Discussion Paper Series

IZA DP No. 18379

February 2026

Gender Gaps Under Comparable Tasks: Evidence from Quasi-Random Assignment

Negar Khaliliaraghi
IFAU

Petter Lundborg
Lund University and IZA@LISER

Johan Vikström
IFAU and Uppsala University

IZA NETWORK @ LISER
ISSN: 2365-9793

Gender Gaps Under Comparable Tasks: Evidence from Quasi-Random Assignment*

Abstract

Gender gaps in earnings persist even among high-skilled workers, partly because men and women often perform different tasks within and across jobs. We study a rare setting in which high-skilled men and women perform the same tasks under comparable conditions, allowing us to assess gender differences in productivity and pay without confounding from task or client allocation. Using administrative data from the Swedish Public Employment Service, we exploit a rotation scheme that quasi-randomly assigns job seekers to employment caseworkers. We find that productivity differences are small: job seekers assigned to female and male caseworkers exit unemployment at similar rates, and hourly wages—conditional on productivity—are nearly identical across genders. Despite this, female caseworkers earn about 8 percent less per year, entirely due to differences in contracted and actual hours worked. We also find suggestive evidence that male caseworkers are more likely to be promoted than equally productive female colleagues. When tasks are standardized and performance is measured objectively, gender differences in productivity and hourly pay are minimal, while gaps in annual earnings and career progression persist.

JEL classification

D84, I12, J12, J21

Keywords

gender gaps, productivity, wages, task allocation

Corresponding author

Petter Lundborg
petter.lundborg@nek.lu.se

* We are grateful to Pascaline Dupas, Heather Sarsons, Petra Persson, Nagore Irreberi, Olle Folke, Eva Ranehill, Karin Hederos, Christina Felfe, Malin Tallås Ahlzén, Therese Nilsson, Ana Rodriguez-Gonzales, Håkan Sehlin, Anna Sjögren, Jonas Cederlöf for their valuable feedback, and to the seminar participants at Stanford University, IFAU, and Lund University for insightful comments and suggestions.

1 Introduction

Gender gaps in earnings remain substantial, despite decades of progress in women's labor market outcomes (Goldin, 2014; Olivetti et al., 2024). Across most Western countries, women now surpass men in educational attainment, yet continue to earn less—particularly in high-paying occupations and senior roles (Blau and Kahn, 2017; Bertrand, 2020). A large literature shows that these gaps are closely linked to differences in occupations, industries, and job tasks, alongside persistent gender differences in career interruptions and caregiving responsibilities.[1]

Even within occupations and firms, however, gender gaps in earnings and career progression often persist. Interpreting these within-job differences is challenging. They may reflect differences in productivity, but they may also arise because men and women are allocated different tasks within the same job titles, face different client assignments or referral networks, or are evaluated differently (e.g., Babcock et al., 2017; Zeltzer, 2020; Sarsons et al., 2021; Sarsons, 2022; Card et al., 2019; Hengel, 2022; Koffi, 2025). In addition, caregiving responsibilities—particularly motherhood—may shape availability and scheduling, which in turn influences task and client assignment, making it difficult to separate productivity from working conditions. (e.g., Bertrand et al., 2010; Lundborg et al., 2017; Barbanchon et al., 2021; Andresen and Nix, 2022; Adams-Prassl et al., 2023; Lundborg et al., 2024). As a result, existing evidence rarely allows for a clean assessment of whether productivity and pay differences persist once task allocation and working conditions are held constant.

This paper addresses this gap by asking a simple but fundamental question: when men and women perform the same tasks under comparable conditions, do productivity and pay still differ? Rather than studying how gender differences in jobs and tasks generate inequality—a question extensively examined in prior work—we study the reverse case. By focusing on an environment where task allocation is effectively neutral, we provide a benchmark for understanding how large gender differences in productivity,

[1] For overviews of the literature, see, e.g., Goldin (2014); Olivetti and Petrongolo (2016); Bertrand (2020); Olivetti et al. (2024).

wages, earnings and promotion opportunities are within identical tasks and how much must instead stem from other margins—such as interruptions, discrimination, subjective evaluation, or bargaining.

We address this gap using a unique real-world setting in Sweden, where high-skilled public sector workers—employment caseworkers at the Swedish Public Employment Service—performed a standardized set of tasks and were effectively randomly assigned to clients (job seekers). The PES implemented a rotation scheme that matched job seekers to caseworkers based on birth dates, generating quasi-random caseworker–client assignments within local offices. This institutional feature provides an unusually clean opportunity to study gender differences in productivity and wages when men and women do the same work for similar clients and under identical production conditions, directly relevant for within-job policy interventions reducing gender bias in task allocation.

In this environment, male and female caseworkers carry out the same core function of helping job seekers find employment. Their tasks are standardized, their client pools are similar, and performance is measured objectively using administrative data on unemployment durations. This setting therefore eliminates confounding from sorting, task differentiation, or subjective performance evaluation—factors that typically obscure productivity comparisons. At the same time, wages are individually negotiated, with substantial variation across workers, creating scope for gender differences in pay to arise through productivity, discrimination, or bargaining behavior.

Our analysis yields four main findings. First, productivity differences between female and male caseworkers are small. If anything, women appear slightly more productive: job seekers randomly assigned to female caseworkers exit unemployment marginally faster than those assigned to male caseworkers. Because client assignment is exogenous, these estimates capture differences in performance rather than differences in client composition or task allocation. The findings are robust across alternative productivity measures and empirical specifications. Our main outcome—the duration of unemployment spells—does not capture job quality, but we find that clients of male and female caseworkers earn similar wages in their first post-unemployment job and display comparable employment rates and earnings five years later. We also find no gender differences in the number of unique job seekers served per contracted hour. Taken together, these results indicate that productivity differences are minimal when men and women perform comparable tasks under similar conditions.

Second, we examine the relationship between gender, productivity, and wages. Even if productivity is equal, wage gaps could arise through discrimination or gender differences in bargaining. If productivity were the primary driver of wage differences, however, small productivity gaps would imply small wage gaps. Consistent with this logic, we find that hourly wage differences are minimal between equally productive male and female caseworkers. This pattern suggests that, in this setting, factors often emphasized in wage determination—such as discrimination or gender differences in bargaining—play at most a limited role in shaping hourly pay.[2]

Third, we turn to annual earnings. Despite similar productivity and nearly identical wages, female caseworkers earn about 8 percent less per year than their male colleagues. Since wages are equal, this gap is entirely driven by differences in hours worked. Roughly half reflects differences in contracted hours, while the remainder reflects differences in actual hours supplied. These findings show that gender gaps in earnings can persist even when productivity and hourly wages are equal, operating instead through differences in time worked.

Fourth, we study promotions, a key channel through which earnings and career gaps can arise (Lazear and Rosen, 1990). Because productivity is measured objectively and task allocation is standardized, we can study promotion differences while holding performance constant. We find some suggestive evidence that male caseworkers are more likely to be promoted than female colleagues, although the estimates are imprecise and do not reach conventional levels of statistical significance. One possible explanation is that women apply for promotions less frequently or decline offers

[2]See Biasi and Sarsons (2021), Dreber et al. (2022), Roussille (2024), and Cortés et al. (2024) for recent evidence on the importance of wage bargaining for gender wage gaps.

more often. However, the promotion gap does not vary systematically with childcare demands, where constraints on long or inflexible hours would be most binding.

Taken together, the findings show that when tasks are standardized and performance is measured objectively, women and men appear similarly productive and receive essentially the same hourly pay. Residual gender differences in earnings arise through hours worked—both contracted and actual—rather than through productivity, while evidence of a promotion gap is, at most, suggestive. These patterns underscore the importance of task allocation as a central margin in shaping gender gaps: policies and organizational practices that reduce discretion or bias in how tasks, clients, and responsibilities are assigned—such as transparent, rule-based, or rotational assignment mechanisms—may therefore be particularly relevant for understanding and addressing persistent gender differences.

Our analysis offers new insights into the sources of gender pay gaps among high-skilled workers. While the setting is specific to employment caseworkers at the Swedish Public Employment Service, this occupation is representative of a large class of high-skilled public-sector jobs with substantial female employment. As in many other high-skilled professions, caseworkers exercise substantial discretion in how they perform their work, and their actions have meaningful consequences for clients (Graversen and van Ours, 2008; Crepon et al., 2013; Schiprowski, 2020; Cederlöf et al., 2025; Humlum et al., 2025).[3] Our findings are therefore directly relevant for understanding gender gaps in productivity, pay, and career progression in the public sector, and they complement existing evidence from high-skilled private-sector and elite professions—including law, medicine, academia, and corporate leadership (Bertrand, 2011; Goldin and Katz, 2016; Azmat and Ferrer, 2017; Sarsons et al., 2021), where gender gaps often arise through different mechanisms such as client acquisition, subjective evaluation, or long and inflexible hours.

[3] While Cederlöf et al. (2025) and Humlum et al. (2025) also exploit date-of-birth–based allocation of job seekers to caseworkers, their focus differs from ours. Cederlöf et al. (2025) study caseworker value-added and its determinants but do not examine wage formation or gender wage gaps. Humlum et al. (2025) analyze the effects of training programs using variation in caseworkers' assignment behavior.

Our paper contributes to several strands of literature. It relates to a growing literature documenting gender differences in task allocation within jobs and organizations. A substantial body of work shows that women are more likely to be allocated tasks with lower promotion potential, fewer learning opportunities, or less visibility, even when holding similar positions, and that such allocation patterns can contribute to gender gaps in earnings and career progression (Benschop and Doorewaard, 1998; Ohlott et al., 1994; De Pater et al., 2010; Babcock et al., 2017). Related evidence highlights gendered demand frictions and referral patterns that affect access to clients, projects, and high-return opportunities (Zeltzer, 2020; Sarsons, 2022). Our contribution is not to identify or explain these allocation mechanisms, but to study a setting in which they are largely absent by design. By exploiting quasi-random assignment of clients to caseworkers, we abstract from endogenous task and client allocation and assess outcomes in an environment where men and women are effectively assigned the same work. In this sense, we complement the task-allocation literature, and show what gender differences in productivity, wages and promotions may look like when allocation itself is neutralized, thereby helping to clarify which gender gaps are likely to originate in allocation processes rather than in differences in performance.

Our paper is also related to Azmat and Ferrer (2017), who document substantial gender differences in performance among young lawyers and show that these differences account for a large share of the gender gaps in earnings and promotion in that setting. One insight is that performance in law is strongly shaped by discretionary effort, client acquisition, and career aspirations, all of which exhibit pronounced gender differences. In contrast, we study a high-skilled occupation in which task assignment and client allocation are exogenous by design and where productivity is weakly linked to discretionary hours or self-selected workloads. This institutional feature allows us to assess gender differences in productivity in a setting that abstracts from many of the mechanisms in the legal profession, such as potentially gendered access to clients, networking opportunities, and returns to long hours. Our results therefore provide a complementary benchmark:

while Azmat and Ferrer demonstrate how performance gaps can arise and matter in high-powered private-sector environments, we show what gender differences in productivity, pay, and advancement look like when allocation and effort margins are largely neutralized. Together, the two settings help clarify which gender gaps are likely to originate in allocation and incentive structures respectively performance differences.[4]

Moreover, we contribute to the literature on gender differences in productivity in more standardized, lower-skilled settings. Cook et al. (2020) and Bolotnyy and Emanuel (2022) document productivity and earnings gaps in transportation and driving occupations, where standardized performance metrics are available. While informative, these settings differ markedly from ours. The tasks we study involve interpersonal skills, counseling, and case management—dimensions of performance that are less relevant in driving or transportation contexts and that may respond differently to institutional constraints.

