CAGE

Pay transparency and gender equality

CAGE working paper no. 482

June 2020
(Revised March 2022)

Emma Duchini
Stefania Simion
Arthur Turrell

Previously titled:
Pay transparency and cracks in the glass ceiling

WARWICK
THE UNIVERSITY OF WARWICK

Economic and Social Research Council

Pay Transparency and Gender Equality*

Emma Duchini† Ştefania Simion‡ Arthur Turrell§

December 2021

Abstract

The 2018 UK transparency policy mandates that firms with at least 250 employees publicly disclose gender equality indicators. A difference-in-differences model exploiting variations in this policy across firm size and time shows that greater transparency closes 14 percent of the gender pay gap by reducing men's wage growth. Additionally, the policy increases transparency in the hiring process, as firms are 10 percent more likely to post wages in their vacancies. Reputation concerns seem to influence employers' reactions, as firms publishing worse equality indicators obtain lower impression scores in YouGov's Women's Rankings. Worse performing firms also improve gender equality the most.

JEL codes: J08, J16, J24.
Keywords: pay transparency; gender equality; firm reputation.

*First draft circulated in November 2019 under the title "Firm response to the public disclosure of gender gaps data". We are grateful to Ghazala Azmat, Manuel Bagues, Mirko Draca, Gabrielle Fack, James Fenske, Libertad Gonzalez, Victor Lavy, and Giuseppe Pratobevera for their valuable advice. We further thank Leonardo Bursztyn, Maia Güell, Muriel Niederle, Bobby Pakzad-Hurson, Joahnna Rickne, Paul Robinson, Daniel Schaefer, Sebastian Seitz, Hans H. Sievertsen, Carl Singleton and Simone Traini for their useful comments. A special thanks goes to Mihir Chandraker, George Taylor, Marios Tsoukis, and Giulia Vattuone for excellent research assistance. We also acknowledge participants in Warwick internal seminars, Brunel research seminar (02/2020), the Productivity Insight Network Workshop (02/2020), Bank of England research seminar (06/2020), the 2020 EALE-SOLE-AASLE conference, 19th IZA/SOLE Transatlantic Meeting for Labor Economists, 2020 World Congress, 2020 EEA conference, 2020 ESCoE Conference on Economic Measurement, MILLS Seminar (12/2020), seminars at Padua, Milano Cattolica, Bristol, Royal Holloway, Essex, 2021 EEA Conference, ONS Women's Network seminar (11/2021) for their constructive suggestions. We finally thank Miranda Kyte for facilitating our access to YouGov data. Duchini gratefully acknowledges financial support from the Productivity Insight Network (PIN), and the Centre for Competitive Advantage in the Global Economy (CAGE). The views expressed are those of the authors and may not reflect the views of the Office for National Statistics and the wider UK government. All mistakes are our own.

† University of Essex, Department of Economics, Wivenhoe Park, Colchester, CO4 3SQ, United Kingdom. *Email:* e.duchini@essex.ac.uk.

‡ University of Bristol, School of Economics, Bristol, BS8 1TU, United Kingdom. *Email:* stefania.simion@bristol.ac.uk.

§ Office for National Statistics, Government Buildings, Cardiff Road, Newport, NP10 8XG, United Kingdom. *Email:* arthur.turrell@ons.gov.uk.

1 Introduction

In recent years, many governments have adopted pay transparency policies with the aim of improving gender equality.[1] By increasing the salience of gender gaps in the labor market, transparency measures are meant to act as an information shock that may alter the relative bargaining power of male and female employees vis-à-vis the firm. Coupled with the potential negative effects of unequal pay on firms' reputation, transparency policies could improve women's relative pay and career outcomes. At the same time, pay transparency may have ambiguous effects on workers' productivity and retention. On the one hand, if firms react to the policy by improving gender equality, pay transparency may boost the productivity and the retention of those employees for whom fairness is important. On the other hand, increased transparency could hurt job satisfaction of those employees who perceive that they are being treated unfairly, with negative knock-on effects for productivity and retention. The magnitude of all of these effects is also likely to depend on how strong and salient the information shock is.

This paper studies the effects of pay transparency in a context where firms' gender equality performance is publicly disclosed. Each year since 2018, UK firms with at least 250 employees have been required to publish a series of gender equality indicators both on their own website and on a dedicated government website. These indicators include percentage gaps in mean and median hourly pay, and the percentage of women in each quartile of the wage distribution.

We begin our analysis by studying the impact of this policy on pay and career outcomes of men and women using the Annual Survey of Hours and Earnings (ASHE), the UK matched employer-employee data set, from 2013 to 2019. To identify causal effects, we adopt a difference-in-differences strategy that exploits the variation across firm size and time in the application of the transparency policy. To avoid capturing any potential impact of this policy on firm size, we define the treatment status based on firms' number of employees prior to the introduction of the mandate.

