Gender Gaps in Productivity, Wages, and Promotions: Evidence from a Random Task Allocation Policy*

Seyedehnegar Khaliliaraghi† Petter Lundborg‡

Johan Vikstrom§

October 18, 2024

Abstract

This paper investigates the longstanding question of whether pay differences between men and women are driven by productivity differences. For this purpose, we leverage the quasi-random assignment of jobseekers to caseworkers at the Swedish Employment Agency, enabling us to capture productivity measures free from gender bias in task allocation. Our findings show that female caseworkers are at least as productive as their male counterparts and earn comparable wages. However, significant gender gaps in promotions persist, with women being substantially less likely to be promoted despite equal productivity. Additionally, female caseworkers tend to manage fewer jobseekers and work fewer hours, which explains the majority of the observed earnings gap.

Keywords: Gender Gaps, Productivity, Wages, Task Allocation
JEL Classification: D84, I12, J12, J21

*We are grateful to Heather Sarsons, Petra Persson, and Karin Hederos for their valuable feedback, and to the seminar participants at Stanford University and Lund University for insightful comments and suggestions.

†Corresponding author. Lund University. (email: seyedehnegar.khaliliaraghi@nek.lu.se)

‡Lund University. (email: petter.lundborg@nek.lu.se)

§IFAU. (email: johan.vikstrom@ifau.uu.se)

1 Introduction

Despite the progress that women have made in the labor market over the last half century, gender gaps in earnings remain pervasive (Goldin, 2014; Cortes and Pan, 2020). The gender gap persists even as women are surpassing men in terms of educational attainment in most Western countries and is even more pronounced in high-status, high-income sectors, occupations, and job positions (Blau and Kahn, 2017; Bertrand, 2020). Understanding the source of the remaining gender gaps is important for the design of family-friendly policies. In this paper, we investigate the longstanding question of whether pay differences between men and women are driven by productivity differences.

While the role of human capital has diminished as an explanation for the remaining gender gaps, as women are better educated than ever, there are several reasons why productivity differences by gender could still arise. Even among high-educated workers, the division of household work typically puts the main burden on women, which may interfere with market productivity. In particular, motherhood may affect productivity through a number of channels such as human capital loss during parental leave, working part-time, and through constraints on working long hours and travels (Bertrand et al., 2010; Lundborg et al., 2017; Kleven et al., 2019; Adams-Prassl et al., 2023). In the presence of such mechanisms, employers may engage in statistical discrimination against women and invest less in their firm-specific training or assigning them to less attractive job tasks with less scope for learning (Correll et al., 2007).

There is still limited understanding of whether gender pay differences are driven by productivity differences, however. One key reason for this gap in knowledge is that workplace performance is rarely directly observed. Even when it is, performance can be confounded with task assignments. If task assignment is not gender-neutral and female workers are allocated to less productive tasks, the resulting productivity measures may not accurately reflect the true productivity differences between men and women (De Pater et al., 2010; Babcock et al., 2017; Zeltzer, 2020).[1] Even with similar job tasks, however, subjective productivity measures can be gender biased if women are held to higher standards during performance evaluations (Goldin and Rouse, 2000; Blau and Devaro, 2007; Card et al., 2019; Beg et al., 2021; Sarsons et al., 2021; Hengel and Moon, 2020; Hengel, 2022).

Our paper overcomes these challenges by using data on performance in a unique real-world setting where high-skilled workers were effectively randomly allocated to homogeneous tasks. The setting is the Swedish Employment Agency, who implemented a program that assigned case-workers to job-seekers based on the job seekers date of birth within a month, effectively randomizing job seekers to case workers. This setting gives us two key advantages. First, because of the random assignment, we can be sure that female and male case-workers were allocated to the same type of tasks and faced similar types of job seekers, meaning that our productivity measures are not confounded by any gender bias in task allocation. Second, the performance data provides us with register-based objective productivity measures, such as the time it takes for the case worker to help a client find a job. Since these productivity measures are not based on subjective evaluations, we can also be sure that the performance measures are not gender-biased themselves. These two features allow us to obtain unusually "clean" measures of gender gaps in productivity.

Our first set of findings reveal small differences in productivity between female and male caseworkers, with females, if anything, being slightly more

[1] Babcock et al. (2017), for instance, show that women are more likely to be asked to volunteer for service tasks that are not valued in promotion processes.

productive than their male counterparts. Job seekers that are randomly allocated to female case-workers find jobs somewhat faster than those allocated to male case workers. These findings, based on unbiased productivity measures, contrast with some recent literature showing that female workers are less productive on average (Azmat and Ferrer, 2017; Cook et al., 2020; Bolotnyy and Emanuel, 2022; Gallen, 2023, forthcoming).

We then go on to explore the role of some factors that are commonly believed to affect productivity. We rule out human capital as an important explanation for any gender productivity gaps. We continue by examining the role of parenthood for productivity differences. While female case workers are more productive on average, it is also well known that motherhood takes a greater toll on the labor market careers of women compared to men. Our results show that parenthood, in general, is not strongly related to productivity and that controlling for it does not affect the small gender productivity gap.

Next, we investigate how work experience relates to the gender productivity gap. Female case-workers have accumulated less work experience on average, due to factors such as maternal leave. To the extent that there is learning on the job, one would therefore expect female productivity to be negatively affected. We show that tenure is positively related to productivity but that accounting for tenure does not meaningfully affect the gender productivity gap.

We continue by examining the relationship between gender, productivity, and wages. If productivity is a key determinant of gender differences in wages, as a recent literature suggests, our results would therefore also predict small differences in wages between female and male caseworkers in the offices that used date of birth to allocate job seekers to caseworkers. This is also what we find; wage differences are minimal across equally productive female and male caseworkers. Furthermore, this suggests that other factors potentially affecting wage setting, such as discrimination or gender differences in wage bargaining, are less important in this setting.[2]

We go on to study gender gaps in annual earnings. We show that female caseworkers earn 7 percent less than their male counterparts on an annual basis. Since productivity and wage differences are small, this gap in earnings must reflect gender gaps in hours worked. We show that this is indeed the case; female caseworkers meet fewer job seekers and differences in contracted hours account for half of the gap, while the remaining part reflect gaps in *effective* hours worked.

Finally, we study gender promotion gaps. Since gender gaps in labor market outcomes may arise through differences in promotions, these gaps would not be revealed by examining pay gaps among workers performing the same task in similar types of positions (Lazear and Rosen, 1990). Again, our setting provides an interesting opportunity to investigate gender promotion gaps, as we can rule out that any such gaps reflect differences in productivity. We show that there exists a substantial gender gap in promotions at offices that adopted the random task allocation policy. Male caseworkers are substantially more likely to be promoted at these offices. While the gap could reflect that female caseworkers apply less for promotions or more often turn down offers of promotion, we continue to find large gender gaps for groups of workers where childcare demands, which may conflict with working long hours, are less binding.

Our results are robust to using alternative productivity measures and alternative specifications of the relationship between productivity and wages. Our main productivity measure is based on the length of the job seekers' unemployment spells and does not incorporate the quality of the job. We show, however, that job seekers with male and female caseworkers exhibit

[2]See Biasi and Sarsons (2021) for recent evidence on the on the significance of wage bargaining in relation to gender wage gaps.

comparable earnings at their first job after leaving unemployment and comparable earnings and rate of employment at 5 years after appearing at the employment office.

The results provide new insights on gender pay gaps and the sources behind them among high-skilled workers. Although our findings pertain specifically to caseworkers at the Swedish Public Employment Agency, we believe they hold broader relevance. Similar patterns of earnings and promotion gaps have been observed in other high-skilled professions such as law, and among CEOs, physicians, pharmacists, and university professors (Jagsi et al., 2006; Bertrand, 2011; Goldin and Katz, 2016; Azmat and Ferrer, 2017; Sarsons et al., 2021). Additionally, like other high-skilled workers, caseworkers have significant discretion in choosing their tools, and their services significantly impact jobseeker outcomes (Graversen and van Ours, 2008; Crepon et al., 2013; Schiprowski, 2020; Cederlöf et al., 2021; Humlum et al., 2023). Focusing on caseworkers allows for straightforward comparisons of performance in similar tasks, which is challenging when comparing across different industries or firms. Moreover, the random allocation of tasks ensures unbiased performance measures, enabling us to study gender productivity gaps accurately.

Our paper relates most closely to the small literature that studies gender differences in productivity among high-skilled workers in specific settings. Azmat and Ferrer (2017) show that male lawyers bill more hours and bring in more client revenue than female lawyers and that these performance differences explain a substantial part of the gender earnings gap. Our paper contributes by leveraging a context where task allocation among high-skilled workers is effectively random, ensuring that female and male workers perform the same type of tasks and that our productivity measures are not confounded by task assignment. Moreover, we provide evidence from a public sector setting, which reflects the work situation of a large share of the female work force in countries with a large public sector such as Sweden.[3]

Additionally, our paper contributes to the literature on gender differences in productivity that have focused on low-skilled settings, where productivity is easy to measure and workers perform homogeneous tasks. Cook et al. (2020) show that male Uber drivers drive faster and pick more lucrative driving spots than their female counterparts. These productivity differences, in turn, explain a large part of the gender earnings gap. Focusing on bus and train operators, Bolotnyy and Emanuel (2022) show that female operators work less overtime and value productivity-related job features such as schedule conventionality, predictability, and controllability more than male operators. The earnings gap is largely explained by these differences. While these settings also allow for the estimation of unbiased productivity measures, as the tasks are relatively standardized, it is less clear whether the results generalise to high-skilled settings.

We also add to the literature studying gender gaps in promotions. While gender gaps in promotion has been commonly documented, the mechanisms behind those have been less clear (see reviews of the literature in Cobb-Clark (1998), Bertrand (2011), Blau and Kahn (2017), and Cortes and Pan (2020)). Some studies have shown that gender promotion gaps exist after accounting for performance (Ginther and Kahn, 2006; Blau and Devaro, 2007; Azmat and Ferrer, 2017; Sarsons et al., 2021). Furthermore, Benson et al. (2024) show that women receive lower ratings of future potential despite receiving higher job performance ratings and that such differences explain a large part of the gender promotion gap. In addition, a recent literature has shown that gender promotion gaps in some settings are to a

[3] A related literature studies the relationship between wage gaps and gaps in marginal product for various observable characteristics. Hellerstein et al. (1999) show that differences in wages based are equal to differences in marginal productivity in the U.S., with the exception of gender. Using Danish register data, Gallen (2023, forthcoming) finds that the earnings gap is equal to the productivity gap.

large extent explained by gender gaps in applying for promotion (Bosquet et al., 2019; Hospido et al., 2022; Fluchtmann et al., 2024; Haegele, 2024) or by gender differences in career aspirations (Azmat, Cuñat and Henry, 2024, forthcoming). We add to this literature by studying gender promotion gaps in a situation where high-skilled female and male workers are assigned the same types of tasks and where we can rule out that objective gender productivity gaps drive the gender gap in promotion.

The remainder of the paper proceeds as follows. Section 2 discusses the institutional in which the random-allocation policy takes place. Section 3 describes the data sources used in the empirical analyses. Section 4 introduces our empirical strategy. Section 5 presents our main set of results. Section 6 concludes.

2 Institutional context and caseworker assignment

In this section we describe the institutional setting and provide details about the policy that randomly allocated case workers at the Swedish Public Employment Service to job seekers.

Case workers at the Swedish Public Employment Service (PES) are responsible for helping job seekers find employment and matching employers with suitable candidates. The agency provides a wide range of services, including career guidance, job matching, skills assessments, and training opportunities. Our paper focuses on the early 2000s, where the PES operated through approximately 300 local offices, each providing assistance to job seekers within their respective areas. Individuals looking for work were obliged to register at their nearest PES office in order to receive unemployment insurance benefits and support from the PES.

Swedish caseworkers have a high degree of flexibility and can decide themselves how often to meet with job seekers, which labor market programs to assign them to, and can refer job seekers to relevant job openings using their connections. The case workers also monitor job seekers' search behavior, making sure it complies with the unemployment insurance requirements.

The caseworkers come from diverse educational and professional backgrounds, as the PES has sought to attract individuals with various skills and experiences. During our study period, caseworkers were required to have at least an upper secondary education degree and three years of work experience. In practice, however, 71 percent of the case workers had completed a university degree, often in the field of human resource management.

2.1 Date-of-birth assignment of caseworkers to job seekers

The managers at local PES offices have the flexibility to tailor their activities and organization to meet local needs. This includes deciding how to allocate job seekers to caseworkers. Some offices match job seekers to the caseworker best suited to support them, while others assign caseworkers who specialize in certain industries or groups. Yet others assign job seekers to caseworkers based on their date-of-birth, as this is perceived as a transparent and way to equalize workload and monitor performance. In these cases, the allocation of job seekers to caseworkers becomes effectively random. We exploit data from the offices using the date-of-birth-allocation rule in our empirical analyses in order to estimate the productivity of the case workers, as further explained below.

