Should I stay or should I go? The unintended consequences of pay gap transparency on female workers*

Tom Kemeny[1], Max Nathan[2,3], Ceren Ozgen[4], Guido Pialli[2,5], Anna C. Rosso[†6], Mateo Seré[2] and Anna Valero[3,7]

[1]University of Toronto, Canada
[2]University College London, United Kingdom
[3]Centre for Economic Performance, United Kingdom
[4]University of Birmingham, United Kingdom
[5]University of Turin, Italy
[6]University of Milan & LdA, Italy
[7]London School of Economics and Political Science, United Kingdom

April 24, 2026

Abstract

This paper examines the effects of the UK's 2017 gender pay transparency policy, which mandated large firms to disclose gender pay gap data. Using a novel employer-employee dataset and a difference-in-differences design around the treatment threshold, we find that the policy led to increased job mobility among women, especially low-skilled, in non-managerial roles and in the lower end of the wage distribution. Our findings suggest that transparency influences workforce composition, contributing to a narrowing of the gender pay gap—primarily through separation from women with weaker bargaining power. We conclude that policymakers should consider complementing pay transparency policies with other tools to avoid unintended consequences from such reforms.

JEL classification: D22, J16, J24, J31, L25

Keywords: pay transparency; gender equality; gender pay gap; workforce composition

*This is a draft paper - final results may change. Please do not cite without permission. Audiences at Collegio Carlo Alberto (Turin, IT) provided helpful comments. Thanks to Rebecca Lee at OpenCorporates, Filipe Mesquita at Diffbot, Paul Longley and Justin van Dijk at UCL for help and advice on data/code, and to our advisory board for constructive feedback. This project uses data from Diffbot, OpenCorporates and Orbis / Orbis Historical. Thanks to all data providers. This research is funded by UKRI Grant ES/W010232/1. The project was reviewed and approved by the UCL Research Ethics Board, application 22883/00. This paper represents the views of the authors, not the advisers, data providers or funders. Other usual disclaimers apply.

†Corresponding author. E-mail: anna.rosso@unimi.it

Author contributions

Tom Kemeny: conceptualisation, validation, writing - reviewing and editing; **Max Nathan:** conceptualisation, funding acquisition, methodology, project administration, software, writing - reviewing and editing; **Ceren Ozgen:** conceptualisation, writing - reviewing and editing; **Guido Pialli:** conceptualisation, data curation, formal analysis, investigation, methodology, software, writing - original draft; **Anna Rosso:** conceptualisation, data curation, formal analysis, investigation, methodology, software, supervision, validation, writing - original draft; **Mateo Seré:** data curation, software; **Anna Valero:** conceptualisation, methodology, writing - reviewing and editing. Authors are listed alphabetically.

1 Introduction

After the UK introduced mandatory gender pay gap reporting in April 2017, the BBC published its own Gender Pay Report. In January 2018, Carrie Gracie resigned as the BBC's China Editor, after learning that the other two male international editors were paid 50% more than she was, and she denounced what she described as a "secretive and illegal" pay culture.[1] This episode shows how newly disclosed pay gaps, once perceived as unfair, can trigger consequential career decisions.

Pay transparency policies have been adopted widely: over half (21 on 38) of the OECD countries now require private sector employers to make available gender pay gap information to the broad audience (OECD, 2023).[2] At the same time, gender pay gaps have been slowly declining. The ONS estimates that the UK's median gender pay gap decreased to 7% in April 2024, from 12.2% in 2009, while OECD data show a fall from 16% in 2005 to 11% in 2023, and from 14.1% to 9% during the same period within the EU-27.[3] Despite this progress, gender pay disparities remain a concern for employers, policymakers, and workers, sustaining an active policy debate about how best to reduce it. Gender pay gap transparency policies are believed to narrow gender pay gaps through exerting public pressure on employers to reduce them or encouraging lower-paid women to bargain for higher wages.

While identifying the first-order effect of pay transparency rules is important, understanding second-order behavioral responses is arguably more consequential from a policy perspective. If firms comply with pay transparency mandates in ways that only minimize their costs, wage inequality may shrink without improving women's earnings or welfare. Ignoring these behavioral adjustments risks overstating the policy's effectiveness and lead policymakers to draw incorrect conclusions about its distributional and efficiency consequences. We identify two main channels through which pay gap disclosure can alter labour market outcomes.

Regarding workers, as in the Carrie Grace case, the disclosure of gender pay disparities generates an information shock in a domain that is typically hidden or opaque: the firm's pay-setting

[1] See https://www.ft.com/content/dff6f1c8-06b0-11e8-9650-9c0ad2d7c5b5 (accessed 22 February 2026).

[2] Cullen (2024) distinguishes between three forms of pay transparency: horizontal pay transparency, where coworkers within the same organization know each other's pay; vertical pay transparency, when workers are informed about the pay of more senior co-workers; cross-firm transparency, where employees get access to the pay of workers in other competing firms.

[3] https://www.oecd.org/en/data/indicators/genderwagegap.html

practices and the extent of gender-related disparities (Cullen and Perez-Truglia, 2023). By revealing information about relative pay and, more broadly, about the firm's compensation culture, disclosure changes women's beliefs about the returns to remaining at the firm. Card et al. (2012) show that random information on pay differences influence both job satisfaction and job market behaviour. Moreover, they also document that employees whose pay falls below the median in their department report lower job satisfaction, and are more likely to search for a new job following the disclosure of salary information. Because women are more likely than men to be on the lower part of the within-firm wage distribution — a stylised fact consistent with the large literature documenting residual gender pay gaps driven by differences in bargaining power, risk preferences, and outside options (Bennedsen et al., 2023; Blau and Kahn, 2017, 2023; Goldin, 2014; Manning and Saidi, 2010) — the informational shock of disclosure falls disproportionately on female workers.

In principle, transparency should empower underpaid female workers to negotiate higher salaries. In practice, however, women typically have weaker bargaining power and less aggressive negotiation behaviour than men (Manning and Saidi, 2010; Dittrich et al., 2014; Leibbrandt and List, 2015). When internal wage renegotiation is constrained by bargaining frictions or pay rigidity, the most salient adjustment margin may be separation rather than wage revision. In this case, transparency affects not only wages, but also retention, by reducing job satisfaction and perceived fairness and by increasing the salience of outside options (Card et al., 2012).

From the firm's perspective, the obligation to publish gender pay gap indicators can be understood as a labor market friction that makes a part of the workforce more rigid and costlier to employ, increasing operational costs and financial leverage (Serfling, 2016; Bena et al., 2022). Pressure from employees, the media and the public to narrow the disclosed gender pay gap raised the implicit cost of maintaining large pay disparities.[4] Firms may thus respond along several margins. First, they may restrain wage growth for high-paid (male) employees, compressing the gender pay gap from above without raising female pay, a mechanism documented by Blundell et al. (2025) for the UK context. Second, and more consequentially for the workers the policy is intended to help, firms may adjust the composition of their workforce. Because low-paid, low-skilled female employees con-

[4] Cullen and Perez-Truglia (2023) show that when salary differences are disclosed, firms are less inclined to offer high wages to underpaid employees, since doing so could trigger pressure on wage increases from other employees. As a result, overall wage levels may decline and wage dispersion may shrink.

tribute most to the measured gender pay gap and have the weakest bargaining power, firms seeking to narrow the reported gap at minimum cost may disproportionately separate from these workers. In this scenario, the gender pay gap declines not because women's wages rise, but because the lowest-paid women exit.

Both channels — workers' exit driven by discontent and firms' compositional adjustment — predict a common empirical pattern: an increase in female separations concentrated among low-paid, low-skilled workers. These two mechanisms are difficult to separate empirically, and they may well operate simultaneously. What matters for policy evaluation is that both imply a reduction in the measured gender pay gap that does not reflect an improvement in women's labour market position.

This paper studies how the gender pay transparency policy introduced in the UK in 2017 affected female workers' labour market behaviour. The policy required companies with more than 250 employees to publish gender pay gap indicators with the aim of decreasing the pay gap in companies where it existed and, potentially, in the economy as a whole. To address this question, we use a newly assembled employer-employee database that combines worker profiles from the web with firm-level financial information. This database and related validation exercises are described in Gray et al. (2025).

Following Blundell et al. (2025), we exploit the firm-size cutoff to provide casual estimates of how the policy affects the gender composition of the workforce. Specifically, we analyze job mobility among individuals employed at firms just above or just below the 250-employee threshold, exploiting variation in policy exposure across firm size and over time. While Blundell et al. (2025) examine the effects of the reform on the gender pay gap, we focus on job mobility, leavers' characteristics and destination jobs/employers. Compared to administrative data, our database allows us to observe complete career histories rather than only job flows across firms in the estimation sample, enabling us to identify where female leavers move and to shed light on mechanisms underlying exit decisions.

Because firms may endogenously adjust headcounts in anticipation of the policy, we define the treatment status in 2015, two years before the policy implementation. We restrict the analysis to firms with +-50 employees from the 250-employee threshold and track workers on these firms before and after the policy. To control for unobserved heterogeneity, we include individual, firm, and gender-region-year fixed effects, which absorb time-invariant individual and firm unobserved heterogeneity, and allow for regional shocks that affect women and men differently.

We report four main findings. First, consistent with Blundell et al. (2025), women in treated firms are more likely to change employer after the policy compared to control firms both in absolute terms and relative to men. The policy led to a significant 1.4 percentage points (about the 10% of the pre-policy probability) increase in women's probability to leave the firm compared to men's probability. This effect is robust to alternative bandwidth selections and specifications.

Second, we find that the policy weakly increased men's probability of being hired, with no impact on women's overall probability of being hired. Relatedly, we find no effect of the policy on the probability that women are promoted to higher paid position or to managerial roles, either across or within firms.

Third, we provide evidence on the mechanisms behind women's separation rates. Specifically, consistent with our assumptions, women who leave are more likely to be in the lower part of the wage distribution — measured by the occupation's average earnings — to work in non-managerial occupations, and to be lower-skilled than men in treated firms. These findings are consistent with both channels outlined above: whether driven by worker discontent or firm adjustments, the departures are concentrated among precisely those women for whom the informational shock is most salient and the bargaining position weakest. To shed more light on which of the two channels above is most likely, we distinguish female leavers based on the unionization rates of their employer's sector. We find that the effect of the policy is stronger for female workers in less-unionized sectors.