[1] Following the recommendations of the European Commission, Austria, Denmark, Italy, and Germany introduced transparency laws, (Aumayr-Pintar et al. 2018). Though pay transparency requirements are less common in the United States, many states have prohibited employers from imposing pay secrecy clauses on their employees (Siniscalco et al. 2017).

To enhance comparability between the treated and control group, in the main specification we restrict the sample to firms with +/-50 employees from the 250 threshold.

Consistent with the hypothesis that pay transparency reduces the relative bargaining power of highly-paid employees (Cullen and Pakzad-Hurson 2021), our results show that the disclosure policy leads to a significant 2.6 percent decrease in men's real hourly pay in treated firms relative to control ones. We also find suggestive evidence that this effect is larger at the top of the wage distribution. In turn, pay compression from the top of the wage distribution results in a 14 percent reduction in the gender hourly pay gap, off a base of £3.

Event-study exercises show that these results are not driven by different pre-policy trends in the outcomes of interest between treated and control groups. Additional robustness checks, such as triple-differences regressions, difference-in-discontinuities specifications, and placebo regressions, exclude the possibility that our estimates capture the impact of time shocks affecting firms above and below the 250-employee threshold differently. Also, our estimates are not sensitive to choices made in the main specification, such as the bandwidth width around the 250 cutoff.

Regarding career outcomes, we do not find any significant effect of the policy on the promotions, recruitment, or retention of either male or female employees in the first two years following its introduction, though point estimates suggest that the probability of separations increases for both men and women in treated firms compared to control firms.

To complement our analysis of career outcomes, we study whether the policy affects firms' recruitment strategies by combining the difference-in-differences strategy with a text analysis of online job vacancies collected by Burning Glass Technologies. We find that, after the introduction of the policy, targeted firms are 10 percent more likely than control firms to post wages in their job ads—especially so in vacancies for the highest paid occupations. In other words, employers increase the transparency of the hiring process, potentially in an effort to attract highly-talented women.

To conclude our difference-in-differences analysis, we use the Business Structure Database, which covers 99 percent of UK firms, and the Annual Business Survey, which covers the production, construction, distribution, and service industry, to study the impact of the policy on labor productivity, as well as on labor costs and firm profits. Consistent with previous results in this literature (Card et al. 2012, Breza et al. 2018, Dube et al. 2019, Bennedsen et al. 2021), point estimates indicate that pay transparency reduces labor productivity, though the effect is marginally insignificant in the UK setting. However, the policy has no significant impact on profits, as the effect on productivity is compensated by a significant reduction of labor costs.

As for the mechanisms behind the main results, we provide descriptive evidence that a reputation motive may be important to understand firms' reactions. First, by linking the data published by targeted businesses with YouGov's Women's Rankings, we find that worse performing firms receive lower reputation scores from women. Second, by linking the equality indicators with ASHE, we find that worse performing firms reduce their gender pay gap the most between 2018 and 2019. To us, this suggests that, by enhancing public scrutiny and enabling comparisons across firms, the public disclosure of the equality indicators magnifies the disciplinary effects of the policy (Perez-Truglia and Troiano 2015, Luca 2018, Johnson 2020).[2]

Our paper contributes to different strands of literature. To begin with, we make several contributions to the growing number of studies analyzing the impact of pay transparency on the gender pay gap and wage inequality more broadly (Card et al. 2012, Mas 2017, Breza et al. 2018, Cullen and Perez-Truglia 2021, Cullen and Perez-Truglia 2018, Baker et al. 2019, Burn and Kettler 2019, Cullen and Pakzad-Hurson 2021, Dube et al. 2019, Bennedsen et al. 2021, Blundell 2021, Gulyas et al. 2020). The closest studies to ours are Baker et al. (2019), Bennedsen et al. (2021), Blundell (2021), and Gulyas et al. (2020). Baker et al. (2019) study the effect on the gender pay gap of a Canadian law requiring public sector organizations to publish employees' salaries above a certain pay threshold, while Bennedsen et al. (2021) and Gulyas et al. (2020) analyze the effect on the gender pay gap of, respectively, a 2006 Danish law and a 2011 Austrian law, mandating that private firms provide their employees with pay data by gender and occupation. Both Baker et al. (2019) and Bennedsen et al. (2021) find that transparency leads to pay compression from the

[2]Using the standard event-study methodology, we also analyze the reaction of the stock market and find that firms targeted by the policy experience a short-lived 35-basis-point drop in their cumulative abnormal returns.

upper part of the wage distribution. In contrast Gulyas et al. (2020) find no impact on individual wages or the gender pay gap, and suggest that this may be due to the fact that, in Austria, the pay information is not disclosed publicly. Relative to these studies, the UK legislation has two unique features that could help improve our understanding of the effects of pay transparency. First, the information disclosed focuses on the gender pay gap rather than pay levels by gender. While in the latter case, workers could react both to cross-gender comparisons and to comparisons with their own gender, in the UK this second channel is shut down. Second, in the UK setting the information is disclosed publicly rather than provided exclusively to employees' representatives, which allows us to study the role of reputation in influencing firms' responses to the policy.