Finally, our findings speak to the literature on gender gaps in promotions. Although promotion gaps are well documented, their underlying mechanisms remain less well understood.[5] Several studies report promotion disparities even after conditioning on performance (Ginther and Kahn, 2006; Blau and Devaro, 2007; Azmat and Ferrer, 2017; Sarsons et al., 2021). Recent work highlights the role of gender differences in evaluations of leadership potential and in promotion applications (Benson et al., 2024; Bosquet et al., 2019; Hospido et al., 2022; Fluchtmann et al., 2024; Haegele, 2024; Azmat et al., 2024). We contribute to this literature by examining promotion outcomes in a setting where men and women perform identical tasks and where productivity differences can be ruled out as a primary explanation.

The remainder of the paper proceeds as follows. Section 2 describes the institutional context. Section 3 outlines our data sources. Section 4

[4]A related literature examines whether wages align with marginal productivity across groups. Hellerstein et al. (1999) find that wages generally track productivity, with gender as a notable exception. Using Danish data, Gallen (2023) show that earnings gaps closely mirror productivity gaps following childbirth.

[5]See, for example, Cobb-Clark (1998), Bertrand (2011), Blau and Kahn (2017), and Cortes and Pan (2020).

presents the empirical strategy. Section 5 reports the main results. Section 6 concludes.

2 Institutional context and caseworker assignment

In this section, we provide background on the Swedish Public Employment Service (PES), describe the role and responsibilities of caseworkers, and outline the institutional features most relevant for our empirical analysis.

The PES is the central public agency responsible for helping job seekers find work and assisting employers in filling vacancies. In the early 2000s, it operated through roughly 300 local offices across the country. Individuals seeking unemployment insurance benefits were required to register with their nearest office, where they were assigned to a caseworker.

Caseworkers are responsible for guiding job seekers back to employment. Their tasks include career counseling, job matching, referrals to labor market programs, and monitoring compliance with unemployment insurance requirements. Once a caseworker is tasked with supporting a job seeker, (s)he has considerable discretion in designing the support: (s)he decides how frequently to meet with clients, which programs to recommend, and which job openings to highlight based on their own networks.

Caseworkers are recruited from diverse educational and professional backgrounds. While the formal requirement during our study period was an upper secondary degree and three years of work experience, in practice 71 percent held a university degree—most commonly in human resource management—indicating that the group we study is highly skilled.

Wages for caseworkers are determined through a combination of collective agreements and individual negotiations. Yearly agreements between trade unions and the government establish the overall framework for wage increases and overtime pay.[6] Within this framework, wages are set through individual bargaining between each caseworker and their local manager. The intention is to create a clear link between performance and pay. By 2010, 78 percent of all wages in the state sector were determined through

[6] Two trade unions, ST and SACO-S, represent caseworkers, while Arbetsgivarverket negotiates on behalf of the government (Samarbetsrådet, 2008).

such individual wage-setting conversations. Even when local unions negotiated with office managers, individual discussions preceded the final agreement, ensuring scope for differentiated wages.[7]

These institutional features are central to our analysis. Caseworkers perform standardized tasks supporting job seekers, yet their discretion allows for meaningful variation in how they assist job seekers. At the same time, the wage-setting system creates substantial scope for individual pay differences, as we further document in Section 3.6. Taken together, these elements make the PES a valuable context for examining gender differences in productivity and pay.

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

Managers at local PES offices have flexibility in how they 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. A subset of offices allocate job seekers to caseworkers based on job seekers' date of birth. According to PES officials, this method is perceived as transparent, administratively simple, and useful for equalizing workloads across caseworkers.[8] When date of birth is used, the allocation of job seekers to caseworkers becomes effectively random, as we show in Section 4.1. Our empirical analysis focuses on these offices.

Although our dataset does not directly identify which offices use date-of-birth allocation, this can often be easily inferred from the data. Figure 1 illustrates this. In Panel A, caseworkers are responsible for clients born on specific dates of the month, producing sharp discontinuities in the distribution of birth dates across caseworkers. In Panel B, where no such rule is in place, birth dates are evenly distributed across caseworkers.

Because not all cases are as clear-cut as those illustrated in Figure 1, we rely on a formal test to classify offices as random or non-random. For

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

[8] Sample statistics in Table 1, further discussed in Section 3.4, reveal no systematic differences between date-of-birth offices and other offices, except that the former tend to be somewhat larger.

each office and year, we regress the job seeker's day of birth on caseworker dummies and conduct an F-test of their joint significance. Offices with $F > 100$ are classified as date-of-birth offices (random), while those with $F < 20$ are identified as non-random offices. Offices with F-statistics between 20 and 100 are excluded from the analysis, as such cases could plausibly be either random or non-random; the thresholds are chosen to select offices that truly are in the respective groups. As we later show, our results are robust to alternative thresholds for defining random offices.[9]

Even in date-of-birth offices, occasional deviations from strict assignment occur—for instance, to handle temporary workload shocks or to direct clients with special needs to designated caseworkers. To address this, we follow Cederlöf et al. (2025) and Humlum et al. (2025) and construct a predicted caseworker for each job seeker: the caseworker most frequently assigned to clients born on the same day of the month within the same office and year. This is the caseworker who would have been assigned if the date-of-birth rule had been strictly followed. Panel A of Figure 1 confirms that this approach captures the assignment rule closely.

3 Data

3.1 Data on caseworkers

We use administrative data from the PES covering caseworkers and their assigned clients between 2003 and 2014. These data allow us to identify date-of-birth offices, link job seekers to caseworkers and offices, and observe caseworker activity. We then connect caseworkers to Statistics Sweden registers containing information on wages, demographics, education, and employment histories.

The PES data record each job seeker's exact date of birth, the case-

[9]Figure A1 shows the distribution of F-statistics across all offices, truncated at 200. Figure A2 shows the prevalence of the date-of-birth rule across offices and years. Around one-quarter to one-third of offices are classified as date-of-birth offices in a typical year, with some fluctuation over time. While the figure reports aggregate shares, it is important to note that offices can also change status, with some switching between policy and non-policy years during the sample period. Further details on the offices are provided in Section 3.

worker to whom they are assigned, and the office they belong to.[10] The data also document caseworker activity, including the number of assigned job seekers and the actions taken for each job seeker.

We link caseworkers to Statistics Sweden's employer–employee register, which covers the universe of matched employment spells, allowing us to identify all PES employees and trace their career histories. The population register provides demographics, education, and annual income. The wage statistics offer yearly information on contracted working hours and monthly wage rates for all caseworkers and other public sector workers.[11] Importantly, these are actual wage rates, not constructed from earnings and hours.

The registers also contain occupational codes, which we use to measure experience and promotions. Experience is defined as the number of years employed as a caseworker at the PES. Promotions are defined as shifts from caseworker to managerial or higher-ranking positions, accompanied by a discrete jump in wages.

3.2 Data on job seekers

We use data on outcomes of all job seekers registered at the PES between 2003 and 2014 to calculate our measures of caseworker performance. These records provide detailed information on unemployment spells, including their start date and duration until exit. Our main measure of caseworker productivity is an indicator for whether a job seeker leaves unemployment within 180 days. In robustness analyses, we also consider alternative measures: the total number of days spent unemployed and indicators for leaving unemployment within 30 and 90 days.

We complement these measures with indicators of job quality. By linking job seekers to matched employer–employee data from Statistics Sweden, we observe earnings and tenure in the first post-unemployment job. These outcomes allow us to assess not only the speed of re-employment but also the quality of jobs.

[10] When caseworkers meet with clients across multiple offices in a year, we assign them to the office with the largest share of their clients.

[11] Wages are recorded as full-time equivalent monthly wage rates.

We also measure the number of unique job seekers handled per contracted hour of work. This captures another dimension of productivity: while some caseworkers may place clients more effectively, they may do so by devoting more time to each case, reducing the number of clients they can serve. Note that our measure is based on contracted rather than actual hours, which may introduce some noise if, for example, female caseworkers work fewer actual hours than male caseworkers.

Taken together, these measures allow us to capture both the speed and quality of job placements, as well as the intensity of caseworker activity.

3.3 Sample restrictions and statistics

To ensure reliable measures of productivity, we exclude small or atypical offices—those with fewer than 200 registered job seekers per year—and caseworkers with fewer than 30 assigned clients in a year. After these restrictions, our dataset includes 8,805 individual caseworkers employed across 257 offices between 2003 and 2014. Together, these caseworkers handled 1,632,001 registered job seekers. Of the total, 5,297 caseworkers worked at least once in an office that applied a date-of-birth allocation rule during our study period. Of the 257 offices classified as random or non-random under the rule described in Section 2.1, 179 implemented the date-of-birth policy at some point.

3.4 Comparison of office-years with and without the date-of-birth policy

Table 1 compares the characteristics of caseworkers across office-years with and without the date-of-birth policy. Some offices never implemented the policy, some always operated under it, and others contribute observations for both types of years. Panel A shows that caseworkers in date-of-birth office-years are on average 46 years old, 63 percent are female, 16 percent are immigrants, and 70 percent hold a university degree. Panels A–B indicate that caseworker and job seeker characteristics are broadly similar across office-years with and without the policy. Some differences are statistically significant but small in magnitude. Panel C shows that offices in policy years are larger in terms of both caseworkers and job seekers, but they have similar caseloads per caseworker and comparable shares of female managers. Overall, apart from their larger size, offices in policy years closely resemble those in non-policy years, and there is little evidence that caseworker or client composition drove adoption of the date-of-birth rule.

3.5 Caseworker characteristics

Table 2 reports descriptive statistics for female and male caseworkers in date-of-birth offices. Female caseworkers are on average 45.5 years old, compared to 46.4 years for male caseworkers, and 67 percent of women hold a university degree, similar to men. The most common fields of study are business management and social work. Overall experience levels are high: 57 percent of caseworkers have more than eight years of tenure, while only 16 percent have less than two years.

3.6 Variation in wages at the Public Employment Service

We next examine the extent of wage variation among caseworkers in our data. As described in Section 2, the institutional framework allows for a large degree of individual wage setting, and here we document how much wages and earnings differ across caseworkers in practice. Figure 2a also shows substantial variation in starting wages among caseworkers, even after accounting for year and office fixed effects. Figure 2b further highlights considerable heterogeneity in annual wage changes, pointing to individualized wage trajectories rather than uniform adjustments.

We then explore whether observable characteristics can explain these differences. Regressions flexibly adjusting for age, education, and experience, as well as year and office fixed effects, account for only 48 percent of the variation in wage changes. Thus, more than half of the variation remains unexplained by standard observables.

These findings and institutional background in Section 2 underscore the substantial scope for other factors—such as bargaining, productivity, or discrimination—to influence wage outcomes in this setting.

4 Empirical strategy

We analyze gender differences in caseworker outcomes in two steps. First, we compare the productivity of male and female caseworkers, using job seeker outcomes as performance measures. Second, we examine gender gaps in wages, earnings, and promotions, and relate those gaps to caseworker characteristics and productivity.

Our identification strategy exploits offices where job seekers are allocated to caseworkers based on date-of-birth rules. This assignment mechanism ensures that, on average, male and female caseworkers are matched with comparable job seekers, making caseworker characteristics orthogonal to client composition. As a result, we can estimate productivity differences. It also allows us to study how gender gaps in pay and career progression relate to productivity. Without the date-of-birth rules, differences in job seeker assignments (task allocation) risk confounding such comparisons.

4.1 Measuring productivity differences

The first part of our empirical analysis focuses on gender differences in productivity. To measure performance, we use outcomes for the job seekers assigned to each caseworker. Because helping job seekers return to work is the primary responsibility of caseworkers, our main productivity measure is whether a client exits unemployment. We also examine job quality and caseload per contracted hour as complementary dimensions of productivity.