Fourth, we document that women are more likely to move into occupations with a lower average earnings than its previous employment spell. In a further version of the paper, we also look at the specific firms where women are going, especially whether these are gender-friendly firms or not.

This paper contributes to a small but growing empirical literature on the labour market effects of pay transparency policies. Bennedsen et al. (2022) study a legislation change in Denmark that mandated firms with more than 35 employees to report salary information separately for men and women. Using employer-employee administrative data, Bennedsen et al. (2022) find a reduction of the gender pay gap, an increasing share of hired women and greater probability of promoting women in higher paid positions in treated firms relative to control firms. Baker et al. (2023) examine the effects of the introduction of public sector disclosure laws across Canadian provinces and find a reduction in the gender gap among full-time academic employees at Canadian universities in treated provinces. However, their findings also suggest that the reduction in the gender pay gap is driven by a slower relative growth of men's salaries. On the contrary, Gulyas et al. (2023) study the effects of the 2011 Austrian pay transparency law and do not find any economically significant effect on the gender pay gap or in separation rates after the reform. They conclude that revealing wages might have actually increased job satisfaction by alleviating previous concerns about unfair salary differences. Gamage et al. (2024) study the impact of a gender pay publication by the Times Higher Education guide in 2007, finding a reduction in the gender pay gap among academics after the policy. The estimated effect is entirely driven by an increase in female wages. By leveraging a German law that allowed employees in large firms to request information on coworkers' salaries, Brütt and Yuan (2022) do not find that the policy significantly affected the gender pay gap, nor it changed the workforce composition of treated firms.Employing proprietary data from Revelio, Dambra et al. (2025) and Liang et al. (2025) study the Pay Ratio Disclosure (PRD) law in the US to show that disclosing firm pay increases job mobility, especially in the lower-end of the wage and rank distribution within firm. Our contribution is to provide causal evidence on the second-order responses — worker exits, their characteristics, and destination employers — that underlie the aggregate effects documented in prior work.

This paper is structured as follows. Section 2 introduces the UK pay transparency policy and highlights its salience in the UK context. Section 3 describes data and empirical design. Section 4 presents the main results. We discuss extensions and robustness checks in Section 5. Finally, Section 6 summarizes the conclusions and discusses policy implications.

2 The UK Gender Pay Gap policy

In 2015, the UK government started a consultation with employers to design a policy reform aimed at enhancing pay transparency. As reported by the Government, the objective was to encourage employers to adopt workplace policies and practices that promote gender equality (Government Equalities Office, 2016). Following positive feedback from the consultation, in February 2017, the

UK Government enacted the Equality Act 2010 (Gender Pay Gap Information) Regulations 2017 No. 172.[5] The regulations mandated all firms registered in Great Britain with at least 250 employees to publish gender equality indicators - both on their own website and on a dedicated portal managed by the Government Equalities Office - by 30 March 2018 for the first reporting year and annually thereafter by the end of the financial year (5 April).[6] The mandated indicators include mean and median gender hourly pay gaps, mean and median gender bonus gaps, the proportion of male and female employees who received bonuses, and the proportion of female employees in each quartile of a company pay distribution.[7] This policy applies to both public and private-sector organizations in England, Wales and Scotland; Northern Ireland is excluded.[8] All companies that are part of a group are required to report indicators individually if they meet thresholds. At the time of the law, 10,500 firms were in the targeted group. Treated firms account for only 0.4% of all UK firms, but 40% of employment and 48% of turnover in the UK economy.[9] Moreover, it is worth to highlight that no other policies targeted this slice of the firm population before the gender pay gap policy.

Importantly, there are no sanctions for firms failing to improve their gender equality indicators over time. However, the Equality and Human Right Commission, responsible for the policy implementation, could issue court orders and impose unlimited fines on firms that failed to report. As reported on the Governemnt's website, 94% of eligible firms complied with the policy by the deadline in the first year. Finally, fewer than 600 non-eligible firms (i.e., those below the 250-employee threshold) voluntarily published the gender equality indicator over the years, representing less than 0.1% of firm in that population, as per the Business Structure Database statistics.

The policy was designed to create an information shock both within and outside organizations. This is also justified by the fact that, according to a survey conducted by the Government Equal-

[5] For the actual legislation, see https://www.legislation.gov.uk/uksi/2017/172/contents/made.

[6] Gender pay gap reports of the covered firms are available at: https://gender-pay-gap.service.gov.uk/search

[7] To compute their gender pay gap indicators, firms must first collect data separately for "relevant employees" and "full-pay relevant employees". Relevant employees are those with a regular employment contract, including workers who are part-time, job sharing, on leave and self-employed. Full-pay employees are those receiving their full basic pay or less than their full basic pay, but not because of leave. For a complete set of instructions for firms to prepare their gender pay gap reports, see: https://www.gov.uk/government/publications/gender-pay-gap-reporting-guidance-for-employers/preparing-your-data.

[8] Public sector entities include most government departments, the armed forces, local authorities, NHS bodies, schools and universities.

[9] See Blundell et al. (2025) for more information on the sample of firms that reported over the years and the evolution of the raw firm-level indicators.

ities Office before the introduction of the policy, firms were not aware of their gender gaps, and employees were discouraged to talk about salaries with their colleagues.

Finally, the policy intended to shame firms into action and it was therefore characterized by substantial salience. Every year, the indicators have received substantial media attention, and reports are posted on firms' main webpages and the government website to ensure accessibility to employees (Blundell et al., 2025). Figure A1 in the Appendix shows two examples of how the gender pay gap indicator is reported in two different firms, one that always complied with the policy and submitted its report on time, while the other one that uploaded its report with delay. Everyone can thus access the website and identify those firms that submitted their reports inaccurately.

Moreover, the policy attracted significant public interest around the time of its implementation. Figure A2 shows Google searches for "gender wage gap" between March 2013 and March 2022. Frequencies are normalized to the peak, which is on April 2018, the time of the first deadline for reporting the gender pay gap indicators. However, the figure also displays a peak of searches in March 2017, the month after the enforcement of the policy.

Figure A3 shows the evolution of two important gender inequality indicators whose publication was imposed by the reform: difference in median hourly wage gap (black line) and percentage of females in the top quartile of the distribution (red line). The graph shows that the level of median pay gap in 2023 was basically the same as in 2018, while the share of female in the top quartile of the wage distribution slightly increased from 37.2% to 38.7%.

3 The effect of the policy

3.1 Data and sample description

To study the effects of the UK gender pay gap policy on individual-level outcomes, we leverage a newly assembled UK employer-employee dataset, which links proprietary individual-level data from Diffbot - a commercial knowledge graph provider - to administrative company-level data. Diffbot extracts individual and company-level profiles from multiple sources across the public web and organises them into a searchable knowledge graph. Individual records include socio-demographic characteristics (gender, year of birth, self-reported languages), education histories, detailed employment histories (employer identifiers, start and end dates, job titles, and brief job descriptions), and a taxonomy of self-reported skills. Diffbot uses supervised learning techniques to assign individual workers to companies based on their career histories. For UK firms, company-level profiles include company names and unique identifiers from Companies House, the UK's Open Companies Register.[10]

We use these company identifiers to merge individual-level data from Diffbot with company-level information from Bureau van Dijk's Orbis Historical, which gives us detailed financial data for companies over a long time frame.[11] Linking these sources yields an employer-employee database covering more than 800,000 individuals and 10,530 'medium' and 'large' companies active between 2007 and 2023.[12] While the Appendix summarises the elements most relevant to our research design, a comprehensive description of data construction and validation is provided in Gray et al. (2025).

To construct the firm sample for our design (see section 3.3), we select organizations around the 250-employee threshold.[13] Then, we compare outcomes for individuals employed in firms whose size is just above the threshold of 250 employees (eligible to treatment) with those in firms just below the threshold. Because consultation on the policy began in 2015 and continued through 2017, we define treatment status using firm size in 2015 to avoid endogenous anticipated responses from those firms seeking to change headcounts to move above or below the threshold. Additionally, to enhance comparability between treatment and control firms, we restrict the estimation sample to firms in a bandwidth of +-50 employees around the 250-threshold.[14]

Our baseline estimation sample includes information for 35,613 employees and 210,450 individual-year observations for the period from 2013 to 2022. We hold the individual sample constant over time, retaining only those individuals observed in all years, thereby forming a balanced panel.

[10]The UK CRN coincides with the ORBIS BvD identifier, which is usually the country-specific registry identifier preceded by the country two-digit code.

[11]Orbis is widely used for firm-level analysis and includes harmonised cross-country financial information on close to 462 million companies worldwide, making it one of the most extensive and reliable sources of corporate financial data.

[12]We classify companies as medium or large following Companies House, which distinguishes 'micro' and 'small' companies entities from medium and large based on thresholds for turnover, assets and employee counts. Further details are available at: https://www.gov.uk/government/publications/life-of-a-company-annual-requirements/life-of-a-company-part-1-accounts.

[13]As it is common in firm-level studies, we exclude firms in agriculture and mining sectors.

[14]We assess the robustness of our results to alternative bandwidth choices. Results are reported in figure A2 in the Appendix.

3.2 Outcome variables

We examine mobility across firms and within firms (via occupational shifts) by assessing how the policy altered workforce composition. In the absence of wage data, we cannot estimate wage effects directly; yet, we can investigate how the policy has reshaped workforce composition and infer wage effects indirectly by examining shifts from lower-paid to better-paid occupations, using occupations' average earnings.

Our first outcome variable is the probability that an individual leaves the firm. Precisely, we build a dummy variable equal to one if an individual leaves the firm the following year and zero otherwise. By construction, this variable is missing for the last year of our sample. We begin with this outcome to align our analysis with prior research. This outcome is salient because pay transparency can affect job satisfaction through bargaining power. If individuals lack the power to renegotiate wages, job satisfaction may fall; conversely, if transparency revises downward workers' priors about their wage gap, job satisfaction may rise. Although our data do not include direct measures of job satisfaction, separation rates can serve as proxies (Card et al., 2012; Gulyas et al., 2023).