While our dataset does not explicitly specify which offices implement a date-of-birth rule or which caseworkers are responsible for different date-of-births, this information can be easily inferred from the available data. Figure 1 provides examples of how job seekers are allocated to caseworkers based on date of birth in two different types of offices: one that uses a date-of-birth allocation mechanism (panel A) and one that uses other allocation mechanisms (panel B).

In panel A, it is evident that caseworkers are primarily responsible for job seekers born on certain dates. For instance, Caseworker 1 primarily handle job seekers born on four specific dates the month, while Caseworker 2 is responsible for another range of dates. Conversely, panel B illustrates the allocation in offices using different mechanisms, where the distribution of birth dates across caseworkers is more uniform.

Although Panel A shows an example of an office where date-of-birth rules are clearly followed when allocating job seekers to caseworkers, it also reveals that exceptions are sometimes made. These exceptions can be due to temporary increases in workload or job seekers with special needs. To address this type of non-random allocation in our empirical analyses, we follow Cederlöf et al. (2021) and define a *predicted* caseworker for each job seeker. The predicted caseworker is the one who would have been assigned if the date-of-birth rule had been strictly followed. Specifically, for each office, year, and day of the month, we identify the predicted caseworker as the one with the largest number of job seekers born on that day of the month.

In order to determine whether an office followed the date-of-birth, we need a formal method to determine whether an office follows this rule. To do this, we conduct an F-test for each office to assess whether the distribution of job seekers' dates of birth is uniform across caseworkers. Specifically, we regress the job seekers' date of birth on caseworker dummies (within each office and year) and test their joint significance using an F-test. We classify offices with F-values above 100 as following a date-of-birth rule for allocating job seekers.[4] In robustness analyses, we test the sensitivity of

[4]Figure A1 illustrates the distribution of F-statistics across all offices, truncated at 200. Figure A2 compares the distribution of offices using date-of-birth rules (F-statistics greater than 100) and those employing other allocation mechanisms across time. See the data section for further details.

our results using alternative thresholds.

3 Data

In our analyses, we leverage data from the Swedish Public Employment Service (PES), containing detailed information about every caseworker employed within the organization and all the job seekers assigned to the caseworkers during 2003-2014.

The caseworkers from PES data are linked to Statistics Sweden (SCB) datasets which include the population register, wage statistics, and the universe of employer-employee matches. The population register contains information on demographics, education level, and annual income of the caseworkers. We utilize this data in addition to employer-employee dataset to identify the individuals who are employed at the unemployment agency and see their employment history.

From these data, we are able to determine their exact tasks, ranging from being a caseworker, to holding more senior positions. With information on occupational codes, we construct a measure of caseworker experience, which reflects years of employment in PES as a caseworker, and a measure of caseworker promotion. We define promotion as instances where the occupational code shifts from a caseworker role to a managerial or higher-ranking position, coupled with an increase in wages. Using the meetings each caseworker had with different job seekers, we identify the exact PES office the caseworker was employed in. Some caseworkers engage in meetings with job seekers across multiple offices within a given year. In such instances, we assume that the caseworker is formally employed by the office where the caseworker holds the majority of allocated job seekers. Furthermore, we utilize wage statistics datasets which have data covering all individuals employed in the public sector, including PES. To construct a variable for wages and hours worked, we obtain data on contracted wage and hours from the wage statistics and link it to the caseworkers.

To build a dataset on the job seekers' information, we first use the PES dataset which includes all of the job seekers who were registered at PES between 2003-2014. This dataset provides detailed insights into job seekers' unemployment spells, including the initiation of their unemployment, duration until leaving unemployment, assigned caseworkers, and the meetings they had with their designated caseworkers. The dataset also contains information on job seeker's date of birth, a crucial component which we use to identify the offices employing the date-of-birth method for allocation of job seekers.

Using data on unemployment spells, we create an indicator for "Leaving employment within 180 days" for each job seeker, which serves as our primary measure of caseworker productivity in our analysis. Additionally, we develop alternative productivity metrics, such as the total number of days spent unemployed and indicators for leaving unemployment within 30 and 90 days. The job seekers from PES dataset are then linked to SCB datasets to obtain information about their demographics, education, and employment. This data in addition to the PES data enables us to construct variables showcasing job seekers' post-unemployment job quality, including future earnings and the tenure of their first job.

In our analyses, we apply several sample restrictions on our population of caseworkers and job seekers. For the calculation of productivity measures, we exclude small and atypical offices–those with fewer than 200 job seekers per year or caseworkers with fewer than 30 assigned job seekers per year. Additionally, we drop job seekers who were registered multiple times within a single year from the sample. These restrictions result in a dataset comprising 8,719 individual caseworkers who work in offices with or without date-of-birth rule.[5] They have meetings with 1,589,249 job seekers registered at these 257 unique offices during 2003-2014. Within this sample, 5,140 caseworkers work in offices which implemented date-of-birth rule in at least one year during 2003-2014.

As statistics illustrate in Table 1, caseworkers employed at DOB offices are on average 46 years old, 63% are female, 16% are immigrants, and almost 70% have a university degree. On average, the random offices are larger than non-random ones, indicating a higher number of caseworker employed and, consequently, a greater number of registered job seekers at these locations. Caseworker and job seeker characteristics are, however, similar across the two types of offices. The small differences in characteristics are sometimes statistically significant but are in all cases small in magnitude.

Table 2 shows descriptive statistics on the caseworkers and job seekers only in date-of-birth rule offices. Female caseworkers are on average 45.5 years old and 66% of them hold a university degree. Male caseworkers are on average somewhat older, at 46.4 years, and have comparable levels of education. The most common degrees among the caseworkers are business management and social work. 54% of the caseworkers have more than 8 years of experience as caseworkers, while only 10 percent have less than 2 years of experience.

3.1 Is there enough variation in wages at the Public employment agency?

One concern with the PES setting is the potential limited variation in wages and earnings in our data. However, it is important to note that the PES employs individually negotiated wages. We next discuss how wages are set

[5] This number excludes caseworkers working in offices where it cannot be definitively determined whether they adhere to the date-of-birth rule or not.

in this context and provide three pieces of evidence showing substantial variation in wages across caseworkers.

Yearly collective agreements between trade unions and the government set the framework for wage bargaining between caseworkers and their managers at the PES (Samarbetsrådet, 2008).[6] These collective agreements regulate the overall scope of wage increases and overtime pay. Based on these agreements, an individual and differentiated salary setting is applied. In practice, this can be implemented in two ways. The primary method is individual bargaining, where wages are set through individual wage negotiation talks between the manager at the local offices and each caseworker. The intention is to achieve a clear link between caseworkers' performance and their wages.

Statistics show that in 2010, 78 percent of all salaries in the state sector were set through this type of individual salary negotiation.[7] The alternative is to set local wages through semi-collective wage bargaining between the local union and the office managers. However, even this arrangement involves individual salary setting. Collective negotiations are preceded by salary discussions between the caseworkers and the local managers. Subsequently, individual salaries are negotiated based on the results of these salary discussions.

Our data supports the existence of individual wage-setting. First, Figure 2a shows substantial variation in starting salaries among caseworkers after accounting for year-specific effects. Second, Figure 2b shows significant variation in wage *changes* across subsequent years, consistent with individual wage-setting. Third, we observe that age, year, and level of education only explains 43 percent of the variation in wages.

[6]For caseworkers, two approximately equally sized trade unions are relevant (ST and SACO-S), while the negotiating party for the government is Arbetsgivarverket.

[7]https://www.arbetsgivarverket.se/statistik-och-analys/staten-i-siffror-loner/staten-i-siffror-loneutveckling/statistik-om-lonesattande-samtal/.

4 Empirical strategy

Our empirical strategy leverages data from Public Employment Service (PES) offices where job seekers are assigned to caseworkers based on date-of-birth rules. This quasi-random assignment ensures that, on average, male and female caseworkers are assigned similar types of job seekers. Consequently, the characteristics of the caseworkers will be uncorrelated with those of the job seekers. This allows us to estimate productivity differences, such as how quickly job seekers find employment and the quality of those jobs, between male and female caseworkers. Without the date-of-birth rules, which effectively randomize the allocation of job seekers to caseworkers, making meaningful productivity comparisons between male and female caseworkers would be challenging due to potential systematic differences in job seeker assignments. Our research design follows the empirical strategy developed by Cederlöf et al. (2021) and also used by Humlum et al. (2023).

Facilitating job seekers in securing employment is the primary responsibility of caseworkers, making it the most obvious reflection of their productivity. Consequently, our main productivity measure is whether the caseworkers' job seekers exit unemployment within 180 days. For robustness analyses, we also consider alternative outcomes related to job quality and various measures of unemployment duration, such as securing employment within 30 days and 90 days.

4.1 Measuring productivity differences

In the first part of our empirical analysis, we focus on gender differences in caseworker productivity. Specifically, we regress job seeker outcomes (primarily leaving unemployment within 180 days) on a dummy for being a female caseworker, $Female^{CW}$, and other caseworkers characteristics,

X^{CW}, using the following model:

$$y_{icpgt} = \alpha + \delta Female_c^{CW} + \beta X_{ct}^{CW} + (\gamma_t \times \theta_p \times \lambda_g) + \epsilon_{icpgt}, \tag{1}$$

where y_{icpt} denotes the outcome of job seeker i who is assigned to caseworker c at office p in age group g in year t. δ captures the productivity difference between male and female caseworkers. The other caseworker characteristics, X_{ct}^{CW}, include, age dummies, level and type education, number of children, immigrant status, and years of experience as a caseworker. Note that some of these caseworkers characteristics may change over time. The model also includes interacted year, γ_t, office, θ_p, and jobseeker-below-25, (λ_g), fixed-effects. The first two fixed effects enable us to leverage the within-office variation created by the date-of-birth rules, while the last one accounts for differences in allocation strategies for job seekers under the age of 25 in various offices.[8]

As explained in Section 2.1, exemptions to the date-of-birth rule for allocating job seekers to caseworkers are sometimes made, potentially introducing non-random sorting. To mitigate this issue, we use an instrumental variable framework, leveraging the predicted caseworker whom the job seeker would have been assigned to according to the date-of-birth rule as an instrument for the actual assigned caseworker. Specifically, we achieve this by instrumenting the characteristics of the actual caseworker by the corresponding characteristics of the predicted caseworker. For instance, the female caseworker dummy is instrumented by a female dummy for the

[8] Cederlöf et al. (2021) notes that it is quite common for offices to use a separate date-of-birth allocation system for youths (aged 24 or younger). Offices often have specific caseworkers dedicated to supporting youths and assign these youths to "youth" caseworkers based on their date of birth, while non-youths are assigned to other caseworkers using a different date-of-birth allocation. For example, one caseworker may support all youths in the office born between the 1st and 15th of the month, while another caseworker may support all older workers born between the 1st and 8th. Therefore, we assign the predicted caseworker separately for youths and non-youths and interact the office and year fixed effects with an age group dummy for being younger than 25 (λ_g).

predicted caseworker. This IV-strategy allow us to exploit only the as-if random variation created by the date-of-birth rules, which assigns all job seekers in an office who are born on the same day of the month to the same caseworker.

To validate that the predicted caseworker characteristics are good predictors of the characteristics of the actual caseworker assigned to the job seeker characteristics, Table A1 reports first-stage estimates, where we regress the female dummy and other characteristics of the predicted caseworker on the corresponding characteristics of the actual caseworker. As illustrated by the regressions, each characteristic of the actual caseworker is highly correlated with the predicted caseworker characteristics.

4.2 Validating the as-if randomization of caseworkers to job seekers

As date-of-birth offices allocate caseworkers to job seekers based on the latter's date of birth (day in the month), the allocation process of (predicted) caseworkers to job seekers should mimic a random matching process. To verify this, we regress the characteristics of the job seekers on those of the predicted caseworkers. If the allocation is truly random, there should be no systematic correlation between the characteristics of job seekers and the predicted caseworkers. To contrast, a similar analysis conducted on data from offices without a date-of-birth rule, where the assignment process is not random, can shed light on whether there is any systematic allocation based on gender or other attributes.

