiosyncratic shock realization:

$$\tilde{\varepsilon}_{p,f} = y_{p,f} - t_{p,f} - \xi_f \tag{15}$$

This leads to a lower selection rate in the second stage of the application process and to on average lower wages for females.

[9] As a third potential mechanism, we could analyze different bargaining powers of men and women. However, we do not have any direct proxy for the level of bargaining power in our dataset. In addition, we show in Appendix B.4 that our results are very similar at firms with and without an institutionalized bargaining agreement (e.g., collective bargaining).

$$w\left(y_{p,f},\varepsilon_{p,f},\xi_f\right)=\alpha\left(y_{p,f}-\varepsilon_{p,f}-t_{p,f}\right)+\left(1-\alpha\right)\xi_f \tag{16}$$

As women anticipate that they will be hired at discriminating firms with a lower probability if they apply and that they will earn lower wages, only a smaller fraction of women will send an application to these firms in the first place, i.e. the cutoff for application is lower:

$$\tilde{e}_{p,f}=E\eta_{p,f}\tilde{w}\left(\varepsilon_{p,f}\right)-\xi_f, \tag{17}$$

Overall, taste-based discrimination at the hiring stage leads to lower female application rates, lower female selection rates, and lower female wages.

2.2.2 (Non)Linear Production Functions and Sorting

Next, we analyze the implications of two types of production function (linear and nonlinear). Let's assume for illustration purposes that there are two types of workers. Type 1 workers are willing/able to provide a larger input, λ_j, than type 2 workers. In addition, we assume that type 1 workers are above the threshold, $\lambda_1>\lambda^*$, while type 2 workers are below, $\lambda_2<\lambda^*$.

Under these assumptions, we obtain four different cutoff points:

$$\tilde{\varepsilon}_{nl,1}=\lambda_1 a_{nl}-\xi_1, \tag{18}$$

$$\tilde{\varepsilon}_{l,1}=\lambda_1 a_l-\xi_1, \tag{19}$$

$$\tilde{\varepsilon}_{nl,2}=\lambda_2\left(1-\delta\right)a_{nl}-\xi_2, \tag{20}$$

$$\tilde{\varepsilon}_{l,2}=\lambda_2 a_l-\xi_2. \tag{21}$$

And under our assumptions, the following ranking holds:

$$\tilde{\varepsilon}_{nl,1}>\tilde{\varepsilon}_{l,1}, \tag{22}$$

and

$$\tilde{\varepsilon}_{l,2}>\tilde{\varepsilon}_{nl,2}. \tag{23}$$

Thus:

$$\eta_{nl,1}>\eta_{l,1}, \tag{24}$$

$$\eta_{l,2} > \eta_{nl,2}. \tag{25}$$

Intuitively, type 1 workers generate the largest production at nonlinear production firms and thereby face the largest selection rate at these firms. By contrast, type 2 workers generate the largest production at firms with linear production functions. The same ranking is true for wages and thereby the probability to apply at the respective firms.

Under certain parametrizations (large differences in production and small dispersion of idiosyncratic application costs), our model generates a complete sorting equilibrium of the following sort:

$$\eta_{nl,1} > \eta_{l,1} = 0 \tag{26}$$

$$\eta_{l,2} > \eta_{nl,2} = 0 \tag{27}$$

In this case, type 1 workers would have no surplus at linear jobs and type 2 workers would have no surplus at nonlinear jobs. Although this example appears to be extreme, it is very useful for illustration purposes.

How could different production functions and input provisions interact with gender? Even nowadays women bear a larger responsibility in terms of childcare and other family-related responsibilities. Thereby, a larger fraction of women may be less flexible in terms of input provision than men (i.e. they may have trouble working long hours, being available on short notice, or doing business travel). Assume that a larger share of men is type 1 workers (compared to women). In this case, we would observe that the average application rate of females at high-productivity firms (those with nonlinear production function) is lower. Note that under complete sorting those women who match at nonlinear firms (only type 1 females) have the same selection rate and the same wage as males.

We are unable to observe type 1 and type 2 persons in the data directly. However, one of the key data innovations is that we have proxies which measure the required flexibility at specific job vacancies (e.g., hours worked or other flexibility requirements) and for measuring how flexible workers can act (e.g. whether women are mothers or not).

2.2.3 Model and Data

Although our theoretical model is too simple to be used for structural model estimations, it provides useful guidance at which outcome variables we should look at. The model provides a roadmap for the empirical analysis.

As we have AKM firm-fixed effects for each firm and we observe the number of applicants for each job, we can calculate the share of female applicants and the probability of being selected (upon application) for jobs with different wage premiums. In a first step,

we will test our hypothesis of taste-based discrimination at the hiring stage by checking whether hiring probabilities for women (upon application) are generally lower than for men (controlling for observables). In addition, we will check whether such a pattern is prevalent in high-wage premium firms. If high-wage firms discriminate more than low-wage firms, this would lead to a gender earnings gap, as women would apply at these firms with lower probability and as they would be selected with lower probability at these firms. Overall, this would depress the share of women in firms with the highest earnings.

In a second step, we will analyze the connection between female application behavior and employer-sided flexibility requirements at the job level. This will help us to understand whether these flexibility requirements (potentially driven by nonlinear production functions) may be an important driver for gender differences. In addition, it will help us to understand whether the share of male applicants may be a suitable proxy for these flexibility requirements.

In a third step, we will analyze whether the share of male applicants matters for the realized wage. We will analyze whether females who match in a pool with a larger share of male applicants earn more than females who match in a pool with a large share of female applicants (controlling for observables).

Finally, we move to the person level and analyze how the share of male applicants is correlated with worker-fixed effects. Under sorting, we expect them to be larger with a larger share of male applicants, as workers at more demanding workplaces generate more output and thereby larger wages. In addition, we will directly check whether having children affects certain outcomes for women. This provides a direct test for the question whether nonlinear production functions and inflexibility for women with children interact.

3 Data

3.1 Data Sources

We use the IAB-Stellenerhebung (IAB Job Vacancy Survey, JVS, see Moczall et al., 2015) as our primary source of data. The JVS covers up to 14,000 establishments per year and is a representative survey among establishments in Germany from all sectors and from all establishment size classes. Each year, the survey collects information on the hiring process of German establishments.[10]

An important component of the JVS is an array of questions about the recruitment process of their most recent new hire. These questions gather information on job characteristics such as the exact job requirements, search channels, search duration, the exact hiring date, individual hire attributes such as gender, age, as well as match-specific char-

[10]We use the information from the 'main' survey, which is conducted in each fourth quarter. For a subset of establishments, there are follow-up questionnaires in the three next quarters.

acteristics like educational qualification, wage bargaining, and, in some waves, the hourly wage. Crucial for our purposes, the JVS asks about details on the pool of applications for the most recent hire. Specifically, employers report the number of applicants, the (self-assessed) number of suitable applicants, the number of invited individuals, and their gender composition.

We complement the JVS data with information from the German social security system. Specifically, we use the method developed by Lochner (2019) to identify establishments' last hires in the administrative records, the Integrated Employment Biographies (IEB). The identification is based on overlapping information such as the hiring date, workers' age, gender, and occupational codes. Using a deterministic matching algorithm, around 70% of the last hires from the JVS can be found in the administrative records.[11] The IEB encompasses labor market information for the majority of workers in Germany.[12] Combining the survey data with the administrative records, hence allows us to observe workers' entire employment and earnings history.

In our baseline specifications in the main part, we restrict the sample to full-time jobs, which we define as a job with more than 25 contractual hours. In Appendix C, we show that all our results are robust when giving up this restriction and taking also part-time jobs into account.

3.2 Administrative Data Linkages and Imputations

The social security data reports the total wage sum over workers' employment spell. These sums are right-censored at the contribution assessment ceiling ('Beitragsbemessungsgrenze'), given by the statutory pension fund. We follow Dustmann et al. (2009) and fit a series of Tobit regression to impute the censored part of the wage distribution.[13] The administrative data lacks details on hours worked, so only wages for full-time workers can be estimated. However, the share of part-time observations with censored wages is negligibly small (less than 1%).

For workers' educational attainment, we construct a variable from information on both schooling and education in terms of the German vocational system. First, we correct for misreporting and inconsistencies using the procedure proposed by Fitzenberger et al. (2006). Then, we build a categorical variable with five distinct values: 1) intermediate school leaving certificate without vocational training, 2) intermediate school leaving certificate with vocational training, 3) upper secondary school leaving certificate without

[11] The algorithm performs several plausibility checks with respect to deviations in the overlapping information. Note that hires with missing information in the key variables are not taken into account.

[12] The IEB covers around 80 percent of the German working population, only excluding civil servants and the self-employed.

[13] First, wages are deflated. Then, Tobit regressions are performed separately for East and West Germany as well as males and females. All regressions control for age and education categories, and all possible interactions.

vocational training, 4) upper secondary with vocational training, 5) college or university degree.

To exploit the role of children, we will use established proxies for motherhood (Mueller and Strauch, 2017).[14] The proxy uses family-related breaks in the employment biography of females to identify childbirth in the administrative data. For identification, the approach uses either employment notifications (maternity allowance payments by the statutory health insurance provider during paid maternal leave) or detailed process data of the Federal Employment Agency (e.g., withdrawal into maternity allowance) about unemployment and benefits. Since the procedure is suitable for all of the administrative data, we can run it on our linked JVS-IEB sample and hence identify women with children among the identified JVS hires.

3.3 Final Sample

For our analysis, we use the JVS from 2010–2016.[15] We then link the administrative data to the survey information. In the end, our estimation sample consists of 21,694 distinct new hires for which we have further information on the recruitment process such as the pool of applicants. Furthermore, we can link workers' full employment history to the new hires data. Table 1 shows descriptive statistics for our main variables, separately for females and males.

Table 1: Main variables by gender

	females		males	
individual characteristics	mean	std. dev.	mean	std. dev.
age	35.86	10.75	36.46	10.91
education (scale 1-5)*	2.39	1.73	2.05	1.66
experience (years)	8.19	8.19	9.67	8.38
match characteristics				
contract hours	34.40	7.69	38.85	4.20
job requirements (scale 1-4) **	2.17	0.61	2.12	0.63
firm size decile	5.47	2.92	5.44	2.88
firm wage premium decile	5.47	2.89	5.58	2.84
log daily wage	4.13	0.47	4.36	0.44
log daily wage if full-time	4.22	0.42	4.37	0.42

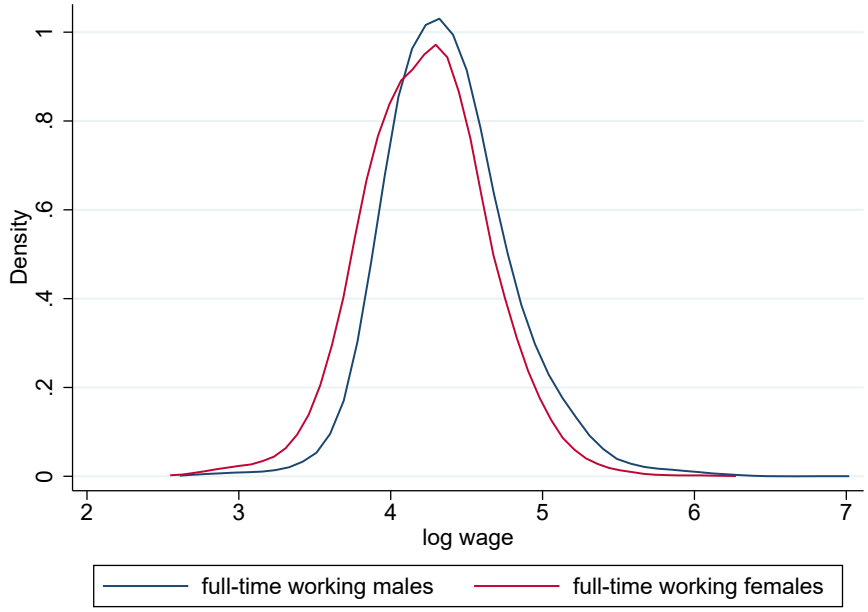
Note: *1) intermediate school leaving certificate without vocational training, 2) as 1) but with vocational training, 3) upper secondary school leaving certificate without vocational training, 4) as 3) but with vocational training, 5) College or university degree; ** 1) missing 2) unskilled 3) vocational training, 4) college or university; Source: JVS, IEB;

[14] The administrative data also allows to use a proxy for marriage (Baechmann et al., 2021). We have experimented with this proxy. However, motherhood appears to be the more meaningful variable to use.

[15] Due to legal reasons, we can only link individual information from the administrative sources to the JVS from 2010 onward.

On average, at the time of being hired, males are about half of a year older and about 1.4 years more experienced than females in our sample. Females are somewhat more educated. Males work on average around 4 hours longer. Females and males are not different with respect to the formal job requirements that are linked to the positions in which they get hired. The same is true with respect to firm size. However, when we look at wage outcomes, we observe large differences. The unconditional difference in the hiring wage amounts to 23 log points on average for all kinds of jobs in our sample.[16] Figure 2 shows the distributions of the hiring wages for females and males in full-time jobs.