To estimate gender differences, we regress these job seeker outcomes on an indicator for whether the caseworker is female, $Female^{CW}$, and a vector of other caseworker characteristics, X^{CW}, using the following model:

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

where y_{icpt} is the outcome of job seeker i assigned to caseworker c at office p in job-seeker-age group g and year t. The coefficient δ captures the productivity difference between male and female caseworkers. The vector of caseworker characteristics, X_{ct}^{CW}, includes age dummies, education level and field, number of children, immigrant status, and years of experience as a caseworker. The specification also includes interacted year fixed effects γ_t, office fixed effects θ_p, and job-seeker-age-group fixed effects λ_g. The year and office effects allow us to exploit the within-office variation created by the date-of-birth rules, while the age-group effect accounts for separate allocation systems often used for job seekers under age 25.[12]

This specification exploits the variation created by the date-of-birth assignment rules to identify gender differences in productivity. As discussed in Section 2.1, offices occasionally deviate from the rule, implying imperfect compliance and scope for non-random sorting. To address this, we adopt an instrumental variables approach that follows the logic in Cederlöf et al. (2025) and Humlum et al. (2025). Specifically, we instrument the characteristics of the actual caseworker with those of the predicted caseworker—that is, the caseworker a job seeker would have been assigned to if the date-of-birth rule had been strictly applied. For example, the indicator for being assigned to a female caseworker is instrumented with an indicator for whether the predicted caseworker is female. The resulting local average treatment effects (LATEs) capture gender differences in productivity for the complier population, defined as job seekers at offices with a date-of-birth rule who are assigned to caseworkers in accordance with that rule.

To validate the IV strategy, we assess the standard assumptions. Table A1 reports first-stage regressions of actual caseworker characteristics on their predicted counterparts. Each actual characteristic is strongly correlated with its predicted value, with F-statistics ranging from 207 to 506, ruling out weak instruments as a concern. At the same time, although predicted and actual caseworker characteristics are strongly related, compliance with the date-of-birth assignment rule is imperfect. For example,

[12]Information from the PES and our data show that many offices run parallel date-of-birth allocation systems for youths (aged 24 or younger). These systems typically assign all youths to designated "youth" caseworkers, while older job seekers are assigned to different caseworkers under a separate rule. For example, one caseworker may handle all youths born between the 1st and 15th of a month, while another manages all older job seekers born between the 1st and 8th. Accordingly, we assign predicted caseworkers separately for youths and non-youths and interact office and year fixed effects with an under-25 indicator.

the rule correctly predicts assignment to a female caseworker for about 46 percent of job seekers. Similarly, the rule correctly predicts whether the assigned caseworker has a university degree for roughly 48 percent of job seekers.

To test *randomization*, we regress job seeker characteristics on predicted caseworker characteristics. Under random assignment, there should be no systematic correlation between the two. Table 3 confirms this prediction. Using data from offices with a date-of-birth rule, and controlling for interacted office and year fixed effects as in model (1), we find no evidence of systematic correlation between predicted caseworker characteristics and job seeker demographics, education, or earnings histories. The only significant coefficient, out of 32 coefficients, is economically negligible: female caseworkers are 0.46 percentage points less likely to be assigned a male job seeker.[13]

As a contrast, we repeat the analysis using data from offices *without* a date-of-birth rule, where assignments are not random.[14] Table 4 shows strong and systematic correlations between actual caseworker characteristics and job seeker characteristics in these offices, confirming that job seekers are non-randomly assigned to caseworkers. This underscores the importance of the date-of-birth allocation for overcoming selection.

Table 4 also illustrates the nature of systematic allocation in offices without random assignment. Female caseworkers are disproportionately assigned to younger job seekers, to women, and to those with lower prior earnings. Such assignment patterns may disadvantage female caseworkers, since these clients often face greater barriers to employment. Differences in caseload composition can then affect professional development and performance evaluations, embedding gender disparities into the allocation process

[13]While our setting provides an objective productivity measure free from gender-based task allocation, gender may still affect productivity through channels such as client behavior or employer responses. For example, female caseworkers may face discrimination from job seekers or employers. Although we find small overall gender gaps, these channels remain important and merit further study. For recent evidence on customer discrimination, see Kelley et al. (2023).

[14]As above, we adjust for office-by-year fixed effects to capture systematic allocation within offices.

itself.[15] These findings echo the broader concern raised in the literature that measured productivity may conflate performance with differences in task assignment (De Pater et al., 2010; Babcock et al., 2017; Zeltzer, 2020). Absent random allocation, observed gender productivity gaps may simply reflect differences in caseloads rather than true performance. This underscores why the date-of-birth allocation mechanism is crucial for isolating genuine productivity differences between male and female caseworkers.

Another key assumption is *monotonicity*, which requires that having a certain predicted caseworker according to the date-of-birth rule should monotonically increase the likelihood of having that caseworker as the actual caseworker. For instance, having a predicted caseworker with a university degree should increase the likelihood of actually being assigned a caseworker with a university degree. Following the arguments in Bhuller et al. (2020), this should hold for all sub-samples of the population. First-stage regressions for various sub-samples (e.g., separately for male and female job seekers) confirm that this is the case, lending support to the monotonicity assumption (results available on request).

A final assumption is the *exclusion restriction*. It requires that the date-of-birth rule affect job-seeker outcomes only through the identity of the assigned caseworker. This is a natural assumption in our setting: birth dates are predetermined, economically irrelevant for labor-market outcomes, and used by the PES only as a mechanical sorting device to rotate clients across caseworkers. There is no plausible channel through which birth date could directly influence job-finding outcomes other than through caseworker assignment.

4.2 Calculating productivity measures

In the second part of the analysis, we study how productivity, gender, and other caseworker characteristics relate to caseworkers' own labor market outcomes. Specifically, we examine gender gaps in wages, earnings, and

[15] Other caseworker attributes are also systematically related to job seeker characteristics. For example, more experienced caseworkers are more likely to be assigned to job seekers who are older, male, disabled, better educated, higher earning, and more likely to have been unemployed.

promotions, and how these gaps change when conditioning on productivity. Restricting to offices with a date-of-birth rule, we estimate:

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

where LM_{cpt}^{CW} denotes caseworker labor market outcomes, including wages, annual earnings, and an indicator for promotion. The coefficient δ captures the difference in outcomes between female and male caseworkers. The vector X_{ct}^{CW} contains standard controls used in studies of gender wage gaps, such as age, education level and field, number of children, and caseworker experience. Year and office fixed effects ensure that identification comes from within-office variation at the office with a date-of-birth rule.

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. (2025) 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 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 that 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 measures of caseworker value-added. For the validity of our empirical design, it is crucial that caseworkers significantly impact job seekers' outcomes. This was confirmed by Cederlöf et al. (2025), who demonstrated substantial heterogeneity in caseworker value-added in Sweden, showing that these terms meaningfully relate to job seeker outcomes.

5 Results

5.1 Productivity differences by gender

We begin by asking whether male and female caseworkers differ in productivity when assigned comparable clients under the date-of-birth rule, using the IV model described in Section 4.1. Our baseline estimates focus on job placement outcomes, with productivity measured as the probability that a job seeker exits unemployment within 180 days. We then add controls for caseworker characteristics such as education, parenthood, and experience. Next, we test robustness with alternative measures of productivity, including unemployment durations of 30 and 90 days, as well as indicators of job quality. Finally, we examine differences in the number of clients served per contracted hour, providing a broader view of productivity across dimensions.

Table 5 presents the results from our baseline specification and shows how the estimated gender gap in caseworker productivity changes as we add controls for relevant background characteristics. Column 1 includes only age fixed-effects, office-by-year-by-job-seeker-below-25 fixed-effects, and immigrant status. We find that female caseworkers are slightly more productive than their male counterparts: job seekers assigned to female caseworkers are 0.81 percentage points more likely to exit unemployment within 180 days. Given a baseline exit rate of 63 percent, this corresponds to a roughly 1 percent productivity advantage. This small but statistically significant difference suggests that productivity gaps are unlikely to explain any broader gender disparities in wages, earnings and/or promotions.

We continue by examining to what extent accounting for additional characteristics affects the gender productivity gap. As shown in Table 2, female and male caseworkers differ across dimensions such as education, number of children, and work experience. If productivity is related to such factors, the gender productivity gap may increase or decrease when accounting for them. Column 2 adds education variables, including measures of 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.

We next examine the role of parenthood for the gender gap. While female caseworkers 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 a larger productivity advantage for female caseworkers. The results in column 3 cast some 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 4, does not alter the small gender productivity gap to any meaningful extent.[16] The exception is those having 3 or more children who are significantly less productive, but this difference is relatively minor, at 1.1 percent of the mean of the outcome.

It may appear surprising that childbearing is largely unrelated to productivity among the caseworkers, given that extended parental leave periods in Sweden may lead to human-capital depreciation. We can think of several reasons why the presence of children is unrelated to productivity. 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 or less for child-rearing

[16] 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.

constitutes a minor fraction of an entire career span. Additionally, flexibility of work arrangements at the Swedish employment agency may facilitate the combination of professional and family responsibilities.

The fourth column examines the influence of years of experience. For this analysis, we add dummies indicating different degrees of experience working at the PES. Here, we drop the childbearing variables as we want to allow any effects of lost experience also through parental leave periods. We find that experience does not markedly alter the gender productivity gap. Furthermore, the coefficients associated with the work experience variables are small and not statistically significant.[17] In the fifth column, we simultaneously incorporate all the controls, which does not change the gender gap.

The results in Table 5 are based on one particular productivity measure; the likelihood that the caseworkers' job seekers find a job within 180 days. As robustness, Table A2 shows results for 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 those in Table 5: job seekers allocated to female caseworkers are significantly more likely to find a job within 90 days and experience shorter unemployment spells, though the magnitude of the differences is small. For instance, job seekers allocated to female caseworkers have 2 percent shorter unemployment durations. For the 30-day follow-up, there is no significant difference, however.

The 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 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 be due to differences in job structure, perhaps reflecting that caseworkers are better able to balance family responsibilities with their professional roles and

[17]It should be noted that the specification includes age fixed-effects which are highly correlated with experience.

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 be less important. These differences between caseworkers and lawyers suggest that even among high-skilled workers, there is substantial variation in gender productivity gaps and in the role of motherhood for productivity.

Our results so far suggest that female caseworkers are slightly more productive in terms of helping their job seekers back to work. This measure of productivity is particularly relevant within the context of the unemployment agency, as it is the most visible indicator of caseworker performance to management. However, the quality of the job is also relevant, as swifter job placement could come at the expense of job quality. To this end, we consider additional indicators of productivity: the duration of first job after exit from unemployment, and long-term earnings and employment status.

Studying job-quality outcomes requires conditioning on job seekers who eventually secure employment, which raises the risk of selection bias if female and male caseworkers differ in their ability to place clients into work in the first place. To address this concern, we begin by testing whether assignment to a female caseworker affects the overall likelihood of ever being observed finding a job in our data. Column 1 of Table 6 shows that it does not: job seekers assigned to female and male caseworkers are equally likely to be employed in the long run.

Having ruled out differential selection into employment, we turn to the quality of the jobs secured. Columns 2–4 of Table 6 reveal no productivity differences between female and male caseworkers. Job seekers assigned to female caseworkers have similar first-job durations (in days; column 2) and comparable long-term earnings and employment outcomes five years after unemployment (columns 3–4). These results demonstrate that the somewhat faster job placements achieved by female caseworkers do not come at the expense of job quality.

So far have focused on job placement outcomes for caseworkers' assigned job seekers, but productivity may also vary in other dimensions. For instance, female caseworkers might handle fewer job seekers per hour worked. To examine this, column 5 uses the monthly number of unique job seekers divided by monthly contracted hours as the outcome variable. However, we find only minor and statistically insignificant differences between male and female caseworkers on this measure. This suggests that female caseworkers' slightly faster job placements do not come at the expense of assisting fewer job seekers. In fact, given that female caseworkers work fewer actual hours per month, as we later discuss, these results suggest that they likely assist slightly more job seekers per actual hour—without any trade-off in terms of placement speed or job quality.