Relative to the work conducted in parallel on the UK's policy (Blundell 2021), our paper combines several sources of data to provide comprehensive evidence on the effects of pay transparency not only on gender pay gaps, but also on career outcomes, labor productivity, and firm-level outcomes, which is necessary to evaluate the effectiveness of this type of policy. In this respect, our paper complements the evidence provided by Bennedsen et al. (2021) on Denmark's transparency policy, who, to the best of our knowledge, are the only ones to study the impact of a transparency policy on firm-level outcomes.

Next, our paper contributes to the increasing number of studies that use job advert data to analyze different dimensions of the labor market (Deming and Kahn 2018, Adams et al. 2020, Azar et al. 2020a, Azar et al. 2020b). In particular, to the best of our knowledge, this is the first paper to document a correlation between firms' hiring practices and the magnitude of firms' gender pay gaps. Moreover, our paper complements contemporary work on the impact of pay history inquiry bans on recruitment practices (Sran et al. 2020), by showing that the content of job vacancies is influenced by broader management practices.

The paper proceeds as follows. Section 2 describes the institutional setting and the UK transparency policy. Section 3 discusses the identification strategy. Section 4 presents the main results. Section 5 illustrates the robustness checks. Section 6 discusses potential mechanisms. Section 7 presents results on firm-level outcomes. Section 8 concludes.

2 Institutional setting

In 2015, the UK government launched a process of consultations with employers to enhance pay transparency. At that time, the average gender pay gap for all employees in the UK stood at 19.1 percent. Moreover, women made up only 34 percent of managers, directors, and senior officials (Government Equalities Office 2015). According to the government's view, "greater transparency will encourage employers and employees to consider what more can be done to close any pay gaps. Moreover, employers with a positive story to tell will attract the best talent" (Government Equalities Office 2015).

In February 2017, this process resulted in the passing of the *Equality Act 2010 (Gender Pay Gap Information) Regulations 2017*. This mandate requires all firms registered in Great Britain that have at least 250 employees to publish gender equality indicators both on their own website and on a dedicated website managed by the Government Equalities Office (GEO hereafter).[3] In the UK these comprise around 10,500 firms each year, representing only 0.4 percent of all UK firms but accounting for 40 percent of employment and 48 percent of turnover (Business Structure Database). To the best of our knowledge, no other substantial law exclusively targeted firms in this size band when the transparency mandate was introduced.[4]

The timing of the publication of the equality indicators works as follows. If a firm has at least 250 employees by April 5th of a certain year (the end of the financial year in the UK), it has to calculate the gender equality indicators as of that date, and publish them by the end of the following financial year. Firms themselves must calculate their number of employees using guidelines provided by the government. Importantly, they have to adopt an extended definition of

[3]The mandate does not apply in Northern Ireland, while in England, Wales, and Scotland, it applies to both private and public sectors. Note also that the public sector in these countries was already subject to some transparency measures. Further details on this are provided on the Equality and Human Right Commission's website: https://www.equalityhumanrights.com/en/advice-and-guidance/public-sector-equality-duty.

[4]Since 2010, employees working in firms with at least 250 employees have the right to request time off for training. Note that, even if this policy affected employees' outcomes differently below and above the 250-employee cutoff, the difference-in-differences strategy would take care of these effects, unless they interacted with the transparency policy. Also, since 2020, publicly listed firms with at least 250 employees have been required to publish pay gaps between the CEO and the median employee. However, note that only 1 percent of businesses with at least 250 employees are publicly listed.

an employee that includes agency workers. Partners of firms are also included in the definition of employee but should not be included in the calculation of the indicators. Finally, part-time workers have the same weight as full-time workers in the calculations.

The indicators that firms have to report include: the gender gap in the median (mean) hourly pay, expressed relative to men's pay; the gender gap in the median (mean) bonus pay, expressed relative to men's bonus pay; the proportion of male and female employees who receive any bonus pay; and the gender ratio in each quartile of a company's wage distribution. Table 1 provides sample means of these indicators for 2017/2018 and 2018/2019. Note that in mid-March 2020 the transparency mandate was temporarily paused due to the Coronavirus outbreak, and firms were only asked to start publishing the equality indicators again in October 2021. Considering this, we only study the effects of the policy in the first two years since its introduction.