Our second outcome relates to the gender composition of new hires. The interest in this outcome comes from the fact that employers could have decided to close the gap by hiring more female workers. In addition, this outcome also allow us to understand how managers become more accountable for new hires and workforce composition, leading them to be more focused on observables (Castilla, 2015). On the other hand, they may hire more women in occupations where they can offer fairer compensations (Bennedsen et al., 2022). We measure this as the probability that an individual is hired by the firm in the current year.

Finally, we study occupational upgrading. In the Danish case, for example, female were less likely to leave treated firms but more likely to be promoted from the bottom of the hierarchy to more senior positions (Bennedsen et al., 2022). To implement a similar strategy, we link 4-digit SOC occupation codes to median occupational earnings by 4-digit SOC code extracted from the aggregated tables of the Annual Survey of Hours and Earnings (ASHE) provided by the UK Office of National Statistics (ONS). Occupational upgrading is measured with a dummy variable equal to one if an individual moves from a lower-paid occupation to a better-paid occupation between $t-1$

and t. Similarly, we also examine promotions into managerial roles by constructing a dummy equal to one if an individual moves from a non-managerial to a managerial position between $t-1$ and t. We classify as managerial all occupations in the aggregated SOC 1-digit category "Managers, directors and senior officials".

3.3 Econometric design

We want to estimate the impact of the UK pay transparency policy on individual labour market outcomes. We follow the identification strategy in Blundell et al. (2025), who estimate the impact of the UK's pay transparency policy on the gender gap in hourly pay and assess whether the effect operates through changes in women's or men's pay. We exploit variation in policy exposure across firm size and time, comparing the evolution of the outcome in firms just above the 250-employee threshold (treated) with firms just below it (controls) after the implementation of the policy. In our baseline analysis, we restrict the sample to firms within a bandwidth of $\pm$ 50 employees.

We employ a triple-difference regression model to estimate the impact of the policy on both men's and women's outcomes. Specifically, the policy effect is estimated for employee i, working in firm f, at time t, for the period from 2013 to 2022:

$$\begin{aligned} Y_{i,f,t} = \beta_1 & \left(Treated_f \times Post_t\right) \\ & + \beta_2 \left(Treated_f \times Post_t \times Female_i\right) \\ & + \beta_3 \left(Treated_f \times Female_i\right) \\ & + \alpha_i + \mu_f + \delta_{j,t} + \chi_{g,r,t} + \epsilon_{i,f,t} \end{aligned} \tag{1}$$

where $Y_{i,f,t}$ is an outcome variable at the individual level, $Post_t$ is a dummy equal to one from 2017 onward and zero before 2017, $Treated_f$ is a dummy equal to one if a firm's number of employees is between 250 and 300 employees in 2015 and 0 if it is between 200 and 249; $Female_i$ is a dummy equal to one if the individual is female (0 otherwise). We control for time-invariant individual characteristics (α_i), such as individual unobserved ability, and for time-invariant firm characteristics (μ_f). We control for time-varying shocks that are common to all firms in the same industry by including 1-digit-industry fixed effects interacted with year fixed effects ($\delta_{j,t}$), as well as for gender-specific regional shocks ($\chi_{g,r,t}$), which capture, for example, time-varying employment practices that are region-and gender-specific. In this baseline specification, we do not include any further controls.

Our parameter of interest is β_2, which identifies how the policy differently affects female outcomes relative to men in treated vis-à-vis control firms. Specifically, it captures the deviation from a parallel trend induced by the policy of the differences between females and males in the outcome variables in treated firms relative to control firms. The coefficient β_1 estimates the effect of the policy on male employees while the sum $\beta_1 + \beta_2$ gives us the effect on female employees. We formally test the effect on female employees by performing a t-test on the sum of the coefficients β_1 and β_2. Finally, the coefficient β_3 tells us the trend in the outcome variable for women in treated firms before the policy implementation. Standard errors are double-clustered at the individual and firm levels.

Our specification includes individual fixed effects to ensure that the estimated impact of the policy is not driven by selection on unobservables— for example, the most or least mobile individuals. By controlling for individual fixed effects, we account for such heterogeneity. Additionally, individual fixed effects also capture the differential sorting of men and women into high-wage firms, given that occupational choice is a major driver of the gender pay gap (Card et al., 2016; Bamieh and Ziegler, 2025). Similarly, we include firm fixed effects to rule out the possibility that our results are driven by differences across firms, such as only firms with highest/lowest employment turnover. This dual control strengthens our confidence that the observed policy effect is not confounded by selection at either the individual or firm-level. Therefore, we identify the effects of the policy within-firm and within-worker, i.e., the additional effect of the reform after controlling for unobserved but time-constant worker and company characteristics. Consequently, our results are not equally driven by individuals with some specific characteristics (i.e. greater preferences for mobility) moving to more flexible firms, which could be different across genders.

The assumption behind our research design is that, in absence of the policy, treated and control firms would have experienced parallel trends in the outcome variables between males and females. To formally test this assumption, we estimate the following dynamic specification:

$$
\begin{aligned}
Y_{i,f,t} = & \sum_{j=2013}^{2021} \beta_j^1 \left(Treated_f \times YEAR_{t=j}\right) \\
& + \sum_{j=2013}^{2021} \beta_j^2 \left(Treated_f \times Female_i \times YEAR_{t=j}\right) \\
& + \alpha_i + \mu_f + \delta_{j,t} + \chi_{g,r,t} + \epsilon_{i,f,t}
\end{aligned} \tag{2}
$$

where $YEAR_{t=j}$ is a binary variable that takes value equal to unity when $t = k$ and zero otherwise. The reference year is 2016, the year before the enforcement of the gender pay gap policy.

3.4 Descriptive statistics

This section provides summary statistics. Table 1 provides summary statistics for individual-level characteristics in our sample for the pre-policy period (2013-2017), separately for treated and control firms, as well as for the full sample. On average, employees in our sample are 44 years old and have 12 years of work experience. The share of female workers is 0.36, while the share of workers with a college degree is 0.69. Therefore, our sample under-represents female workers relative to UK statistics, while it over-represents high-skilled workers. The selection on these characteristics reflects both the paper's focus on large firms, which usually over-represent male high-skilled workers, on average, and the nature of our employer-employee database, as highlighted in Gray et al. (2025).

Several features are worth to notice. The profile of workers in treated and control firms is basically similar yet some slight differences emerge. Control firms have a slightly lower share of college-degree workers, younger workforce, more female and workers are, on average, less experienced. These small differences in firm-level variables are partly captured in our analysis by the inclusion of individual fixed effects. Results from t-tests on the differences in observable characteristics between control and treated firms indicate that college degree share, age, worker experience, promotion to manager and entering the company are the variables for which we observe statistically significant differences between the two groups.

Additionally, we run a linear probability model where the probability to be treated is regressed against the same observable characteristics in Table 1 averaged over the pre-policy period 2013-3017. Results from this test are presented in Table A1 in the Appendix. While age, being a new hire and being promoted to a managerial position are significantly related to the probability of being treated, the overall F-test on these observable characteristics is not statistically significant at standard confidence levels.

Table 1. Summary statistics, pre-policy period (2013-2017)

	Control		Treated		All (Treated + Control)			t-test
	Mean	SD	Mean	SD	Mean	SD	Obs.	p-value
College degree (share)	0.68	0.47	0.70	0.46	0.69	0.46	24,138	0.045
Migrant (share)	0.16	0.36	0.16	0.37	0.16	0.37	20,991	0.451
Age (years)	43.57	10.05	44.10	9.97	43.81	10.02	19,787	0.000
Female (share)	0.37	0.48	0.36	0.48	0.36	0.48	31,461	0.181
Worker experience (years)	12.29	7.91	12.56	7.93	12.41	7.92	31,461	0.002
Separation at $t+1$ (share)	0.13	0.21	0.13	0.20	0.13	0.21	31,461	0.989
Promotion (share)	0.14	0.24	0.14	0.24	0.14	0.24	31,461	0.174
Promotion to manager (share)	0.02	0.09	0.01	0.08	0.02	0.09	31,461	0.000
Enter the company (share)	0.22	0.29	0.21	0.28	0.22	0.28	31,461	0.000

Notes: This table reports mean and standard deviation of covariates, separately for individuals in treatment and control firms, before the implementation of the policy (2013-2017). Variables are averaged over the period 2013-2017, so that there is a single observation for each individual. Treatment status is defined based on firm size in 2015. The sample includes all firms with an employment level of +- 50 around the policy threshold, as defined in the baseline regression.

4 Baseline results

This section presents the baseline results. Section 4.1 discusses the effects of pay gap disclosure on separation rates by gender. Section 4.2 reports the effects of the policy on the gender composition of new hires.

Table 2. Impact of the policy on the likelihood of leaving the firm

	(1) Leaving firm at t+1	(2) Leaving firm at t+1	(3) Leaving firm at t+1	(4) Leaving firm at t+1
Treated firm × post	-0.00534*	-0.00622**	-0.00408	-0.00411
	(0.00276)	(0.00304)	(0.00309)	(0.00312)
Treated firm × post × female	0.0109***	0.0134***	0.0133**	0.0138***
	(0.00389)	(0.00515)	(0.00521)	(0.00523)
Treated firm × female	-0.00395	-0.00441	0.0807	0.0766
	(0.0753)	(0.0752)	(0.0631)	(0.0634)
Firm FE	Yes	Yes	Yes	Yes
Year FE	Yes	Absorbed	Absorbed	Absorbed
Individual FE	Yes	Yes	Yes	Yes
Gender-Year FE	No	Yes	Absorbed	Absorbed
Gender-Year-Region FE	No	No	Yes	Yes
Industry-Year FE	No	No	No	Yes
Observations	210,450	210,450	210,450	210,450
Adj. R-squared	0.230	0.230	0.232	0.233
Diff. in women coeff	0.00560	0.00717	0.00924	0.00973
p-value women coeff	0.119	0.0844	0.0275	0.0216
Men's pre-policy mean	0.14	0.14	0.14	0.14
Women's pre-policy mean	0.16	0.16	0.16	0.16

Notes: This table reports different specifications of the model presented in equation (1). The dependent variable is a dummy equal to one if the individual leaves the firm at $t + 1$. The estimation sample includes individuals working in firms between 200 and 300 employees. A treated firm is defined as having at least 250 employees in 2015. Post is a dummy equal to one from 2017 onward. The p-value reported at the bottom of the table refers to the t-test of the two reported coefficients in the Table, corresponding to the impact of the policy on female employees. The pre-policy mean represents the mean of the dependent variable for the treated group between 2013 and 2017 separately for women and men. Standard errors double-clustered at the individual and firm levels are reported in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

4.1 Impact of the policy on separation rates

This section presents the results from estimating the effect of the policy on the likelihood of leaving the firm at $t + 1$. Specifically, Table 2 shows the results from alternative specifications of equation (1). At the bottom of the table, we report the p-value of the t-test on the sum of the two reported coefficients β_1 and β_2, which captures the effect of the policy on female employees. The baseline specification in column (1) includes individual, firm and year fixed effects. Additional sets of fixed effects are introduced sequentially. Column (2) adds gender-by-year fixed effects, which capture time-varying gender-specific factors. We include gender-region-year fixed effects in column (3). Finally, column (4) further controls for time-varying industry-specific shocks.