Finally, note that our instrumental-variables estimates identify local average treatment effects for the complier population—job seekers whose caseworker assignments follow the date-of-birth rule. As a result, these estimates do not capture productivity differences for job seekers whose assignments are overridden. A potential concern is therefore that discretionary deviations from the rule may be more common for particularly difficult cases, which could be systematically assigned to caseworkers of one gender rather than the other in ways not captured by the complier estimates. We therefore examine heterogeneity in productivity by client difficulty by splitting the sample into job seekers with predicted unemployment durations below and above the median, and estimating gender differences in productivity separately for these two groups. We find no evidence that gender differences in productivity vary systematically with client difficulty, suggesting that our main findings are not driven by the exclusion of harder, non-complier cases.[18]

[18] Predicted unemployment duration is constructed using a regression model based on the predetermined job seeker characteristics reported in Table 1. The estimated effect of assignment to a female caseworker on the likelihood to exit unemployment within 180 days is 0.0079 (s.e. = 0.0030) below the median of predicted duration and 0.0073 (s.e. = 0.0040) above the median.

5.2 Gender, productivity, and wages

Having shown that male and female caseworkers are similarly productive, we now turn to the question of whether this parity in productivity is reflected in wages. If productivity is a key determinant of gender differences in wages, as some recent literature suggests, our results would therefore also predict small differences in wages between female and male caseworkers. However, if gender gaps in wages mainly reflect other factors than productivity, such as discrimination or gender differences in wage bargaining, gender gaps could still arise. 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.

We proceed in three steps. First, we validate that our productivity measure is related to wages by estimating how wage levels vary with different specifications of caseworker productivity. Second, we analyze whether any wage difference between male and female caseworkers persists after conditioning on productivity. Third, we study how gender wage gaps vary by experience, which may be relevant for wage bargaining.

To validate the productivity measures, Table 7 presents results from various specifications of the productivity-wage relationship, using offices with a date-of-birth rule, and adjusting for age fixed-effects and office-by-year fixed effects. Column 1 shows that the estimated productivity measure, entered without any further transformation, has no statistically significant relationship with wages. The same holds when we account for non-linearities by adding a squared term and when we used lagged productivity at t-1 (columns 2 and 3). However, if managers at the PES offices classify their caseworkers into different productivity categories when setting wages, a more relevant measure may be the within-office productivity quintile each caseworker falls into. Column 4 also shows that this is the case. Adding the within-office productivity quintile as a continuous variable gives a significantly and positive correlation with the wage rate. The relationship is also present when using the same measure at t-1, indicating that the previous year's productivity is likely an important factor in determining current year wages (column 5).

The coefficient indicates that moving up one quintile in the productivity distribution within an office at t-1 is associated with a 0.3 percent increase in monthly wages. Next, Figure A3 plots average wages by productivity quintile. It indicates a clear breakpoint at the median: caseworkers at or above the median have higher wages than those below it. This is confirmed in columns 6–7 of Table 7, where we divide the caseworkers between above/below median productivity and show that caseworkers above the median have significantly higher wages.

Having established that our productivity measures are closely linked to wages, Table 8 turns to the relationship between gender, productivity, and wages. Column 1 examines the gender–wage gap, while controlling for education and experience. The estimated gap is small and statistically insignificant. In Column 2, we additionally control for contracted hours to capture potential penalties from working shorter hours. The results remain virtually unchanged: the gender gap continues to be small and insignificant.

In Column 3, we introduce productivity as a control. Given that gender differences in productivity were found to be small, this adjustment has, as expected, only a minor effect on the wage gap. Including the number of children as an additional control leaves the estimated gap unchanged (column 4).[19] The coefficient on the female dummy is also stable when alternative productivity measures are applied or when yearly wage growth is used as the outcome (Tables A3 and A4). Likewise, varying the thresholds for the F-statistics that define a date-of-birth office does not materially affect the results (Table A5).

How do these results square with other recent findings 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

[19] Productivity is here defined as being above the office-level median.

explained the former (Azmat and Ferrer, 2017; Bolotnyy and Emanuel, 2022; Cook et al., 2020). 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 promotions. We return to this below.

Our findings do differ from other studies, however, that identify a gender pay gap in settings with individual wage-bargaining, even after controlling for 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 between newly hired and experienced caseworkers. Table A6 therefore stratifies the sample by experience level, but the gender wage gap is small and insignificant for both inexperienced and experienced caseworkers. This suggests that gender differences in wage requests may be relatively small in the Swedish context, as also observed by Säve-Söderbergh (2019).

5.3 Gender and earnings

We next shift our focus from wages to annual earnings, which combine both pay and time spent working. Given the small gender differences in both hourly wages and productivity, this earnings gap must be driven by other factors. A natural explanation is that female caseworkers tend to work fewer hours—either due to lower contracted hours or because of more frequent absences, for instance, when taking care of sick children.

To explore this in more detail, Table 9 decomposes the sources of the gender earnings gap. In Column 1, we estimate the annual log earnings gap, showing that female caseworkers earn about 8 percent less than their male colleagues. Column 2 adds controls for the monthly wage rate, which has little effect on the earnings gap, confirming that it is not driven by wage differences but by time worked. Once we control for contracted hours, the earnings gap shrinks by about half to 3.5 percent (Column 3). Column 4 further adds controls for individual productivity, measured by caseworker value added, but this has virtually no effect on the remaining earnings gap, consistent with productivity differences playing little role. This remaining gap likely reflects differences in actual hours worked, which we cannot observe directly.

To investigate whether family-related absences contribute to the remaining gap, Column 5 adds indicators for parenthood. While having children significantly reduces earnings, the gender earnings gap remains largely unchanged.[20] This points to unobserved factors—such as sick leave or unpaid leave—as likely drivers of the residual difference. Combined with the earlier finding that female caseworkers assist as many—or more—job seekers per contracted hour, this also implies that female caseworkers likely see somewhat more job seekers per actual hour worked.

The fact that female caseworkers are at least as productive as male caseworkers despite more frequent absences may reflect the nature of the caseworker profession itself. The profession may have a high degree of substitutability—colleagues can temporarily step in—making it easier to balance family obligations without a productivity penalty (Goldin, 2014). Consistent with this view, Azmat et al. (2022) show that women are more likely to work in firms that allow for such flexibility by providing better access to substitutes. This implies that the residual gender earnings gap we observe may be larger in settings where substitutability is lower and temporal flexibility more constrained.

In summary, this section shows that gender differences in annual earnings arise not from differences in pay or productivity, but from differences in actual hours worked. Female caseworkers perform just as well as male caseworkers on an hourly basis, but earn less over the year due to more frequent time away from work—likely reflecting their greater share of family responsibilities.

[20] Interacting parenthood variables with gender shows, as expected, that the earnings impact of parenthood is stronger for female caseworkers. However, this does not explain the gender earnings gap in Table 9.

5.4 Gender and promotions

While our previous results show small or no gender gaps in productivity and wages, gender disparities in promotions may still exist. Such differences would not be visible when comparing workers in similar current roles but could arise through biased promotion processes. Importantly, our setting allows us to test whether equally productive male and female caseworkers are promoted at similar rates, helping to isolate the role of non-productivity-related factors such as discrimination, working hours, or application behavior.

In Table 10, we analyze the gender promotion gap, defining a promotion as a transition from caseworker to a senior or managerial position, based on occupational codes in the registers. The estimation framework follows the same model as in the previous analysis.

We start in Column 1 by estimating the promotion gap without controlling for productivity.[21] The results suggest a sizable gender promotion gap: female caseworkers are 44 percent less likely to be promoted. The estimate is statistically insignificant, however, which likely reflects the small number of observed promotions—only about 50 in total.

Column 2 adjusts for productivity to rule out the possibility that small productivity differences are driving the observed gap. As anticipated, the gender promotion gap remains unchanged. In column 3, we add the number of contracted hours worked, as women working fewer than full-time hours might hinder promotion to a managerial role. However, when doing so, the gender promotion gap remains largely unaffected.[22]

We next examine whether there are indications that females may apply for promotions less frequently, a mechanism emphasized in recent work (Bosquet et al., 2019; Hospido et al., 2022; Fluchtmann et al., 2024; Haegele,

[21]The estimation sample in Table 10 is somewhat smaller than in the rest of the analysis (4,852 versus 5,172 observations), since promotions cannot be observed for 2014, the final year in the data. The results in other parts of the analysis remain robust when restricting the sample accordingly.

[22]One possible concern is that the promotion gap could partly reflect gender differences in turnover—if, for example, female caseworkers are more likely to leave the PES before being considered for promotion, or if lower promotion prospects induce women to exit. However, we find no evidence of gender differences in the probability of leaving PES offices (Table A7).

2024). A potential reason is that roles requiring long hours and limited flexibility are harder to reconcile with current or future childcare responsibilities (Bertrand et al., 2010). While our data do not contain information on promotion applications, we can test this indirectly: if family responsibilities reduce the likelihood that female caseworkers apply to or accept promotion offers, the effect should be most pronounced among those with children. To investigate this, columns 4–5 report separate regressions for caseworkers with and without children, controlling for productivity in both cases. For neither group do we find a significant promotion gap.

Overall, we find suggestive evidence of a gender gap in promotions. However, the small number of promotion events limits statistical power, and the estimates do not reach conventional levels of statistical significance.[23]

6 Conclusions

This paper examines gender differences in productivity, pay, and career progression in a setting where men and women perform the same tasks under comparable conditions. Exploiting quasi-random assignment of job seekers to caseworkers at the Swedish Public Employment Service, we study an environment in which task and client allocation are effectively neutralized and performance can be measured objectively. This design allows for a clean assessment of gender differences that is rarely possible in observational data.

Three core findings emerge. First, productivity differences between female and male caseworkers are small. If anything, job seekers assigned to female caseworkers exit unemployment marginally faster, with no evidence of trade-offs in job quality or caseload per contracted hour. Second, both unconditional and conditional on productivity, wage differences are small and statistically insignificant, leaving limited scope for gender differences in pay driven by discrimination or bargaining in this setting. Third, we find suggestive evidence of a promotion gap: male caseworkers appear more likely to advance into managerial roles, although estimates are imprecise.

[23] Using alternative definitions of promotion—for example, a wage increase exceeding 10 percent—yields qualitatively similar but likewise imprecise estimates.

Taken together, the results show that when men and women are assigned the same tasks and evaluated under the same conditions, gender differences in productivity and hourly pay largely disappear. This highlights task allocation as a central margin for understanding gender inequality in the workplace. In more typical labor market settings—where task assignment is discretionary and opaque—gender gaps may therefore reflect differences in how work, clients, and advancement opportunities are allocated rather than differences in underlying performance.

While our setting is specific to a high-skilled public-sector occupation with standardized tasks, the findings provide a benchmark that is informative beyond this context. They suggest that organizational practices governing task and client assignment play a potentially important role in shaping gender gaps. Policies or institutional designs that reduce discretion or bias in allocation—such as transparent, rule-based, or rotational assignment mechanisms—may therefore be an effective complement to interventions aimed at reducing gender inequality in pay and career outcomes.

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* **107**: 1–23.

Andresen, M. E. and Nix, E. (2022). What causes the child penalty? evidence from adopting and same-sex couples, *Journal of Labor Economics* **40**(4): 971–1004.

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

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). Workplace presenteeism, job substitutability and gender inequality, *Journal of Human Resources* . Forthcoming.

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.

Barbanchon, T. L., Rathelot, R. and Roulet, A. (2021). Gender differences in job search: Trading off commute against wage, *The Quarterly Journal of Economics* **136**(1): 381–426.