Table 3 uses data from the offices *without* a date-of-birth rule, and shows regression estimates for the relationship between various caseworker and job seeker characteristics, where the latter includes age, gender, immigrant status, the presence of a disability, education, recipient of welfare benefits, and earnings and unemployment in the previous year.[9] The tables show that female caseworkers are disproportionately assigned to younger job seekers, those who are also female, and those with lower labor earnings, suggesting that female caseworkers may be at at an disadvantage. They are more likely to be allocated job seekers who, on average, may have more barriers to employment. These differences in caseload composition could impact the professional development and performance evaluations of female caseworkers adversely, highlighting an inherent gender disparity within the allocation system of caseworkers to job seekers. We also see that other caseworker characteristics, such as age, education, and experience are systematically related to job seeker characteristics.[10]

In Table 4, we turn to the offices *with* a date-of-birth rule. Since our analyses exploits variation based on the predicted caseworker, we now correlated the characteristics of the predicted caseworker with those of the actual caseworker. In contrast to offices without a date-of-birth rule, we find no evidence of any systematic correlation between job seeker and predicted caseworker characteristics, very much in line with as-if random allocation procedure implied by the date-of-birth rules.[11] This implies that we can use the data from the date-of-birth offices to obtain unbiased measures of productivity differences by gender.

4.3 Calculating productivity measures

In the second part of our analysis, we focus on the role of productivity, gender, and other caseworker characteristics for labor market outcomes of

[9] As before, we control for the interaction between office and year fixed effects to focus on the systematic allocation of caseworkers within an office.

[10] For instance, more experienced caseworkers are more likely to be assigned job seekers who are older, male, disabled, and who are better educated, earn more, and are more likely to have been unemployed.

[11] The only significant estimate suggests that female caseworkers are slightly less likely to be allocated a male job seekers, but the estimate is small at -0.004 percentage points.

the caseworkers. That is, we examine how productivity relates to gender gaps in wages, earnings, and promotions. Thus, we are interested in the following model:

$$LM_{cpt}^{CW} = \alpha + \delta Female_c^{CW} + \beta X_{ct}^{CW} + \kappa Prod_{ct}^{CW} + (\gamma_t \times \theta_p) + \epsilon_{ct}, \quad (2)$$

where LM_{cpt}^{CW} is caseworker labor market outcomes such as wages, earnings, and an indicator for getting promoted to senior positions. We primarily focus on the difference in labor market outcomes between female and male caseworkers, captured by coefficient δ. The other caseworker characteristics, X_{ct}^{CW}, include variables often used when studying gender wage gaps, including age, level and type of education, number of children, and years of experience as a caseworker. As before, year and office fixed-effects are included to exploit the within-office variation created by the date-of-birth rules. Our analysis includes only offices with date-of-birth assignment of job seekers, excluding the other offices. This approach allows us to identify the influence of random task assignment and productivity on gender disparities in the labor market.

Estimating equation (2) requires a productivity measure, $Prod_{ct}^{CW}$, representing the estimated productivity of caseworker c in year t. To this end, we follow Cederlöf et al. (2021) and exploit the date-of-birth allocation as before, but now we estimate caseworker fixed-effects as measures of the productivity of each caseworker. As in equation (2), we use job seeker outcomes to (leaving unemployment within 180 days) capture productivity differences and estimate:

$$y_{icpgt} = \alpha + \mu_{ct} + (\gamma_t \times \theta_p \times \lambda_g) + \epsilon_{icpgt}, \quad (3)$$

where y_{icpt} denotes whether the job seeker i found a job within 180 days. Importantly, μ_{ct} represents the fixed-effect of caseworker c in year t. The fixed effects are allowed to vary across years to capture changes to experience, childbearing, and other things that may change over time. These fixed effects, and various transformations of those, are then used as measures of caseworker productivity, $Prod_{ct}^{CW}$, in equation (2).

When estimating equation (3) we, again, need to adjust for selective exemptions by using the predicted caseworker as an instrument. Previously, we used the characteristics of the predicted caseworker as instruments for the characteristics of the actual caseworker. Here, we use indicators for each predicted caseworker as instruments for the caseworker fixed-effects. That is, for each endogenous variable (caseworker fixed effect) we have one instrument (indicator for the predicted caseworker). Note that we can only estimate fixed effects for caseworkers from whom we have an instrument, i.e, for those who at some point are a predicted caseworker.

The estimated fixed effects can be interpreted as caseworker value-added terms. For the validity of our empirical design, it is crucial that caseworkers significantly impact job seekers' outcomes. This was confirmed by Cederlöf et al. (2021), who demonstrated substantial heterogeneity in caseworker value-added, showing that these terms meaningfully relate to job seeker outcomes. For example, having a caseworker who is one standard deviation better in terms of value-added increased the probability of leaving unemployment within 90 days by 9 percent of a standard deviation. Similar effects were observed for leaving unemployment within 180 days and the overall duration of unemployment.

5 Results

5.1 Productivity differences by gender

We begin our empirical analysis by investigating gender differences in productivity, using data from the offices that randomly allocated case workers to job seekers. In Table 5, the columns show how the (potential) gender gap in caseworker productivity is affected by accounting for various productivity-related factors, such as education, the presence of children, and labor market experience. All regressions include age fixed effects and year-by-office-by-jobseeker-25 fixed effects. As described in Section 4.1, caseworker characteristics are instrumented by predicted caseworker characteristics.

Column 1 in Table 5 presents the gender productivity gap while only controlling for age fixed effects, year-by-office-by-jobseeker-25 fixed effects, and immigrant status. The findings show that females caseworkers are slightly *more* productive compared to their male counterparts. Job seekers assigned to female caseworkers have a statistically significant 0.71 percentage points higher likelihood of exiting unemployment within 180 days. Given a baseline rate of 63 percent, this corresponds to 1 percent higher productivity among female caseworkers. This finding suggests that gender-based productivity differences are small and are unlikely to account for any large disparities in wages or promotion opportunities that may exist between male and female caseworkers.

We continue by examining to what extent accounting for additional controls affects the gender productivity gap. As shown in Table 1 and 2, female and male caseworkers differ across dimensions such as education, number of children, and work experience. To the extent that productivity is related to such factors, the gender productivity gap may therefore increase or decrease when accounting for them.

Column 2 adds education variables as controls. They measure if the caseworker has a university or secondary school degree and whether the degree is in business or social sciences, or other disciplines. Adding these education controls hardly affects the gender productivity gap at all. Furthermore, these human capital measures are, in themselves, not strongly related to productivity.

We next examine the role of parenthood for the gender productivity gap. While female case workers are more productive on average, it is also well known that parenthood takes a greater toll on the labor market careers of women compared to men. To the extent that motherhood is associated with lower productivity, accounting for it could potentially lead to an even greater productivity difference in favor of female caseworkers. The results in column 3 cast doubt on lower productivity being an important mechanism behind the child penalty, however. Accounting for the number of children, and separately accounting for having children below the age of 5, does not alter the small gender productivity gap to any important extent.[12] The exception is those having 3 or more children who are significantly less productive. The difference is relatively minor, however, at 2 percent.

It may appear surprising that childbearing is largely unrelated to productivity among the caseworkers, given that the extended parental leave periods in Sweden may lead to human capital depreciation. We can think of several reasons why the presence of children does not significantly impact the productivity of caseworkers. The duration of parental leave may simply be too short to experience a substantial loss of human capital. Even in Sweden, where maternity leave periods are notably generous, taking one year off for child-rearing constitutes a minor fraction of an entire career span. Additionally, flexibility of work arrangements at the Swedish employment agency may facilitate the combination professional and family responsibilities.

In the fourth column, we examine the influence of years of experience as a caseworker on the gender productivity gap. As shown in Table 1, female caseworkers, on average, possess less labor market experience, limiting their

[12] Additionally, we conducted separate regressions by gender to assess whether the relationship between parenthood and productivity is stronger for female caseworkers. The results provide no evidence to support this.

opportunities for on-the-job learning and career advancement.[13] Accounting for experience could therefore be expected to increase the productivity gap in favor of women, if anything. For this analysis, we add dummies indicating different degrees of experience working at the PES. We drop the dummies indicating child-bearing, however, as we want to allow any effects of lost experience to run through parental leave periods. The findings show that the inclusion of experience in the regression does not markedly alter the gender productivity gap. Furthermore, the coefficients associated with the work experience variables are small and not statistically significant, suggesting that any lost experience due to caregiving responsibilities does not impact productivity levels. In the fifth column, we simultaneously incorporate all the controls. This does not significantly alter the minor gender productivity gap.

The results in Table 5 are based on one particular productivity measure; the likelihood of finding a job within 180 days. To assess the robustness of our results, Table A2 show regression results where we employ alternative productivity measures, including the probability of finding a job within 30 days, 90 days, and the (log) duration of unemployment. The results are similar to the ones in Table 5; job seekers allocated to female caseworkers are significantly more likely to find a job within 90 days and have significantly shorter unemployment duration. In the latter case, job seekers allocated to female caseworkers have 2 percent shorter unemployment durations. At the 30-day follow-up, there is no significant difference, however.

Our findings showing an absence of gender productivity differences contrasts to some results in the recent literature. Azmat and Ferrer (2017) found large gender differences in productivity within the legal profession in the United States. The observed difference was partly attributed to the

[13] The difference in experience is also present conditional on age.

impact of parenthood on the productivity of female lawyers, which appears to be more pronounced than its impact on female caseworkers. The different findings may stem from differences in job structure, perhaps reflecting that caseworkers are better able to balance family responsibilities with their professional roles and their more limited opportunities to work long hours. Additionally, gender differences in career ambitions were found to account for a substantial portion of the productivity gap among lawyers. Given that the wage premium for career advancement is smaller for caseworkers than for lawyers, gender disparities in career aspirations are likely to have a less important role for the productivity within this group. The comparison of results between caseworkers and lawyers suggests that even among high-skilled workers, there are substantial variation in gender productivity gaps and in the role of motherhood for productivity.

Our results so far suggest that female caseworkers exhibit slightly higher productivity in terms of the time it takes to transition their assigned job seekers out of unemployment. This measure of productivity is particularly relevant within the context of the unemployment agency, where it stands as the most visible indicator of caseworker performance to management. However, the correlation between faster job placement and the quality of labor market matches is less clear. A swifter job placement could potentially come at the expense of job quality.

To explore this further, we consider additional indicators of productivity: the duration of employment at the first job following unemployment, initial earnings at this first job, and long-term earnings and employment status. Since such an analysis needs to condition on getting a job during the *entire* observation window, this could cause a selection problem if female caseworkers are more productive also in this regard. Column 1 of Table 6 shows that this worry is unfounded as there is no difference in the long-run likelihood of getting their clients find a job between female and male caseworkers.

Columns 2-4 of Table 6 present findings from these additional productivity measures. Columns 2 demonstrates the absence of long-run productivity differences between female and male caseworkers in terms of the likelihood of the first job lasting at least two years. Moreover, columns 3 and 4 reinforce this finding by showing no significant differences in job seekers' long-term earnings and employment status five years later between those allocated to female or male caseworkers. These findings show that the somewhat faster job placements of female caseworkers does not compromise job quality.

5.2 Gender, productivity, and wages

Our results so far reveal small differences in productivity between male and female caseworkers. If productivity is a key determinant of gender differences in wages, as a recent literature suggests, our results would therefore also predict small differences in wages between female and male caseworkers in the offices that used date of birth to allocate job seekers to caseworkers. If gender gaps in wages prevail, however, this would suggest that other factors than productivity affects wage setting, such as discrimination or gender differences in wage bargaining. Our setting, where we can rule out any important differences in productivity between female and male workers, or control for them using clean productivity measures, thereby provides an attractive setting to test for this.

Before analyzing the relationship between gender, productivity, and wages, we first confirm that our productivity measure is related to the wages of caseworkers. Initially, it is unclear which functional form best describes the relationship between productivity and wages, as this may vary across institutional settings. Table 7 presents results from various specifications of the productivity-wage relationship, using data from offices that apply the date-of-birth rule. The regressions control for age fixed effects and office-by-year fixed effects.

The results in column 1 show that the estimated productivity variable, entered without any further transformation, has no statistically significant relationship with wages. Additionally, when we account for non-linearities by adding a squared term, or measure productivity at t-1 (columns 2 and 3), we also obtain small and insignificant estimates.

If managers at PES offices classify their caseworkers into different productivity categories when setting wages, a more relevant measure might be the within-office productivity category each caseworker falls into. Column 4 presents results where workers are divided into productivity deciles based on our estimated productivity measure, and where the decile is entered as a continuous variable. This within-office decile measure is significantly and positively related to wages. The relationship becomes even stronger when using the same measure at t-1, indicating that the previous year's productivity is likely a crucial factor in determining current year wages (column 5). The coefficient indicates that moving up one decile in the productivity distribution within an office is associated with a 0.2 percent increase in monthly wages. Consequently, a person moving from the first to the highest decile would earn approximately 2 percent more per month. Instead of assuming a linear relationship between the deciles and wages, Figure A3 plots the estimates of each productivity deciles when entered as separate categories. The results indicate a clear breakpoint at the median: caseworkers at or above the median earn more than those below it. Consequently, in columns 6 and 7, we classify workers in an office based on whether their productivity is above or below the median. As expected, workers above the median have significantly higher wages.