The first row of Table 1 shows that the median gender pay gap is just below 12 percent in 2017/2018, and increases slightly in the following year.[5] The mean gender pay gap is around 14 percent in 2017/2018, and instead decreases a little the year after. Both median and mean bonus gaps tend to be smaller than pay gaps but it is worth noting from the large standard deviation that some firms mistakenly reported their level gap rather than a percentage, making it difficult to interpret these bonus gaps.[6] The proportion of women receiving bonus pay is smaller than for men in each year, and both increase slightly in the second year. The gender ratio along the wage distribution is balanced at the bottom, but less than 40 percent of employees in the upper part of the wage distribution are women. Also, the proportion of women in each quartile of the wage distribution increases slightly over the two years. From now on, we will refer to these data as the GEO data.

Three other features of this policy are important to understand the UK context. First, the policy does not impose sanctions on firms that do not improve their gender equality indicators over

[5]When asked to comment on their figures, "Ashfords, a law firm whose pay gap went up from 15.8 per cent to 39.4 per cent, said the figures were driven by a large number of junior women joining the company, which offset the effect of women being promoted to senior positions." (*Financial Times* 2019). In Section 4 we will provide some suggestive evidence supporting this mechanism.

[6]When excluding the bottom and top 1 percent, the median (mean) bonus gap stands at 13.14 (23.56) percent in 2017/18 and 12.35 (23.46) percent in the second year.

time. However, the Equality and Human Rights Commission, the enforcement body responsible for this regulation, can issue court orders and unlimited fines for firms that do not publish these indicators. As of 2020, all firms targeted by the law were deemed to have complied. Panel A of Figure 1 reports the distribution of submission dates for the first two years that the mandate was in place. While some firms do not meet the deadline, the majority publish their data in the last month before it. Note also that only 235 firms with fewer than 250 employees published gender equality indicators in 2018. These represent less than 0.1 percent of active UK firms with fewer than 250 employees in 2018 (Business Structure Database). This tiny percentage is consistent with the hypothesis that firms are reluctant to disclose information on employees' pay if they are not forced to do so (Siniscalco et al. 2017). It is also important to take into account this figure when thinking about the potential general equilibrium effects of this policy.

Second, according to a survey conducted on behalf of GEO prior to the introduction of this policy, out of 855 private and non-profit firms with at least 150 employees, only one third of firms had ever computed their gender pay gap, and just 3 percent had made these figures publicly available. Moreover, up to 13 percent declared that staff were discouraged from talking about their pay with colleagues and 3 percent reported that their contracts included a clause on pay secrecy (Downing et al. 2015). These figures suggest that the transparency policy is likely to represent an information shock both inside and outside the firm.

Finally, this policy is salient. Not only are the figures publicly available via a government website, but they also receive extensive media attention each year when they are published (*BBC* 2018, *The Guardian* 2018, *Financial Times* 2018, *Financial Times* 2019). Importantly, Panel B of Figure 1 shows that Google searches for the term "gender pay gap" spike around each year's deadline, indicating that this policy attracts significant public interest.

3 Identification strategy

To identify the impact of the UK transparency policy on pay, career outcomes and firm-level outcomes, we exploit the variation in its implementation across firm size and over time. Specifically, we estimate a difference-in-differences model that compares the evolution of the outcomes of interest in firms whose size is slightly above (treated group) or below (control) the 250-employee cutoff. As firm size can be endogenously determined, we define treatment status based on firm size in 2015, prior to the start of the consultation process to implement the mandate.[7] Moreover, in the main specification, we consider firms with $+/-50$ employees from the 250 threshold, to enhance comparability between treatment and control group. As both choices could be considered to be arbitrary, we show in the next section that our results are robust both to the use of a different year to define the treatment status and to changes in the bandwidth used to construct the estimation sample. When studying employees' outcomes, our baseline regression model is as follows:

$$Y_{ijt} = \mu_i + \alpha_j + \theta_t + \beta\,(TreatedFirm_j * Post_t) + X'_{it}\pi + Z'_{jt}\delta + u_{ijt}, \tag{1}$$

where i is an employee working in a firm j, which has between 200 and 300 employees, in year t, running between 2013 and 2019.[8] The outcome Y_{ijt} is either a measure of pay (hourly or weekly pay, contractual pay or bonuses), hours worked, promotion, occupation held, or job mobility. As for the regressors, μ_i are individual fixed effects that capture the impact of individual-specific time-invariant characteristics such as motivation or agreeableness; α_j are firm fixed effects that capture the impact of firm-specific time-invariant characteristics such as industry, or firm culture;[9] θ_t are year fixed effects that control for time shocks common to all firms such as electoral cycles; $TreatedFirm_j$ is a dummy equal to one if a firm has at least 250 employees in 2015, and $Post_t$ is

[7] Appendix Figure A1 shows the distribution of firms around the 250-employee cutoff in each year since the introduction of the mandate. Data are drawn from the Business Structure Database, which we further describe in Section 7. While a McCrary test performed separately for each year does not reject the null that there is no jump at the cutoff, it seems cautious to define treatment status based on pre-policy firm size.