Benschop, Y. and Doorewaard, H. (1998). Covered by equality: The gender subtext of organizations, *Organization Studies* **19**(5): 787–805.

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.

Bhuller, M., Dahl, G. B., Løken, K. V. and Mogstad, M. (2020). Incarceration, recidivism, and employment, *The Journal of Political Economy* **128**(4): 1269–1324.

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. (2025). What makes a good caseworker?, *Technical report*.

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.

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

Cortés, P., French, J., Pan, J. and Zafar, B. (2024). Gender differences in negotiations and labor market outcomes: Evidence from an information intervention with college students, *Working Paper 32154*, 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.

Dreber, A., Heikensten, E. and Säve-Söderbergh, J. (2022). Why do women ask for less?, *Labour Economics* **78**(C).

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): 182–219.

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

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.

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.

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. (2025). What works for the unemployed? evidence from quasi-random caseworker assignments, *Technical report*, NBER Working Paper 33807.

Kelley, E. M., Lane, G. V., Pecenco, M. and Rubin, E. A. (2023). Customer discrimination in the workplace: Evidence from online sales, *Working Paper 31998*, National Bureau of Economic Research.
URL: *http://www.nber.org/papers/w31998*

Koffi, M. (2025). Innovative ideas and gender (in)equality, *American Economic Review* **115**(7): 2207–2236.

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 Rasmussen, A. (2024). Is there really a child penalty in the long run? New evidence from ivf treatments, *SSRN Electronic Journal* .

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.

Ohlott, P. J., Ruderman, M. N. and McCauley, C. D. (1994). Gender differences in managers' developmental job experiences, *Academy of Management Journal* **37**(1): 46–67.

Olivetti, C., Pan, J. and Petrongolo, B. (2024). Chapter 8 - the evolution of gender in the labor market, *in* C. Dustmann and T. Lemieux (eds), *Handbook of Labor Economics, Volume 5*, Vol. 5 of *Handbook of Labor Economics*, Elsevier, pp. 619–677.

Olivetti, C. and Petrongolo, B. (2016). The evolution of gender gaps in industrialized countries, *Annual Review of Economics* **8**(Volume 8, 2016): 405–434.
URL: *https://www.annualreviews.org/content/journals/10.1146/annurev-economics-080614-115329*

Roussille, N. (2024). The role of the ask gap in gender pay inequality, *The Quarterly Journal of Economics* **139**(3): 1557–1610.

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. (2022). Interpreting signals in the labor market: Evidence from medical referrals. Working paper.

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.

Säve-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): 169–97.

Figures

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

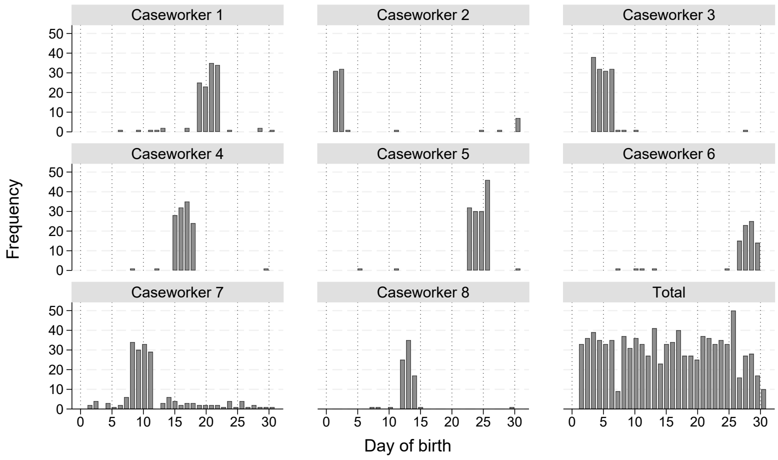

(a) Office with date of birth rule

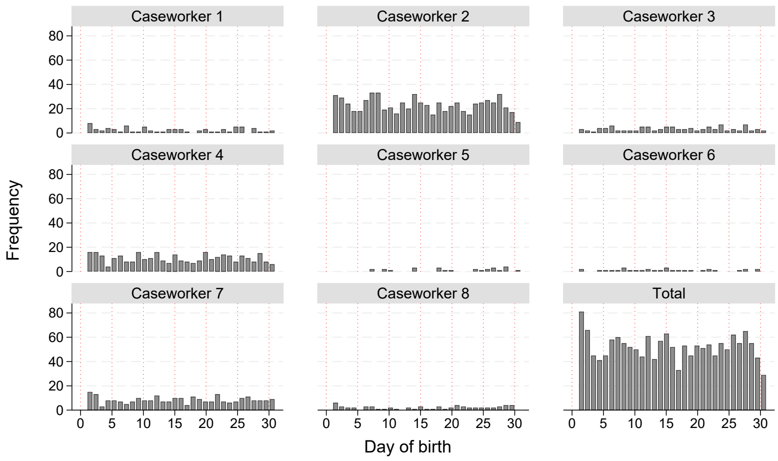

(b) Office without date of birth rule

Notes: The figure shows the allocation of job seekers to caseworkers by day of birth (day of the month). Each panel plots the distribution of job seekers' days of birth for a given caseworker; the "Total" panel aggregates across caseworkers within the office. Panel (a) shows an office that applies a date-of-birth assignment rule, resulting in pronounced clustering by caseworker and a strong first stage (F-statistic = 814.39). Panel (b) shows an office without the rule, where days of birth are approximately uniformly distributed and the first stage is weak (F-statistic = 2.29). F-statistics are computed using one year of data.

Figure 2: Distribution of starting wages and changes in wages for the caseworrkers

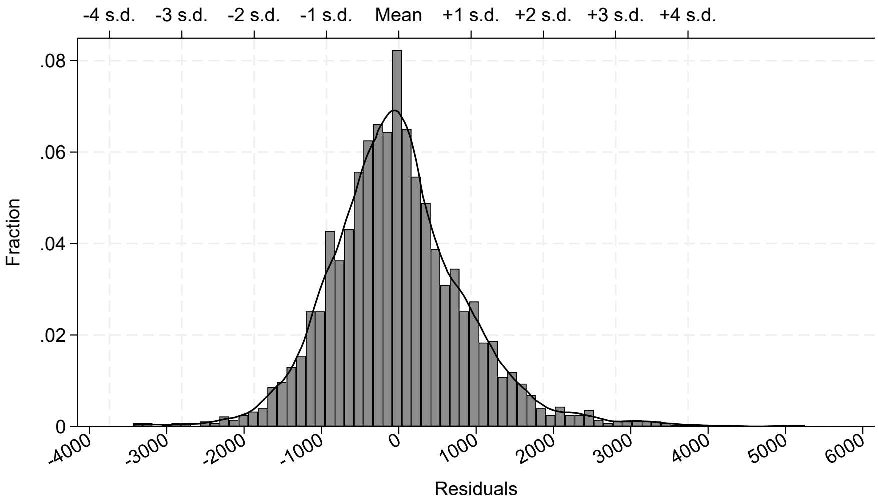

(a) Residual starting wages

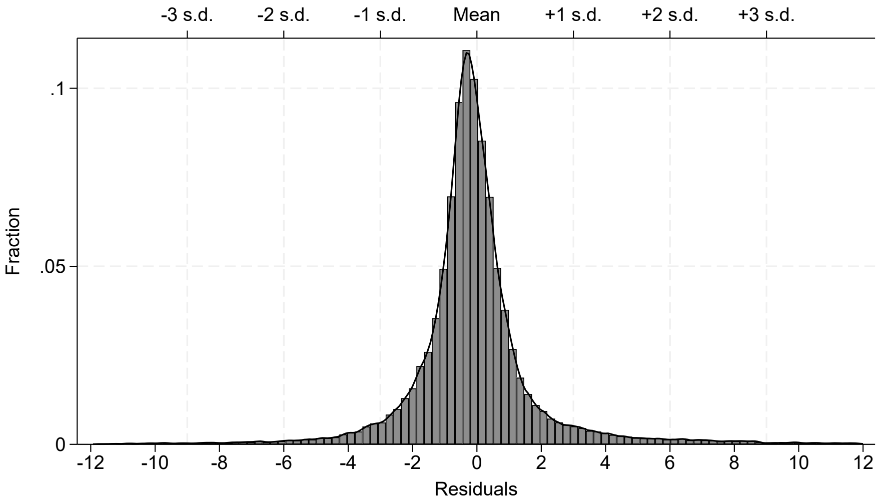

(b) Residual wage increases

Notes: The figure shows the distribution of residualized caseworker wages. Panel (a) plots deviations in starting monthly wages from office-year means, measured in Swedish kronor (SEK) (SEK 100 corresponds to EURO 8 as of June 2024). Panel (b) plots deviations in the annual percentage change in wages from office-year means, measured in percentage points. In both panels, the horizontal axis reports deviations from the office-year mean, and the vertical axis reports the fraction of observations.

Tables

Table 1: Caseworker, job seeker and office statistics by type of office

	Panel A: Caseworker characteristics					
	Random offices		Non-random offices			
	Mean	SD	Mean	SD	Diff	p-value
Age	46.35	10.55	46.49	10.45	0.14	0.22
Female	0.63	0.48	0.64	0.48	0.00	0.76
Immigrant	0.16	0.37	0.16	0.36	-0.01	0.21
Married	0.61	0.49	0.62	0.49	0.01	0.35
Number of children below 16	0.64	0.93	0.64	0.94	0.00	0.83
University Degree	0.70	0.46	0.67	0.47	-0.03***	0.00
Secondary Degree	0.28	0.45	0.30	0.46	0.03***	0.00
Business Degree	0.28	0.45	0.29	0.46	0.01*	0.05
Social Degree	0.18	0.38	0.17	0.38	-0.01	0.07
Log(earnings)	12.41	0.27	12.42	0.25	0.01*	0.03
Log(wage)	10.01	0.10	10.00	0.10	-0.01***	0.00
Experience						
0-2 years	0.16	0.36	0.14	0.35	-0.01**	0.00
2-4 years	0.13	0.33	0.14	0.34	0.01*	0.03
4-6 years	0.09	0.28	0.09	0.28	0.00	0.43
6-8 years	0.06	0.24	0.06	0.25	0.00	0.68
8-10 years	0.06	0.24	0.06	0.24	0.00	0.99
+10 years	0.51	0.50	0.51	0.50	0.00	0.88
# of observations	12,806		20,289		32,813	
# of observations (unique)	5,297		7,015		8,805	

	Panel B: Job seeker characteristics					
	Random offices		Non-random offices			
	Mean	SD	Mean	SD	Diff	p-value
Age	31.90	12.35	32.48	12.54	0.58***	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
Immigrant	0.24	0.43	0.25	0.43	0.01***	0.00
Disabled	0.04	0.20	0.05	0.21	0.01***	0.00
Eligible for UI	0.65	0.48	0.66	0.47	0.00***	0.00
Earnings, $t-1$	99072	121028	101193	126624	2121***	0.00
Secondary Degree	0.54	0.50	0.53	0.50	-0.00***	0.00
University Degree	0.26	0.44	0.25	0.43	-0.02***	0.00
Unemployment duration	264.62	422.77	272.57	442.60	7.95***	0.00
# of observations	1,371,307		992,982		2,364,289	
# of observations (unique)	1,016,644		759,866		1,632,001	

	Panel C: Office 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 reports descriptive statistics for caseworkers (Panel A), job seekers (Panel B), and offices (Panel C), separately for offices that assign job seekers to caseworkers using a date of birth rule ("random offices") and offices without such a rule ("non-random offices"). The column labeled "Diff" reports the difference in means between non-random and random offices (non-random minus random), and the p-value column reports two-sided p-values from tests of equality of means. The sample includes all job seekers and caseworkers 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 for female and male caseworkers