Figure 2: Hiring wage distribution by gender



Note: Estimates for full-time workers only. Source: JVS, IEB;

In contrast to most other datasets, the IAB Job Vacancy Survey contains information on the pool of applicants for a particular hire. Specifically, firms report the number of male and female applicants for their most recent hire. Hence, we can calculate the share of male/female applications. Table 2 shows the distribution of the share of male applications for different occupations.[17] Women are for example more likely to apply in health care related occupations than men, while the opposite is the case in occupations related to construction and architecture. Table A.1 in the Appendix shows similarly distinct application patterns across industry sectors. The share of male applicants is for example much larger in manufacturing than in certain service sectors (e.g. related to

[16] We define the hiring wage as the wage within the first employment spell in the administrative data that refers to the new hire.

[17] Note that the shares of female and male applications always sum up to one for each hire and thereby also for each occupation.

Table 2: Share of male/female hires and applicants across occupations

KldB2010 1-digit	total hires	share of hires		share of applicants	
		males (%)	females (%)	males (%)	females (%)
1 Occupations in agriculture, forestry, farming, and gardening	701	68.47	31.53	66.46	33.54
2 - Occupations in production of raw materials/ goods/ manufacturing	4,785	84.91	15.09	82.65	17.35
3 - Occupations in construction, architecture, surveying, technical building services	1,601	90.82	9.18	88.90	11.10
4 - Occupations in natural sciences, geography, informatics	894	77.85	22.15	75.80	24.20
5 - Occupations in traffic, logistics, safety, security	1,910	80.00	20.00	76.16	23.84
6 - Occupations in commercial services, trading, sales, the hotel business/ tourism	1,814	40.24	59.76	40.26	59.74
7 - Occupations in business organisation, accounting, law, administration	5,643	30.96	69.04	34.82	65.18
8 - Occupations in health care, the social sector, teaching, education	3,679	17.75	82.25	19.44	80.56
9 - Occupations in philology, literature, humanities, social sciences, economics, media, art, culture, design	574	41.99	58.01	43.48	56.52
Total	21,604	53.67	46.33	58.66	41.34

education).[18]

4 Empirical Results

4.1 Application and Selection Patterns at the Firm Level

We start by looking at the application and selection behavior at particular firms through the lens of our theoretical model. For this purpose, we use the information on the pool of applicants for different jobs from the IAB-Job Vacancy Survey. We know the gender-composition of application pools, i.e. the number of male and female applicants. However, we do not know any further characteristics of these applicants. In later steps, we will also use information on the characteristics of the person who was actually hired and the characteristics of the job.

In the theoretical model, higher firm-specific wages may either be driven by a larger output per worker or wage formation. As we do not have any value added or sales information in the IAB Job Vacancy Survey, we analyze how the gender-specific application behavior differs across firm fixed effects from a two-way fixed effects regressions as described in Bellmann et al. (2020) and Lochner et al. (2020).[19]

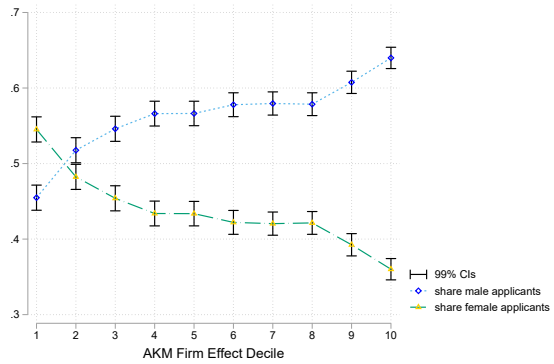
Panel (a) of Figure 3 shows the share of male and female applicants for each of these firms, ranked according to AKM firm-fixed effect deciles (with the firms that pay the

[18] In line with results by Gomes and Kuhn (2019), female application rates are much larger in the public sector than in the rest of the economy. See Appendix.

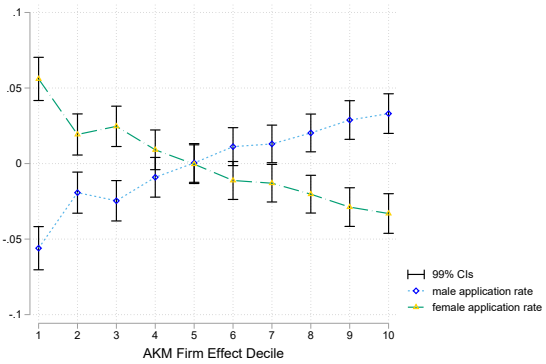
[19] These authors run an AKM wage regression on the universe of German administrative data for 2010-2017 in the spirit of Abowd et al. (1999). These effects imply firm-specific wage premiums (or discounts), often associated with rent-sharing, efficiency wages, or strategic wage posting behavior (see among others Card et al., 2013; Postel-Vinay and Robin, 2002; Burdett and Mortensen, 1998)

Figure 3: Application and selection rate by gender and AKM firm effect deciles

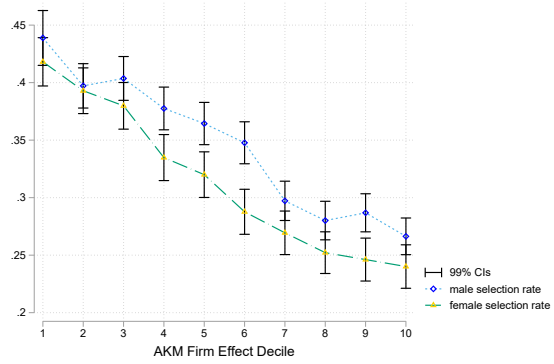
(a) Share of female/male applications



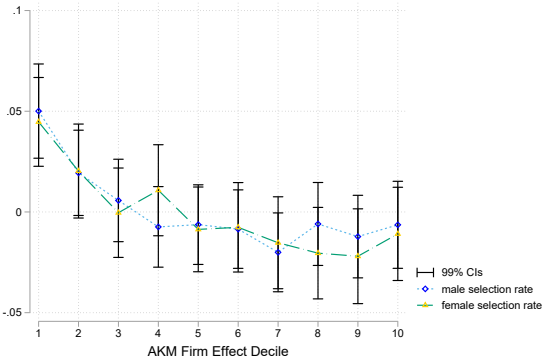
(b) Residualized share of female/male applications



(c) Female/male selection rate



(d) Residualized female/male selection rate



Note: Full-time jobs only. Variables are defined as follows: a) number of female (male) appl./ number of all applications; b) 1/ number of all applicants if male/female hired; Control variables: Industry categories (Nace Rev 2), firm size categories, occupation categories (5 digits); Source: JVS, IEB;

largest average discount on the left-hand side, and the firms with the largest premium on the right-hand side). Men have a more than 20 percentage point larger probability than women to apply at firms which pay the highest earnings premiums. At the bottom of the earnings premium distribution, the opposite is true. Women have a 15 percentage point higher probability to apply at firms which pay the lowest premiums than men.

A sizeable part of these patterns may be driven by women and men choosing different sectors, occupations, or even firms of different sizes, as visible in Table 2 and Table A.1. Therefore, we control for occupation, industry, and firm size in panel (b) of Figure 3, as these are variables that are typically controlled for in Mincer wage regressions. Although the differences between male and female application behavior are quantitatively less pronounced when adding controls, the striking insight is that a substantial gap in the application behavior remains. There is a roughly 6 percentage points larger probability for men to apply at the highest-wage firms and a 10 percentage points larger probability for women to apply at the lowest-wage firms. Through the lens of our model, this large difference in the gender-specific application behavior may either be driven by taste-based discrimination at the hiring stage or by different production functions at different jobs.

To further open up the black box of matching, we use a proxy for the selection behavior of firms conditional on application. Following Hochmuth et al. (2021) and Lochner et al. (2021), we define the selection rate as the inverse of the number of applicants and interpret it as a measure of employer selectiveness. If a firm selects a small fraction from the pool of applicants, it is considered as selective. In terms of our model this means that only a small fraction of applicants is hired.

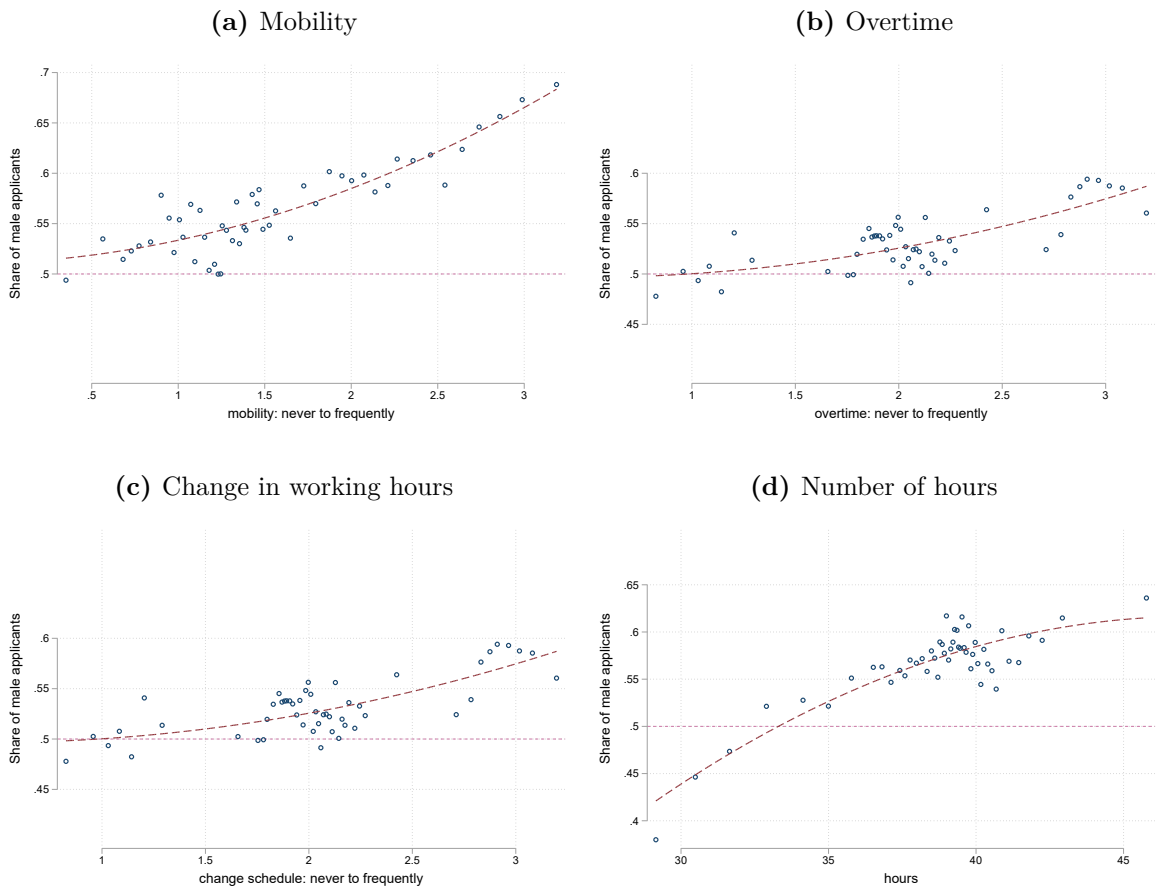
Panel (c) of Figure 3 shows that the (uncontrolled) selection rate of men and women differs across AKM-deciles. However, at firms with the highest wage premiums, the probability of men and women getting hired/selected (conditional on applying) is almost the same (with confidence bands overlapping). When we control for sector, occupation, and firm size in panel (d), male and female selection rates are almost the same in all deciles.

Given the stark differences in gender-specific application rates and the strong similarities in selection rates across AKM-deciles, the model mechanism that high-paying firms discriminate more strongly against women than low-paying firms (and thereby drive up the earnings gap) is not supported by the empirical gender-specific selection patterns. By contrast, the patterns are reconcilable with the second hypothesis that high-paying firms offer different jobs (namely, nonlinear jobs) and attract workers that are willing to provide the necessary flexibility. Thereby, women who sort into these high-paying firms may have the same probability as men to be selected. We will analyze this hypothesis in more detail in the next subsections.

4.2 Application Behavior and Firm-Sided Flexibility Requirements

While our previous analysis was at the firm level, we now move to the job level. The IAB Job Vacancy Survey offers several proxies for firm-sided flexibility requirements. They serve as proxy for Goldin (2014)'s hypothesis of different production functions.

Figure 4: The share of male applicants and flexibility requirements



Note: Figures show binscatters with 50 bins and quadratic fit lines. To residualize the x-variable and y-variables, we regress each variable on the controls, generate the residuals, and add the sample mean of each variable back to its residuals. We then group the x-axis variable into equal-sized bins, compute the mean of the x-axis and y-axis variables within each bin, and create a scatterplot of these data points. Control variables: Industry categories (Nace Rev 2), firm size categories, occupation categories (5 digits); Source: JVS, IEB; Full-time jobs only; Sources: JVS, IEB.