Having established that productivity relates to wages, we next examine the relationship between gender, productivity, and wages. The results are presented in Table 8. Here, we control for productivity in the final column of the table, in order to account for any small gender productivity gaps. We start by analyzing how the gender wage gap evolves when we zoom in on male and female case workers who perform increasingly similar work tasks. This can potentially shed some light on the how much of the wage gap that is explained by differences in task allocation. With our data, we can move from the typical case observed in register data, where tasks are unobserved and non-randomly allocated, to the case where female and male case workers perform the same type of tasks, due to the policy of randomly allocating job seekers to clients.

In column 1 of Table 8, we present the gender wage gap offices where the random task allocation policy was implemented but where the sample includes all case workers at these offices, including the ones who not exclusively work with job seekers and who have other tasks as well. This sample thus consists of workers in the same occupational category, where tasks are *not* equalized as we do not account for the type of tasks performed. As revealed in column 1, the gender gap in wages is small and insignificant in this group of workers, however.

In column 2 we turn to the gender wage gap among case workers working specifically with job seekers in the offices that have implemented the random-allocation policy. In this sample, we know from our previous analysis in Section 5.1 that female and male caseworkers are equally productive, with a slight advantage for the female ones. If anything, the gender wage gap now becomes even smaller and turns slightly in favor of female case workers.

We next examine the subset of caseworkers for whom we have productivity measures. The sample size is now smaller, as we can only measure productivity for part of the original group, as explained in the data section. Column 3 first explores the relationship between gender and wages in this subset without controlling for productivity, while Column 4 includes productivity as a control. As expected, adding productivity to the regression has minimal impact on the gender wage gap, since productivity differences were initially small.[14]

How do these results square with other recent ones in the literature? Whereas the absence of gender productivity differences contrasts to some recent studies, our results for the gender wage gap are consistent. We find no evidence of gender differences in wages in a situation where gender productivity differences are absent, whereas previous studies on lawyers, Uber drivers, and bus and metro drivers found wage differences in situations with large productivity differences by gender, and where the latter fully explained the former. In both cases, this leaves little room for wage discrimination or gender differences in wage bargaining as important phenomena in these contexts. Of course, these results are for a given task and do not rule out gender differences in promotion to different and higher paid job tasks and positions. We return to this in Section 5.4.

Our findings differ from other studies that identify a gender pay gap in negotiable wage settings, even after controlling for later productivity (e.g., (Azmat and Petrongolo, 2014; Biasi and Sarsons, 2021)). One potential explanation for this discrepancy is that our results may obscure heterogeneity in the relationship between gender and wages. For example, newly hired caseworkers may lack the precise productivity data necessary to effectively negotiate wages, while more experienced workers benefit from clearer productivity metrics. Since prior research often suggests that women are less effective in wage negotiations, it is possible that the gender pay gap differs

[14]Productivity is measured here as being an above-median productive worker within the office. The coefficient for the female dummy remains largely unchanged when alternative definitions of productivity are used or when wage changes are examined as the outcome (see Tables A3 and A4). Similarly, varying the thresholds for the F-statistics used to classify an office as a date-of-birth office does not significantly affect the results (see Table A5)

between newly hired and experienced caseworkers. To investigate this, Table A6 stratifies the sample by experience level and re-estimates the model from column 4 of Table 8. The results indicate no significant gender pay gaps for either inexperienced or experienced caseworkers. This finding may suggest that gender differences in salary requests are relatively small in the Swedish context, as observed by Save-Söderbergh (2019).

5.3 Gender and earnings

We next turn to the gender gap in annual earnings. As revealed in Table 2, there exists a sizable annual earnings gap of about 8 percent between male and female case workers. Since wage and productivity differences between male and female caseworkers are small, the source of this gap must be related to other factors. Female caseworkers could earn less on an annual basis because of working fewer contracted hours. Alternatively, female caseworkers may work fewer *effective* hours. This could be the case, for instance, if they take the responsibility of caring for sick children more often than male caseworkers.

In Table 9, we analyze the sources of the gender earnings gap, focusing on workers for whom complete productivity data is available. Column 1 presents the gender gap in the annual number of clients handled, while controlling for factors such as education, experience, age, and office-by-year fixed effects. On average, female caseworkers see 15 fewer clients per year, which can be attributed to their tendency to work fewer hours compared to their male counterparts. This represents approximately a 10 percent reduction in the number of clients seen.

In line with this, column 2 reveals that female caseworkers earn approximately 18,000 SEK less annually than male caseworkers, corresponding to a 7 percent gender earnings gap. Given that the monthly wage rate is similar for both genders, controlling for wage rate in column 3 has little effect on the earnings gap. This indicates, again, that the gender earnings gap is driven primarily by differences in hours worked.

Indeed, column 3 demonstrates that adding contracted hours as a control substantially reduces the earnings gap, though a significant 4 percent gap remains. This residual gap likely stems from differences in actual hours worked, which we are unable to observe directly in the data.

In column 4, we explore whether this remaining gap could be explained by time taken off to care for sick children by including indicators for the presence of children. While having children significantly affects annual earnings, the gender earnings gap remains largely unchanged. This suggests that unobserved factors, such as sick leave or unpaid leave, may be driving the remaining disparity.

It is interesting to note that the higher frequency of absences among female caseworkers apparently does not lead to any differences in productivity compared to their male counterparts. One possible interpretation of this finding is that the caseworker profession exhibits a high degree of "substitutability," meaning that workers in this field can easily be replaced by others, allowing for greater temporal flexibility (Goldin, 2014). Azmat, Hensvik and Rosenqvist (2022, forthcoming) demonstrate that women are more likely to choose firms where they have greater access to substitutes, which mitigates the productivity impact of unexpected absences. Consequently, we would expect the remaining earnings gap we find to be larger in firms with fewer available substitutes.

In summary, the findings in this section indicate that, in a context where task allocation is effectively randomized between male and female workers, the gender earnings gap does not arise from differences in productivity or wages, but rather from disparities in effective hours worked. In simpler terms, female caseworkers perform equally well in their roles but have a lower total output due to more frequent absences. This may be partially attributed to family responsibilities, such as caring for sick children.

5.4 Gender and promotions

Our findings so far reveal that wage and productivity differences among caseworkers performing the same tasks are small. Gender gaps in labor market outcomes may also arise through differences in promotions, however, which would not be revealed by examining caseworkers performing the same task in similar types of positions. Again, our setting provides an interesting opportunity to investigate promotion gaps. If female caseworkers are less likely to get promoted, despite being equally productive as their male counterparts, it must thus be due to other factors than productivity, such as discrimination, hours worked, or gender differences in applying for promotion.

In Table 10, we estimate the gender promotion gap, where a promotion is defined as switching from being a caseworker to having a senior and managerial role, according to occupational codes in the registers. We include all caseworkers in random offices with productivity measures, and control for education, immigrant status, and tenure, as well as age and year-by-office fixed effects. The results in column (1) reveal a large and significant gender promotion gap, where females are 1.9 percentage points, or 86 percent, less likely to get promoted compared to their male counterparts.

In Column 2, we introduce productivity as a control variable to rule out the possibility that small productivity differences are driving the observed gap. As anticipated, the gender promotion gap remains unchanged, suggesting that other factors are at play. One potential factor could be the number of contracted hours, as working fewer than full-time hours might hinder promotion to a managerial role. However, when contracted hours are added to the regression model (Column 3), the gender promotion gap remains largely unaffected. Contracted hours alone show a marginally significant positive relationship with the likelihood of promotion.

Several other factors may explain the gender promotion gap between equally productive workers. One explanation is that female caseworkers apply for promotions less frequently, as demonstrated by a recent literature (Bosquet et al., 2019; Hospido et al., 2022; Fluchtmann et al., 2024; Haegele, 2024). A potential reason for this lower application rate could be that roles requiring long hours and limited flexibility are harder to reconcile with current or future childcare responsibilities (Bertrand et al., 2010).

We do not observe applications for promotions in our data but if family responsibilities prevent some females to apply for promotions, or accept promotion offers, we should observe the effects primarily among workers who have small children. To investigate this, columns (3) and (4) show separate regressions for caseworkers with and without children, while in both cases controlling for productivity differences. Female caseworkers without children are still significantly less likely to be promoted, but the gap is about half in magnitude compared to that of female caseworkers with children. In the former group, female caseworkers are 1.4 percentage points less likely to get promoted. In the latter group, the corresponding effect is 2.7 percentage points.

The findings reveal that female caseworkers with children are less likely to get promoted compared to their childless counterparts and that this cannot be explained by the former group being less productive. But female caseworkers without children are also less likely to be promoted, suggesting that the gap is not entirely due the fertility channel. What could generate these patterns? One explanation would be that managers simply discriminate against female applicants. Managers may also believe that current performance is a poor predictor of future performance and that female workers holds less potential as managers or other higher positions (Benson et al., 2024). Further research should aim at gaining a deeper understanding of the reasons behind the gender promotion gap.

6 Conclusions

We explore the longstanding question of whether pay disparities between men and women stem from differences in productivity. It is commonly believed that women's more frequent absences from the labor market-often due to childbearing and sick leave-lead to productivity losses that contribute to the gender pay gap. Documenting productivity differences between female and male workers and their impact on gender pay gaps poses significant challenges, however. In high-skill professions, productivity measures are often unavailable, and when they are accessible, they may be biased due to potentially gender-biased task assignments. If female workers are assigned less productive tasks, performance measures become confounded with task assignments, obscuring genuine productivity differences between genders. To address this issue, we utilized a natural experiment at the Swedish Public Employment Service, where caseworkers were randomly assigned to job seekers based on their date of birth. This randomization effectively eliminated task assignment biases, enabling us to obtain unbiased and objective measures of productivity differences between men and women.

We present three main findings. First, gender gaps in productivity are minimal. Female caseworkers are at least as productive as their male counterparts in facilitating job placements. If anything, job seekers assigned to female caseworkers tend to find jobs slightly faster, without any compromise in job quality. This contrasts with recent literature that identifies significant gender productivity gaps in both high- and low-skill occupations. Additionally, parenthood showed only a weak association with productivity.

Second, wage differences between female and male caseworkers are small and statistically insignificant. Given the minimal productivity differences, traditional explanations for gender wage disparities, such as discrimination and differences in wage bargaining, hold little relevance in this context. The annual earnings gap observed between female and male caseworkers is primarily due to differences in effective hours worked.

Third, our results underscore the significant role of gender promotion gaps in contributing to gender pay gaps. Despite being at least as productive as their male counterparts, female caseworkers are much less likely to be promoted to managerial positions. This disparity could be due to fewer applications for promotions or a higher likelihood of turning down promotion offers among female caseworkers. However, we also observe substantial gender gaps in promotion among groups of workers with less stringent childcare demands, which are less likely to interfere with working long hours. These findings contrast with those from studies of U.S. lawyers, where performance accounted for most of the gender gap in promotions (Azmat and Ferrer, 2017).

Our findings contribute new insights into the role of productivity in gender disparities within the labor market. Although female caseworkers are as productive as their male counter and earn similar wages, a significant promotion gap persists, highlighting an ongoing issue of gender inequality. This suggests that other factors, such as potential discrimination or gender-specific barriers in promotion processes, continue to impede gender equality in career advancement. Understanding the underlying causes of these promotion disparities is an important area for future research.

References

Adams-Prassl, A., Hara, K., Milland, K. and Callison-Burch, C. (2023). The Gender Wage Gap in an Online Labor Market: The Cost of Interruptions, *The Review of Economics and Statistics* pp. 1–23.

Azmat, G., Cuñat, V. and Henry, E. (2024, forthcoming). Gender promotion gaps and career aspirations, *Management Science* .

Azmat, G. and Ferrer, R. (2017). Gender gaps in performance: Evidence from young lawyers, *Journal of Political Economy* **125**(5): 1306–1355.

Azmat, G., Hensvik, L. and Rosenqvist, O. (2022, forthcoming). Workplace presenteeism, job substitutability and gender inequality, *Journal of Human Resources* .

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

Babcock, L., Recalde, M. P., Vesterlund, L. and Weingart, L. (2017). Gender differences in accepting and receiving requests for tasks with low promotability, *American Economic Review* **107**(3): 714–47.

Beg, S., Fitzpatrick, A. and Lucas, A. M. (2021). Gender bias in assessments of teacher performance, *AEA Papers and Proceedings* **111**: 190–95.

Benson, A., Li, D. and Shue, K. (2024). Potential and the gender promotions gap, *SSRN Electronic Journal* .

Bertrand, M. (2011). Chapter 17 - new perspectives on gender, Vol. 4 of *Handbook of Labor Economics*, Elsevier, pp. 1543–1590.

Bertrand, M. (2020). Gender in the twenty-first century, *AEA Papers and Proceedings* **110**: 1–24.

Bertrand, M., Goldin, C. and Katz, L. F. (2010). Dynamics of the gender gap for young professionals in the financial and corporate sectors, *American Economic Journal: Applied Economics* **2**(3): 228–55.