	Female caseworkers		Male caseworkers			
	Mean	SD	Mean	SD	Diff	p-value
Age	45.56	10.46	46.43	10.51	0.87**	0.00
Immigrant	0.15	0.35	0.16	0.36	0.01	0.38
Secondary Degree	0.31	0.46	0.27	0.45	-0.03**	0.01
University Degree	0.67	0.47	0.70	0.46	0.03*	0.03
Married	0.60	0.49	0.56	0.50	-0.04**	0.01
Number of children below 16	0.68	0.95	0.62	0.93	-0.06*	0.02
Business Degree	0.34	0.47	0.23	0.42	-0.11***	0.00
Social Degree	0.16	0.37	0.18	0.38	0.01	0.17
Log(earnings)	12.39	0.23	12.47	0.17	0.07***	0.00
Log(wage)	10.00	0.09	10.00	0.09	-0.00	0.05
Experience						
0-2 years	0.16	0.36	0.16	0.37	0.00	0.71
4-6 years	0.08	0.28	0.09	0.28	0.00	0.74
6-8 years	0.06	0.24	0.06	0.23	-0.01	0.23
8-10 years	0.06	0.23	0.06	0.23	-0.00	0.65
+10 years	0.52	0.50	0.49	0.50	-0.03*	0.03
# of observations	3,637		2,023		5,660	
# of observations (unique)	1,885		1,055		2,939	

Notes: The table reports descriptive statistics for female and male caseworkers in offices that allocate job seekers to caseworkers using job seekers' date of birth. Columns report group means and standard deviations. The column labeled "Diff" reports the difference in means between male and female caseworkers. The p-value column reports two-sided p-values from tests of equality of means. 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 tests for offices *with* date-of-birth rule

	Dependent variables: job seeker 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	0.00076	0.00009	-3.19264*	6.72615	-0.00007	-0.00007	-0.00003	0.00000
	(0.00089)	(0.00012)	(1.35356)	(12.99831)	(0.00005)	(0.00006)	(0.00002)	(0.00003)
Predicted caseworker gender (=1 if female)	-0.00763	0.00178	-39.49797	-323.02994	0.00461***	-0.00118	0.00067	0.00047
	(0.01764)	(0.00225)	(26.45859)	(251.97594)	(0.00104)	(0.00119)	(0.00041)	(0.00069)
Predicted caseworker education	0.00343	-0.00027	8.62406	-165.55461	-0.00018	0.00050	0.00006	0.00026
	(0.00950)	(0.00119)	(14.69381)	(132.04255)	(0.00054)	(0.00066)	(0.00021)	(0.00035)
Predicted caseworker experience	0.00006	0.00025*	-2.57780	7.24748	0.00006	-0.00009	0.00001	0.00002
	(0.00090)	(0.00012)	(1.39931)	(13.10474)	(0.00005)	(0.00006)	(0.00002)	(0.00003)
# of observations	1,371,307	1,371,307	1,371,307	1,371,307	1,371,307	1,371,307	1,371,307	1,371,307

Notes: The table above shows the estimates of the relationship between various job seeker characteristics and predicted caseworker characteristics in offices that use the date-of-birth rule in allocating job seekers to caseworkers (random offices with F-stat above 100, see section 2.1). Predicted caseworker characteristics are the characteristics of the caseworker that the job seeker would have been assigned to if the date-of-birth rule was fully applied without exceptions. Age and experience (tenure at the PES) are measured in years. Caseworker's education variable categorizes caseworkers by highest level of completed education as follows (1 = Pre-high school [< 9 years], 2 = Pre-high school [9(10) years], 3 = High school , 4 = University or College [< 2 years], 5 = University or College [2 years], 6 = PhD). All models include interacted year, office and job seeker age under 25 fixed effects. Standard errors, clustered at the caseworker level in parentheses. * p<0.1, ** p<0.05, *** p<0.01.

Table 4: Balancing tests for offices *without* date-of-birth rule

	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	0.18873***	0.00543***	16.65898***	851.89685***	-0.00115***	0.00131***	0.00104***	0.00114***
	(0.01369)	(0.00068)	(4.00231)	(84.43477)	(0.00030)	(0.00022)	(0.00013)	(0.00010)
Caseworker's gender (=1 if female)	-1.20519***	-0.02512	90.75665	-11773.31421***	0.05870***	-0.00336	0.00332	-0.00627**
	(0.31961)	(0.01537)	(85.73899)	(1921.96024)	(0.00657)	(0.00494)	(0.00273)	(0.00214)
Caseworker's education	-0.02475	0.01474	92.08661	-2521.75377*	0.01835***	0.00846***	0.00398**	-0.00222*
	(0.16236)	(0.00774)	(54.59081)	(1031.58339)	(0.00341)	(0.00245)	(0.00133)	(0.00112)
Caseworker's experience	0.13946***	0.00434***	5.65814	668.56147***	-0.00064	0.00091***	0.00065***	0.00091***
	(0.01419)	(0.00073)	(4.83767)	(94.80253)	(0.00040)	(0.00021)	(0.00013)	(0.00010)
# of observations	992,981	992,981	992,981	992,981	992,981	992,981	992,981	992,981

Notes: The table above shows the estimates of the relationship between various job seeker characteristics and caseworker characteristics in offices that do not use a date-of-birth rule in allocating job seekers to caseworkers (non-random offices with F-stat below 20, see section 2.1). Age and experience (tenure at the PES) are measured in years. Caseworker's education variable categorizes caseworkers by highest level of completed education as follows (1 = Primary school [$<$ 9 years], 2 = Primary school [9(10) years], 3 = High school, 4 = University or College [$<$ 2 years], 5 = University or College [$\geq$ 2 years], 6 = PhD). All models include interacted year, office and job seeker age under 25 fixed effects. Standard errors, clustered at the caseworker level in parentheses. * $p<0.1$, ** $p<0.05$, *** $p<0.01$.

Table 5: Determinants of caseworker productivity: exits from unemployment

	Productivity measure: job seeker exiting unemployment within 180 days				
	(1)	(2)	(3)	(4)	(5)
Female	0.0081***	0.0079***	0.0081***	0.0075***	0.0077***
	(0.0028)	(0.0028)	(0.0029)	(0.0029)	(0.0029)
University degree		0.0029	0.0032	0.0036	0.0040
		(0.0087)	(0.0089)	(0.0087)	(0.0089)
Secondary degree		0.0082	0.0082	0.0084	0.0083
		(0.0089)	(0.0091)	(0.0089)	(0.0091)
Business degree		0.0000	0.0003	-0.0004	-0.0001
		(0.0033)	(0.0034)	(0.0034)	(0.0034)
Social degree		0.0014	0.0012	0.0017	0.0014
		(0.0040)	(0.0040)	(0.0040)	(0.0040)
At least one child < 4 years			-0.0008		-0.0005
			(0.0045)		(0.0045)
1 Child			0.0028		0.0025
			(0.0038)		(0.0038)
2 Children			-0.0060		-0.0068
			(0.0043)		(0.0043)
+3 Children			-0.0109*		-0.0112*
			(0.0064)		(0.0064)
Experience					
2-4 years				-0.0045	-0.0035
				(0.0050)	(0.0050)
4-6 years				-0.0028	-0.0007
				(0.0058)	(0.0058)
6-8 years				0.0035	0.0057
				(0.0063)	(0.0063)
8-10 years				-0.0005	0.0013
				(0.0064)	(0.0064)
+10 years				0.0023	0.0041
				(0.0052)	(0.0052)
Immigrant	-0.0059	-0.0058	-0.0057	-0.0055	-0.0053
	(0.0040)	(0.0040)	(0.0040)	(0.0040)	(0.0040)
Mean	0.63	0.63	0.63	0.63	0.63
# of observations (job seekers)	1,371,307	1,371,307	1,371,307	1,371,307	1,371,307
# of observations (caseworkers)	5,297	5,297	5,297	5,297	5,297
Caseworker Age FEs	Yes	Yes	Yes	Yes	Yes
Office×Year× Job seeker Age<25 FEs	Yes	Yes	Yes	Yes	Yes

Notes: The table reports instrumental-variables estimates of the relationship between actual caseworker characteristics and the probability that a job seeker exits unemployment within 180 days. Identification exploits as-if random assignment of job seekers to caseworkers based on date of birth in the offices that use this allocation (random offices). Actual caseworker characteristics refer to the characteristics of the caseworker to whom the job seeker is ultimately assigned. Predicted caseworker characteristics correspond to the caseworker the job seeker would have been assigned to under a strict application of the date-of-birth assignment rule, without exceptions, and are used as instruments for actual 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, 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 (days)	(3) Earnings, $t+5$	(4) Employment status, $t+5$	(5) Number of unique job seekers per monthly contracted hours
Female	0.001	-1.145	-1908.134	-0.000	0.003
	(0.004)	(7.110)	(1167.776)	(0.004)	(0.003)
Mean	0.6	667.2	159233.6	0.8	0.2
# of observations	1,371,307	290,431	290,431	290,431	11,223
Caseworker Age FEs	Yes	Yes	Yes	Yes	Yes
Office×Year× Job seeker Age<25 FEs	Yes	Yes	Yes	Yes	Yes

Notes: The table reports instrumental-variables estimates of the relationship between actual caseworker gender and job seeker outcomes (columns 1–4), and OLS estimates of the relationship between caseworker gender and the number of job seekers per contracted monthly hour (column 5). In columns 1–4, identification exploits as-if random assignment of job seekers to caseworkers based on date of birth in offices that use this allocation rule. Actual caseworker gender refers to the gender of the caseworker to whom the job seeker is ultimately assigned. Predicted caseworker gender is defined as the gender implied by a strict application of the date-of-birth assignment rule without exceptions and is used as an instrument for actual caseworker gender. Predicted caseworker age fixed effects are used as instruments for actual caseworker age in columns 1–4. All specifications include fixed effects for interactions between year, office, and an indicator for job seeker age below 25. Column 1 includes all job seekers registered in date-of-birth offices; columns 2–4 restrict the sample to job seekers who exit PES into employment; column 5 includes caseworkers who met at least one job seeker in each month of September, October, and November. Earnings are measured in Swedish kronor (SEK), where SEK 100 equals approximately EUR 8 as of June 2024. Standard errors, clustered at the caseworker level, 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)					
Productivity^2		-0.000					
		(0.000)					
Productivity at $t-1$			-0.000				
			(0.000)				
Productivity quintile				0.001**			
				(0.001)			
Productivity quintile at $t-1$					0.003***		
					(0.001)		
Productivity (=1 if above median)						0.004**	
						(0.002)	
Productivity at $t-1$ (=1 if above median)							0.006*
							(0.004)
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	5,660	5,660	2,744	5,660	2,744	5,660	1,774
Caseworker Age Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Office × Year Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes

Notes: The table reports estimates of the relationship between caseworkers' full-time–equivalent monthly wage rates and various measures of productivity. Productivity measure is constructed as described in section 4.2. The sample is restricted to caseworkers employed in offices that assign job seekers using the date-of-birth allocation rule, which are the offices for which productivity can be measured. Columns 1, 2, 4, and 6 use the full sample of caseworkers with productivity observed in at least one year, while columns 3, 5, and 7 restrict the sample to caseworkers with productivity observed in two consecutive years. All models include interactions between year and office fixed effects, as well as caseworker 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 8: Gender log wage gap among caseworkers