We use four different flexibility requirements from the IAB Job Vacancy Survey that are asked for the last hire, namely the necessity to change the workplace (e.g. due to business travels), the necessity to work overtime, the necessity to change working hours on short notice, and the number of hours worked.[20] In Figure 4, we plot these four employer-sided flexibility requirements against the (residualized) share of male applicants.[21] In line

[20]Employers answer whether these flexibility requirements happen "often," "rarely," or "never" for the last hire. We experimented with further questions from the survey. These four selected dimensions seem to reflect the flexibility dimension best.

[21]Both the horizontal and the vertical axis are residualized by sector, occupation, and firm size.

with the second model hypothesis that there are different types of jobs, all four flexibility requirements comove positively with the share of male applicants for these particular jobs. Thus, these figures show that higher employer-sided flexibility requirements are associated with a larger share of male applicants.

In reality, flexibility requirements are multidimensional. Although the survey questions in the IAB Vacancy are much more detailed in this dimension than in many other surveys, we believe that employer-sided flexibility requirements can only be captured partially.[22] Given the strong connection between observed flexibility requirements and the share of male applicants, we consider the share of male applicants as a suitable proxy for multidimensional flexibility requirements. We will use this proxy for our further empirical analysis. In the next step, we will analyze how the residual gender earnings gap is affected by the gender-specific application behavior.

4.3 Residual Gender Earnings Gap

We start by estimating standard Mincer regressions, where we control for a rich set of observables. In addition, we add a dummy for females to estimate the size of the residual gender earnings gap. Recall that we observe new hires, hence we estimate the gap in hiring earnings without potential gender-specific tenure effects. In a second step, we add our proxy for firm-sided flexibility requirements, namely the share of male applicants. This variable is absent in standard datasets. Thereby, we can check how much of the residual gender earnings gap is due to an omitted variable bias.

Our benchmark Mincer regression looks as follows:

$$Log\ wage_{i,t} = \alpha\ gender_{i,t} + \gamma\ controls_{i,t} + error_{i,t}, \tag{28}$$

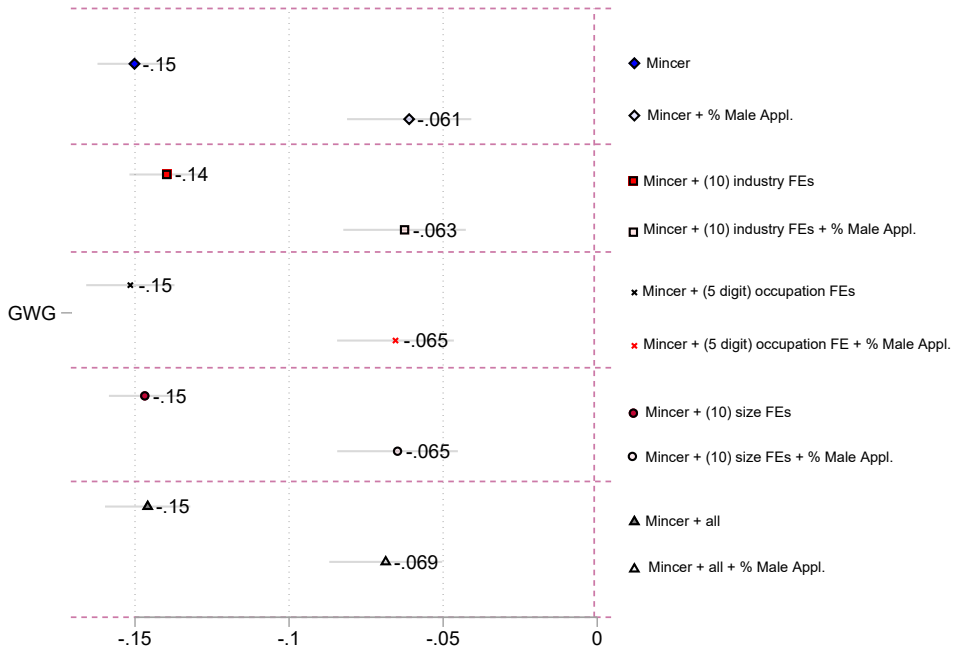
where i is the recruitment from the cross-sectional JVS in year t (2010 to 2016), and *gender* is a dummy with male as the reference group. In our benchmark specification, the set of *controls* includes the total number of applicants, worker age fully interacted with education attainment (measured by five categories), experience in years as well as its squared term, an indicator variable for the previous labor market status (non-employed, unemployed, employed), the contractual hours of the new job, formal job requirements (four categories), and year dummies. We estimate equation (28) for various specifications, which include additional controls. Specifically, we subsequently add a full set of dummies for industries, occupations, and establishment size deciles.

The left-hand side of Figure 5 shows the estimated α-coefficients for different regression specifications as laid out in the Figure legend. The estimated gender gap in the hiring

[22] This idea follows the train of thought in Goldin (2014, p.1104): "By job flexibility I mean a multitude of temporal matters including the number of hours, precise times, predictability and ability to schedule one's own hours."

earnings is 15% in our benchmark specification. Including a set of industry or occupation categories, or establishment size dummies to the control variables barely changes this pattern. Even if we add all these additional controls at once to the benchmark specification, the gender-gap in the hiring earnings is in the same ballpark. This is the same order of magnitude as in the existing literature for Germany (see for example Fuchs et al., 2019).

Figure 5: The gender hiring wage gap



Note: The Figure shows the estimates for the gender gap (α) in the hiring wage as specified in equation 29. Dependent variable: imputed log daily earnings. Default independent variables: gender dummy, the total number of applicants, worker age fully interacted with education attainment (measured by five categories), experience in years as well as its squared term, an indicator variable for the previous labor market status (non-employed, unemployed, employed), the contractual hours of the new job, formal job requirements (four categories), and year dummies. Estimates for full-time workers only. Source: JVS, IEB;

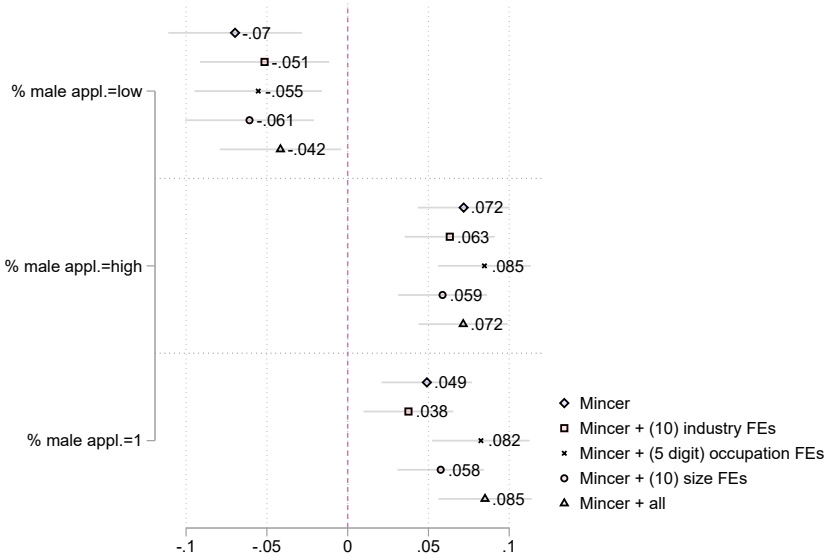
In the second step, we add the share of male applicants as an additional explanatory variable to control for the flexibility requirements of different jobs:

$$Log\ wage_{i,t} = \alpha\ gender_{i,t} + \beta\ share\ male\ appl._{i,t} + \gamma\ controls_{i,t} + error_{i,t}. \tag{29}$$

The right-hand side of Figure 5 shows that adding the gender-share of applicants reduces the gap in the hiring earnings to 6.1% (a reduction of 59%) in the benchmark specification. The same pattern holds in all other specifications. When adding the share of male applicants to the regressions, the residual earnings gap drops substantially

(although the precise magnitude depends on the specification).[23]

Figure 6: Coefficients for categories of share of male applicants, **male** hires



Note: The Figure shows the coefficients for the share of male applicants (β) as specified in equation 29. Dependent variable: imputed log daily earnings; Default independent variables: the total number of applicants, worker age fully interacted with education attainment (measured by five categories), experience in years as well as its squared term, an indicator variable for the previous labor market status (non-employed, unemployed, employed), the contractual hours of the new job, formal job requirements (four categories), and year dummies. Five categories for the number of male appl. (only females, low male share, medium male share (reference), high male share, only males); Estimates for full-time male workers only. Source: JVS, IEB.

Under the second theoretical hypothesis, jobs with a high share of male applicants are different from those with a lower share of male applicants. Both men and women (and not only men) should earn more than men and women with comparable observable characteristics. To be able to test this further, we construct a categorical variable instead of the continuous share of male applicants. We distinguish five categories: one if a vacancy has only female applications, five if there are only male applications, and two, three, and four in between.[24] Two refers to a low, three to a medium, and four to a high share of male applicants. We choose a medium share of male applicants as the reference group, which allows us to compare the coefficients across genders.

Figure 6 shows the estimated coefficients for the categorical variable for hired men only.[25] Conditional on the large set of control variables, men who match at a job with

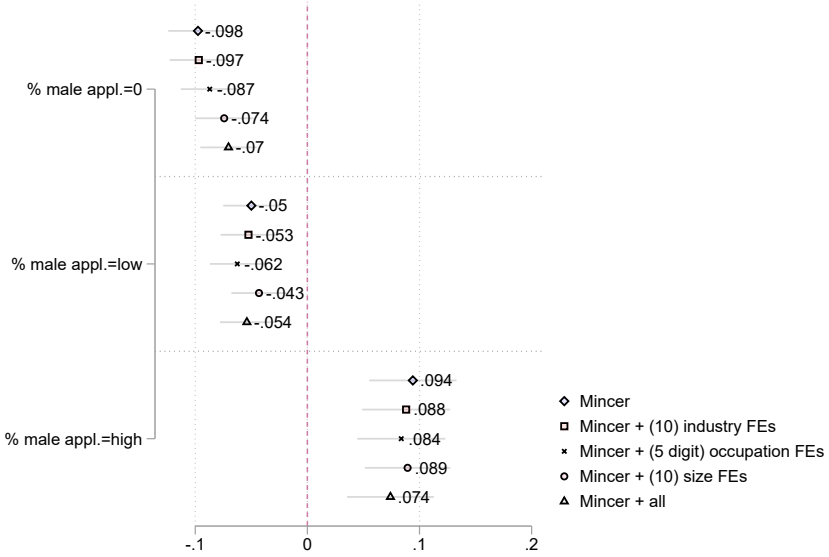
[23] In further robustness checks, we restricted our sample to only female-dominated jobs and used an alternative occupational classification. The pattern that the residual gender earning gap drops substantially when adding the share of male applicants holds in all specifications. Results are available on request.

[24] Figure A.1 in the Appendix shows the categories. We divide the inner part of the distribution into three parts. In the first part the mean of male applicants is 21%, in the second it is 48%, and in the third it is 80%)

[25] We again focus on full-time workers. Figure B.11 and B.12 show that all our findings are qualitatively

a high share of male applicants earn 5.9 to 8.5 percentage points higher earnings as compared to those who match with a medium share.

Figure 7: Coefficients for categories of share of male applicants, **female** hires



Note: The Figure shows the coefficients for the share of male applicants (β) as specified in equation 29. Dependent variable: imputed log daily earnings; Default independent variables: the total number of applicants, worker age fully interacted with education attainment (measured by five categories), experience in years as well as its squared term, an indicator variable for the previous labor market status (non-employed, unemployed, employed), the contractual hours of the new job, formal job requirements (four categories), and year dummies. Five categories for the number of male appl. (only females, low male share, medium male share (reference), high male share, only males); Estimates for full-time female workers only. Sources: JVS, IEB.

Figure 7 shows the estimated coefficients for the categorical variable from regressions for hired females only. In line with the second hypothesis in our theoretical model, the coefficients are increasing in the share of male applicants. We observe large effects in all our regressions. For instance, in our benchmark specification, a female recruitment with a low share of male applicants on average results in 4.3 to 6.2 percentage points lower earnings as compared to a female recruitment where there was a medium share of male applicants. On the other hand, depending on the exact specification, a female recruitment with high share of male applicants on average results in 7.4 to 9.4 percentage points higher earnings as compared to a female recruitment where there was a medium share of male applicants. These patterns in the data provide further evidence for the hypothesis that jobs with a larger share of male applicants are different from those with a low share of male applicants. Employers appear to provide compensating differentials for the higher degree of employer-sided flexibility requirements.