[8] As explained in Section 4.1, we choose this time window because it is the maximum number of years over which we observe all outcomes of interest.

[9] Both employees' personality traits and firm culture can change over time. Nevertheless, it seems plausible to assume that these characteristics will be constant over the period of time considered.

a dummy equal to one from 2018 onward. Note that, in our data, we observe outcomes over fiscal years, which is also the time span over which firms have to publish equality indicators. Hence, for instance, 2018 refers to the fiscal year 2017/18 that starts in April 2017 and ends in March 2018. The vectors X_{it} and Z_{jt} include, respectively, time-varying individual and firm controls. In our main specification we only include firm-region specific time shocks, though we present robustness checks in Section 5 where we alternatively include industry-specific time shocks, or individual controls such as age and age squared. Our main coefficient of interest is β which, conditional on the validity of this identification strategy, should capture any deviation from a parallel evolution in the outcomes of interest between the treatment and the control group due to the introduction of the mandate. In all regressions, we use UK Labor Force Survey (LFS) weights, though in Section 5 we show that our results do not depend on this choice. Standard errors are clustered at the firm level, though in Section 5 we show that the significance of our results does not change when clustering them at the level of firm size times industry. Finally, as our hypothesis is that this policy will affect men and women differently, we will estimate each regression separately by gender. All regression tables will also report the p-value of the t-test on the equality of coefficients for men and women. In Section 5, we complement this strategy by estimating a triple-differences model using gender as the third dimension of treatment variation.

The validity of our identification strategy depends on three assumptions. First, it has to satisfy the parallel-trend assumption, that is, prior to the introduction of the policy, the evolution of the outcomes of interest must be comparable in treated and control firms. Second, our estimates should not capture the effect of other time shocks that coincide with the introduction of pay transparency and affect firms on either side of the 250-employee cutoff differently. Third, the results should not depend on the size of the bandwidth considered around the policy cutoff, nor should they depend on the year chosen to define the treatment status.

To support the validity of the parallel-trend assumption and study the dynamic impact of pay transparency, we will open the discussion of our main findings by illustrating the results of the following event-study exercise:

$$Y_{ijt} = \mu_i + \alpha_j + \theta_t + \sum_{k=2013}^{2019} \beta_k (TreatedFirm_j * \mathbf{1}[t = k]) + X'_{it}\pi + Z'_{jt}\delta + u_{ijt}, \quad (2)$$

where $\mathbf{1}[t = k]$ is an indicator variable that takes value 1 when $t = k$ and 0 otherwise. In what follows, we take 2017, the year prior to the introduction of pay transparency, as the reference year.

Next, in Section 5, we will provide evidence that the other two conditions necessary for the validity of the identification are also met.

4 Impact on employees' outcomes

4.1 Data

To measure employees' outcomes, we use the Annual Survey of Hours and Earnings (ASHE). ASHE is an employer survey covering 1 percent of the UK workforce that is conducted every year and is designed to be representative of the employee population.[10] The ASHE sample is drawn from National Insurance records for working individuals, and the selected workers' employers are required by law to complete the survey. Specifically, ASHE asks employers to report data on gender, pay, hours, and tenure for the selected employees, using a snapshot in April each year. Information on age, occupation, and industrial classification is also available. Once workers enter the survey, they are followed even when changing employer, though individuals are not observed when unemployed or out of the labor force. In practice, ASHE is an unbalanced panel data set at the employee level. While the main limitation of ASHE is its small sample size, crucially it provides the number of employees in a firm and year, which allows us to define the treatment status in our identification strategy.[11] From ASHE, we create the following variables:

[10] Office for National Statistics. (2019). Annual Survey of Hours and Earnings, 1997-2019: Secure Access. [data collection]. 14th Edition. UK Data Service. SN: 6689, http://doi.org/10.5255/UKDA-SN-6689-13.

[11] When none of the employees of a firm is interviewed in ASHE in the year used to define the treatment status, we recover the information on firm size from the Business Structure Database. This concerns 27 percent of firms in our sample. In Section 5 we show that our results are not affected if these firms are excluded from the estimation sample.

Pay measures. Our main variable of interest is log real hourly pay, including bonuses but excluding overtime pay; we also separately consider log basic real hourly pay and bonus payments. To study the impact of the policy on bonuses, we use the inverse hyperbolic sine transformation to take into account the fact that many workers do not receive these payments. Finally, we consider log real weekly pay, and weekly hours worked. All monetary values are deflated using the ONS 2015 CPI Index.

Career outcomes. First, we are interested in the impact of the policy on employees' probability of being promoted; for this, we follow the ONS definition of a promotion and construct a dummy equal to one if an employee has experienced at least a 30 percent increase in his/her hourly pay since the previous year, and/or has acquired managerial responsibility since then. Next, we construct a dummy equal to one if a worker is employed in top-paid occupations, i.e., technical, professional, and managerial occupations. We then use months of tenure in a firm to study mobility into that firm; this variable is missing for around 2 percent of the estimation sample. Finally, we construct a dummy variable that is equal to one if the employee leaves the firm in $t + 1$, either by moving to a different a firm, or by leaving the labor market.