	Log(Wage)			
	(1)	(2)	(3)	(4)
Female	-0.002	0.001	0.001	0.001
	(0.002)	(0.002)	(0.002)	(0.002)
University Degree	0.020***	0.020***	0.020***	0.020***
	(0.006)	(0.006)	(0.006)	(0.006)
Secondary Degree	0.010	0.009	0.009	0.010*
	(0.006)	(0.006)	(0.006)	(0.006)
Business degree	0.001	0.001	0.001	0.001
	(0.002)	(0.002)	(0.002)	(0.002)
Social degree	-0.002	-0.002	-0.002	-0.002
	(0.002)	(0.002)	(0.002)	(0.002)
Immigrant	-0.010***	-0.011***	-0.011***	-0.011***
	(0.002)	(0.002)	(0.002)	(0.002)
Experience				
2-4 years	0.018***	0.019***	0.019***	0.018***
	(0.003)	(0.003)	(0.003)	(0.003)
4-6 years	0.042***	0.044***	0.044***	0.043***
	(0.003)	(0.003)	(0.003)	(0.003)
6-8 years	0.065***	0.066***	0.066***	0.066***
	(0.004)	(0.004)	(0.004)	(0.004)
8-10 years	0.078***	0.081***	0.081***	0.080***
	(0.004)	(0.004)	(0.004)	(0.004)
+10 years	0.101***	0.103***	0.103***	0.103***
	(0.003)	(0.003)	(0.003)	(0.003)
Hours Worked (per week)		0.002***	0.002***	0.002***
		(0.000)	(0.000)	(0.000)
Productivity (=1 if above median)			0.003*	0.003*
			(0.001)	(0.001)
At least one child < 4 years				0.007***
				(0.003)
1 Child				-0.002
				(0.002)
2 Children				0.003
				(0.003)
3+ Children				0.003
				(0.004)
Mean log(wage)	10.0	10.0	10.0	10.0
Mean (wage)	22,087.0	22,087.0	22,087.0	22,087.0
# of observations	5,660	5,660	5,660	5,660
Caseworker Age Fixed Effects	Yes	Yes	Yes	Yes
Office × Year Fixed Effects	Yes	Yes	Yes	Yes

Notes: The table reports estimates of the relationship between caseworkers' full-time-equivalent log monthly wages and caseworker characteristics including productivity. The sample is restricted to caseworkers employed in offices that use date-of-birth allocation and for whom productivity measures can be constructed as described in section 4.2. All models incorporate interactions between year and office fixed effects, as well as caseworker age fixed effects. Earnings 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 9: Gender log earnings gap among caseworkers

	Log(earnings)				
	(1)	(2)	(3)	(4)	(5)
Female	-0.078***	-0.075***	-0.035***	-0.035***	-0.039***
	(0.006)	(0.005)	(0.004)	(0.004)	(0.004)
University Degree	0.004	-0.025	-0.016	-0.016	-0.018
	(0.021)	(0.018)	(0.014)	(0.014)	(0.014)
Secondary Degree	0.012	-0.002	-0.002	-0.002	-0.006
	(0.021)	(0.018)	(0.014)	(0.014)	(0.014)
Business degree	-0.004	-0.005	-0.001	-0.001	-0.000
	(0.007)	(0.006)	(0.005)	(0.005)	(0.005)
Social degree	-0.006	-0.003	-0.001	-0.001	-0.000
	(0.008)	(0.007)	(0.006)	(0.006)	(0.006)
Immigrant	0.002	0.016**	0.007	0.007	0.007
	(0.008)	(0.007)	(0.006)	(0.006)	(0.006)
Experience					
2-4 years	0.007	-0.019**	-0.008	-0.008	-0.006
	(0.010)	(0.010)	(0.008)	(0.008)	(0.008)
4-6 years	-0.007	-0.068***	-0.032***	-0.032***	-0.020**
	(0.012)	(0.012)	(0.010)	(0.010)	(0.009)
6-8 years	0.043***	-0.052***	-0.016	-0.016	-0.010
	(0.012)	(0.012)	(0.011)	(0.011)	(0.010)
8-10 years	0.036***	-0.077***	-0.021**	-0.021**	-0.017*
	(0.014)	(0.014)	(0.011)	(0.011)	(0.010)
+10 years	0.036***	-0.110***	-0.044***	-0.044***	-0.043***
	(0.010)	(0.012)	(0.010)	(0.010)	(0.010)
Log(wages)		1.450***	1.220***	1.220***	1.246***
		(0.058)	(0.043)	(0.043)	(0.044)
Hours worked (per week)			0.026***	0.026***	0.025***
			(0.001)	(0.001)	(0.001)
Productivity (=1 if above median)				0.001	0.001
				(0.004)	(0.004)
At least one child < 4 years					-0.085***
					(0.008)
1 Child					-0.018***
					(0.006)
2 Children					-0.036***
					(0.007)
3+ Children					-0.050***
					(0.011)
Mean	12.4	12.4	12.4	12.4	12.4
# of observations	5,660	5,660	5,660	5,660	5,660
Caseworker Age Fixed Effects	Yes	Yes	Yes	Yes	Yes
Office× Year Fixed Effects	Yes	Yes	Yes	Yes	Yes

Notes: The table reports estimates of the relationship between caseworkers' log annual earnings and caseworker characteristics including productivity. The sample is restricted to caseworkers employed in offices that use date-of-birth allocation and for whom productivity measures can be constructed as described in section 4.2. All models incorporate interactions between year and office fixed effects, as well as caseworker age fixed effects. Earnings 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 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	-0.004	-0.004	-0.003	-0.003	-0.004
	(0.003)	(0.003)	(0.003)	(0.004)	(0.008)
University Degree	0.007	0.008	0.008	0.013	0.009
	(0.005)	(0.005)	(0.005)	(0.008)	(0.009)
Secondary Degree	0.002	0.002	0.002	0.009	0.002
	(0.005)	(0.005)	(0.005)	(0.008)	(0.010)
Business degree	0.004	0.n004	0.004	0.002	0.003
	(0.003)	(0.003)	(0.003)	(0.004)	(0.008)
Social degree	0.002	0.002	0.002	-0.000	0.004
	(0.005)	(0.005)	(0.005)	(0.005)	(0.010)
Immigrant	-0.005*	-0.006*	-0.006*	-0.000	-0.011
	(0.003)	(0.003)	(0.003)	(0.005)	(0.007)
Experience					
2-4 years	0.022***	0.022***	0.022***	0.018*	0.036**
	(0.008)	(0.008)	(0.008)	(0.010)	(0.015)
4-6 years	0.015*	0.015*	0.015**	0.011	0.031**
	(0.007)	(0.007)	(0.008)	(0.011)	(0.015)
6-8 years	0.010	0.011	0.011	0.024**	0.004
	(0.008)	(0.008)	(0.008)	(0.012)	(0.012)
8-10 years	0.020**	0.020**	0.021**	-0.001	0.039**
	(0.009)	(0.009)	(0.009)	(0.010)	(0.016)
+10 years	0.014***	0.014***	0.015***	0.015**	0.019*
	(0.005)	(0.005)	(0.005)	(0.007)	(0.010)
Productivity (=1 if above median)		-0.004	-0.004	-0.002	-0.014**
		(0.003)	(0.003)	(0.003)	(0.007)
Hours worked (per week)			0.000*	0.000	0.001
			(0.000)	(0.000)	(0.001)
Mean	0.009	0.009	0.009	0.006	0.015
# of observations	5,274	5,274	5,274	3,184	1,900
# of unique caseworkers	2,784	2,784	2,784	1,750	1,124
Caseworker Age Fixed Effects	Yes	Yes	Yes	Yes	Yes
Office× Year Fixed Effects	Yes	Yes	Yes	Yes	Yes

Notes: The table presents estimates of the relationship between caseworker promotions to senior positions and caseworker characteristics including productivity and hours worked. The sample is restricted to caseworkers employed in offices that use date-of-birth allocation and for whom productivity measures can be constructed as described in section 4.2. Columns 4 and 5 separate the sample into caseworkers with and without children. All models include interactions between year and office fixed effects, as well as caseworker age fixed effects. Robust standard errors are shown in parentheses. * p<0.1, ** p<0.05, *** p<0.01.

Appendix: Additional figure and tables

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

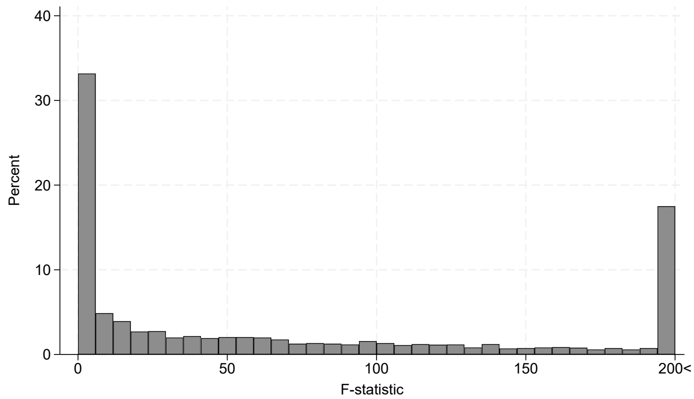

Notes: The figure reports F-statistics from joint significance tests of caseworker dummies in regressions of job seekers' day of birth within each office and year. An F-statistic above 100 indicates the presence of a date-of-birth rule in the allocation of job seekers.

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

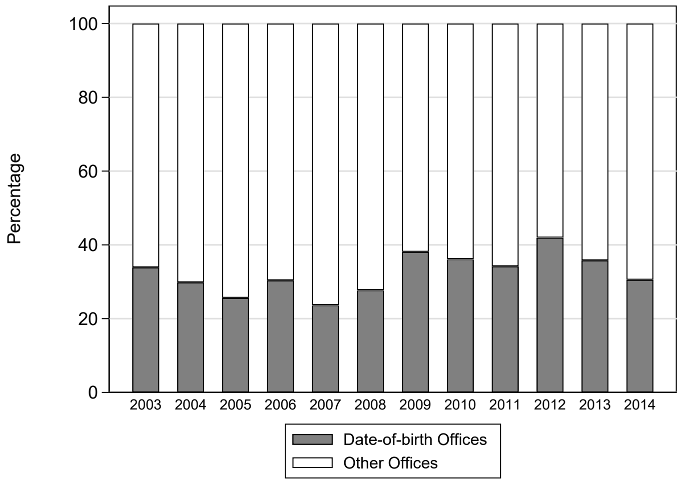

Notes: The figure shows the fraction of date-of-birth allocation offices (F-statistic ≥ 100) and other offices (F-statistic < 100) over time.

Figure A3: Productivity quintiles and log wages

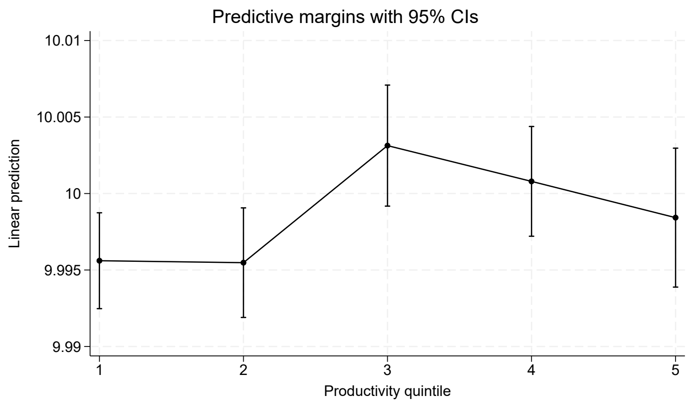

Notes: The figure plots estimates of the relationship between caseworker productivity quintiles (represented as dummies) and caseworkers' log wages. The analysis controls for age fixed effects and interactions between year and office fixed effects, using the sample of caseworkers in date-of-birth rule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

Table A1: First stage estimates for predicted and actual caseworker characteristics

	Dependent variables: Actual caseworker characteristics					
	(1)	(2)	(3)	(4)	(5)	(6)
	Age	Female	Immigrant	University degree	Secondary degree	Experience Experience
Predicted caseworker age	0.4478***	0.0001	0.0003	0.0006	-0.0006	-0.0113
	(0.0111)	(0.0005)	(0.0004)	(0.0004)	(0.0004)	(0.0086)
Predicted caseworker female	-0.0594	0.4561***	-0.0038	0.0022	-0.0005	-0.0250
	(0.1410)	(0.0083)	(0.0052)	(0.0066)	(0.0064)	(0.1276)
Predicted caseworker immigrant	0.0857	-0.0115	0.4720***	0.0041	-0.0039	0.0017
	(0.1960)	(0.0102)	(0.0137)	(0.0086)	(0.0084)	(0.1830)
Predicted caseworker university degree	-0.0765	0.0246	-0.0043	0.4810***	0.0028	0.2471
	(0.5919)	(0.0265)	(0.0181)	(0.0320)	(0.0203)	(0.8188)
Predicted caseworker secondary degree	-0.2873	0.0224	-0.0063	0.0065	0.4776***	0.2009
	(0.5961)	(0.0268)	(0.0185)	(0.0328)	(0.0218)	(0.8391)
Predicted caseworker tenure	-0.0130	-0.0001	0.0001	0.0003	-0.0002	0.4419***
	(0.0106)	(0.0005)	(0.0004)	(0.0005)	(0.0005)	(0.0150)
# of observations	1,371,307	1,371,307	1,371,307	1,371,307	1,371,307	1,371,307
F-statistic	412.714	506.475	207.054	370.538	345.477	239.577

Notes: The table reports first-stage estimates relating actual caseworker characteristics to predicted caseworker characteristics implied by the date-of-birth assignment rule. Each column corresponds to a different actual caseworker characteristic used as the dependent variable. Predicted caseworker characteristics are constructed under a strict application of the date-of-birth assignment rule without exceptions. The sample includes caseworkers employed in offices that allocate job seekers using the date-of-birth rule. All models incorporate interactions between year, office, and age (below 25) fixed effects. Standard errors, clustered at the caseworker level, in parentheses. * p<0.1, ** p<0.05, *** p<0.01.