Across all specifications in Table 2, the differential effect of the policy on women (the coefficient on "Treated firm × post × female") is positive and statistically significant. The coefficient from the full specification in column (4) suggests that the policy leads to a nearly 1.4% increase in women's probability to leave relative to men's. Relative the pre-policy separation rate for men of 0.14, the coefficient of 0.014 estimated in column (4) corresponds to a 10% increase in women's relative probability of leaving. In contrast, the coefficient on "Treated firm x post", which reports the effect of the policy on men in treated vs. control firms, does not reach statistical significance in our preferred full specification shown in column (4). On the contrary, the p-value reported at the bottom of Table 2 indicates that effect of the policy is driven by a higher probability of women of leaving treated firms after the policy implementation. Although the sample and estimation strategy differ somewhat from ours, this result is in line with Blundell et al. (2025), who find that the policy increases women's probability of leaving by 16% relative to men.[15] Finally, the coefficient on the interaction "Treated firm × female" is not statistically significant, suggesting that there is no differential pre-policy trend in female separation between treated and control firms.

To verify the parallel trends assumption, we run the event-study analysis illustrated in equation (2). Figure 1 shows the results from the event study on the gender gap in leaving rates, with the 2016 as the benchmark year. To minimize differences across firms, the analysis is carried out on firms with +-20 employees around the 250-threshold. Figure 1 shows that the evolution of the outcome variable is comparable across treatment and control groups for both female and male employees. After 2016, there is an immediate increase in women's leaving rates that persists until 2020.

We next report a series of robustness checks on the effect of the policy on women's separation rates. First, we show that our results vary little when we change the bandwidth around the 250-employee cutoff. Figure A2 reports the estimates of β_2 from equation (1) for different bandwidths

[15] See column (2), Table 4, in Blundell et al. (2025). It is worth to highlight, however, that their result does not pass all the robustness checks they perform. Moreover, in their specification, they do not include worker fixed effects. Therefore, their effects also capture individual sorting.

Figure 1. Event study for the probability of leaving the firm

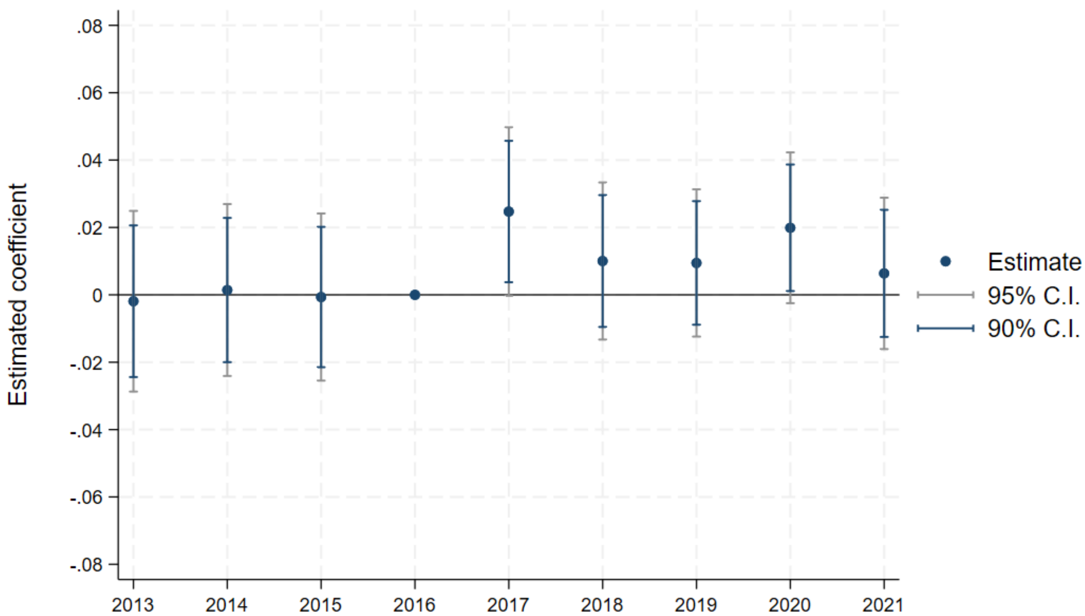

Notes: The graph reports the estimates of the coefficients β_k^1 of equation 2. These coefficients capture the leads and lags of the differential effect of the policy on women's probability of leaving relative to men's. The sample is restricted to firms within a bandwidth of +-20 around the cutoff. The graph also reports confidence intervals at the 90% and 95% levels.

ranging from 100 to 20 around the 250-employee cutoffs. Precisely, the estimated coefficient of "Treated x post x female" is only marginally insignificant when we consider bandwidths of +-100 (p-value 0.141) and +-70 (p-value 0.109), in which, however, treated and control firms might become less comparable.

Second, we run a set of placebo regressions using cutoffs other than the 250-employee threshold set by the policy. A positive and statistically significant effect in these regressions would raise concerns about the presence of contemporaneous shocks affecting our dependent variable, beyond the pay transparency law. Specifically, we pretend that the cutoff was at 300, 200 or 150 employees, and include all firms within a bandwidth +-50 employees around each cutoff, as in our main specification. However, as shown in figure A5, the estimates are not statistically significant at any of these different cutoffs.

Third, we check the robustness of our results to additional specifications. Figure A6 shows that the estimated effect of the policy on the female probability to leave is robust to restricting the sample to individuals who are older than 25 years, considering only workers with more than 8 years of experience (corresponding to the 25^{th} percentile of the distribution), excluding sectors for which our sample is less representative than the BPE (sectors G, I and Q), or the over-represented Information and Technology sector J. Finally, our results are robust to the inclusion of time-varying gender-industry specific shocks. However, our results are not robust to considering the year 2018 as the year of the reform, albeit the sign of the coefficient is still positive.

4.2 Impact of the policy on hiring rates

We now turn to the effects of the policy on hiring and the gender composition of new hires. Precisely, as discussed in section 3.2, we want to assess whether the policy has increased hiring of new employees or altered gender composition of new hires. Our dependent variable equals one if the individual is hired in the current year and zero otherwise. We show the estimates of the effect of the policy on new hires in Table 3. We replicate the same specifications of Table 2. Column (1) uses individual, firm and year fixed effects. Column (2) adds gender-year fixed effects. We also introduce a full set of gender-region-year fixed effects in column (3). Finally, column (4) includes industry-year fixed effects. The coefficient of "Treated firm × post × female" is negative in almost all specifications but statistically significant only for column (1). On the contrary, the coefficient "Treated firm × post", which yields the effect of the policy on male hiring, is positive and statistically significant across all columns. Column (4) indicates that men are more likely to be hired by about 0.9 percentage points in treated firms relative to control firms after the policy. As indicated by the p-value of the t-test on the sum of the reported coefficients, the effect on female only is not statistically significant at standard confidence levels. Therefore, we conclude that the pay transparency policy has changed the gender composition of new hires towards a greater share of men employees, but it did not have any effect on female employees. These findings seem to be consistent with a standard firm microeconomic reaction. As female labour becomes more costly relative to male labour, firms respond by reducing the demand for female labour, the costlier factor input, towards a greater share of male labour.

5 Mechanisms

5.1 Heterogeneity effects on leavers' attributes

This section shows selected heterogeneous effects of the policy. Precisely, we want to understand whether the policy had differential effects by gender on selected categories of workers. Our conceptual background predicts that the policy affects separation rates of women through greater job dissatisfaction following pay disparities revelations. Prior works show that job dissatisfaction would affect more workers in the lower wage distribution or with weaked bargaining power (Card et al., 2012). To account for such differences, we assess the differential effect of the policy on different categories of workers depending on their job ranking (i.e., managerial vs non-managerial jobs), position in the wage distribution (i.e., being below the median or above the median of the wage distribution), level of qualification (high-skilled vs low-skilled) and work experience (i.e., having a below-median or above- median level of experience).