In the empirical analysis, we use data over the period 2013–2019. We start from 2013, as we can observe all outcomes since then, and stop in 2019, as we prefer not to include years affected by the pandemic. However, note that we use information from 2012 and 2020 to construct, respectively, the promotion dummy and the indicator for leaving a firm in $t + 1$. In terms of sample restrictions, we drop individuals with missing id or missing firm id (0.5 percent of the sample); we drop secondary jobs (3 percent); we drop individuals who work at least once more than 100 hours per week and those with an hourly pay greater or equal to £1000 (0.2 percent). Our resulting sample is formed of 5,165 men and 4,489 women, for a total of 15,899 individual-year observations for men and 13,568 for women. We observe men across 2,940 firms and women across 2,652 firms.

Table 2 provides summary statistics for the main outcomes measured in the pre-policy period, 2013–2017. Several features are worth noting. First, the profile of workers in treated and control firms is remarkably similar. Second, focusing on the treatment group (columns 1 and 3), the unconditional hourly pay gap amounts to £3.19, or 19 percent of men's pay. There is also a large gender gap in the probability of receiving bonuses (45 percent), and a very large gap in the amount received (62 percent). While both men and women have a 7 percent chance of getting promoted each year, there is an 8 percent gender gap in favor of men in the probability of working in top-paid occupations. Men are also more likely to stay longer in a firm than women, and to work in the private sector—though this share is already around 80 percent for women, which prevents us from studying heterogeneous effects between public and private-sector employees. Finally, note that among both men and women, only one third of workers are covered by a collective agreement, which similarly limits our ability to study heterogeneous effects between unionized and non-unionized workers.

4.2 Findings

This section presents the impact of the pay transparency policy on employees' outcomes. Figure 2 shows the event studies for the main variable of interest, log hourly pay, separately for men and women. The event studies for the other outcomes are reported in Appendix A. The graph in Panel A refers to men, the one in Panel B to women. From these figures, we observe, first, that the evolution of the outcome in the pre-policy period seems to be comparable across treatment and control groups, both for male and female employees. Second, we see that men's real hourly pay decreases in treated firms relative to control firms after the introduction of the mandate, with this effect becoming significant at the 5 percent level in 2019. In contrast, the policy does not appear to have any visible impact on women's pay.

Table 3 reports estimates of the corresponding average effects of the policy on the main outcomes of interest. Panel A refers to men, while Panel B refers to women. Each column shows a different outcome. At the bottom of the table, we report the p-value of the t-test on the equality of coefficients for men's and women's outcomes and the pre-policy mean for the treated group calculated over the period 2013–2017. Consistent with the dynamics shown in the event study, Column 1 shows that the policy leads to a 2.6 percent decrease in men's real hourly pay in treated firms relative to control firms, with this effect being significant at 5 percent. In contrast, the policy does not seem to have a significant effect on women's pay. In turn, as indicated by the p-value of the t-test on the equality of coefficients, these effects translate into a significant reduction in the gender pay gap, amounting to 14 percent of the pre-policy mean,[12] which is very similar to the effect found by Bennedsen et al. (2021).

In Columns 2 to 7, we unpack these wage effects into different pay components. Columns 2 and 3 suggest that the effect on hourly pay comes from wages rather than hours worked, though the impact on men's weekly pay is just marginally insignificant. Moreover, the event study reported in Appendix Figure A2 shows a significant drop in men's weekly pay in the second year of the transparency regime. Note also that women's hours worked and weekly pay seem to be negatively affected by the policy, consistent with a scenario in which firms become more inclined to accommodate women's requests for flexible work in an effort to retain them. Columns 4 and 5 of Table 3 complement this result by showing that the effect on men's pay is driven by the contractual part of pay, while the impact on bonuses is negative but not significant, nor significantly different from the effect for women. Note also that the (insignificant) decline in women's bonuses is consistent with the negative coefficient on hours worked, considering that, as shown in Appendix Figure A4, and documented by literature (Azmat and Ferrer 2017), incentive payments tend to be positively correlated with hours worked.

In Table 4, we complement these results by showing the impact of the policy on career outcomes. First, point estimates in Column 1 suggests that the policy leads to fewer promotions for men (a decrease of 7 percent relative to the pre-policy mean) and more for women (13 percent increase relative to the pre-policy mean), though the coefficients are imprecisely estimated. Note also that in the event study of men's promotions reported in Appendix Figure A3, the negative effect of the policy becomes significant in 2019. At the same time, point estimates in Column 2 of Table 4 suggest that the policy decreases the probability that women work in top-paid occupations.