Table A2: Gender gaps using alternative measures of productivity related to exits from unemployment

	(1) Leave unemployment within 30 days	(2) Leave unemployment within 90 days	(3) log (days in unemployment)
Female	0.0019	0.0064**	-0.0200**
	(0.0018)	(0.0029)	(0.0080)
University degree	-0.0089*	-0.0061	0.0157
	(0.0052)	(0.0095)	(0.0256)
Secondary degree	-0.0052	-0.0009	-0.0040
	(0.0053)	(0.0096)	(0.0264)
Business degree	-0.0002	-0.0014	0.0052
	(0.0022)	(0.0035)	(0.0094)
Social degree	-0.0013	0.0010	-0.0020
	(0.0025)	(0.0042)	(0.0113)
Any child < 4 years old	0.0000	0.0001	-0.0056
	(0.0029)	(0.0047)	(0.0125)
1 Child	0.0007	0.0004	0.0010
	(0.0025)	(0.0040)	(0.0107)
2 Children	-0.0034	-0.0070	0.0219*
	(0.0027)	(0.0045)	(0.0122)
+3 Children	-0.0114***	-0.0078	0.0400**
	(0.0040)	(0.0067)	(0.0180)
Experience			
2-4 years	0.0008	-0.0052	0.0127
	(0.0033)	(0.0054)	(0.0136)
4-6 years	0.0068*	0.0031	0.0064
	(0.0038)	(0.0062)	(0.0159)
6-8 years	0.0054	0.0065	-0.0111
	(0.0042)	(0.0068)	(0.0177)
8-10 years	0.0074*	0.0018	-0.0126
	(0.0042)	(0.0069)	(0.0177)
+10 years	0.0026	0.0041	-0.0030
	(0.0034)	(0.0055)	(0.0146)
Immigrant	-0.0018	-0.0024	0.0139
	(0.0023)	(0.0041)	(0.0111)
Mean	0.1	0.4	4.8
# of observations	1,371,307	1,371,307	1,371,307
Caseworker Age Fixed Effects	Yes	Yes	Yes
Office×Year× Job seeker Age<25 Fixed Effect	Yes	Yes	Yes

Notes: The table reports instrumental-variables estimates of the relationship between actual caseworker characteristics and job seeker outcomes. The outcomes are exiting unemployment within 30 days (column 1), exiting unemployment within 90 days (column 2), and the log number of days the job seeker remains unemployed (column 3). Identification exploits as-if random assignment of job seekers to caseworkers based on date of birth in the offices that use this allocation (random offices). Actual caseworker characteristics refer to the characteristics of the caseworker to whom the job seeker is ultimately assigned. Predicted caseworker characteristics correspond to the caseworker the job seeker would have been assigned to under a strict application of the date-of-birth assignment rule, without exceptions, and are used as instruments for actual 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, in parentheses. * p<0.1, ** p<0.05, *** p<0.01.

Table A3: Gender, wages, and productivity - alternative specifications of productivity (1)

	(1) Log(wage)	(2) Log(wage)	(3) Log(wage)	(4) Percentage change in wages	(5) Log(wage)	(6) Log(wage)
Female	-0.002	-0.002	-0.000	0.115	-0.002	-0.000
	(0.002)	(0.002)	(0.003)	(0.108)	(0.002)	(0.003)
Productivity	0.000	0.001				
	(0.000)	(0.001)				
Productivity^2		-0.000				
		(0.000)				
Productivity at $t-1$			0.000	0.003		
			(0.000)	(0.010)		
Productivity quintile					0.001**	
					(0.001)	
Productivity quintile at $t-1$						0.003***
						(0.001)
Mean	10.0	10.0	10.0	2.5	10.0	10.0
# of observations	5,660	5,660	2,744	2,744	5,660	2,744
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 quintiles are calculated within each year and office. All models include the control variables discussed in section 4 and interactions between year and office fixed effects, as well as caseworker 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 (2)

	Log(wage)			
	(1)	(2)	(3)	(4)
Female	-0.002	-0.002	-0.002	-0.002
	(0.002)	(0.002)	(0.002)	(0.002)
Productivity 30 (=1 if above median)	0.000			
	(0.001)			
Productivity 90 (=1 if above median)		0.002		
		(0.001)		
Productivity 180 (=1 if above median)			0.003**	
			(0.001)	
Productivity log(duration) (=1 if above median)				-0.001
				(0.001)
Mean	10.0	10.0	10.0	10.0
# of observations	5,660	5,660	5,660	5,660
Controls	Yes	Yes	Yes	Yes
Caseworker 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 and interactions between year and office fixed effects, as well as caseworker 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

	Outcome: Log(wage)		
	(1)	(2)	(3)
	F-stat≥ 100	F-stat≥ 50	F-stat≥ 200
Female	-0.002	-0.001	-0.006**
	(0.002)	(0.001)	(0.002)
Productivity	0.003**		
	(0.001)		
Productivity 50		0.002	
		(0.001)	
Productivity 200			0.005**
			(0.002)
Mean	10.0	10.0	10.0
# of observations	5,660	7,834	3,242
# of observations (unique)	2,939	3,667	1,956
Controls	Yes	Yes	Yes
Caseworker 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 and include interactions between year and office fixed effects, as well as caseworker 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

	Log(Wage)			
	(1) 0-2 years	(2) 3-5 years	(3) 6-10 years	(4) +11 years
Female	0.001	-0.001	0.010*	-0.004
	(0.003)	(0.005)	(0.005)	(0.003)
Productivity (=1 if above median)	0.002	-0.003	-0.004	0.005*
	(0.002)	(0.006)	(0.006)	(0.003)
University Degree	0.063**	0.001	0.023	0.014*
	(0.026)	(0.013)	(0.015)	(0.008)
Secondary Degree	0.060**	0.009	0.019	0.005
	(0.027)	(0.014)	(0.015)	(0.008)
Immigrant	-0.006**	-0.012**	0.010	-0.021***
	(0.003)	(0.006)	(0.008)	(0.005)
Business degree	-0.003	0.006	0.013*	0.001
	(0.003)	(0.006)	(0.007)	(0.003)
Social degree	-0.004	-0.001	-0.005	0.013*
	(0.003)	(0.006)	(0.008)	(0.007)
At least one child $<$ 4 years	0.008*	-0.000	0.007	0.005
	(0.005)	(0.008)	(0.008)	(0.008)
1 Child	0.001	-0.012	-0.009	-0.006
	(0.004)	(0.008)	(0.008)	(0.005)
2 Children	-0.002	-0.014	0.004	0.004
	(0.005)	(0.011)	(0.009)	(0.006)
3+ Children	0.001	0.006	0.016	-0.012
	(0.007)	(0.013)	(0.012)	(0.009)
Mean	9.9	10.0	10.0	10.0
# of observations	728	577	603	2,593
Caseworker Age Fixed Effect	Yes	Yes	Yes	Yes
Workplace $\times$ Year Fixed Effect	Yes	Yes	Yes	Yes

Notes: The table presents estimates of the gender log wage gap in different brackets of tenure, using a sample of caseworkers from date-of-birth rule offices for whom the productivity measure is available. All models include interactions between year and office fixed effects, as well as caseworker 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 A7: Gender exit gap

	Exiting PES				
	All caseworkers	All caseworkers	All caseworkers	Caseworkers without children	Caseworkers with children
	(1)	(2)	(3)	(4)	(5)
Female	0.001	0.001	-0.001	0.003	-0.018
	(0.006)	(0.006)	(0.007)	(0.008)	(0.015)
University Degree	0.002	0.001	0.001	0.001	0.012
	(0.018)	(0.018)	(0.018)	(0.023)	(0.044)
Secondary Degree	0.003	0.002	0.003	-0.006	0.028
	(0.018)	(0.018)	(0.018)	(0.023)	(0.043)
Business degree	-0.006	-0.007	-0.007	0.006	-0.029**
	(0.007)	(0.007)	(0.007)	(0.008)	(0.014)
Social degree	-0.001	-0.001	-0.001	-0.006	0.003
	(0.010)	(0.010)	(0.010)	(0.013)	(0.019)
Immigrant	-0.026***	-0.025***	-0.025***	-0.003	-0.074***
	(0.009)	(0.009)	(0.009)	(0.012)	(0.018)
Experience					
2-4 years	-0.031**	-0.031**	-0.032**	-0.036*	-0.003
	(0.016)	(0.016)	(0.016)	(0.020)	(0.027)
4-6 years	-0.027	-0.027	-0.028	-0.041	-0.047
	(0.019)	(0.019)	(0.019)	(0.028)	(0.029)
6-8 years	-0.056***	-0.056***	-0.058***	-0.052*	-0.067**
	(0.017)	(0.017)	(0.017)	(0.028)	(0.027)
8-10 years	-0.061***	-0.061***	-0.063***	-0.066**	-0.064**
	(0.017)	(0.017)	(0.017)	(0.026)	(0.027)
+10 years	-0.064***	-0.064***	-0.066***	-0.080***	-0.059**
	(0.014)	(0.014)	(0.015)	(0.021)	(0.025)
Productivity		0.005	0.005	0.004	0.004
(=1 if above median)		(0.006)	(0.006)	(0.007)	(0.012)
Hours Worked			-0.001	-0.002**	0.001
(per week)			(0.001)	(0.001)	(0.002)
Mean	0.049	0.049	0.049	0.041	0.064
# of observations	4,852	4,852	4,852	2,924	1,740
# of unique caseworkers	2,574	2,574	2,574	1,614	1,043
Caseworker Age Fixed Effect	Yes	Yes	Yes	Yes	Yes
Workplace × Year Fixed Effect	Yes	Yes	Yes	Yes	Yes

Notes: The table presents estimates of the relationship between caseworkers' exit from PES and caseworker characteristics, productivity, and hours worked. The sample includes only caseworkers for whom a productivity measure can be calculated. Columns 4 and 5 separate the sample into caseworkers with and without children. All models include interactions between year and office fixed effects, as well as caseworker age fixed effects. Robust standard errors are shown in parentheses. * p<0.1, ** p<0.05, *** p<0.01.