To examine these differences, we estimate the following quadruple-differences-in-differences model:

Table 3. Impact of the policy on the likelihood of entering the firm

	(1) New hire	(2) New hire	(3) New hire	(4) New hire
Treated firm × post	0.0171***	0.0101**	0.00808*	0.00891**
	(0.00392)	(0.00426)	(0.00432)	(0.00436)
Treated firm × post × female	-0.0247***	-0.00541	-0.00603	-0.00840
	(0.00553)	(0.00744)	(0.00754)	(0.00756)
Treated firm × female	0.0547	0.0484	0.0405	0.0408
	(0.0671)	(0.0672)	(0.0660)	(0.0660)
Firm FE	Yes	Yes	Yes	Yes
Year FE	Yes	Absorbed	Absorbed	Absorbed
Individual FE	Yes	Yes	Yes	Yes
Gender-Year FE	No	Yes	Absorbed	Absorbed
Gender-Year-Region FE	No	No	Yes	Yes
Industry-Year FE	No	No	No	Yes
Observations	210,450	210,450	210,450	210,450
Adj. R-squared	0.150	0.151	0.152	0.154
diff women coeff	-0.00762	0.00466	0.00206	0.000508
p-value women coeff	0.136	0.445	0.739	0.935
Men's pre-policy mean	0.16	0.16	0.16	0.16
Women's pre-policy mean	0.19	0.19	0.19	0.19

Notes: The dependent variable is a dummy equal to one if the individual is a new hire. The estimation sample includes individuals working in firms between 200 and 300 employees. A treated firm is defined as having at least 250 employees in 2015. Post is a dummy equal to one from 2017 onward. The p-value reported at the bottom of the table refers to the t-test of the two reported coefficients in the Table, corresponding to the differential impact of the policy on female employees. The pre-policy mean represents the mean of the dependent variable for the treated group between 2013 and 2017. Standard errors double-clustered at the individual and firm levels are reported in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

$$
\begin{aligned}
Y_{i,f,t} = {} & \beta_1 \left(Treated_f \times Post_t \right) \\
& + \beta_2 \left(Treated_f \times Post_t \times Female_i \right) \\
& + \beta_3 \left(Treated_f \times Post_t \times Group_i^k \right) \\
& + \beta_4 \left(Treated_f \times Post_t \times Female_i \times Group_i^k \right) \\
& + \Gamma \mathcal{L}_{i,f,t} + \alpha_i + \mu_f + \delta_{j,t} + \chi_{g,r,t} + \varepsilon_{i,f,t}
\end{aligned} \tag{3}
$$

where Y_{ift} denotes the outcome of individual i, employed by firm f, at time t. Relative to (1), we add $(Group_i^k)$, which is an indicator variable for membership in subgroup k, where $k \in \{$Manager, High-skilled, Above-median wage, Above-median experience$\}$. The term $\mathcal{L}_{ift}$ collects all lower-order terms implied by the full interaction of $Treated_f$, $Post_t$, $Female_i$, and $Group_i^k$. As per the baseline model in (1), the specification further includes individual fixed effects α_i, firm fixed effects μ_f, industry-by-year fixed effects δ_{jt}, and gender-by-region-by-year fixed effects χ_{grt}.

Equation (3) is a fully saturated difference-in-differences specification. The coefficient β_1 captures the effect of the policy for the reference category, namely men in the baseline subgroup $(Female_i = 0,\ Group_i^k = 0)$. The coefficient β_2 measures how the policy effect for women differs from that for men within the baseline subgroup. The coefficient β_3 measures how the policy effect for individuals in subgroup $k = 1$ differs from that for individuals in the reference subgroup $k = 0$ among men. β_4 captures whether the gender-specific effect of the policy varies across subgroup k. Equivalently, β_4 measures whether the female–male policy gap differs between individuals in subgroup $k = 1$ and those in subgroup $k = 0$.

Under this parameterization, the policy effect is given by β_1 for men in the reference subgroup, by $\beta_1 + \beta_2$ for women in the reference subgroup, by $\beta_1 + \beta_3$ for men in subgroup $k = 1$, and by $\beta_1 + \beta_2 + \beta_3 + \beta_4$ for women in subgroup $k = 1$. Accordingly, β_2 identifies the female–male difference in the policy effect within subgroup $k = 0$, β_3 identifies the subgroup-k differential in the policy effect among men, and β_4 identifies whether that subgroup differential is itself different for women.

All remaining lower-order interactions are included to absorb baseline differences across treated and untreated units, gender, and subgroup membership. These terms do not affect the interpretation of the policy effect, which is identified by the interactions involving $Treated_f \times Post_t$, therefore they are omitted from the following discussion.

Table 4 shows the results of the interactions from estimating equation (3). The dependent variable is an indicator equal to one if the individual leaves the firm at $t + 1$.

In Panel A, where the heterogeneity dimension is managerial status, the policy significantly increases turnover among non-manager women. The estimated effect is 0.0122 and statistically significant at the 1 percent level. Moreover, among non-managers, the effect of the policy is significantly larger for women than for men: the female-male differential is 0.0156 and precisely estimated. By contrast, among managers, the female-male differential is small and statistically indistinguishable from zero. Consistent with this pattern, the estimated policy effect for manager women is close to zero and not statistically significant. The difference in the policy effect between manager and non-manager women is also negative, though imprecisely estimated. Taken together, these results suggest that the policy's impact on women's separations is concentrated among non-managerial workers, with no evidence of a comparable effect among managers.

Panel B looks at the heterogeneity by skill level. The policy has a positive and statistically significant effect on turnover among low-skilled women, with an estimated coefficient of 0.0205. The female-male differential within the low-skilled group is also positive and marginally significant, indicating that the policy raises exits more for low-skilled women than for comparable men. Among high-skilled workers, the female-male differential remains positive and statistically significant at conventional levels ($p < 0.05$), although the estimated policy effect for high-skilled women is smaller and not statistically distinguishable from zero. The difference between high-skilled and low-skilled women is negative and imprecisely estimated. Overall, the evidence points to stronger effects among low-skilled women.

In Panel C, which splits workers by their position in the wage distribution, the policy again has a positive and statistically significant effect on women below the median wage distribution. The corresponding estimate is 0.0148, and the female-male differential within the below-median group is also positive and significant ($p < 0.05$). For women above the median, by contrast, the estimated effect is close to zero and statistically insignificant, and the female-male differential is not significant. The difference between above-median and below-median women is negative, but not precisely estimated. This suggests that the policy disproportionately affects women located in the lower part of the wage distribution.

Finally, Panel D shows heterogeneity by work experience. The policy significantly increases turnover among below-median-experience women, with an estimated effect of 0.0181. The female-male differential is also positive and statistically significant in this group ($p < 0.05$), implying that the policy has a stronger effect on less experienced women than on comparable men. In contrast, the estimated policy effect for above-median-experience women is essentially zero. The difference between more and less experienced women is negative and statistically significant ($p < 0.01$), indicating that the policy effect is concentrated among workers with lower experience. Among the four heterogeneity dimensions considered in the table, this is the clearest evidence that the effect declines systematically across groups.

Taken together, the results indicate that the policy increased women's probability of leaving the firm primarily among workers in more junior or less advantaged positions— non-managers, low-skilled workers, below-median wage workers, and less experienced workers. By contrast, there is little evidence that the policy had a meaningful effect on separations among women in managerial, higher-wage, or more experienced positions. This pattern is consistent with the job satisfaction channel that the policy's impact was strongest among women with weaker bargaining positions or lower attachment to firm-specific career ladders.

5.2 Impact of the policy on occupational upgrading

So far, we have shown that, following the policy implementation, women in treated firms are more likely to leave than those in control firms.

We now examine the policy's effects on occupational shifts. The results are reported in Table 5. We first examine the effects of the policy on within-firm occupational upgrading. In this case, our estimation strategy also controls for firm $\times$ individual fixed effects, hence studying the effects of the policy on individuals who remain in the same firm. We define occupational upgrading as the move between t-1 and t from a lower-paid to a better paid occupation. Column (1) shows the results

Table 4. Effects by gender and occupational group

Effect	Coefficient	Std. error	T-stat	p-value
Panel A. Managers vs non-manager				
Non-manager females	0.0122***	0.0045	2.71	0.007
Females vs males non-manager	0.0156***	0.0055	2.83	0.005
Females vs males manager	0.0074	0.0163	0.45	0.649
Manager females	-0.0038	0.0132	-0.29	0.771
Manager vs. Non-manager females	-0.0160	0.0139	-1.15	0.248
Panel B. High-skilled vs low-skilled				
Low-skilled females	0.0205***	0.0086	2.38	0.017
Females vs males low-skilled	0.0191*	0.0109	1.75	0.079
Females vs males high-skilled	0.0166**	0.0085	1.97	0.049
High-skilled females	0.0108	0.0069	1.57	0.117
Low vs. high-skilled females	-0.0097	0.0109	-0.89	0.373
Panel C. Below vs above-median wage				
Below-median females	0.0148***	0.0052	2.84	0.005
Females vs males below-median	0.0175**	0.0069	2.55	0.011
Females vs males above-median	0.0076	0.0082	0.93	0.351
Above-median females	0.019	0.0070	0.27	0.787
Above-median females vs. below-median	-0.0129	0.0086	-1.50	0.134
Panel D. Sector's unionization rate				
Below-median females	0.0181***	0.0059	3.04	0.002
Females vs males below-median	0.0244***	0.0077	3.17	0.001
Females vs males above-median	0.0021	0.0081	0.26	0.795
Above-median females	-0.0000	0.0068	-0.00	0.998
Above-median females vs. below-median	-0.0181***	0.0089	-2.0189	0.04

Notes: This table reports results of the interaction model presented in equation (2) for each occupational category. The dependent variable is a dummy equal to one if the individual leaves the firm at $t + 1$. The estimation sample includes individuals working in firms between 200 and 300 employees. A treated firm is defined as having at least 250 employees in 2015. Post is a dummy equal to one from 2017 onward. The p-value reported at the bottom of the table refers to the t-test of the two reported coefficients in the Table, corresponding to the impact of the policy on female employees. Standard errors double-clustered at the individual and firm levels are reported in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

of this specification. We cannot find any effect of the policy on the likelihood of being promoted towards a better-paid occupation for either females or males. In column (2), we look at the

Columns (1) and (3) use the full specification with individual, firm, year, and gender–region–year fixed effects, capturing the policy's impact on the probability of promotion both within and across firms. By contrast, Columns (2) and (4) include interacted individual–firm fixed effects, isolating the policy's impact on within-firm promotions. As shown in Table 5, we find no statistically significant effects for either men or women. None of the coefficients reaches statistical significance at standard confidence levels. We thus conclude that the policy has no distinguishable effects on promotions—either within firms or across firms—for men or women. These findings are consistent with Blundell et al. (2025) for the UK policy.