[12]According to the estimates shown in Table 3, the transparency policy reduces men's pay by 2.6 percent relative to a pre-policy mean of £17.07, that is £0.44. The pre-policy gender hourly pay gap amounts to 3.19 £. Thus, the policy leads to a reduction of 0.44/3.19 or 14 percent.

While this may appear surprising, it is consistent with businesses' claims that they have tried to hire more women in junior roles, as reported in Section 1. Results in Column 3 show that tenure in the firm is barely affected by the policy (the point estimates represent an effect of at most 0.001 percent relative to the pre-policy means). In contrast, the policy seems to decrease retention of both male and female employees in treated firms relative to control firms, though the coefficients in Column 4 are not precisely estimated.

Overall, these results show that the transparency policy leads to a slowdown of men's pay growth that is concentrated in the contractual part of pay, and that is partly driven by a decrease in the opportunities for promotion.

The next section is dedicated to showing that these results are not driven by time shocks that affect treated and control firms differently, and that they are robust to the use of different models and changes in the regression specification. Following this, Section 6 will explore the contribution of different channels in explaining these results.

5 Robustness checks

This section presents two sets of robustness checks. First, we show that our results are unlikely to be driven by contemporaneous shocks to the policy that have heterogeneous effects across treated and control firms. Second, we show that our results do not depend on the choices made in the main specification, in particular in terms of the size of the bandwidth around the policy cutoff, and the year used to define the treatment status. To summarize all these results, we visually represent them in Figure 3, where Panel A refers to men, and Panel B refers to women. Appendix Section A reports detailed regression tables. We then conclude this section by discussing potential general equilibrium effects and the external validity of our results.

Contemporaneous shocks. To make sure that our estimates do not capture the effect of other events occurring at the same time as the introduction of pay transparency requirements that could affect treated and control firms differently, we perform three robustness checks. First, we run a series of placebo tests pretending that the mandate binds at different firm size thresholds. Figure 3 reports the estimates of the main placebo regression for both men and women, and Appendix Table A1 displays the corresponding detailed regression results. The estimation sample includes firms with $+/-50$ employees from the 150 threshold. Reassuringly, this placebo mandate does not appear to have an impact on hourly pay for either men or women. This exercise should help exclude the possibility that our estimates capture the impact of time shocks that happen at the same time as the mandate and affect larger firms differently to smaller firms.[13] As a second robustness check, we estimate the following triple-differences model with the gender dimension as the third difference:

$$\begin{aligned} Y_{ijt} &= \mu_i + \alpha_j + \theta_t + \beta\,(TreatedFirm_j * Post_t) \\ &\quad + Fem_i[\gamma_0 + \gamma_1 TreatedFirm_j + \gamma_2 Post_t + \gamma_3 (TreatedFirm_j * Post_t)] \\ &\quad + X'_{it}\pi + Z'_{jt}\delta + u_{ijt}, \end{aligned} \tag{3}$$

where Fem_i is a dummy variable that is equal to one if i is a woman, and all other variables are defined as in regression 1. As such, this alternative specification controls for within-group time shocks that are common to male and female employees. Reassuringly, the estimates from this triple-differences model, reported in Figure 3 and Appendix Table A2, are practically indistinguishable from those of the double-differences model. As a third robustness check, we estimate the following difference-in-discontinuities model:

$$\begin{aligned} Y_{ijt} &= \mu_i + \alpha_j + X'_{it}\pi \\ &\quad + Post_t[\delta_0 + \delta_{reg} + \delta_1 FirmSize_{j2015} + TreatedFirm_j(\beta_0 + \beta_1 FirmSize_{j2015})] \\ &\quad + u_{ijt}, \end{aligned} \tag{4}$$

[13] As for larger placebo cutoff values considered in Appendix Table A1, it should be noted that these regressions include all treated firms. The fact that the magnitude of the effects are non-zero may simply point to heterogeneous effects of the policy across firm size, consistent with the idea that larger firms are more exposed to public scrutiny.

where δ_{reg} are region fixed effects and $FirmSize_{j2015}$ is a continuous variable measuring the number of employees in firm j in 2015 relative to the 250-employee cutoff. The main difference between our main specification and this one is that the difference-in-discontinuities model takes into account the possibility that firms with a different number of employees are on different trends (Grembi et al. 2016). Though our event studies seem to exclude that this is the case, this exercise should further support this assumption. Results are reported in Figure 3 and Appendix Table A3. While the impact on men's pay is only significant at 10 percent in this specification, the point estimate changes little compared to our preferred specification.