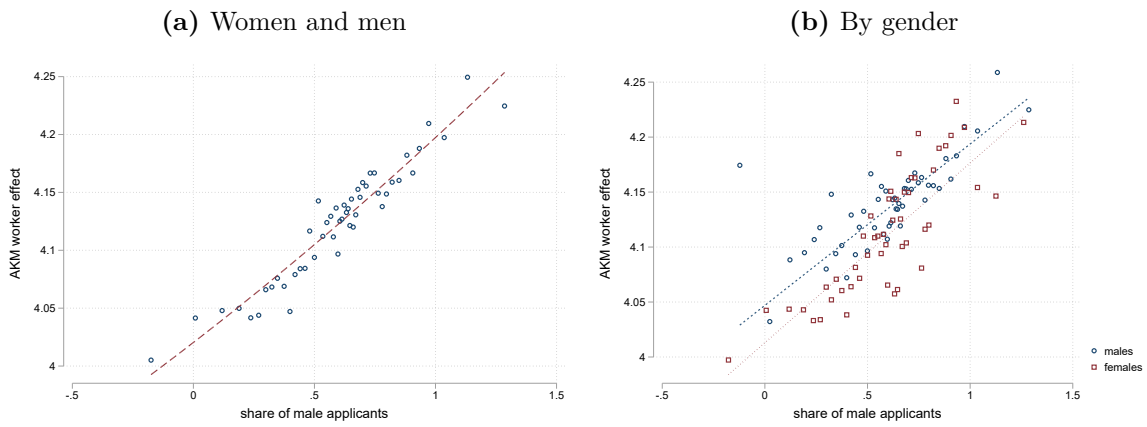
unaltered once we include part-time workers.

4.4 Evidence for Nonlinear Jobs on the Person Level

In our final step, we analyze the interaction of the share of male applicants with characteristics of the person who matched. More precisely, we analyze the connection between the share of male applicants and the worker fixed effect from the two-way fixed effects regressions.[26] In addition, we check how being a mother affects the residual gender earnings gap and how this interacts with the share of male applicants.

Figure 8a shows the residualized share of male applicants and the residualized worker fixed effect of the hired workers from the AKM two-way fixed effects regression. A larger share of male applicants is associated with a larger AKM worker fixed effect. Through the lens of our model, workers that are willing/able to provide a high input and who are hired in a nonlinear job will produce more than hires in linear jobs. A certain fraction of this higher production will be passed on in form of higher wages (under Nash bargaining or any other wage formation where wages depends on produced output) and show up as larger worker-specific wage premiums. Figure 8b shows the relation between the AKM person effects of hired workers and the share of male applicants separately for hired males and females. A higher share of male applicants is associated with higher AKM worker fixed effects both for men and women. Thus, higher flexibility requirements at certain jobs are associated with higher worker fixed effects for both genders.

Figure 8: AKM Person effects and the gender distribution of the application pool



Note: Figures show binscatters with 50 bins and linear fit lines. To residualize the x-variable and y-variables, we regress each variable on the controls, generate the residuals, and add the sample mean of each variable back to its residuals. We then group the x-axis variable into equal-sized bins, compute the mean of the x-axis and y-axis variables within each bin, and create a scatterplot of these data points. Control variables: Industry categories (Nace Rev 2), firm size categories, occupation categories (5 digits); Source: JVS, IEB; Full-time jobs only; Sources: JVS, IEB.

Obviously, the connection between the share of male applicants and AKM worker fixed effects cannot be interpreted causally. The worker fixed effects capture unobserved worker

[26] As our data is a cross-section of hires, we cannot estimate person fixed effects directly. However, we can use the worker fixed effects that were estimated on the universe of German administrative data and link it to our cross section.

heterogeneity and differences in the worker fixed effect in Figure 8a may therefore (partly) be driven by ex-ante worker ability. Against this background, the positive correlation is in line with the result by Lamadon et al. (2022) who show that compensating differentials are larger for high-ability workers and smaller for low-ability workers. As shown before, applicant pools with a larger share of male applicants can be found at firms with higher firm fixed effects and are thus associated with higher pay (i.e., a compensating differential for higher employer-sided flexibility requirements).[27]

The quantitative difference between these two estimated curves in Figure 8b is relatively small (i.e. an order of magnitude smaller than the gender earnings gap when not controlling for the share of male applicants). Thus, the connection between the share of male applicants and AKM worker fixed effects is on average very similar for men and women.

So far, our empirical results suggest that high-flexibility jobs (i.e., those with a larger share of male applicants) are associated with a disamenity and thereby pay compensating differentials. At the person level, we can also test the hypothesis whether these patterns are driven by different production functions. Assume that a person that is unable to provide high-flexibility matches at a firm with nonlinear production function. In this case, our model would predict low output at this job and a particularly large wage discount for the matched person. Although we do not have any information about the degree of flexibility that a person can provide, we consider motherhood as a suitable proxy. Mothers in Germany still bear a larger fraction of childcare than fathers and thereby tend to be less flexible.

Therefore, we use the established proxy for being a mother in the administrative data (Mueller and Strauch, 2017). Based on this proxy, we estimate the residual gender earnings gap relative to males for female mothers and for childless females. Column (1) of Table 3 shows that the residual gender earnings gap is about 6 percentage points larger (-20 vs. -14 percent) for mothers compared to childless women. When we add our proxy for firm-sided flexibility requirements (i.e., the share of male applicants) to the regression in column (2), the gap between female mothers and childless females remains similarly large (-12 vs. -6 percent). Overall, this exercise shows that mothers face a larger hiring earnings discount in the labor market than women without children.

Finally, we interact the share of male applicants with dummies for mothers and childless females. Figure 9 shows the (predicted) wage discount for mothers and childless females (relative to men) split up according to the shares of male applicants at the respective jobs (from 0.1 to 0.9, as share of 0 and 1 have to be excluded as only one gender matches at those jobs).[28] When mothers match at a job with a 90 percent share of male

[27] Note that worker ex-ante heterogeneity is absent in our model and therefore the model is silent on this issue.

[28] We include an interaction term of the share of male applicants as a continuous variable with a dummy variable which has distinct values for mothers and childless females relative to males in our regression.

	(1) log earnings	(2) log earnings
mother	-0.2024***	-0.1231***
(male=reference)	(0.0142)	(0.0158)
childless female	-0.1389***	-0.0637***
(male=reference)	(0.0071)	(0.0093)
Observations	12,945	11,631
Adjusted R^2	0.6038	0.6126

Table 3: Estimates for full-time workers only; Standard errors in parentheses; Controls: total number of applicants, a set of worker age dummies fully interacted with education dummies, experience in years as well as its squared term, a dummy for the previous labor market status (non-employed, unemployed, employed), the hours of the new contract, dummies for formal job requirements, year dummies, industry categories, occupation categories, and establishment size deciles; Columns (2) additionally control for the share of male applicants; * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

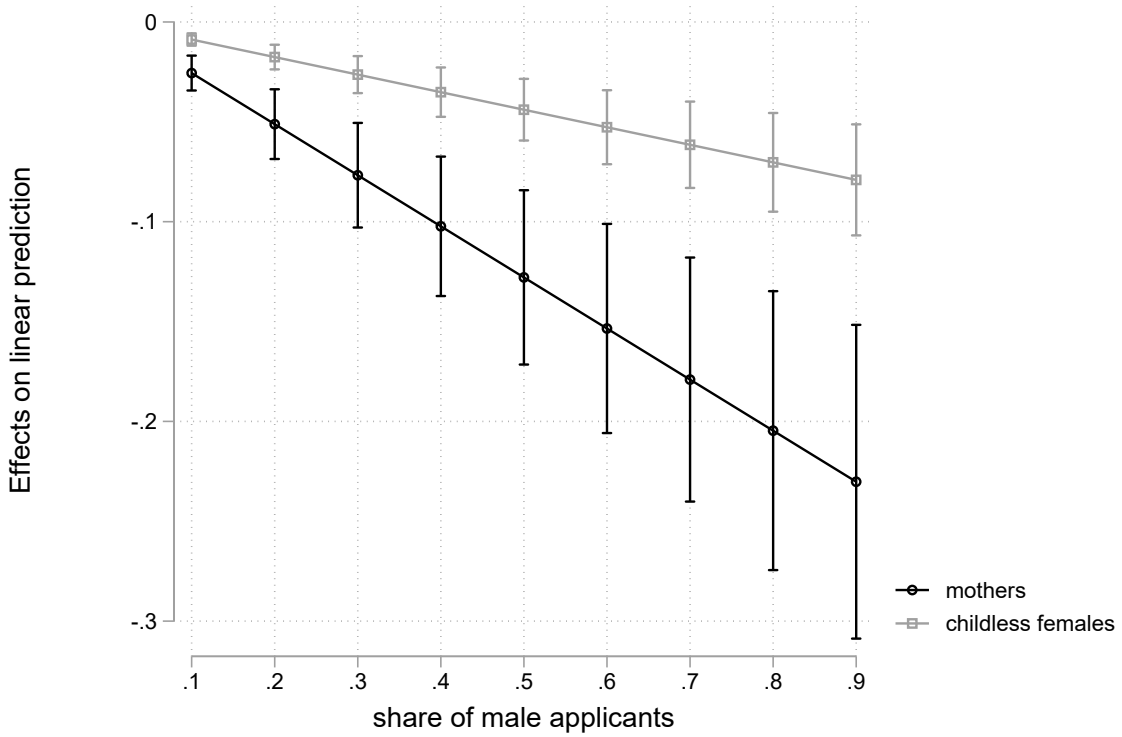
applicants, they face a more than 20 percent residual gender wage gap relative to men, while this number is very small at low shares of male applicants. Note that the weighted average of these estimates corresponds to the point estimates in column (2) of Table 3. This is in line with our interpretation that jobs with high share of male applicants tend to be nonlinear jobs. As mothers are unable to provide the employer-sided (desired) flexibility, they produce less and thereby face a large wage discount. It is also striking that the wage discount differential between mothers and childless females increases with the share of male applicants. While the differences in the point estimates are economically very small for matches with small shares of male applicants, it is more than 15 percentage points for matches with 90 percent male applicants.[29]

It is also worthwhile discussing that the wage discount for childless females also increases in the share of male applicants. First, the economic differences between the highest and lowest share of male applicants are relatively small. Second, having children is an incomplete proxy for the ability and willingness of women to provide flexibility. It is for example well known that women also bear a larger burden of care responsibilities that are not related to children (e.g. elderly care). Thus, even women without children may on average be less flexible than men. This may be one of the drivers for the small differences across the share of male applicant categories.

Based on this regression, we then calculate marginal effects over a grid of values of the share of male applicants.

[29]The confidence bands are larger for larger share of male applicants as the number of observations is small. This is due to two reasons. First, due to the matching of the IAB Job Vacancy and administrative data, the sample size is reduced. Second, by definition at jobs with a larger share of male applicants, the absolute number of females and even more so for mothers is small.

Figure 9: Mothers and Women without Children



Note: Figure shows the earnings gap (marginal effects) for mothers and childless females compared to males as a reference group at various levels of the share of male applicants. Controls: the share of male applicants interacted with a dummy for mothers and childless females (male=reference), the total number of applicants, a set of worker age dummies fully interacted with education dummies, experience in years as well as its squared term, a dummy for the previous labor market status (non-employed, unemployed, employed), the hours of the new contract, dummies for formal job requirements, year dummies, industry categories, occupation categories, and establishment size deciles; Full-time jobs only; Sources: JVS, IEB.

5 Conclusion

This paper shows that gender-specific application behavior is key for understanding hiring wage differences. Even within industries, firm size categories, and occupations, women are 10 percentage points more likely to apply at the lowest-wage firms than men. Our theoretical labor market flow model rationalizes this behavior based on different production functions at different jobs, where the highest paying jobs are nonlinear in input, as defined by Goldin (2014).

We show that the share of applicants is a positive function of various measurable dimensions of employer-sided flexibility requirements. Therefore, we consider it as a suitable proxy for multidimensional flexibility requirements at the job-level. Once we include this proxy into standard Mincer regressions (beyond standard observable variables such as occupations, sectors, and worker characteristics), the residual gender-wage gap drops by 50-60 percent. This illustrates that the gender-specific application behavior is an important explanatory variable that is typically omitted in Mincer-type wage regressions, as it is not contained in standard datasets.

Our paper combines information from the IAB Job Vacancy Survey with administrative information on the last hire. This combination allows us to use a proxy whether women have children. We show that wage discounts are particularly large for women with children. This wage discount increases in our proxy for employer-sided flexibility. Again, this is in line with the nonlinear jobs hypothesis. When women with children match at nonlinear jobs, they are less able to provide a high-degree of employer-sided flexibility and thereby face a large wage discount.

Our paper offers variable policy-relevant lessons. Policy interventions that allow women to get access to jobs with high-flexibility requirements (such as better access to childcare or incentives for different intra-family sharing of care responsibilities) will change their application behavior and thereby can reduce the gender-earnings gap. Furthermore, the Covid-19 pandemic has shown that a different organization of work is possible (e.g., more working from home arrangements). Only future research will show whether this new work environment will stick and whether it will boost women's possibilities to get better access to jobs with high-flexibility requirements.

References

Abowd, John M., Francis Kramarz, and David N. Margolis (1999) "High Wage Workers and High Wage Firms," *Econometrica*, Vol. 67, pp. 251–333.

Azmat, Ghazala and Rosa Ferrer (2017) "Gender Gaps in Performance: Evidence from Young Lawyers," *Journal of Political Economy*, Vol. 125, pp. 1306–1355.

Baechmann, Ann-Christin, Corinna Frodermann, Benjamin Lochner, Michael Oberfichtner, and Simon Trenkle (2021) "Identifying Couples in Administrative Data for the years 2001–2014," FDZ-Methodenreport 202103 (en), Institute for Employment Research, Nuremberg, Germany.