Table 5. Alternative outcome variables

	(1) Within-firm Upgrading	(2) Within-firm Manager	(3) Leaving Upgrading	(4) Leaving Downgrading
Treated firm × post	0.00219	0.000266	0.000152	-0.00426**
	(0.00259)	(0.00115)	(0.00228)	(0.00213)
Treated firm × post × female	-0.00291	0.00212	0.00584	0.00800**
	(0.00451)	(0.00204)	(0.00380)	(0.00357)
Treated firm × female			0.0392	0.0374
			(0.0493)	(0.0722)
Firm FE	Absorbed	Absorbed	Yes	Yes
Individual FE	Absorbed	Absorbed	Yes	Yes
Individual-Firm FE	Yes	Yes	No	No
Gender-Year-Region FE	Yes	Yes	Yes	Yes
Industry-Year FE	Yes	Yes	Yes	Yes
Observations	201,201	201,201	201,450	201,450
Adj. R-squared	0.0929	0.0626	0.232	0.235
diff women coeff	-0.000717	0.00239	0.00599	0.00373
p-value women coeff	0.847	0.157	0.0518	0.197

Notes: The dependent variables are: 1) a dummy equal to one if the individual is promoted from a lower-paid to a better-paid occupation; 2) a dummy equal to one if the individual is promoted from a lower-paid to a better-paid occupation within the firm; 3) a dummy equal to one if the individual is promoted from a non-managerial to a managerial occupation; 4) a dummy equal to one if the individual is promoted from a non-managerial to a managerial occupation within the firm. The estimation sample includes individuals working in firms between 220 and 280 employees. A treated firm is defined as having at least 250 employees in 2015. Post is a dummy equal to one from 2017 onward. The p-value reported at the bottom of the table refers to the t-test of the two reported coefficients in the Table, corresponding to the differential impact of the policy on female employees. Standard errors double-clustered at the individual and firm levels are reported in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

6 Conclusion

This paper examines the second-order labour market consequences of the UK's 2017 gender pay transparency policy, which mandated firms with more than 250 employees to publish gender pay gap indicators. While prior work has largely focused on whether such policies reduce the measured gender pay gap, we focus on the behavioural responses of female workers and firms to these changes. Using a novel employer-employee dataset that combines individual career profiles from Diffbot with firm-level financial data from Orbis, and exploiting variation in policy exposure around the 250-employee threshold in a difference-in-differences design, we provide causal evidence on how pay gap transparency affects female workers' employment.

Our findings reveal four key results. First, the policy led to a significant increase in women's probability of leaving the firm relative to men's — approximately 1.4 percentage points, or about 10% of the pre-policy separation rate. This effect is robust to alternative bandwidth selections and specifications, and is supported by an event-study analysis confirming the parallel trends assumption. Second, the policy weakly increased the hiring of men in treated firms while having no effect on women's hiring probability. Third, women who leave are disproportionately those in the lower part of the wage distribution, in non-managerial positions, with lower skills and in less-unionized sectors — precisely the workers for whom the informational shock is most salient and bargaining power weakest. This pattern is consistent with both worker-driven exit due to job discontent following gender pay gap reporting and firm-side compositional adjustment aimed at reducing pay gaps and alleviate the consequences of greater labour frictions due to employees' and public pressure. Fourth, we find no evidence that the policy promoted women into better-paid or managerial occupations, either within or across firms. Instead, women who leave tend to move into occupations with lower average earnings than their previous employment spell.

Taken together, these results paint a nuanced picture of the effects of the UL pay transparency policy. While the policy may contribute to narrowing the measured gender pay gap, our evidence suggests that this reduction does not necessarily reflect an improvement in women's labour market position. Rather, the decline in the reported gap may depend on the exit of the lowest-paid women from treated firms. This represents an unintended consequence that should stimulate a debate about the welfare implications of transparency-based policies.

These findings carry important policy implications. Pay transparency is a necessary condition for addressing gender pay disparities, but it is not sufficient on its own. Our results suggest that the reduction in gender pay gap may be masked by the (un)intended separation from a part of the female workforce. Therefore, policymakers should consider complementing disclosure requirements with mechanisms that increase bargaining leverage of underpaid employees.

This paper opens several avenues for future research. In a forthcoming version, we will extend the analysis by examining the specific firms where female leavers relocate. Our employer-employee dataset, which provides complete career histories rather than only job flows within the estimation sample, uniquely enables us to track the destination employers of women who exit treated firms. This extension will allow us to assess whether women move to firms with smaller gender pay gaps, stronger female representation in senior positions, or more equitable pay cultures — or, conversely, whether they sort into firms offering lower pay or fewer advancement opportunities. Understanding whether the exits induced by the policy represent moves toward better-matched employers or reflect downward mobility is crucial for a complete welfare assessment of pay transparency mandates.

References

Arora, A., and Dell, M. (2023). Linktransformer: A unified package for record linkage with transformer language models. *arXiv preprint arXiv:2309.00789*.

Baker, M., Halberstam, Y., Kroft, K., Mas, A., and Messacar, D. (2023). Pay transparency and the gender gap. *American Economic Journal: Applied Economics*, *15*(2), 157–183.

Bamieh, O., and Ziegler, L. (2025). Can wage transparency alleviate gender sorting in the labour market? *Economic Policy*, *40*(122), 401–426.

Bena, J., Ortiz-Molina, H., and Simintzi, E. (2022). Shielding firm value: Employment protection and process innovation. *Journal of Financial Economics*, *146*(2), 637–664.

Bennedsen, M., Larsen, B., and Wei, J. (2023). Gender wage transparency and the gender pay gap: A survey. *Journal of Economic Surveys*, *37*(5), 1743-1777.

Bennedsen, M., Simintzi, E., Tsoutsoura, M., and Wolfenzon, D. (2022). Do firms respond to gender pay gap transparency? *The Journal of Finance*, *77*(4), 2051-2091.

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

Blau, F. D., and Kahn, L. M. (2023). The gender pay gap. *Journal of Economic Perspectives*, *37*(4), 3–28.

Blundell, J., Duchini, E., Simion, Ş., and Turrell, A. (2025). Pay transparency and gender equality. *American Economic Journal: Economic Policy*, *17*(2), 418–445.

Brütt, K., and Yuan, H. (2022). Pitfalls of pay transparency: Evidence from the lab and the field. *Tinbergen Institute Discussion Paper*.

Card, D., Cardoso, A. R., and Kline, P. (2016). Bargaining, sorting, and the gender wage gap: Quantifying the impact of firms on the relative pay of women. *The Quarterly Journal of Economics*, *131*(2), 633–686.

Card, D., Mas, A., Moretti, E., and Saez, E. (2012). Inequality at work: The effect of peer salaries on job satisfaction. *American Economic Review*, *102*(6), 2981–3003.

Castilla, E. J. (2015). Accounting for the gap: A firm study manipulating organizational accountability and transparency in pay decisions. *Organization Science*, *26*(2), 311-333.

Cullen, Z. (2024). Is pay transparency good? *Journal of Economic Perspectives*, *38*(1), 153–180.

Cullen, Z., and Perez-Truglia, R. (2023). The salary taboo privacy norms and the diffusion of information. *Journal of Public Economics*, *222*, 104890.

Dambra, M., Khavis, J., Lin, Z., and Suk, I. (2025). Labor market consequences of pay transparency: Evidence from the initial pay ratio disclosure. *Available at SSRN 4826506*.

De Loecker, J., Obermeier, T., and Van Reenen, J. (2024). Firms and inequality. *Oxford Open Economics*, *3*(Supplement_1), i962-i982.

Dittrich, M., Knabe, A., and Leipold, K. (2014). Gender differences in experimental wage negotiations. *Economic Inquiry*, *52*(2), 862–873.

Gamage, D. K., Kavetsos, G., Mallick, S., and Sevilla, A. (2024). Pay transparency intervention and the gender pay gap: Evidence from research-intensive universities in the UK. *British Journal of Industrial Relations*, *62*(2), 293–318.

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

Government Equalities Office. (2016). *Closing the gender pay gap: Government consultation.* United Kingdom Government Equalities Office.

Gray, S., Kemeny, T., Nathan, M., Ozgen, C., Pialli, G., Rosso, A. C., … Valero, A. (2025). Graph-ee: Building employer - employee panels from a knowledge graph and company microdata. *Working Paper*.

Gulyas, A., Seitz, S., and Sinha, S. (2023). Does pay transparency affect the gender wage gap? Evidence from Austria. *American Economic Journal: Economic Policy*, *15*(2), 236–255.

Kalemli-Özcan, Ş., Sørensen, B. E., Villegas-Sanchez, C., Volosovych, V., and Yeşiltaş, S. (2024). How to construct nationally representative firm-level data from the orbis global database: New facts on smes and aggregate implications for industry concentration. *American Economic Journal: Macroeconomics*, *16*(2), 353–374.

Leibbrandt, A., and List, J. A. (2015). Do women avoid salary negotiations? Evidence from a large-scale natural field experiment. *Management Science*, *61*(9), 2016–2024.

Liang, C., Lourie, B., Qu, H., and Wang, F. Y. (2025). Pay ratio disclosure and employee turnover. *Available at SSRN 4702852*.

Manning, A., and Saidi, F. (2010). Understanding the gender pay gap: What's competition got to do with it? *Industrial and Labor Relations Review*, *63*(4), 681–698.

OECD. (2023). *Reporting gender pay gaps in OECD countries: Guidance for pay transparency implementation, monitoring and reform.* Gender Equality at Work, OECD Publishing, Paris. doi: https://doi.org/10.1787/ea13aa68-en

Serfling, M. (2016). Firing costs and capital structure decisions. *The Journal of Finance*, *71*(5), 2239–2286.

Appendices

A Data description

This section describes the data used in this paper, along with sources and variables.

A.1 Diffbot

The empirical analysis of this work leverages a novel employer-employee database for UK companies. This section provides a more detailed description of these data. We invite the reader to refer to Gray et al. (2025) for full details.

Individual-level data are sourced from Diffbot, the world's leading commercial knowledge graph database of the public web. At the start of 2025, the graph included 3.6m active UK companies with CRN and 10.4m workers. This compares to totals of 5.5m active firms and 33.9m workers aged 16+ from the 2025 Business Population Estimates.

Diffbot constructs a knowledge graph by scraping the web continually, using a combination of tools ranging from image recognition, natural language processing and supervised learning. Then, it organizes such data into a graph of entities, including individuals, organizations, places, which are linked to each other. This database has several desirable properties that we exploit to build an employer-employee database. First, the graph includes data on UK companies, including the Company Registration Number (CRN). This information allows us to link companies in Diffbot to other company-level data like Orbis. Second, Diffbot includes extensive information on individuals, including gender, age, detailed education and employment histories, skills and job titles, among others. All these information allows us to assemble an employer-employee database for a set of UK companies.