Biasi, B. and Sarsons, H. (2021). Flexible Wages, Bargaining, and the Gender Gap, *The Quarterly Journal of Economics* **137**(1): 215–266.

Blau, F. D. and Devaro, J. (2007). New evidence on gender differences in promotion rates: An empirical analysis of a sample of new hires, *Industrial Relations: A Journal of Economy and Society* **46**(3): 511–550.

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

Bolotnyy, V. and Emanuel, N. (2022). Why do women earn less than men? evidence from bus and train operators, *Journal of Labor Economics* **40**(2): 283–323.

Bosquet, C., Combes, P.-P. and Garcia-Penalosa, C. (2019). Gender and promotions: Evidence from academic economists in france, *The Scandinavian Journal of Economics* **121**(3): 1020–1053.

Card, D., DellaVigna, S., Funk, P. and Iriberri, N. (2019). Are Referees and Editors in Economics Gender Neutral?*, *The Quarterly Journal of Economics* **135**(1): 269–327.

Cederlöf, J., Söderström, M. and Vikström, J. (2021). What makes a good caseworker?, *Technical report*, IFAU Working Paper 2021:9.

Cobb-Clark, D. (1998). Getting Ahead: the Determinants of and Payoffs to Internal Promotion for Young Men and Women, *CEPR Discussion Papers 395*, Centre for Economic Policy Research, Research School of Economics, Australian National University.

Cook, C., Diamond, R., Hall, J. V., List, J. A. and Oyer, P. (2020). The Gender Earnings Gap in the Gig Economy: Evidence from over a Million Rideshare Drivers, *The Review of Economic Studies* **88**(5): 2210–2238.

Correll, S., Benard, S. and Paik, I. (2007). Getting a job: Is there a motherhood penalty?, *American Journal of Sociology* **112**(5): 1297–1338.

Cortes, P. and Pan, J. (2020). Children and the remaining gender gaps in the labor market, *Working Paper 27980*, National Bureau of Economic Research.

Crepon, B., Duflo, E., Gurgand, M., Rathelot, R. and Zamora, P. (2013). Do Labor Market Policies have Displacement Effects? Evidence from a Clustered Randomized Experiment, *The Quarterly Journal of Economics* **128**(2): 531–580.

De Pater, I. E., Van Vianen, A. E. M. and Bechtoldt, M. N. (2010). Gender differences in job challenge: A matter of task allocation, *Gender, Work & Organization* **17**(4): 433–453.

Fluchtmann, J., Glenny, A. M., Harmon, N. A. and Maibom, J. (2024). The gender application gap: Do men and women apply for the same jobs?, *American Economic Journal: Economic Policy* **16**(2): 182â219.

Gallen, Y. (2023, forthcoming). Motherhood and the gender productivity gap, *Journal of the European Economic Association* .

Ginther, D. and Kahn, S. (2006). Women's careers in academic social science. progress, pitfalls and plateaus. Working Paper.

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

Goldin, C. and Katz, L. F. (2016). A most egalitarian profession: Pharmacy and the evolution of a family-friendly occupation, *Journal of Labor Economics* **34**(3): 705–746.

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

Graversen, B. K. and van Ours, J. C. (2008). How to help unemployed find jobs quickly: Experimental evidence from a mandatory activation program, *Journal of Public Economics* **92**(10): 2020–2035.

Haegele, I. (2024). The broken rung: gender and the leadership gap. Working Paper.

Hellerstein, J. K., Neumark, D. and Troske, K. R. (1999). Wages, productivity, and worker characteristics: Evidence from plant level production functions and wage equations, *Journal of Labor Economics* **17**(3): 409–446.

Hengel, E. (2022). Publishing While Female: are Women Held to Higher Standards? Evidence from Peer Review, *The Economic Journal* **132**(648): 2951–2991.

Hengel, E. and Moon, E. (2020). Gender and quality at top economics journals, *Working Papers 202001*, University of Liverpool, Department of Economics.

Hospido, L., Laeven, L. and Lamo, A. (2022). The Gender Promotion Gap: Evidence from Central Banking, *The Review of Economics and Statistics* **104**(5): 981–996.

Humlum, A., Munch, J. R. and Rasmussen, M. (2023). What works for the unemployed? evidence from quasi-random caseworker assignments, *Technical report*, IZA Discussion Paper 16033.

Jagsi, R., Guancial, E. A., Worobey, C. C., Henault, L., Chang, Y., Starr, R., Tarbell, N. J. and Hylek, E. (2006). The "gender gap" in authorship of academic medical literature–a 35-year perspective., *The New England Journal of Medicine* .

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

Lazear, E. P. and Rosen, S. (1990). Male-female wage differentials in job ladders, *Journal of Labor Economics* **8**(1, Part 2): S106–S123.

Lundborg, P., Plug, E. and Würtz Rasmussen, A. (2017). Can women have children and a career? iv evidence from ivf treatments, *American Economic Review* **107**(6): 1611–37.

Samarbetsrådet (2008). Stöd för lokal lönebildning. råd till lokala parter vid tillämpningen av rals 2007-2010, *Technical report*, Samarbetsrådet för det statliga avtalsområdet.

Sarsons, H., Gërxhani, K., Reuben, E. and Schram, A. (2021). Gender differences in recognition for group work, *Journal of Political Economy* **129**(1): 101–147.

Save-Söderbergh, J. (2019). Gender gaps in salary negotiations: Salary requests and starting salaries in the field, *Journal of Economic Behavior Organization* **161**: 35–51.

Schiprowski, A. (2020). The role of caseworkers in unemployment insurance: Evidence from unplanned absences, *Journal of Labor Economics* **38**(4): 1189–1225.

Zeltzer, D. (2020). Gender homophily in referral networks: Consequences for the medicare physician earnings gap, *American Economic Journal: Applied Economics* **12**(2): 169â97.

Figures

Figure 1: Job seeker allocation to caseworkers based on day of birth

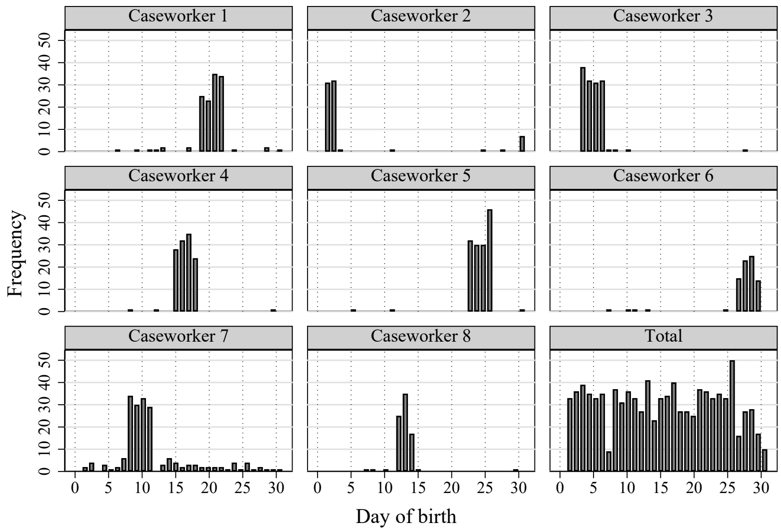

(a) Office with date of birth rule

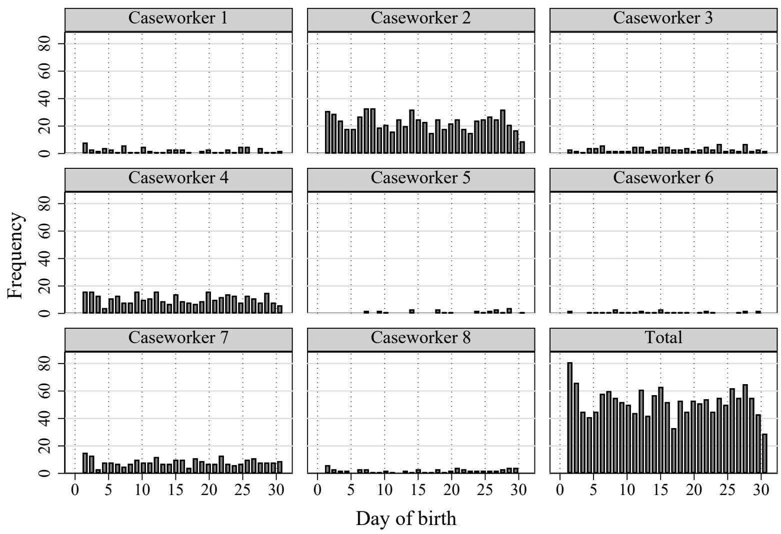

(b) Office without date of birth rule

Notes: The figure illustrates the allocation of job seekers to caseworkers based on the job seekers' birth dates in an office with a date-of-birth rule and with an F-stat=814.39 (panel a) and in an office without the date-of-birth rule with F-stat=2.29 (panel b). See text for details on the F-statistics.

Figure 2: Distribution of starting wages and changes in wages

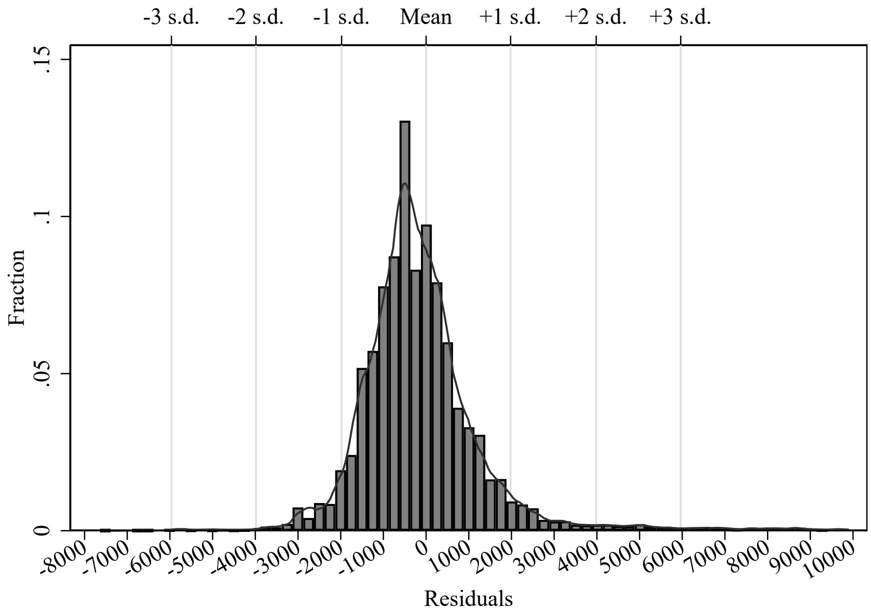

(a) Variation in starting salary of the caseworkers

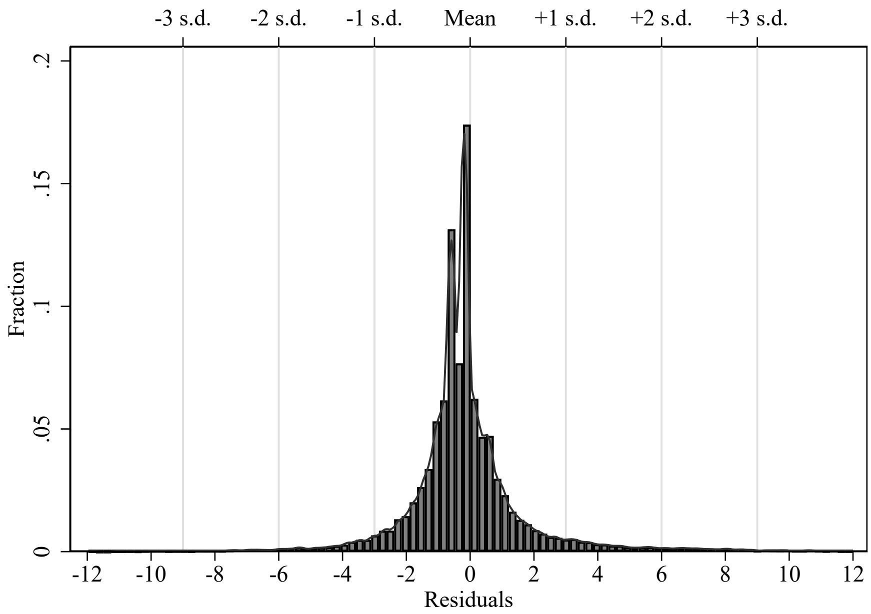

(b) Variation in changes in the wages of the caseworkers

Notes: The figures illustrates the distribution of the residuals of starting salaries (panel a) and percentage changes in salaries (panel b) after taking out the year fixed effects.