Barrero, Jose Maria, Nicholas Bloom, and Steven Davis (2020) "Why Working From Home Will Stick," Discussion Paper 174, Becker-Friedman-Institute, University of Chicago.

Bellmann, Lisa, Benjamin Lochner, Stefan Seth, and Stefanie Wolter (2020) "AKM Effects for German Labour Market Data," FDZ Methodenreport 01/2020, IAB Nuremberg.

Blau, Francine D. and Lawrence H. Kahn (2017) "The Gender Wage Gap: Extent, Trends, and Explanations," *Journal of Economic Literature*, Vol. 55, pp. 789–865.

Bolotnyy, Valentin and Natalia Emanuel (2019) "Why Do Women Earn Less Than Men? Evidence from Bus and Train Operators," *Harvard University, mimeo.*

Brown, Alessio, Christian Merkl, and Dennis Snower (2016) "An Incentive Theory of Matching," *Macroeconomic Dynamics*, Vol. 131, pp. 633–686.

Budig, Michelle J. and Melissa J. Hodges (2010) "Differences in Disadvantage: Variation in the Motherhood Penalty across White Women's Earnings Distribution," *American Sociological Review*, Vol. 75, pp. 705–728.

Burdett, Kenneth and Dale T. Mortensen (1998) "Wage Differentials, Employer Size, and Unemployment," *International Economic Review*, Vol. 39, pp. 257–273.

Card, David, Ana Rute Cardoso, and Patrick Kline (2016) "Bargaining, Sorting, and the Gender Wage Gap: Quantifying the Impact of Firms on the Relative Pay of Women," *Quarterly Journal of Economics*, Vol. 19, pp. 643–686.

Card, David, Jörg Heining, and Patrick Kline (2013) "Workplace Heterogeneity and the Rise of West German Wage Inequality," *The Quarterly Journal of Economics*, Vol. 128, pp. 967–1015.

Carrillo-Tudela, Carlos, Hermann Gartner, and Leo Kaas (2020) "Recruitment Policies, Job-Filling Rates and Matching Efficiency," Discussion Paper 13240, IZA.

Chugh, Sanjay K. and Christian Merkl (2016) "Efficiency and Labor Markt Dynamics in a Model of Labor Selection," *International Economic Review*, Vol. 57, pp. 1371–1404.

Cook, Cody, Rebecca Diamond, Jonathan Hall, John A. List, and Paul Oyer (2021) "The Gender Earnings Gap in the Gig Economy: Evidence from over a Million Rideshare Drivers," *Review of Economic Studies*, Vol. 88, pp. 2210–2238.

Cortés, Patricia, Jessica Pan, Laura Pilossoph, and Basit Zafar (2021) "Gender Differences in Job Search and the Earnings Gap: Evidence from Business Majors," Working Paper 28820, National Bureau of Economic Research.

Dustmann, Christian, Johannes Ludsteck, and Uta Schönberg (2009) "Revisiting the German Wage Structure," *The Quarterly Journal of Economics*, Vol. 124, pp. 843–881.

Faberman, R. Jason, Andreas I. Mueller, Ayşegül Şahin, and Giorgio Topa (2017) "Job Search and the Gender Wage Gap," *Society for Economic Dynamics, Meeting Papers*.

Fitzenberger, Bernd, Aderonke Osikominu, and Robert Völter (2006) "Imputation Rules to Improve the Education Variable in the IAB Employment Subsample," *Schmollers Jahrbuch : Journal of Applied Social Science Studies / Zeitschrift für Wirtschafts- und Sozialwissenschaften*, Vol. 126, pp. 405–436.

Fluchtmann, Jonas, Anita M. Glenny, Nikolaj Harmon, and Jonas Maibom (2020) "The Gender Application Gap: Do men and women apply for the same jobs?" , unpublished manuscript.

Fuchs, Michaela, Anja Rossen, Antje Weyh, and Gabriele Wydra-Somaggio (2019) "Unterschiede in der Lohnlücke erklären sich vor allem durch die Betriebslandschaft vor Ort," *IAB Kurzbericht*, Vol. 10.

Goldin, Claudia (2014) "A Grand Gender Convergence: Its Last Chapter," *American Economic Review*, Vol. 104, pp. 1091–1119.

Gomes, Pedro and Zoe Kuhn (2019) "You're the One that I Want! Public Employment and Women's Labor Market Outcomes," IZA Discussion Paper 12702.

Hochmuth, Brigitte, Britta Kohlbrecher, Christian Merkl, and Hermann Gartner (2021) "Hartz IV and the Decline of German Unemployment: A Macroeconomic Evaluation," *Journal of Economic Dynamics and Control*, p. 104114.

Illing, Hannah, Johannes F Schmieder, and Simon Trenkle (2021) "The Gender Gap in Earnings Losses after Job Displacement," Working Paper 29251, National Bureau of Economic Research.

Lamadon, Thibaut, Magne Mogstad, and Bradley Setzler (2022) "Imperfect Competition, Compensating Differentials, and Rent Sharing in the US Labor Market," *American Economic Review*, Vol. 112, pp. 169–212.

Le Barbanchon, Thomas, Roland Rathelot, and Alexandra Roulet (2020) "Gender Differences in Job Search: Trading off Commute against Wage," *The Quarterly Journal of Economics*, Vol. 136, pp. 381–426.

Lochner, Benjamin (2019) "A Simple Algorithm to Link "Last Hires" from the Job Vacancy Survey to Administrative Records," *FDZ-Methodenreport*, Vol. 6.

Lochner, Benjamin, Christian Merkl, Heiko Stüber, and Nicole Gürtzgen (2021) "Recruiting Intensity and Hiring Practices: Cross-sectional and Time-series Evidence," *Labour Economics*, Vol. 68, p. 101939.

Lochner, Benjamin, Stefan Seth, and Stefanie Wolter (2020) "Decomposing the Large Firm Wage Premium in Germany," *Economics Letters*, Vol. 194, p. 109368.

Moczall, Andreas, Anne Müller, Martina Rebien, and Kurt Vogler-Ludwig (2015) "The IAB Job Vacancy Survey. Establishment Survey on Job Vacancies and Recruitment Processes. Waves 2000 to 2013 and Subsequent Quarters from 2006," , FDZ Datenreport 04/2015, Institut für Arbeitsmarkt- und Berufsforschung (IAB), Nürnberg.

Mueller, Dana and Katharina Strauch (2017) "Identifying mothers in administrative data," FDZ-Methodenreport 201713 (en), Institute for Employment Research, Nuremberg, Germany.

Postel-Vinay, Fabien and Jean-Marc Robin (2002) "Equilibrium Wage Dispersion with Worker and Employer Heterogeneity," *Econometrica*, Vol. 70, pp. 2295–2350.

Sorkin, Isaac (2018) "Ranking Firms Using Revealed Preference," *The Quarterly Journal of Economics*, Vol. 133, pp. 1331–1393.

Taber, Christopher and Rune Vejlin (2020) "Estimation of a Roy/Search/Compensating Differential Model of the Labor Market," *Econometrica*, Vol. 88, pp. 1031–1069.

Wiswall, Matthew and Basit Zafar (2017) "Preference for the Workplace, Investment in Human Capital, and Gender," *The Quarterly Journal of Economics*, Vol. 133, pp. 457–507.

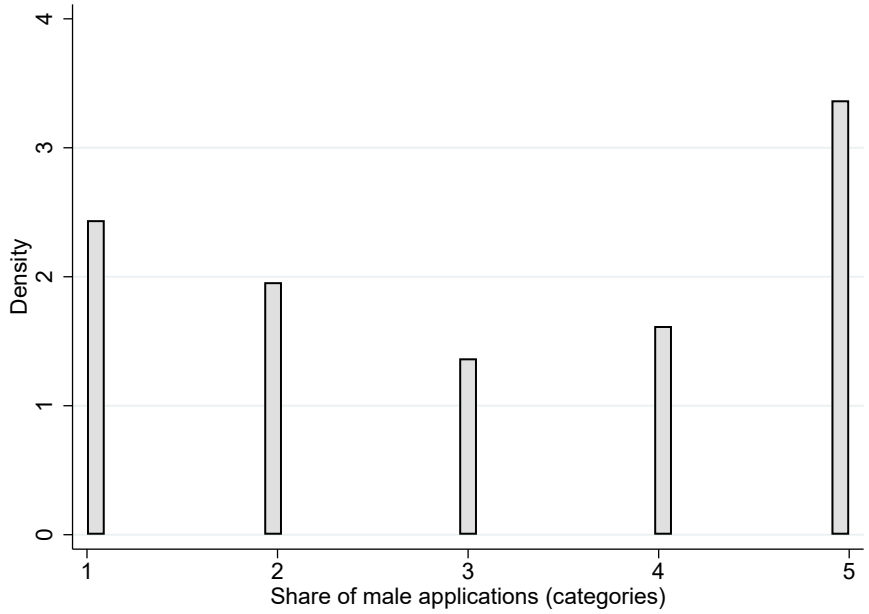
Xiao, Pengpeng (2021) "Wage and Employment Discrimination by Gender in Labor Market Equilibrium," *VATT Working Papers, 144*.

A Data Appendix

Table A.1: Share of male/female hires and applicants across industries

		Share of hires		Share of applicants	
NACE Rev. 2	total hired	males (%)	females (%)	males (%)	females (%)
A - Agriculture, forestry and fishing	941	67.59	32.41	66.13	33.87
B - Mining and quarrying					
C - Manufacturing	4,952	72.70	27.30	70.72	29.28
D - Electricity, gas, steam and air cond. supply	1,579	68.84	31.16	69.65	30.35
E - Water supply, sewerage, waste management					
F - Construction	826	87.89	12.11	85.00	15.00
G - Wholesale and retails trade, repair of motor	1,613	68.20	31.80	65.36	34.64
H - Transportation and storage					
I - Accomomodation and food service activities	664	41.27	58.73	39.03	60.97
J - Information and communication	4,470	52.24	47.76	52.99	47.01
K - Financial and insurance activities					
L - Real estate activities					
M - Professional, scientific and technical activities					
N - Administrative and support service activities					
O - Public administration/defence, comp. social sec.	1,860	34.68	65.32	37.17	62.83
P - Education	4,789	26.12	73.88	27.87	72.13
Q - Human health and social work activities					
R - Arts, entertainment and recreation					
S - Other service activities					
T - Activities of households as employers					
U - Activities of extraterritorial organisations/bodies					
Total	21,694	53.72	46.28	57.10	42.90

Figure A.1: Share of male applicants: categories



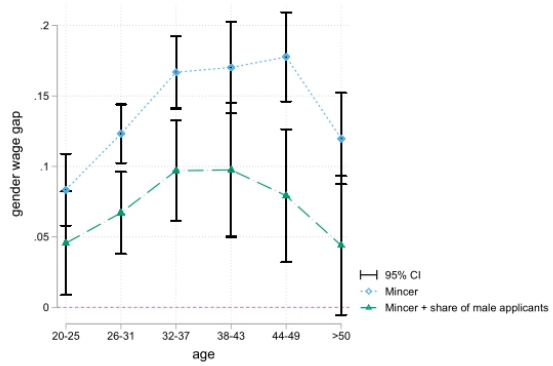
Source: JVS, IEB;