We construct the employer-employee database in three steps. First, we assembled a dataset with fixed individual characteristics, including the year of birth, languages and gender. Second, we built a panel dataset with education histories, including information on highest qualification, field and school/university. We use the country of the earliest available educational qualification achieved to proxy for the migrant status of an individual. For example, if someone's lowest level of qualification reported on the web is high-school, and this qualification has been obtained in the UK, we assume that the individual is British. Additionally, we also use information on education histories to construct an individual's highest level of education attained based on a set of keywords. Further details on these procedures and validation checks are reported in Gray et al. (2025).

Third, we compile full employment histories, including the start and end dates of each employment spell, employers, job titles, and brief job descriptions. Finally, we merge these three components (individual fixed characteristics, education and employment histories) into an employer-employee panel dataset, converting education variables into time-invariant characteristics while preserving the time dimension of employment histories. Diffbot assigns to each entity a unique identifier, allowing us to study job flows across organizations.

Diffbot also provides brief job-title description but does not link them to standard occupational classifications. Therefore, we use LinkTransformer (Arora and Dell, 2023), an open-source LLM-based text classifier, to map the universe of Diffbot job-title descriptions in our sample to the UK 2020 Standard Occupational Classification (SOC2020). We experiment with multiple levels of SOC granularity and obtain the best performance at the 4-digit level.

A.2 Company-level data

As explained in Gray et al. (2025), the employer-employee database is build on a tractable sample of firms. Relying on extensive checks, we find that web information on entities is available especially for large organizations. Additionally, as it is well-known, financial and employment information in databases such as Orbis is more complete and higher quality for medium and large firms (Kalemli-Özcan et al., 2024). Therefore, we first preselect a sample of 'medium' and 'large' firms active at some point between 2007 and 2023. Companies House distinguishes 'micro' and 'small' companies entities from medium and large companies based on turnover, assets and employee thresholds. We apply these post-2016 thresholds backwards to ensure a time-consistent sampling frame. Companies above this threshold need to provide complete, audited annual accounts. This means that company-level information is most complete and highest quality for this sample. Applying these definitions to Orbis Historical, using data from unconsolidated balance sheets only, we obtain a sample of 55,775 companies. We then match this sample of companies to Companies House (through OpenCorporates), which leaves us with a sample of 55,187 companies.

We follow the cleaning procedures for Orbis data documented in Kalemli-Özcan et al. (2024)

and De Loecker et al. (2024). We keep only firm-year observations for which financial variables are expressed in GBP pounds. We use the account closing date to determine the calendar year. If the closing date is after or on June 1st, we assign it to the current year. If it is before June 1st, we assign it to the previous year. At this stage, Orbis may contain multiple annual observations for some firms. We design a routine of sequential steps to remove firm-year observations duplicates, similar to De Loecker et al. (2024):

- We keep the annual report values when both the annual report and local registry filing are present, and the annual report values are non-missing.
- When annual report and local registry filing values are not the same (and both are non-missing), we prefer annual report values.
- When annual report values are missing, and local registry filings are non-missing, we keep local registry filings.
- After the selection above, we prefer consolidated accounts to unconsolidated accounts.
- We remove remaining duplicates by taking the observations with fewer missing values for the number of employees, EBITDA and costs of employees.

A.3 Tables

Table A1. Probability to be treated on observable characteristics

	(1)
College or higher degree	-0.00493
	(0.0348)
Migrant	0.0317
	(0.0401)
Age	0.00212*
	(0.00122)
Female	0.00396
	(0.0319)
Worker experience	0.00101
	(0.00145)
Leaving company	0.0251
	(0.0413)
New hire	-0.0925**
	(0.0409)
Promotion	0.0226
	(0.0348)
Managerial promotion	-0.117*
	(0.0696)
Observations	9,011
Adj. R-squared	0.00612
F-test p-value	0.114

Notes: This table reports a linear probability model of the probability to be in a treated firm against a set of observable characteristics. Variables are averaged over the pre-policy period 2013-2017. Standard errors double-clustered at the individual and firm levels are reported in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

A.4 Figures

Figure A1. Gender policy reports uploaded to the Government website

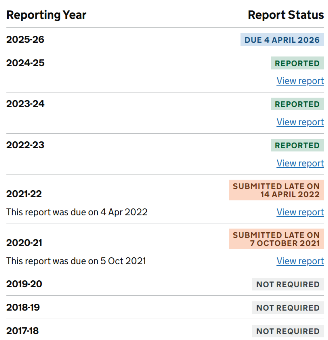

Reporting Year	Report Status
2025-26	DUE 4 APRIL 2026
2024-25	REPORTED View report
2023-24	REPORTED View report
2022-23	REPORTED View report
2021-22 This report was due on 4 Apr 2022	SUBMITTED LATE ON 14 APRIL 2022 View report
2020-21 This report was due on 5 Oct 2021	SUBMITTED LATE ON 7 OCTOBER 2021 View report
2019-20	NOT REQUIRED
2018-19	NOT REQUIRED
2017-18	NOT REQUIRED

(a) Firm uploading its gender report with delays

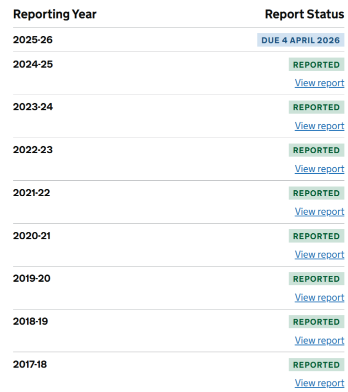

Reporting Year	Report Status
2025-26	DUE 4 APRIL 2026
2024-25	REPORTED View report
2023-24	REPORTED View report
2022-23	REPORTED View report
2021-22	REPORTED View report
2020-21	REPORTED View report
2019-20	REPORTED View report
2018-19	REPORTED View report
2017-18	REPORTED View report

(b) Firm uploading its gender report on time

Notes: The figure shows how gender policy reports are reported on the Government website.

Figure A2. Google searches for ’gender wage gap’

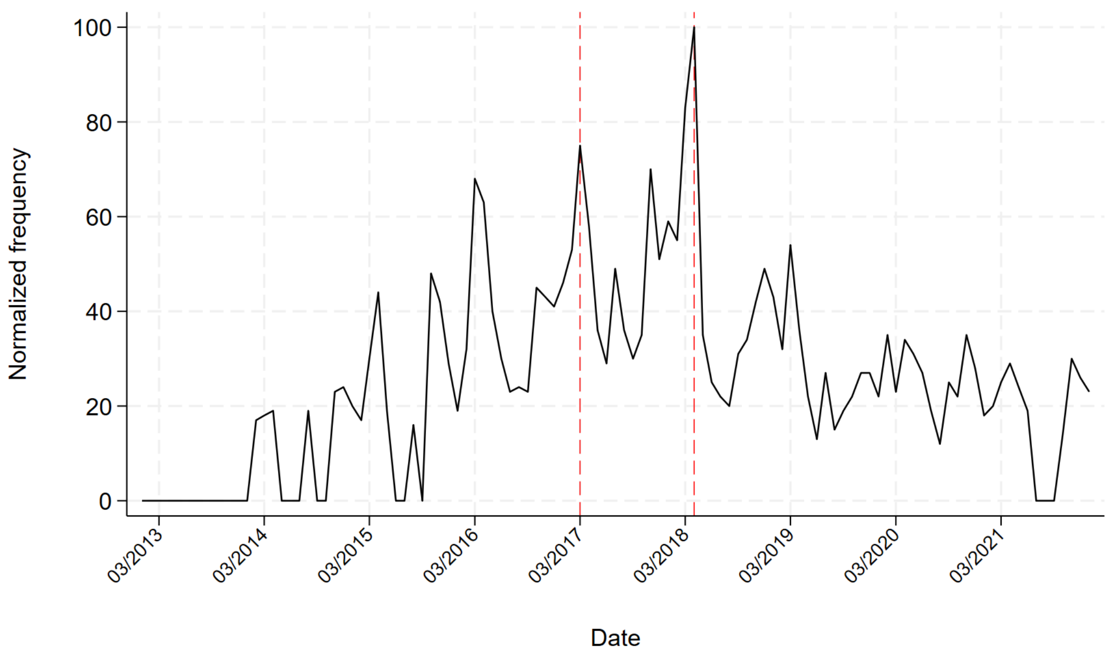

Notes: The figure shows the relative Google searches for ”gender wage gap” in the UK between March 2013 and March 2022. Numbers are expressed as relative to the peak in April 2018.

Figure A3. Evolution of gender inequality indicators for reporting firms

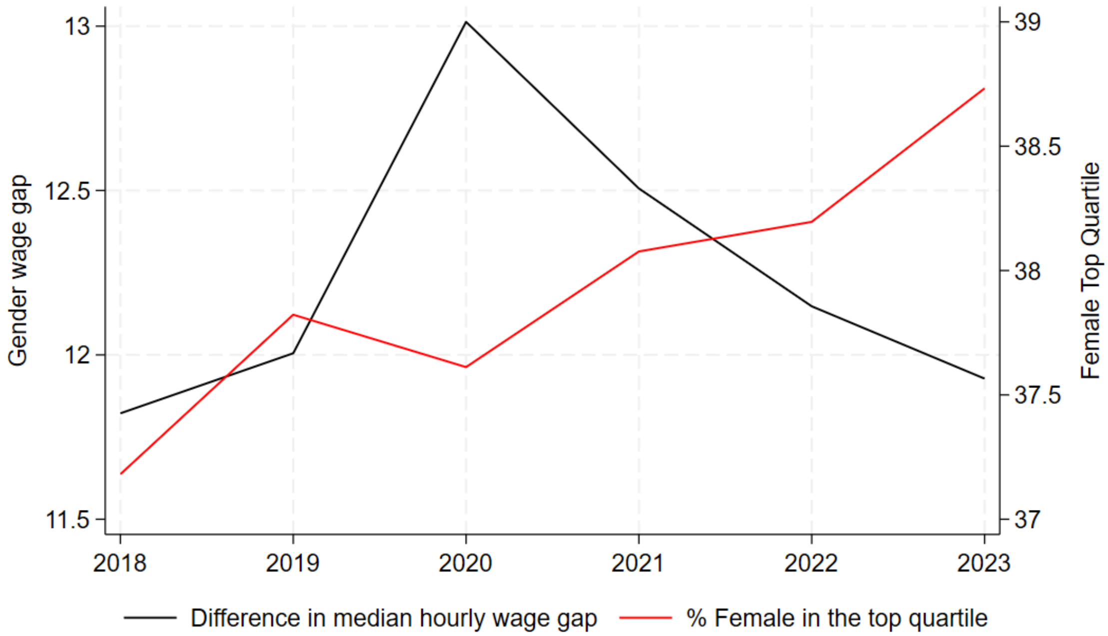

Notes: The figure reports the evolution of the difference in median hourly wage gap between males and females (black line) and the share of female in the firm's top quartile distribution (red line) for the period from 2018 to 2023.