Tables

Table 1: Descriptive statistics by type of office

	Panel A: Caseworkers' characteristics					
	Random offices		Non-random offices			
	Mean	SD	Mean	SD	Diff	p-value
Age	46.47	10.40	46.57	10.27	0.10	0.40
Female	0.64	0.48	0.64	0.48	0.00	0.91
Immigrant	0.16	0.37	0.16	0.36	-0.00	0.30
Married	0.61	0.49	0.62	0.49	0.01	0.39
Number of children below 16	0.63	0.93	0.64	0.94	0.01	0.49
University Degree	0.69	0.46	0.67	0.47	-0.03***	0.00
Secondary Degree	0.28	0.45	0.31	0.46	0.03***	0.00
Business degree	0.29	0.45	0.30	0.46	0.01	0.06
Social degree	0.18	0.38	0.17	0.38	-0.01	0.13
Log(earnings)	12.43	0.24	12.43	0.23	0.00	0.47
Log(wage)	10.01	0.10	10.00	0.10	-0.01***	0.00
Experience						
0-2 years	0.10	0.31	0.09	0.29	-0.01**	0.00
2-4 years	0.18	0.38	0.19	0.39	0.01	0.17
4-6 years	0.11	0.32	0.11	0.32	0.00	0.59
6-8 years	0.10	0.30	0.10	0.29	-0.01	0.11
8-10 years	0.13	0.34	0.12	0.33	-0.01	0.11
+10 years	0.37	0.48	0.39	0.49	0.01**	0.01
# of observations	12,151		19,444		31,595	
# of observations (unique)	5,140		6,880		8,719	

	Panel B: job seekers' characteristics					
	Random offices		Non-random offices			
	Mean	SD	Mean	SD	Diff	p-value
Age at inflow	31.87	12.35	32.44	12.55	0.57***	0.00
Female	0.47	0.50	0.46	0.50	-0.01***	0.00
Married	0.23	0.42	0.24	0.43	0.01***	0.00
At least one child	0.35	0.48	0.36	0.48	0.01***	0.00
Born outside Sweden	0.24	0.43	0.25	0.43	0.01***	0.00
= 1 if reg. as disabled	0.04	0.20	0.05	0.21	0.01***	0.00
= 1 Eligiable for UI	0.65	0.48	0.66	0.47	0.01***	0.00
Earnings 1 year before	98885	120911	100880	126457	1994***	0.00
High-school	0.54	0.50	0.54	0.50	-0.01***	0.00
University	0.26	0.44	0.25	0.43	-0.02***	0.00
Unemployment duration	264.09	422.03	271.68	442.03	7.59***	0.00
# of observations	1,317,963		960,718		2,278,681	
# of observations (unique)	984,404		741,072		1,589,249	

	Panel C: Offices' characteristics					
	Random offices		Non-random offices			
	Mean	SD	Mean	SD	Diff	p-value
Number of caseworkers	16.27	11.59	10.68	10.63	-5.59***	0.00
Number of job seekers	1819.97	1566.72	994.66	1098.50	-825.31***	0.00
# of observations	818		1,102		1,920	
# of observations (unique)	179		210		257	

Notes: The table illustrates means, standard deviations, and t-test results of mean comparisons of caseworker, job seeker, and office characteristics by type of office. The sample includes job seekers registered at PES during 2003-2014. Earnings are measured in Swedish SEK (SEK 100 corresponds to EURO 8 as of June 2024).Standard errors in parentheses. * p<0.1, ** p<0.05, *** p<0.01.

Table 2: Descriptive statistics by caseworkers' gender

	Panel A: Caseworkers					
	Female caseworkers		Male caseworkers			
	Mean	SD	Mean	SD	Diff	p-value
Age	45.55	10.50	46.45	10.56	0.91**	0.00
Immigrant	0.15	0.36	0.16	0.37	0.01	0.17
Secondary Degree	0.31	0.46	0.27	0.44	-0.04**	0.00
University Degree	0.66	0.47	0.70	0.46	0.04*	0.01
Married	0.61	0.49	0.57	0.50	-0.04*	0.01
Number of children below 16	0.69	0.95	0.61	0.92	-0.08**	0.00
Business degree	0.33	0.47	0.23	0.42	-0.10***	0.00
Social degree	0.16	0.37	0.18	0.38	0.01	0.25
Log(earnings)	12.39	0.24	12.47	0.17	0.08***	0.00
Log(wage)	10.00	0.09	10.00	0.09	-0.00	0.43
Experience						
0-2 years	0.10	0.29	0.10	0.30	0.00	0.58
4-6 years	0.10	0.30	0.11	0.32	0.01	0.24
6-8 years	0.09	0.28	0.07	0.26	-0.01	0.07
8-10 years	0.12	0.32	0.10	0.30	-0.02	0.08
+10 years	0.42	0.49	0.42	0.49	0.01	0.71
# of observations	3,365		1,807		5,172	
# of observations (unique)	1,752		954		2,705	
	Panel B: job seekers					
	Female caseworkers		Male caseworkers			
	Mean	SD	Mean	SD	Diff	p-value
Age	31.71	12.34	32.14	12.37	0.43***	0.00
Female	0.48	0.50	0.45	0.50	-0.03***	0.00
Married	0.23	0.42	0.23	0.42	0.01***	0.00
Having children	0.35	0.48	0.34	0.48	-0.00***	0.01
Immigrant	0.24	0.43	0.24	0.42	-0.00***	0.00
Disabled	0.04	0.20	0.04	0.19	-0.00***	0.00
Eligible for UI	0.65	0.48	0.66	0.47	0.02***	0.00
Earnings, $t-1$	96,469	120,113	102,800	122,091	6,330***	0.00
Secondary degree	0.54	0.50	0.54	0.50	0.00***	0.00
University degree	0.27	0.44	0.26	0.44	-0.01***	0.00
Unemployment duration	261.82	419.77	267.78	425.65	5.96***	0.00
# of observations	815,011		502,952		1,317,963	
# of observations (unique)	660,168		425,138		984,404	

Notes: The table illustrates means, standard deviations, and t-test results of mean comparisons of caseworker and job seekers' characteristics by caseworker gender in date-of-birth rule offices. The sample includes job seekers registered at PES during 2003-2014. Earnings are measured in Swedish SEK (SEK 100 corresponds to EURO 8 as of June 2024). Standard errors in parentheses. * p<0.1, ** p<0.05, *** p<0.01.

Table 3: Balancing test in non-date-of-birth-rule offices

	Dependent variables: job seekers' characteristics							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	Age	Education	Welfare benefits 1 year before	Earnings 1 year before	Female	Born outside Sweden	= 1 if disabled	Unemployed 24 months before
Caseworker's age	0.18818***	0.00546***	17.69384***	841.09986***	-0.00115***	0.00134***	0.00107***	0.00119***
	(0.01424)	(0.00070)	(4.05833)	(86.60513)	(0.00030)	(0.00024)	(0.00012)	(0.00010)
Caseworker's gender (=1 if female)	-1.27070***	-0.02931	94.02833	-12363.54031***	0.05899***	-0.00285	0.00292	-0.00653**
	(0.32158)	(0.01550)	(85.32112)	(1916.13699)	(0.00652)	(0.00501)	(0.00273)	(0.00218)
Caseworker's education	0.00991	0.01407	99.42217	-2322.34808*	0.01739***	0.00861***	0.00429**	-0.00209
	(0.16400)	(0.00784)	(55.80317)	(1033.68825)	(0.00337)	(0.00252)	(0.00134)	(0.00115)
Caseworker's experience	0.27008***	0.00641***	9.70791	1473.76543***	-0.00369***	0.00135**	0.00154***	0.00201***
	(0.03059)	(0.00152)	(9.63809)	(201.45116)	(0.00074)	(0.00048)	(0.00028)	(0.00023)
# of observations	960,718	960,718	960,718	960,718	960,718	960,718	960,718	960,718

Notes: The table presents OLS estimates of the relationship between various job seeker characteristics and caseworker characteristics in offices without a date-of-birth rule. All models include fixed effects for interactions between year, office, and job seeker age under 25. Standard errors, clustered at the caseworker level, are shown in parentheses. * p<0.1, ** p<0.05, *** p<0.01.

Table 4: Balancing test in date-of-birth-rule offices

	Dependent variables: job seekers' characteristics							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	Age	Education	Welfare benefits 1 year before	Earnings 1 year before	Female	Born outside Sweden	= 1 if disabled	Unemployed 24 months before
Predicted caseworker age	0.00099	0.00003	-3.02003*	9.59215	-0.00007	-0.00008	-0.00003	0.00001
	(0.00093)	(0.00012)	(1.38185)	(13.25088)	(0.00005)	(0.00006)	(0.00002)	(0.00003)
Predicted caseworker gender (=1 if female)	-0.00605	0.00200	-31.01948	-428.96232	0.00510***	-0.00065	0.00065	0.00074
	(0.01778)	(0.00227)	(26.56603)	(252.71025)	(0.00106)	(0.00119)	(0.00042)	(0.00070)
Predicted caseworker education	0.00210	0.00061	10.70709	-173.90337	-0.00014	0.00054	-0.00005	0.00010
	(0.00959)	(0.00117)	(15.06751)	(132.77378)	(0.00055)	(0.00066)	(0.00022)	(0.00036)
Predicted caseworker experience	0.00192	0.00009	-5.60934*	9.11548	-0.00010	0.00000	0.00008	0.00004
	(0.00197)	(0.00023)	(2.79157)	(27.64615)	(0.00011)	(0.00012)	(0.00005)	(0.00008)
# of observations	1,317,963	1,317,963	1,317,963	1,317,963	1,317,963	1,317,963	1,317,963	1,317,963

Notes: The table presents OLS estimates of the relationship between job seeker characteristics and predicted caseworker characteristics in offices with a date-of-birth rule. All models include fixed effects for interactions between year, office, and job seeker age under 25. Standard errors, clustered at the predicted caseworker level, are shown in parentheses. * p<0.1, ** p<0.05, *** p<0.01.

Table 5: Determinants of productivity

	Leave unemployment within 180 days				
	(1)	(2)	(3)	(4)	(5)
Female	0.0072**	0.0071**	0.0073**	0.0071**	0.0074**
	(0.0028)	(0.0028)	(0.0029)	(0.0028)	(0.0029)
University degree		0.0030	0.0034	0.0038	0.0044
		(0.0087)	(0.0089)	(0.0086)	(0.0088)
Secondary degree		0.0075	0.0076	0.0074	0.0077
		(0.0089)	(0.0091)	(0.0089)	(0.0090)
Business degree		-0.0003	0.0000	-0.0009	-0.0006
		(0.0034)	(0.0034)	(0.0034)	(0.0034)
Social degree		0.0001	-0.0003	0.0002	-0.0001
		(0.0041)	(0.0041)	(0.0041)	(0.0041)
At least one child < 4 years			-0.0005		0.0004
			(0.0045)		(0.0045)
1 Child			0.0022		0.0018
			(0.0038)		(0.0039)
2 Children			-0.0068		-0.0080*
			(0.0042)		(0.0043)
+3 Children			-0.0122*		-0.0128**
			(0.0063)		(0.0063)
Experience					
2-4 years				-0.0054	-0.0054
				(0.0053)	(0.0053)
4-6 years				-0.0032	-0.0018
				(0.0058)	(0.0059)
6-8 years				0.0025	0.0041
				(0.0060)	(0.0061)
8-10 years				0.0023	0.0037
				(0.0060)	(0.0061)
+10 years				0.0069	0.0080
				(0.0060)	(0.0061)
Immigrant	-0.0060	-0.0059	-0.0058	-0.0052	-0.0050
	(0.0041)	(0.0041)	(0.0041)	(0.0041)	(0.0041)
Mean	0.63	0.63	0.63	0.63	0.63
# of observations (job seekers)	1,317,963	1,317,963	1,317,963	1,317,963	1,317,963
# of observations (caseworkers)	5,140	5,140	5,140	5,140	5,140
Age Fixed Effect	Yes	Yes	Yes	Yes	Yes
Office×Year×Age<25 Fixed Effect	Yes	Yes	Yes	Yes	Yes

Notes: The table presents the estimates of the relationship between caseworker characteristics and the likelihood of job seekers exiting unemployment within 180 days in offices with a date-of-birth rule. All actual caseworker characteristics are instrumented using predicted caseworker characteristics. All models include fixed effects for interactions between year, office, and job seeker age under 25. Standard errors, clustered at the actual caseworker level, are shown in parentheses. * p<0.1, ** p<0.05, *** p<0.01.