B Additional Empirical Results

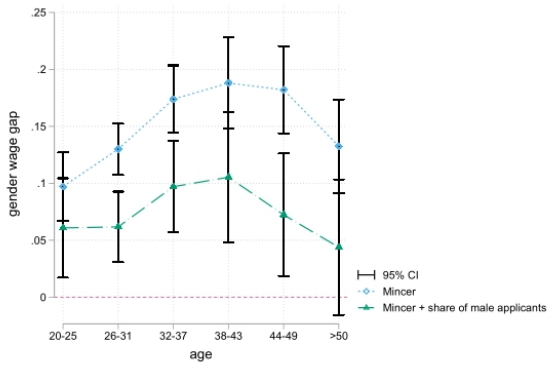
B.1 Age Cohorts

Figure B.1: GWG estimates by 5 year cohorts

(a) Full- and part-time workers



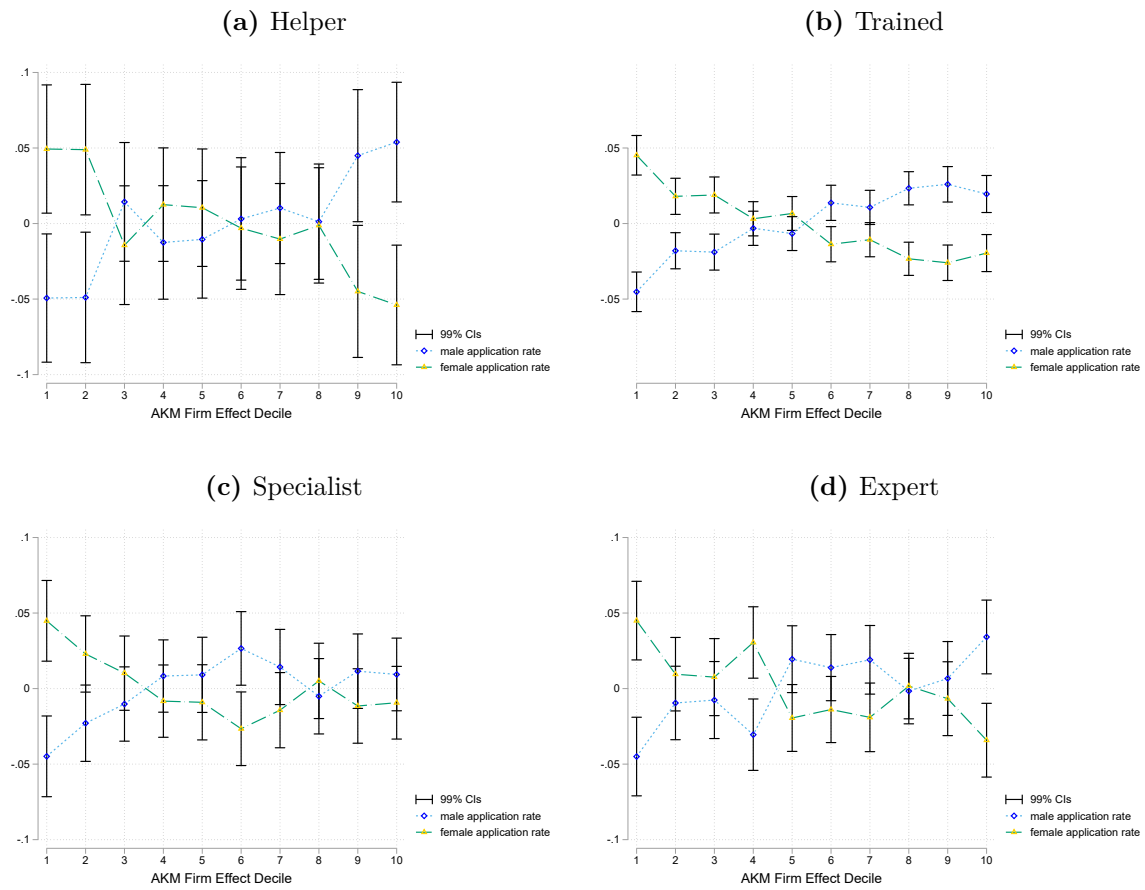
(b) Full-time workers only



Note: Figure shows the estimates for the gender gap in the hiring wage by age groups as laid out on the x-axis. Dependent variable: imputed log daily earnings. Default independent variables: gender dummy, the total number of applicants, worker age fully interacted with education attainment (measured by five categories), experience in years as well as its squared term, an indicator variable for the previous labor market status (non-employed, unemployed, employed), the contractual hours of the new job, formal job requirements (four categories), and year dummies. Source: JVS, IEB;

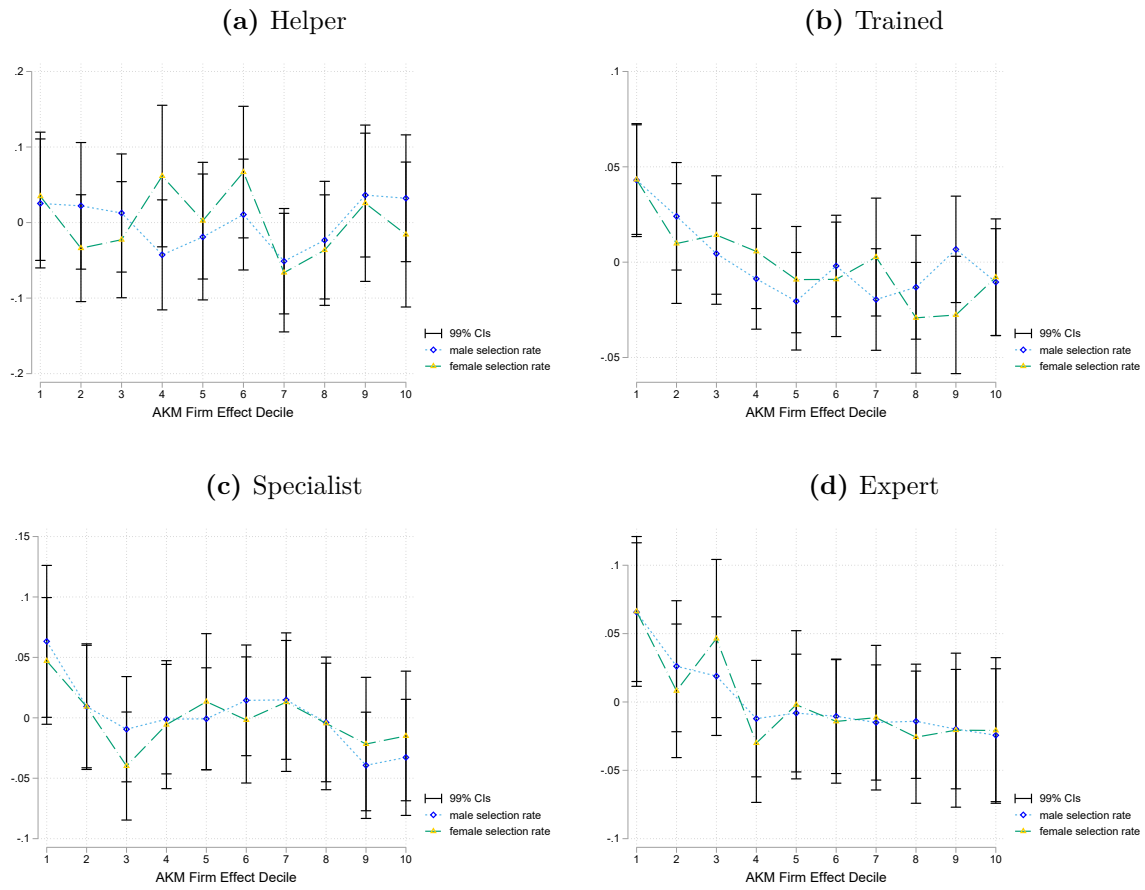
B.2 Application and Selection Behavior within Task Complexities

Figure B.2: Residualized share of male applicants over grid of AKM firm effect deciles by task complexity



Note: Full-time jobs only; Variables are defined as follows: a) number of female (male) appl./ number of all applications; b) 1/ number of all applicants if male/female hired; Control variables: Industry categories (Nace Rev 2), firm size categories, occupation categories (5 digits); Source: JVS, IEB;

Figure B.3: Residualized selection rates over grid of AKM firm effect deciles by job level



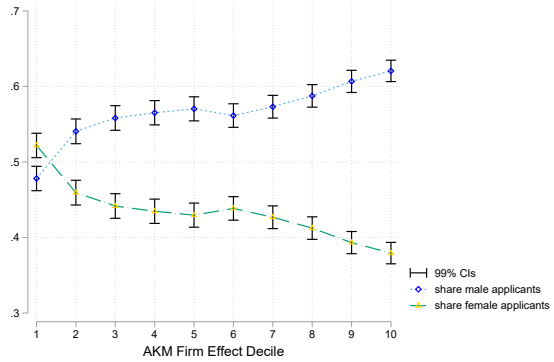
Note: Full-time jobs only; Variables are defined as follows: a) number of female (male) appl./ number of all applications; b) 1/ number of all applicants if male/female hired; Control variables: Industry categories (Nace Rev 2), firm size categories, occupation categories (5 digits); Source: JVS, IEB;

B.3 Application and Selection Behavior with Alternative Firm Fixed Effects

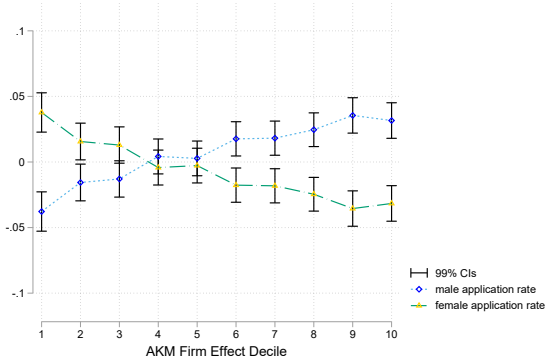
Figures B.4 and B.5 show the patterns in the data with differently estimated firm-fixed effects. In this case, the firm-fixed effect is estimated separately for men and women (i.e., each firm has two wage premia: one for men and one for women).

Figure B.4: Application and selection rate by gender and AKM firm effect deciles (estimated from males only)

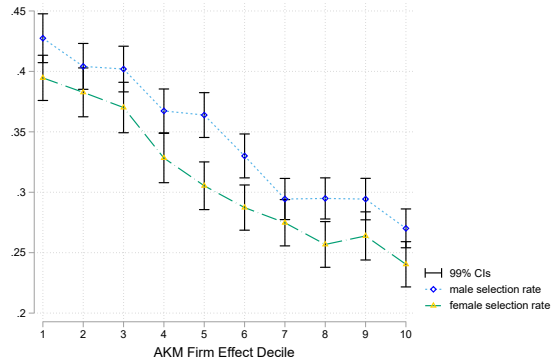
(a) Share of female/male applications



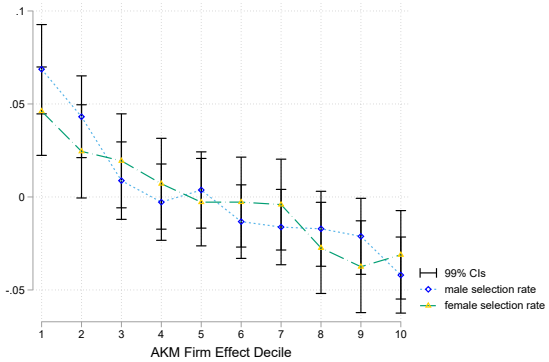
(b) Residualized share of female/male applications



(c) Female/male selection rate



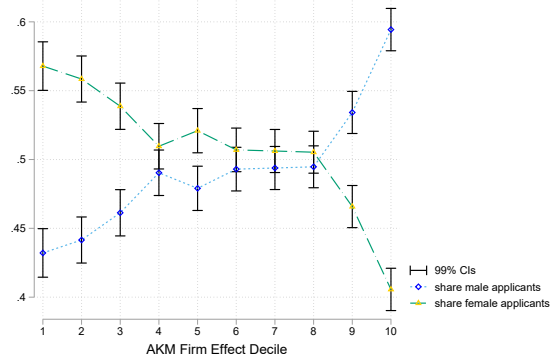
(d) Residualized female/male selection rate



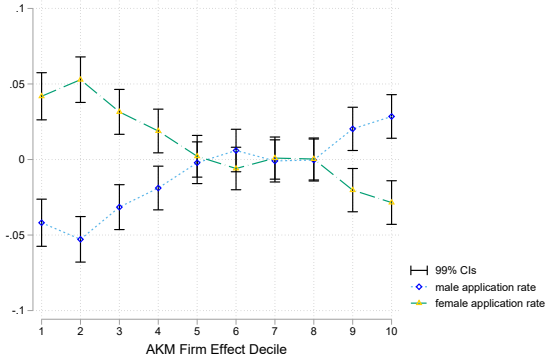
Note: Full-time jobs only. Firm effects estimates for males only. Variables are defined as follows: a) number of female (male) appl./ number of all applications; b) 1/ number of all applicants if male/female hired; Control variables: Industry categories (Nace Rev 2), firm size categories, occupation categories (5 digits); Source: JVS, IEB;

Figure B.5: Application and selection rate by gender and AKM firm effect deciles (estimated from females only)

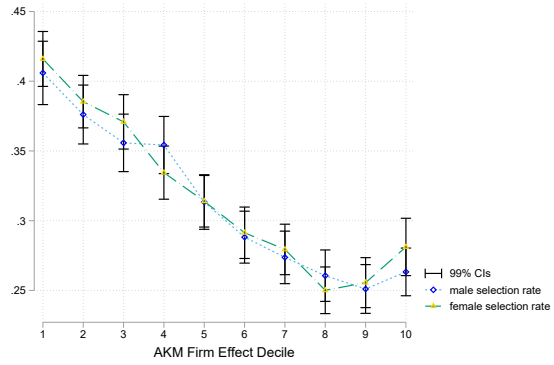
(a) Share of female/male applications



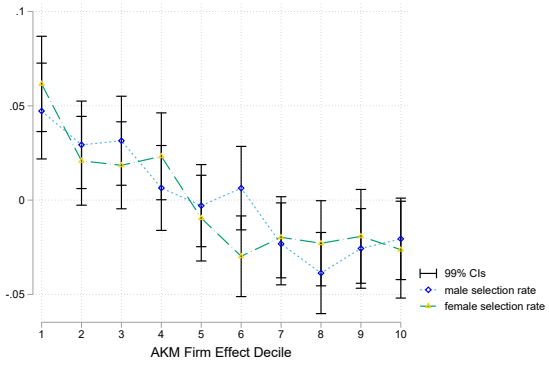
(b) Residualized share of female/male applications



(c) Female/male selection rate



(d) Residualized female/male selection rate



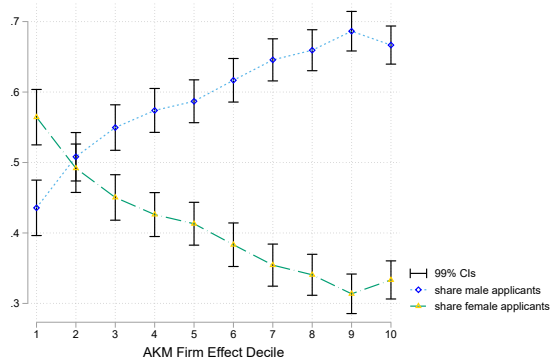
Note: Full-time jobs only. Firm effects estimates for females only. Variables are defined as follows: a) number of female (male) appl./ number of all applications; b) 1/ number of all applicants if male/female hired; Control variables: Industry categories (Nace Rev 2), firm size categories, occupation categories (5 digits); Source: JVS, IEB;

B.4 Application and Selection Behavior and Bargaining

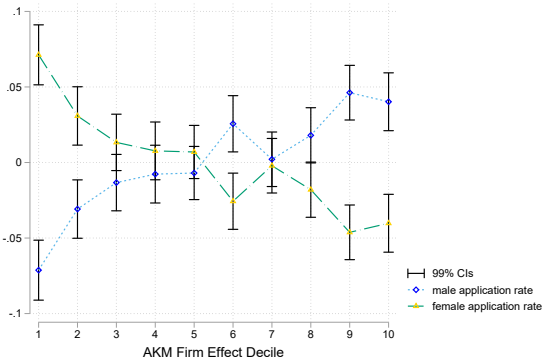
Figures B.6 and B.7 show the application and selection behavior across AKM firm effect deciles, separately for firms that are inside a collective or firm-level bargaining agreement and those that are not, respectively. Although the application rates differ somewhat in the raw data, once we control for our full set of controls, the quantitative results are very similar to our baseline sample.