Figure A4. Differential effect of the policy on females relative to males for alternative bandwidths

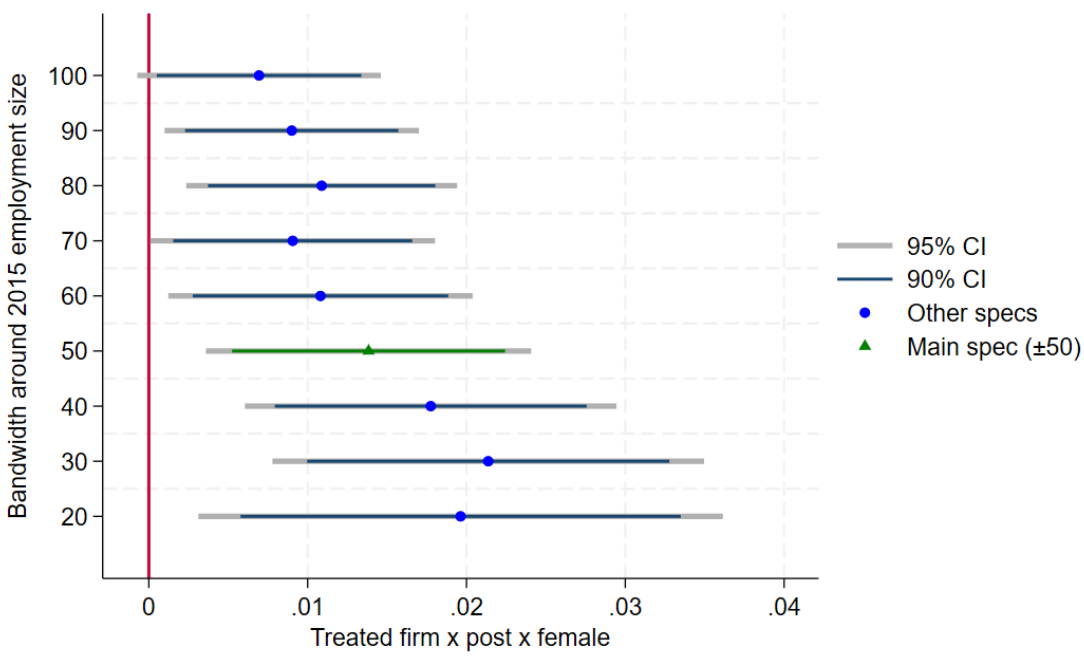

Notes: The figure shows the estimated coefficient of Treated firm $\times$ post $\times$ female for alternative bandwidths.

Figure A5. Robustness to alternative treatment years

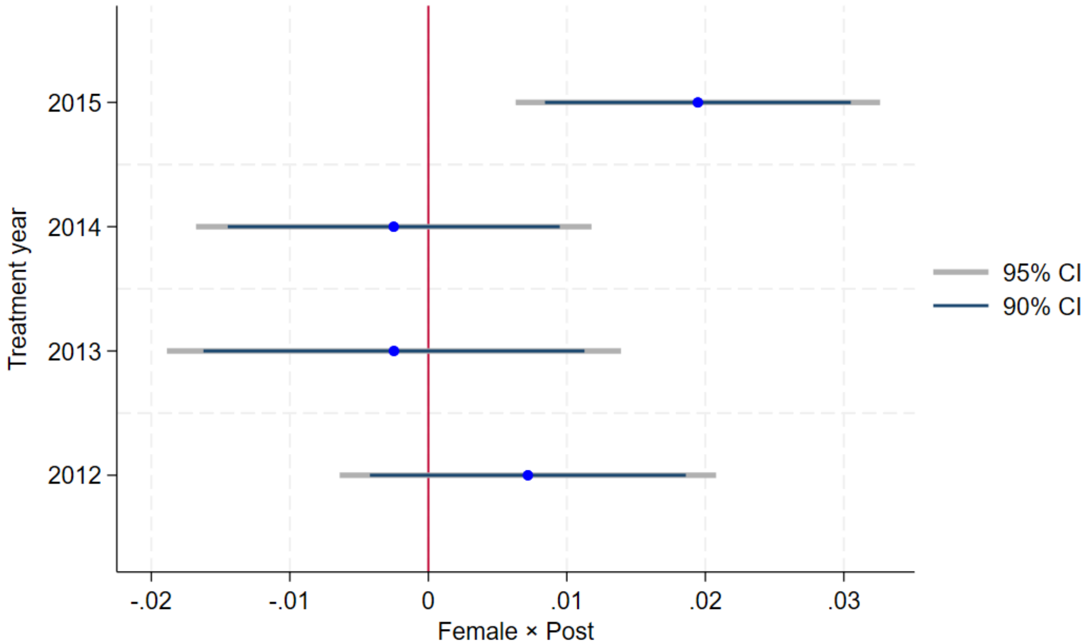

Notes: The figure shows the estimated coefficient of Treated firm × post × female for alternative treatment years.

Figure A6. Placebo regressions

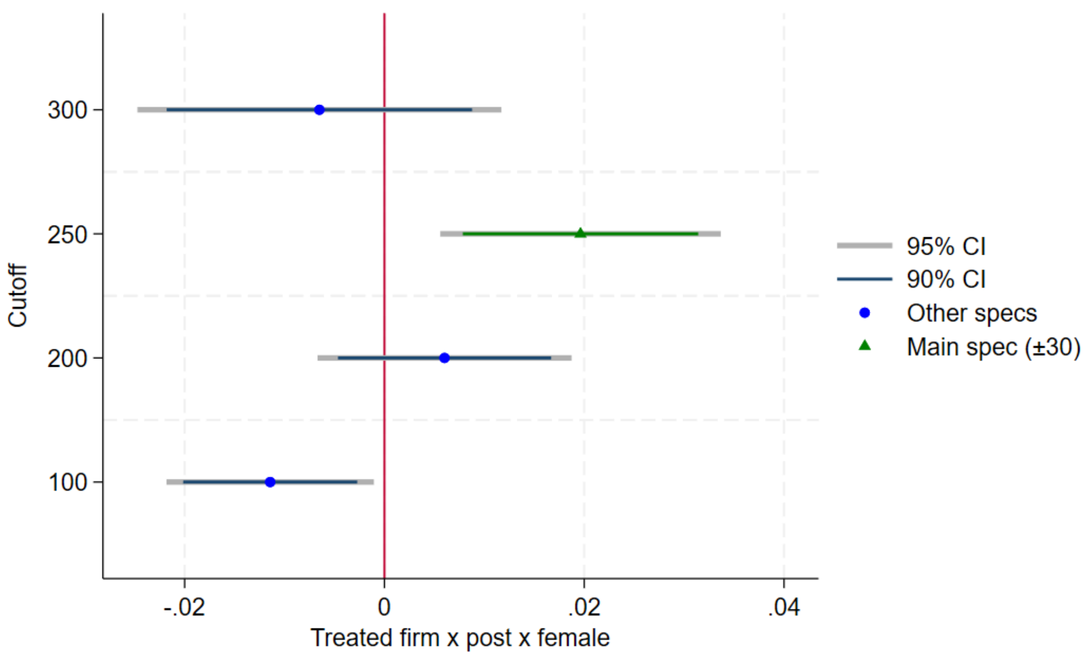

Notes: The figure shows the estimated coefficient of Treated firm × post × female for alternative placebo cutoffs.

Figure A7. Additional robustness checks

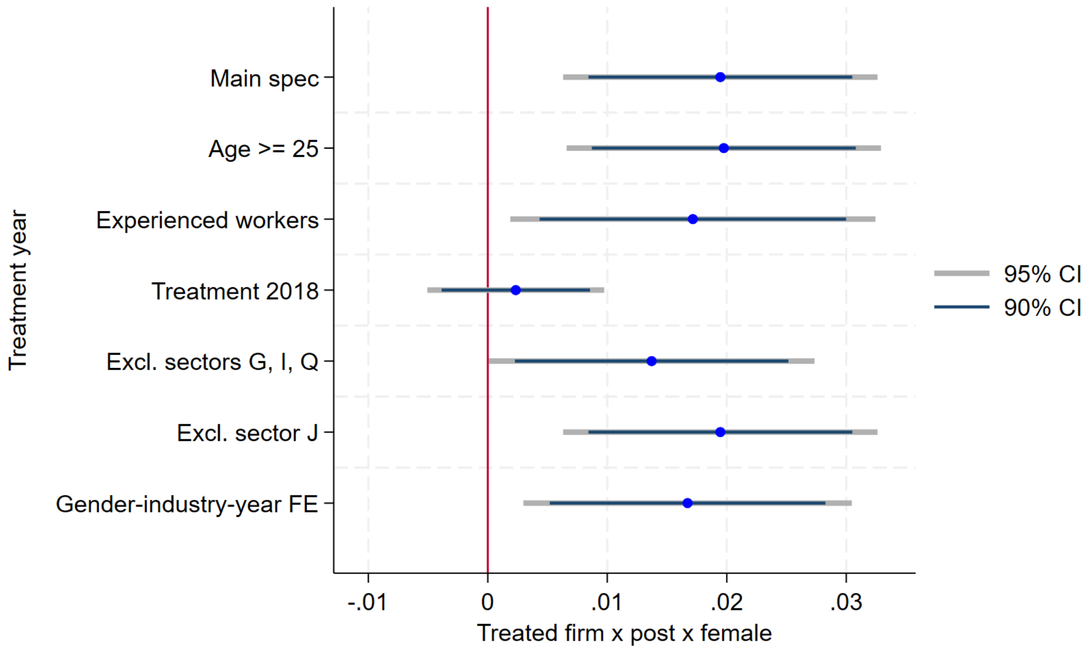

Notes: The figure shows the estimated coefficient of Treated firm × post × female for alternative treatment years.