Table 6: Other dimensions of caseworker productivity: job-quality and long-run earnings of job seekers

	(1) Exit to employment	(2) First job duration (=1 if at least 2 years)	(3) Earnings, $t+5$	(4) Employment status, $t+5$
Female	-0.000	-1.450	-1721.979	0.000
	(0.003)	(7.030)	(1170.005)	(0.004)
Mean	0.6	666.2	159109.3	0.8
# of observations	1,317,963	279,866	279,866	279,866
Office×Year×Age<25 Fixed Effect	Yes	Yes	Yes	Yes

Notes: The table presents estimates of the relationship between job seekers' outcomes and caseworkers' gender in offices with a date-of-birth rule. Actual caseworker gender and age are instrumented using predicted caseworker gender and age. In addition to gender and age, all models include fixed effects for interactions between year, office, and job seeker age under 25. Column 1 includes all job seekers registered in DOB offices, while the remaining columns include only those who left PES with a job. Earnings are measured in Swedish SEK (SEK 100 is equivalent to EUR 8 as of June 2024). Standard errors, clustered at the caseworker level, are shown in parentheses. * p<0.1, ** p<0.05, *** p<0.01.

Table 7: Wages and different definitions of productivity

	Log(wage)						
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Productivity	-0.000	0.000					
	(0.000)	(0.002)					
Productivityˆ2		0.000					
		(0.000)					
Productivity at t-1			-0.000				
			(0.001)				
Productivity decile				0.001***			
				(0.000)			
Productivity decile at t-1					0.002***		
					(0.000)		
Productivity (=1 if above median)						0.006**	
						(0.002)	
Productivity at t-1 (=1 if above median)							0.011***
							(0.003)
Mean wage	22,056.4	22,056.4	22,501.6	22,056.4	22,501.6	22,056.4	22,501.6
Mean log(wage)	10.0	10.0	10.0	10.0	10.0	10.0	10.0
# of observations	5,172	5,172	2,523	5,172	2,523	5,172	2,523
# of observations (unique)	2,705	2,705	1,397	2,705	1,397	2,705	1,397
Age Fixed Effect	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Office × Year Fixed Effect	Yes	Yes	Yes	Yes	Yes	Yes	Yes

Notes: The table presents estimates of the relationship between caseworkers' wages and productivity, using various specifications of the productivity variable within the sample of caseworkers in date-of-birth rule offices. All models include interactions between year and office fixed effects, as well as age fixed effects. Robust standard errors are shown in parentheses. * p<0.1, ** p<0.05, *** p<0.01.

Table 8: Gender wage gap among caseworkers

	Log(Wage)			
	All caseworkers in DOB offices	Caseworkers in DOB offices doing the same task	Caseworkers in DOB offices doing the same task	Caseworkers in DOB offices doing the same task
	(1)	(2)	(3)	(4)
Female	-0.001	0.001	-0.001	-0.001
	(0.001)	(0.001)	(0.002)	(0.002)
University Degree	0.019***	0.017***	0.021***	0.021***
	(0.004)	(0.005)	(0.006)	(0.006)
Secondary Degree	0.003	0.002	0.011*	0.011*
	(0.004)	(0.005)	(0.006)	(0.006)
Immigrant	-0.012***	-0.011***	-0.011***	-0.011***
	(0.001)	(0.002)	(0.002)	(0.002)
Business degree	-0.003**	-0.005***	-0.005**	-0.005**
	(0.001)	(0.002)	(0.002)	(0.002)
Social degree	-0.002	-0.001	-0.001	-0.001
	(0.001)	(0.002)	(0.002)	(0.002)
Experience				
2-4 years	0.016***	0.015***	0.015***	0.015***
	(0.001)	(0.002)	(0.003)	(0.003)
4-6 years	0.052***	0.053***	0.051***	0.051***
	(0.002)	(0.002)	(0.003)	(0.003)
6-8 years	0.080***	0.076***	0.073***	0.073***
	(0.002)	(0.003)	(0.004)	(0.004)
8-10 years	0.094***	0.093***	0.092***	0.092***
	(0.002)	(0.003)	(0.004)	(0.004)
+10 years	0.116***	0.114***	0.110***	0.110***
	(0.002)	(0.002)	(0.003)	(0.003)
At least one child < 4 years	0.004***	0.006***	0.009***	0.009***
	(0.001)	(0.002)	(0.003)	(0.003)
1 Child	-0.002	-0.004**	-0.003	-0.003
	(0.001)	(0.002)	(0.003)	(0.003)
2 Children	-0.001	-0.002	0.001	0.001
	(0.001)	(0.002)	(0.003)	(0.003)
3+ Children	0.002	-0.000	0.002	0.002
	(0.002)	(0.003)	(0.004)	(0.004)
Productivity (=1 if above median)				0.004*
				(0.002)
Mean log(wage)	10.0	10.0	10.0	10.0
Mean wage	22,370.9	22,190.5	22,056.4	22,056.4
# of observations	20,605	10,237	5,172	5,172
# of observations (unique)	7,231	4,481	2,705	2,705
Age Fixed Effect	Yes	Yes	Yes	Yes
Office × Year Fixed Effect	Yes	Yes	Yes	Yes

Notes: The table presents estimates of the relationship between caseworkers' wages, caseworker gender, and productivity across different samples of caseworkers. All models include interactions between year and office fixed effects. Wages are measured in Swedish SEK (with SEK 100 equivalent to EUR 8 as of June 2024). Robust standard errors are shown in parentheses. * p<0.1, ** p<0.05, *** p<0.01.

Table 9: Gender log earnings gap among caseworkers

	Number of job seekers	Log(earnings)				
	(1)	(2)	(3)	(4)	(5)	(6)
Female	-14.617***	-0.083***	-0.081***	-0.039***	-0.039***	-0.043***
	(2.476)	(0.006)	(0.005)	(0.005)	(0.005)	(0.004)
University Degree	1.167	0.004	-0.024	-0.013	-0.013	-0.015
	(6.988)	(0.021)	(0.018)	(0.014)	(0.014)	(0.014)
Secondary Degree	8.632	0.006	-0.009	-0.004	-0.004	-0.008
	(7.456)	(0.021)	(0.019)	(0.014)	(0.014)	(0.014)
Business degree	-2.901	-0.011	-0.005	-0.002	-0.002	-0.001
	(2.951)	(0.008)	(0.007)	(0.005)	(0.005)	(0.005)
Social degree	-1.412	-0.005	-0.004	-0.002	-0.002	0.000
	(3.734)	(0.009)	(0.008)	(0.007)	(0.007)	(0.006)
Immigrant	1.010	0.003	0.019**	0.008	0.008	0.007
	(3.275)	(0.008)	(0.008)	(0.006)	(0.006)	(0.006)
Experience						
2-4 years	34.133***	0.035***	0.013	0.033***	0.033***	0.035***
	(4.631)	(0.013)	(0.013)	(0.011)	(0.011)	(0.011)
4-6 years	34.204***	0.047***	-0.025*	0.017	0.017	0.023**
	(5.260)	(0.014)	(0.014)	(0.012)	(0.012)	(0.012)
6-8 years	34.925***	0.061***	-0.040**	0.015	0.015	0.018
	(5.658)	(0.016)	(0.016)	(0.013)	(0.013)	(0.013)
8-10 years	34.154***	0.070***	-0.057***	0.003	0.003	-0.002
	(5.687)	(0.016)	(0.017)	(0.015)	(0.015)	(0.014)
+10 years	41.089***	0.082***	-0.069***	0.003	0.003	-0.003
	(5.272)	(0.015)	(0.016)	(0.013)	(0.013)	(0.013)
Log(wage)			1.378***	1.155***	1.156***	1.197***
			(0.062)	(0.046)	(0.046)	(0.047)
Hours worked (per week)				0.026***	0.026***	0.025***
				(0.001)	(0.001)	(0.001)
Productivity (=1 if above median)					-0.004	-0.005
					(0.005)	(0.005)
At least one child < 4 years						-0.092***
						(0.009)
1 Child						-0.018***
						(0.006)
2 Children						-0.039***
						(0.008)
3+ Children						-0.061***
						(0.012)
Mean	148.9	12.4	12.4	12.4	12.4	12.4
# of observations	5,172	5,172	5,172	5,172	5,172	5,172
Age Fixed Effect	Yes	Yes	Yes	Yes	Yes	Yes
Office× Year Fixed Effect	Yes	Yes	Yes	Yes	Yes	Yes

Notes: The table presents estimates of the relationship of the caseworkers' annual earnings and number of unique job seekers in a year with their characteristics and productivity in date-of-birth rule offices. The sample includes only caseworkers for whom a productivity measure can be calculated. All models incorporate interactions between year and office fixed effects. Earnings are measured in Swedish SEK (with SEK 100 equivalent to EUR 8 as of June 2024). Robust standard errors are shown in parentheses. * p<0.1, ** p<0.05, *** p<0.01.

Table 10: Gender promotion gap among caseworkers

	Promotion				
	All caseworkers	All caseworkers	All caseworkers	Caseworkers without children	Caseworkers with children
	(1)	(2)	(3)	(4)	(5)
Female	-0.017***	-0.017***	-0.015***	-0.011	-0.027**
	(0.005)	(0.005)	(0.005)	(0.007)	(0.011)
University Degree	0.024**	0.024**	0.024**	0.031*	0.017
	(0.010)	(0.010)	(0.010)	(0.016)	(0.012)
Secondary Degree	0.040***	0.040***	0.040***	0.054***	0.013
	(0.011)	(0.011)	(0.011)	(0.017)	(0.013)
Business degree	-0.007	-0.007	-0.006	-0.002	-0.004
	(0.006)	(0.006)	(0.006)	(0.008)	(0.011)
Social degree	-0.009	-0.009	-0.009	0.003	-0.009
	(0.006)	(0.006)	(0.006)	(0.009)	(0.009)
Immigrant	0.018***	0.018***	0.018***	0.039***	-0.010
	(0.007)	(0.007)	(0.007)	(0.011)	(0.009)
Experience					
2-4 years	0.006	0.006	0.007	0.016	-0.000
	(0.006)	(0.006)	(0.006)	(0.010)	(0.011)
4-6 years	0.020**	0.020**	0.021**	0.026*	0.023
	(0.008)	(0.008)	(0.008)	(0.014)	(0.015)
6-8 years	0.031***	0.031***	0.032***	0.048**	0.020
	(0.010)	(0.010)	(0.010)	(0.019)	(0.015)
8-10 years	0.034***	0.034***	0.036***	0.033**	0.045**
	(0.010)	(0.010)	(0.010)	(0.015)	(0.018)
+10 years	0.035***	0.034***	0.036***	0.052***	0.011
	(0.008)	(0.008)	(0.008)	(0.013)	(0.014)
Productivity		0.003	0.002	0.007	-0.001
		(0.005)	(0.005)	(0.007)	(0.009)
Hours worked (per week)			0.001*	0.002***	-0.001
			(0.001)	(0.001)	(0.001)
Mean	0.022	0.022	0.022	0.024	0.018
# of observations	5,310	5,310	5,310	3,273	2,037
Office× Year Fixed Effect	Yes	Yes	Yes	Yes	Yes

Notes: The table presents estimates of the relationship between caseworker promotions to senior positions and caseworker characteristics, productivity, and hours worked. The sample consists of caseworkers in date-of-birth rule offices. Columns 4 and 5 separate the sample into caseworkers with and without children. All models include interactions between year and office fixed effects. Robust standard errors are shown in parentheses. * p<0.1, ** p<0.05, *** p<0.01.

Appendix

Figures

Figure A1: Distribution of F-statistics for testing the presence of date-of-birth rule offices

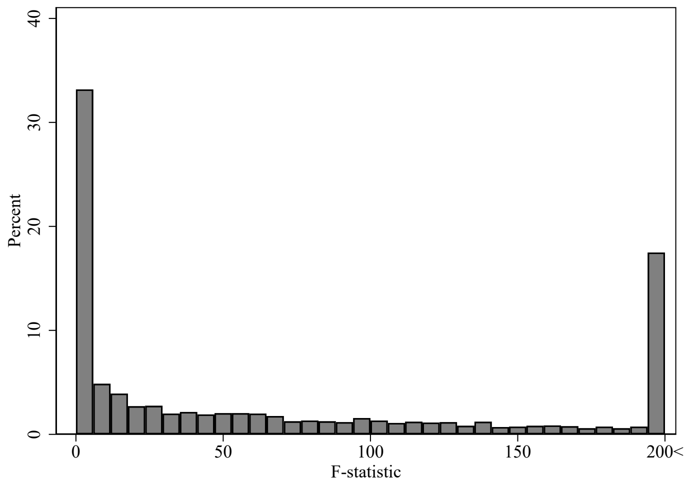

Notes: The figure shows the F-statistics from regressing job seekers' day of birth on caseworker dummies within each office and year. An F-statistic above 100 indicates the presence of a date-of-birth rule in the allocation of job seekers.