Figure B.6: Application and selection rate by gender and AKM firm effect deciles, with collective bargaining

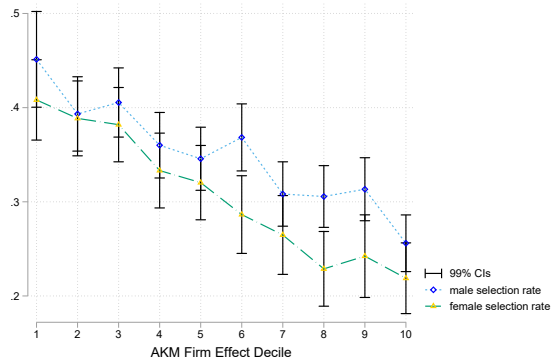
(a) Share of female/male applications



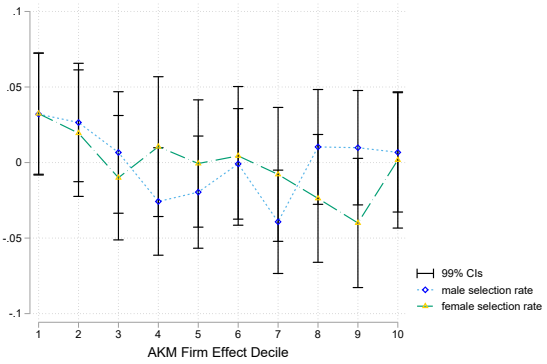
(b) Residualized share of female/male applications



(c) Female/male selection rate



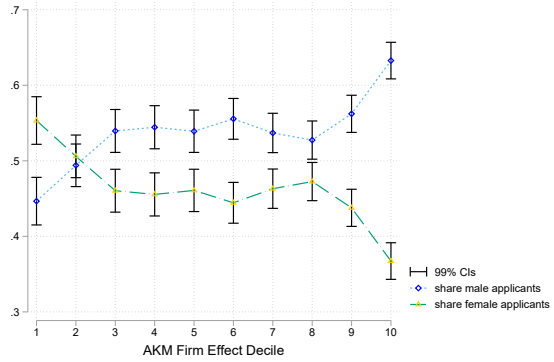
(d) Residualized female/male selection rate



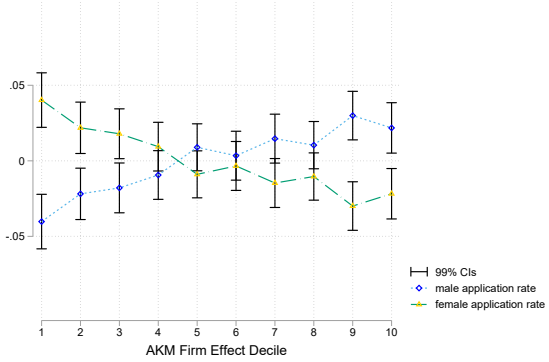
Note: Full-time jobs with collective bargaining only. Firm effects estimates for females only. Variables are defined as follows: a) number of female (male) appl./ number of all applications; b) 1/ number of all applicants if male/female hired; Control variables: Industry categories (Nace Rev 2), firm size categories, occupation categories (5 digits); Source: JVS, IEB;

Figure B.7: Application and selection rate by gender and AKM firm effect deciles, without collective bargaining

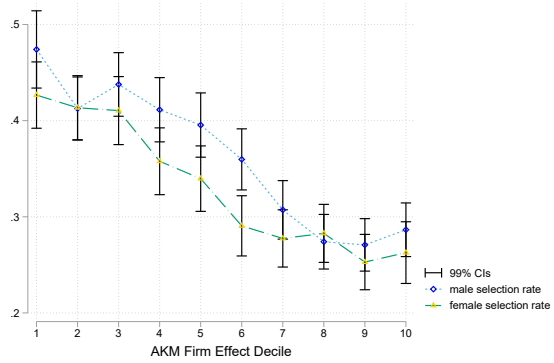
(a) Share of female/male applications



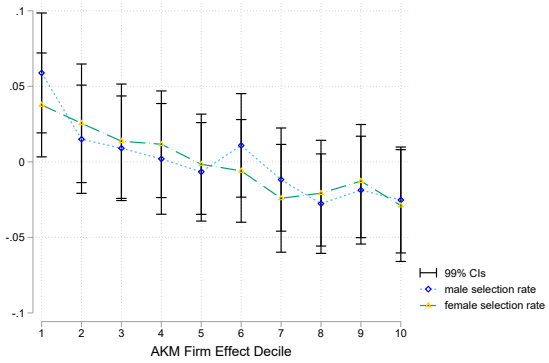
(b) Residualized share of female/male applications



(c) Female/male selection rate



(d) Residualized female/male selection rate



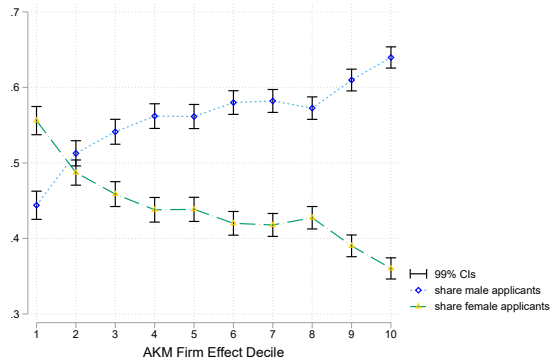
Note: Full-time jobs without collective bargaining only. Firm effects estimates for females only. Variables are defined as follows: a) number of female (male) appl./ number of all applications; b) 1/ number of all applicants if male/female hired; Control variables: Industry categories (Nace Rev 2), firm size categories, occupation categories (5 digits); Source: JVS, IEB;

C Alternative Sample Restriction

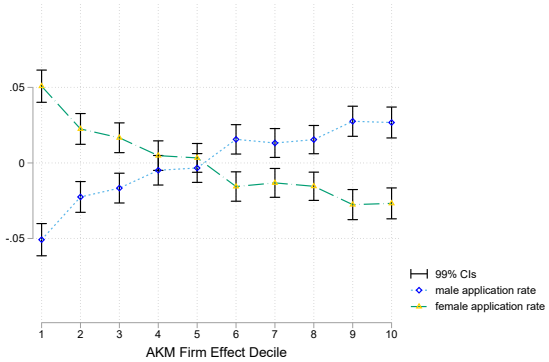
This Appendix replicates all main results, without imposing the full-time restriction (i.e., only workers with more than 25 hours working time). All our key insights are unaffected by the chosen sample restrictions, although the quantitative numbers slightly differ.

Figure B.8: Application and selection rate by gender and AKM firm effect deciles

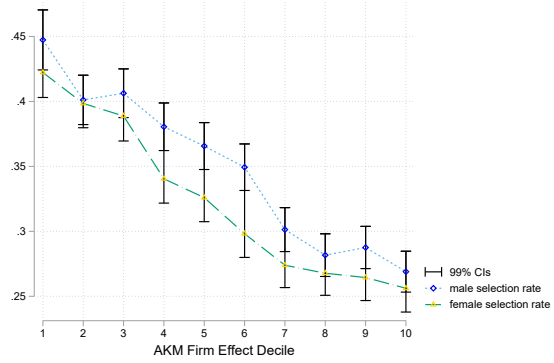
(a) Share of female/male applications



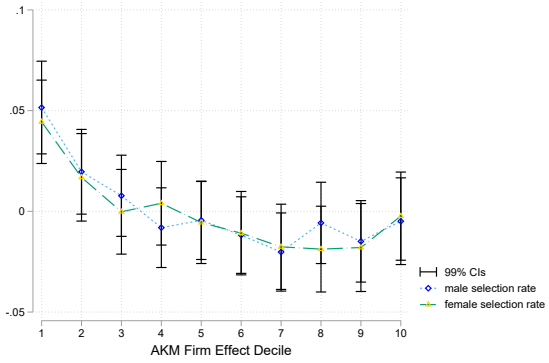
(b) Resiudalized share of female/male applications



(c) Female/male selection rate

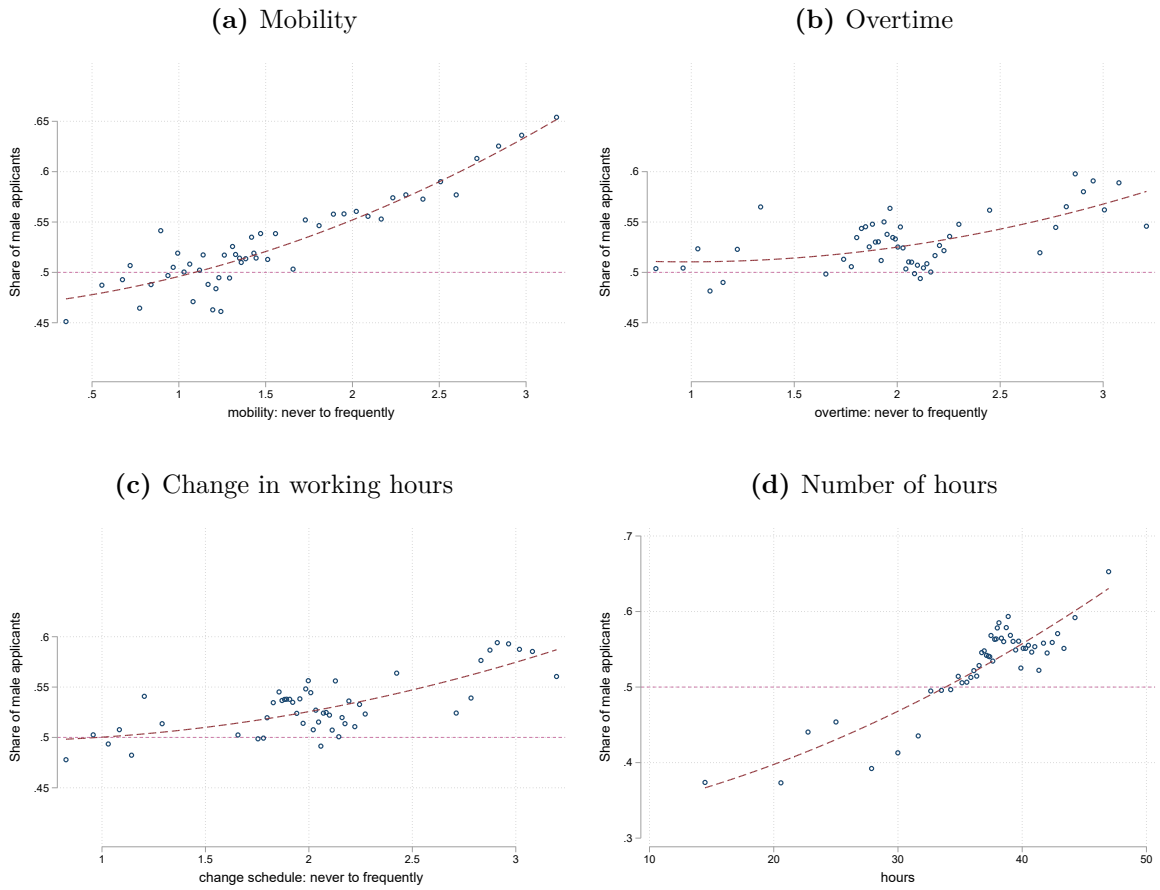


(d) Resiudalized female/male selection rate



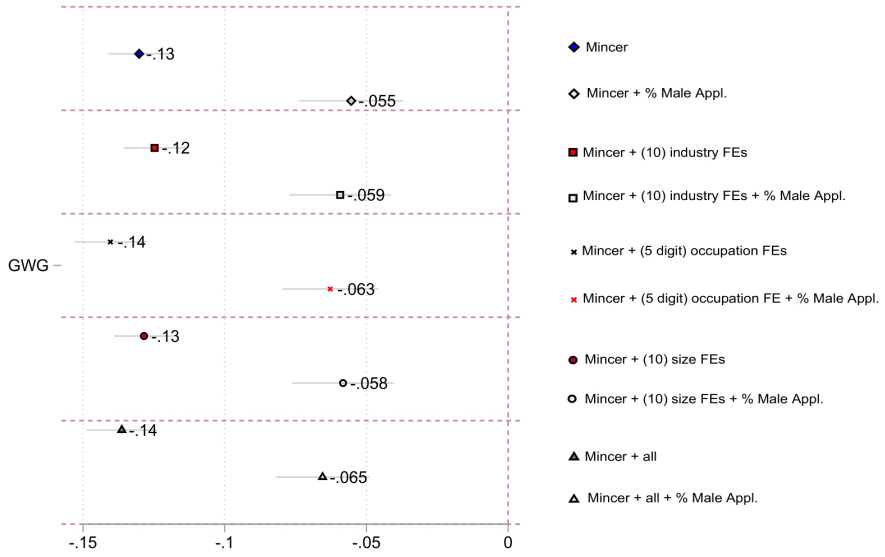
Note: Full-time and part-time jobs. Variables are defined as follows: a) number of female (male) appl./ number of all applications; b) 1/ number of all applicants if male/female hired; Control variables: Industry categories (Nace Rev 2), firm size categories, occupation categories (5 digits); Source: JVS, IEB;

Figure B.9: The share of male applicants and flexibility requirements



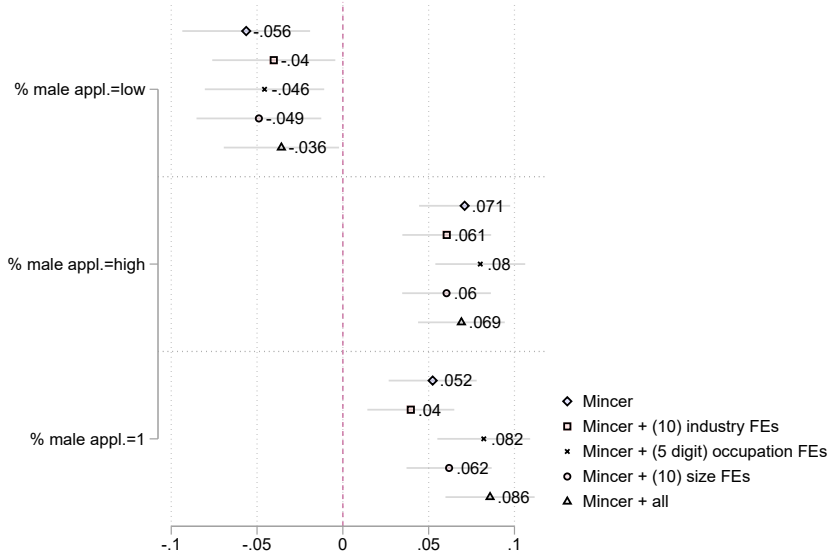
Note: Figures show binscatters with 50 bins and quadratic fit lines. To residualize the x-variable and y-variables, we regress each variable on the controls, generate the residuals, and add the sample mean of each variable back to its residuals. We then group the x-axis variable into equal-sized bins, compute the mean of the x-axis and y-axis variables within each bin, and create a scatterplot of these data points. Control variables: Industry categories (Nace Rev 2), firm size categories, occupation categories (5 digits); Source: JVS, IEB; Full-time and part-time jobs; Sources: JVS, IEB.

Figure B.10: The gender hiring wage gap



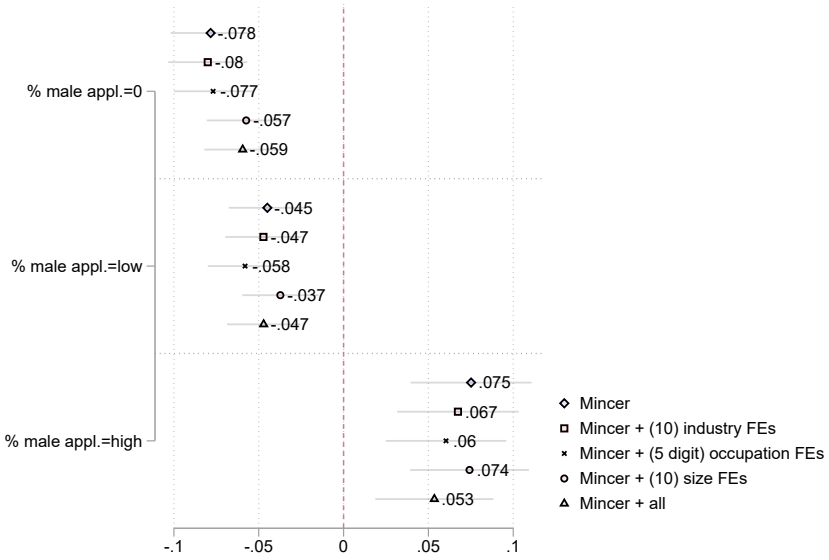
Note: The Figure shows the estimates for the gender gap (α) in the hiring wage as specified in equation 29. Dependent variable: imputed log daily earnings. Default independent variables: gender dummy, the total number of applicants, worker age fully interacted with education attainment (measured by five categories), experience in years as well as its squared term, an indicator variable for the previous labor market status (non-employed, unemployed, employed), the contractual hours of the new job, formal job requirements (four categories), and year dummies. Estimates for full-time and part-time workers. Source: JVS, IEB;

Figure B.11: Coefficients for categories of share of male applicants, **male** hires



Note: The Figure shows the coefficients for the share of male applicants (β) as specified in equation 29. Dependent variable: imputed log daily earnings; Default independent variables: the total number of applicants, worker age fully interacted with education attainment (measured by five categories), experience in years as well as its squared term, an indicator variable for the previous labor market status (non-employed, unemployed, employed), the contractual hours of the new job, formal job requirements (four categories), and year dummies. Five categories for the number of male appl. (only females, low male share, medium male share (reference), high male share, only males); Estimates for full-time and part-time male workers. Source: JVS, IEB.

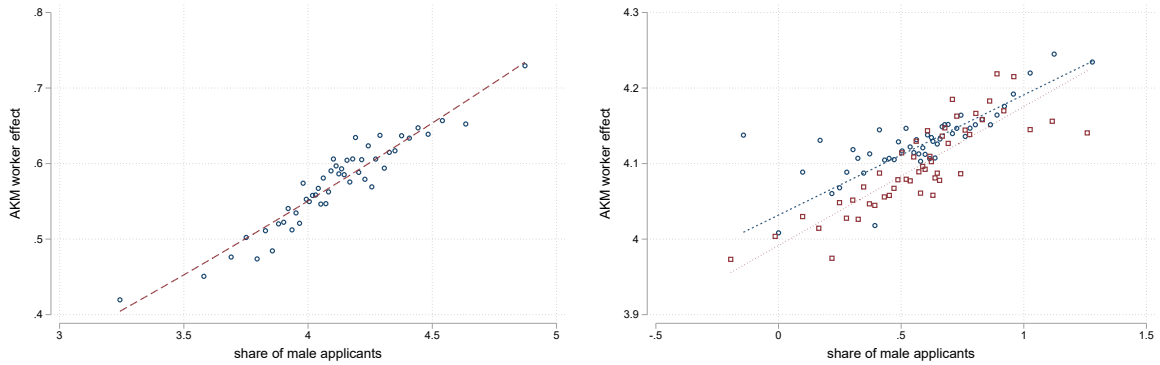
Figure B.12: Coefficients for categories of share of male applicants, **female** hires



Note: The Figure shows the coefficients for the share of male applicants (β) as specified in equation 29. Dependent variable: imputed log daily earnings; Default independent variables: the total number of applicants, worker age fully interacted with education attainment (measured by five categories), experience in years as well as its squared term, an indicator variable for the previous labor market status (non-employed, unemployed, employed), the contractual hours of the new job, formal job requirements (four categories), and year dummies. Five categories for the number of male appl. (only females, low male share, medium male share (reference), high male share, only males); Estimates for full-time and part-time female workers. Sources: JVS, IEB.

Figure B.13: AKM Person effects and the gender distribution of the application pool

(a) Women and men **(b)** By gender

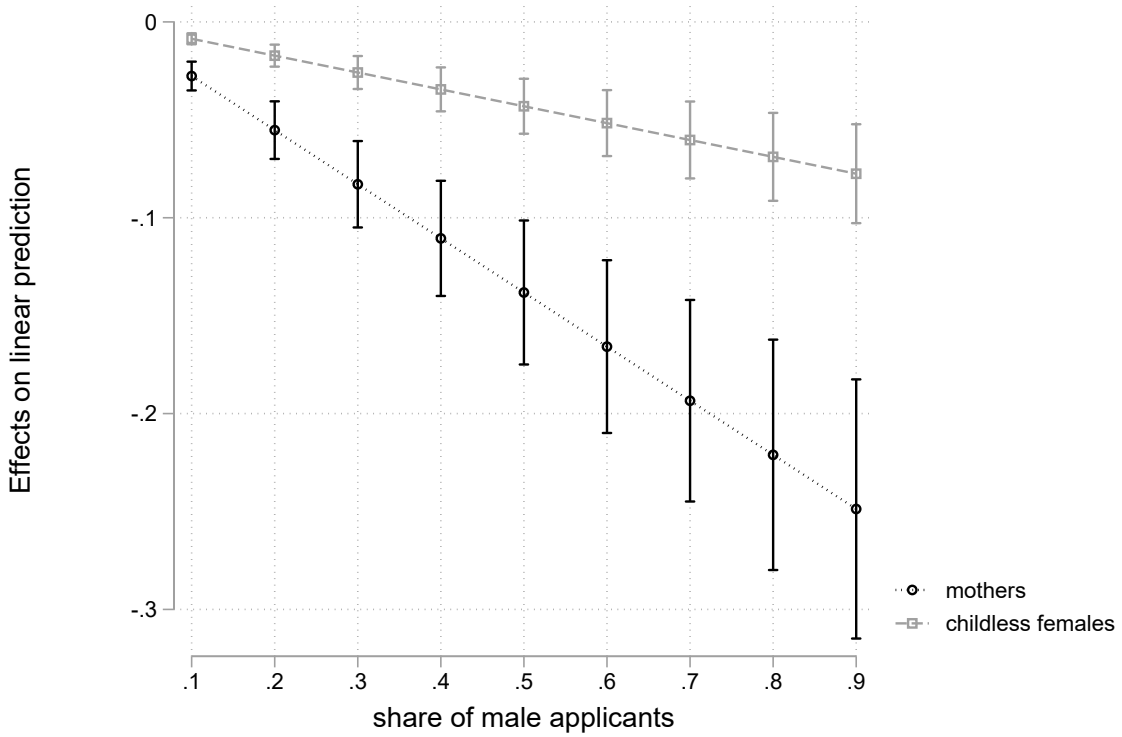


Note: Figures show binscatters with 50 bins and linear fit lines. To residualize the x-variable and y-variables, we regress each variable on the controls, generate the residuals, and add the sample mean of each variable back to its residuals. We then group the x-axis variable into equal-sized bins, compute the mean of the x-axis and y-axis variables within each bin, and create a scatterplot of these data points. Control variables: Industry categories (Nace Rev 2), firm size categories, occupation categories (5 digits); Source: JVS, IEB; Full-time and part-time jobs; Sources: JVS, IEB.

	(1) log earnings	(2) log earnings
mother	-0.1877***	-0.1150***
(male=reference)	(0.0114)	(0.0129)
childless female	-0.1300***	-0.0607***
(male=reference)	(0.0063)	(0.0084)
Observations	18,324	16,390
Adjusted R^2	0.6417	0.6498

Table B.1: Estimates for full-time and part-time workers; Standard errors in parentheses; Controls: total number of applicants, a set of worker age dummies fully interacted with education dummies, experience in years as well as its squared term, a dummy for the previous labor market status (non-employed, unemployed, employed), the hours of the new contract, dummies for formal job requirements, year dummies, industry categories, occupation categories, and establishment size deciles; Columns (2) additionally control for the share of male applicants; * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Figure B.14: Mothers and Childless Females



Note: Figures show the earnings gap (marginal effects) for mothers and childless females compared to males as a reference group at various levels of the share of male applicants. Controls: the share of male applicants interacted with a dummy for mothers and childless females (male=reference), the total number of applicants, a set of worker age dummies fully interacted with education dummies, experience in years as well as its squared term, a dummy for the previous labor market status (non-employed, unemployed, employed), the hours of the new contract, dummies for formal job requirements, year dummies, industry categories, occupation categories, and establishment size deciles; Full-time and part-time jobs; Sources: JVS, IEB.