Figure A2: Prevalence of date-of-birth rule usage in offices over time

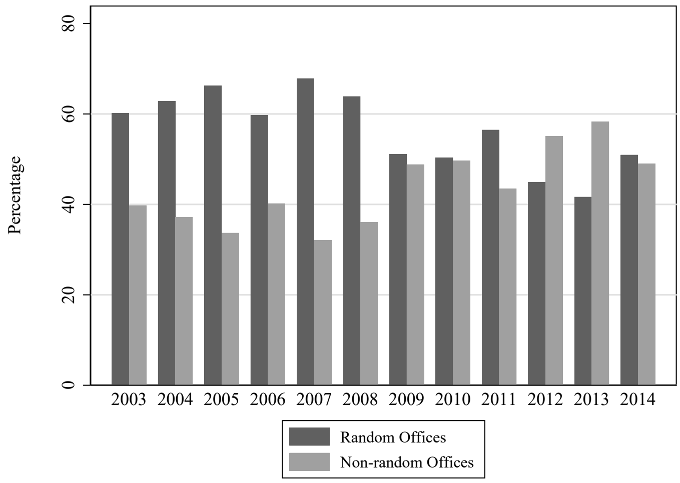

Notes: The figure shows the number of date-of-birth allocation offices (with F-statistics above 100) and non-date-of-birth allocation offices (with F-statistics below 20) over time.

Figure A3: Productivity deciles and log wages

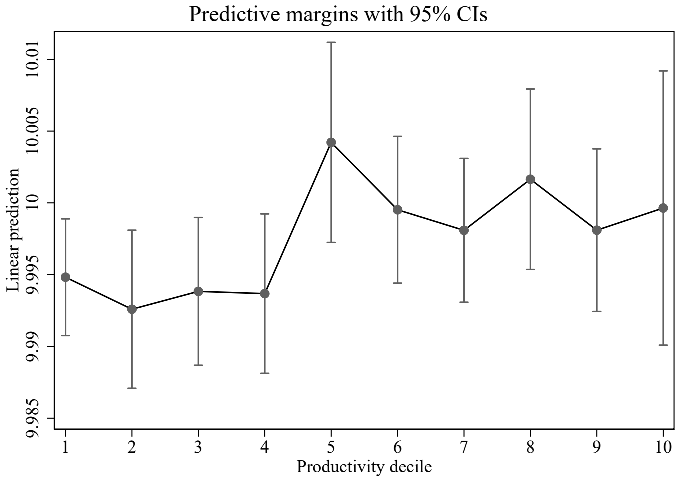

Notes: The figure plots estimates of the relationship between caseworker productivity deciles (represented as dummies) and caseworkers' log wages. The analysis controls for age fixed effects and interactions between year and office fixed effects, using the sample of caseworkers in date-of-birth rule offices.

Tables

Table A1: First stage results

	Dependent variables: Actual caseworker characteristics					
	(1) Age	(2) Female	(3) Immigrant	(4) University degree	(5) Secondary degree	(6) Experience
Predicted caseworker age	0.4623***	-0.0001	0.0002	0.0004	-0.0004	0.0002
	(0.0107)	(0.0004)	(0.0003)	(0.0004)	(0.0004)	(0.0038)
Predicted caseworker female	-0.1360	0.4653***	-0.0037	0.0034	-0.0014	-0.0224
	(0.1377)	(0.0085)	(0.0052)	(0.0067)	(0.0065)	(0.0685)
Predicted caseworker immigrant	0.0235	-0.0125	0.4825***	0.0056	-0.0050	0.0527
	(0.1948)	(0.0104)	(0.0141)	(0.0089)	(0.0086)	(0.0942)
Predicted caseworker university degree	-0.0500	0.0236	-0.0028	0.4890***	0.0004	0.1930
	(0.5845)	(0.0268)	(0.0182)	(0.0318)	(0.0202)	(0.2710)
Predicted caseworker secondary degree	-0.2339	0.0204	-0.0054	0.0073	0.4821***	0.1578
	(0.5840)	(0.0270)	(0.0183)	(0.0324)	(0.0219)	(0.2734)
Predicted caseworker experience	-0.0358**	0.0002	0.0006	0.0013	-0.0008	0.4331***
	(0.0173)	(0.0009)	(0.0007)	(0.0008)	(0.0008)	(0.0108)
# of observations	1,317,963	1,317,963	1,317,963	1,317,963	1,317,963	1,317,963

Notes: The table presents first-stage estimates of the relationship between actual and predicted caseworker characteristics. The sample includes caseworkers in date-of-birth rule offices. All models incorporate interactions between year, office, and age (below 25) fixed effects. Standard errors, clustered at the caseworker level, are shown in parentheses. * p<0.1, ** p<0.05, *** p<0.01.

Table A2: Gender gaps in alternative measures of productivity

	(1) Leave unemployment within 30 days	(2) Leave unemployment within 90 days	(3) log (days of unemployment
Female	0.0022	0.0067**	-0.0203**
	(0.0018)	(0.0029)	(0.0080)
University degree	-0.0065	-0.0046	0.0121
	(0.0052)	(0.0095)	(0.0256)
Secondary degree	-0.0030	-0.0005	-0.0051
	(0.0054)	(0.0097)	(0.0264)
Business degree	-0.0012	-0.0022	0.0076
	(0.0021)	(0.0035)	(0.0095)
Social degree	-0.0010	0.0006	-0.0002
	(0.0024)	(0.0042)	(0.0113)
Any child < 4 years old	-0.0003	0.0019	-0.0052
	(0.0029)	(0.0047)	(0.0126)
1 Child	0.0004	0.0005	0.0007
	(0.0025)	(0.0040)	(0.0108)
2 Children	-0.0037	-0.0075*	0.0239**
	(0.0027)	(0.0044)	(0.0120)
+3 Children	-0.0111***	-0.0083	0.0425**
	(0.0041)	(0.0067)	(0.0178)
Experience			
2-4 years	-0.0021	-0.0111**	0.0193
	(0.0033)	(0.0054)	(0.0142)
4-6 years	0.0004	-0.0055	0.0147
	(0.0036)	(0.0060)	(0.0159)
6-8 years	0.0035	0.0024	-0.0030
	(0.0039)	(0.0064)	(0.0168)
8-10 years	0.0034	-0.0010	-0.0001
	(0.0038)	(0.0063)	(0.0165)
+10 years	-0.0008	0.0034	-0.0059
	(0.0036)	(0.0061)	(0.0165)
Immigrant	-0.0023	-0.0030	0.0139
	(0.0023)	(0.0041)	(0.0112)
Mean	0.13	0.41	4.79
# of observations	1,317,963	1,317,963	1,317,963
Age Fixed Effects	Yes	Yes	Yes
Office×Year× Age<25 Fixed Effect	Yes	Yes	Yes

Notes: The table presents estimates of the relationship between caseworker characteristics and various job seeker outcomes in date-of-birth rule offices. All actual caseworker characteristics are instrumented using predicted caseworker characteristics. All models include fixed effects for interactions between year, office, and job seeker age under 25. Standard errors, clustered at the actual caseworker level, are shown in parentheses. * p<0.1, ** p<0.05, *** p<0.01.

Table A3: Gender, wages, and productivity - alternative specifications of productivity

	(1) Log(wage)	(2) Log(wage)	(3) Log(wage)	(4) Percentage change in wages	(5) Log(wage)	(6) Log(wage)
Female	-0.001	-0.001	0.001	0.088	-0.001	0.000
	(0.002)	(0.002)	(0.003)	(0.117)	(0.002)	(0.003)
Productivity	-0.000	0.001				
	(0.000)	(0.001)				
Productivityˆ2		0.000				
		(0.000)				
Productivity at $t-1$			0.000	0.008		
			(0.000)	(0.011)		
Productivity decile					0.001***	
					(0.000)	
Productivity decile at $t-1$						0.002***
						(0.000)
Mean	10.0	10.0	10.0	2.5	10.0	10.0
# of observations	5,172	5,172	2,523	2,523	5,172	2,523
Controls	Yes	Yes	Yes	Yes	Yes	Yes
Age Fixed Effect	Yes	Yes	Yes	Yes	Yes	Yes
Office × Year Fixed Effect	Yes	Yes	Yes	Yes	Yes	Yes

Notes: The table presents estimates of the relationship between caseworker log wages and caseworker gender and productivity, using a sample of caseworkers from date-of-birth rule offices. Productivity deciles are calculated within each year and office. All models include the control variables discussed in section 4.3 and interactions between year and office fixed effects, as well as age fixed effects. Wages are measured in Swedish SEK (with SEK 100 equivalent to EUR 8 as of June 2024). Robust standard errors are shown in parentheses. * p<0.1, ** p<0.05, *** p<0.01.

Table A4: Gender, wages, and productivity - alternative measures of productivity

	Log(wage)			
	(1)	(2)	(3)	(4)
Female	-0.001	-0.001	-0.001	-0.001
	(0.002)	(0.002)	(0.002)	(0.002)
Productivity 30 (=1 if above median)	0.001			
	(0.002)			
Productivity 90 (=1 if above median)		0.002		
		(0.002)		
Productivity 180 (=1 if above median)			0.004*	
			(0.002)	
Productivity log(duration) (=1 if above median)				-0.002
				(0.002)
Mean	10.0	10.0	10.0	10.0
# of observations	5,172	5,172	5,172	5,172
Controls	Yes	Yes	Yes	Yes
Age Fixed Effect	Yes	Yes	Yes	Yes
Office × Year Fixed Effect	Yes	Yes	Yes	Yes

Notes: The table presents estimates of the relationship between caseworker log wages and caseworker gender and productivity, using a sample of caseworkers from date-of-birth rule offices. Productivity metrics (30, 90, and 180) are calculated using dummies for leaving unemployment within 30, 90, and 180 days, respectively. Productivity (log duration) is calculated using the logarithm of the total unemployment duration for each job seeker. All models include the control variables discussed in section 4.3 and interactions between year and office fixed effects, as well as age fixed effects. Wages are measured in Swedish SEK (with SEK 100 equivalent to EUR 8 as of June 2024). Robust standard errors in parentheses. * p<0.1, ** p<0.05, *** p<0.01.

Table A5: Gender, wages, and productivity using various definitions of date-of-birth-rule offices

	Log(wage)		
	(1)	(2)	(3)
Female	-0.001	-0.002	-0.005**
	(0.002)	(0.002)	(0.002)
Productivity F-stat$\geq$ 100	0.004*		
	(0.002)		
Productivity F-stat$\geq$ 50		0.003*	
		(0.002)	
Productivity F-stat$\geq$ 200			0.005*
			(0.003)
Mean	10.0	10.0	10.0
# of observations	5,172	5,207	3,009
# of observations (unique)	2,705	2,718	1,792
Controls	Yes	Yes	Yes
Age Fixed Effect	Yes	Yes	Yes
Office × Year Fixed Effect	Yes	Yes	Yes

Notes: The table presents estimates of the relationship between caseworker log wages and caseworker gender and productivity, using a sample of caseworkers from date-of-birth rule offices and applying various F-stat thresholds for defining these offices (see text for details). All models include the control variables discussed in section 4.3 and include interactions between year and office fixed effects, as well as age fixed effects. Wages are measured in Swedish SEK (with SEK 100 equivalent to EUR 8 as of June 2024). Robust standard errors are shown in parentheses. * p<0.1, ** p<0.05, *** p<0.01.

Table A6: Gender wage gap and tenure

	Wage			
	(1) 0-2 years	(2) 3-5 years	(3) 6-10 years	(4) +11 years
Female	59.714	-57.710	47.893	-67.222
	(49.925)	(136.937)	(118.785)	(98.290)
Productivity (=1 if above median)	-19.213	99.983	180.616	46.845
	(61.790)	(172.696)	(123.668)	(104.480)
University Degree	924.200**	594.266*	830.223***	196.808
	(458.552)	(356.786)	(197.463)	(259.565)
Secondary Degree	866.953*	541.363	618.594***	-91.959
	(461.657)	(328.898)	(215.735)	(256.634)
Immigrant	-101.330*	-192.798*	-340.076***	-584.496***
	(55.448)	(103.405)	(131.096)	(156.858)
Business degree	31.993	-235.684	-219.402*	-34.557
	(69.048)	(183.983)	(125.272)	(102.947)
Social degree	-37.055	-128.875	116.072	325.010
	(54.849)	(127.065)	(148.232)	(216.494)
At least one child $<$ 4 years	151.766*	83.230	425.123***	-112.909
	(79.757)	(190.475)	(163.905)	(295.434)
1 Child (below 16)	-92.586	46.674	10.814	-164.901
	(83.145)	(205.130)	(144.711)	(134.445)
2 Children (below 16)	-29.003	-184.607	-4.587	45.866
	(86.567)	(191.203)	(151.071)	(160.590)
3+ Children (below 16)	1.568	37.346	145.597	-155.718
	(143.803)	(214.229)	(204.566)	(337.424)
Mean	20808.8	21432.9	21523.2	23413.2
# of observations	876	744	1,148	1,717
Age Fixed Effect	Yes	Yes	Yes	Yes
Workplace $\times$ Year Fixed Effect	Yes	Yes	Yes	Yes