Specification. Our second set of robustness checks aims to verify above all that our results are robust to the choice of the bandwidth around the 250-employee cutoff, do not depend on the year we use to define the treatment status, and are not sensitive to the data we use to define treatment status. Figure 3 shows that the estimates of β from equation 1 change very little when restricting or enlarging the bandwidth around the 250-employee cutoff. Appendix Table A4 displays the corresponding detailed regression results and provides estimates using further bandwidth choices. Specifically, estimates of the impact of the policy on men's wages only become marginally insignificant when estimating the model using small samples (bandwidths between 30 and 40), while they are significant and of similar magnitude when enlarging the sample. Conversely, estimates of the effect on women's pay are always close to zero and insignificant, with the estimated zero effect becoming more precisely estimated as we enlarge the sample.

Finally, Figure 3 and Appendix Table A5 show that our main findings change little when: changing the year used to define the treatment status; including industry-specific time-shocks in place of region-specific time shocks; adding age controls; not using LFS weights; restricting the sample to either workers aged 16-65 or those aged 25+; considering only full-time employees; restricting the sample to firms for which we can use only ASHE-based information on the number of employees to define the treatment status; and clustering standard errors at firm-size times industry.

To sum up, our estimates are very stable across different specifications and sample sizes, which strongly supports the validity of our identification strategy.

General equilibrium effects and external validity. One of the most striking results in Table 3 is the absence of effects on women's pay. In principle, this could be due to the fact that both treated and control firms have raised women's pay as they compete for the same workers. To explore this possibility, in Appendix Figure A5 we plot the evolution of both men's and women's pay separately by treatment group. On the one hand, the effect on men's pay clearly comes from a decrease of real wages in treated firms. As for women, their wages increase in both treated and control firms, but we do not see any sharp increase after the introduction of pay transparency in either of the two group. In other words, it seems implausible that general equilibrium effects could completely explain the null effect on women's pay.

To provide some insights regarding the external validity of our estimates, Appendix Figure A6 compares the occupational and industry distribution of men and women in the estimation sample to that of the entire ASHE population, over the period studied. The occupational distribution displayed in Panel A is remarkably similar in the two samples both for men and women, with only some under-representation of sales occupations in women's estimation sample. As for the industry distribution, Panel B of Figure A6 shows that, with the exception of the manufacturing sector being over-represented in men's estimation sample, the distribution matches well across the two samples. Taken together, these figures suggest that, in the absence of large equilibrium effects, the estimated effects can hold across the firm size distribution.

6 Mechanisms

Our results indicate that the UK transparency policy led targeted firms to slow down the wage growth of male employees relative to similar firms that were not subject to this mandate. In this section, we exploit different sources of data to explore which employees are most affected, which firms react the most, and whether a reputation motive helps explain firm reactions.

Which employees. Many studies show that the gender pay gap is larger at the top of the wage distribution, a stylized fact that contributes to the broader phenomenon of the "glass ceiling" (Bertrand et al. 2010, Cortes and Pan 2019). The policy may therefore push firms to act mostly on the upper part of the wage distribution to more effectively close the gap. Using ASHE, Appendix Table A6 lends some support to this hypothesis, by comparing the effect of the policy on employees working in top-paid occupations, i.e., managerial, professional, and technical occupations, versus other professions.[14] While the results are not statistically different across subgroups, point estimates in Column 2 and 3 of Panel A suggest that the slowdown in men's pay is larger in top-paid occupations compared to other professions. Point estimates in Column 2 of Panel B also suggest that women in top-paid occupations experience pay increases relative to their counterparts in the control group, while the effect is close to zero in other professions. Note that, taken at face value, these estimates imply an 18 percent reduction of the gender pay gap in top-paid occupations.[15]

Which firms. One of the most innovative features of the UK transparency policy as compared to the mandates introduced in other countries is that firms have to make the equality indicators publicly available. The behavioral economics literature suggests that when individuals receive information on their relative performance, those performing worst improve the most afterwards (Allcott and Kessler 2019), and the same may be true of firms. Unfortunately, using ASHE, we cannot use the difference-in-differences design to study whether firms react in this way as we cannot compute the firm-level gender pay gap pre-policy.[16] However, we explore this mechanism descriptively in Figure 4. Panel A shows the evolution of firms' (publicly available) median gender hourly pay gap relative to their 2018 gender pay gap. Consistent with the behavioral literature, firms with a worse performance at baseline reduce the gender pay gap the most between the two

[14]Note that, when analyzing the impact of the policy by subgroup, we exclude region-specific time shocks from the regression to avoid running out of degrees of freedom.

[15]The 2.9 percent decrease in men's pay in top-paid occupations from a base of £23.51 corresponds to £0.68. The 1.1 increase in women's pay in these occupations compared to a pre-policy mean of £18.73 corresponds to £0.18. Relative to a pre-policy gender gap of £4.78 in these occupations, the total £0.86 effect corresponds to an 18 percent decrease in the pay difference.

[16]Note that ASHE does not provide information on all employees in a firm.

years. By merging the publicly available data with ASHE,[17] we also find that this seems to come both from a relatively lower growth in men's real pay (left-hand side graph of Panel B) and a higher increase in women's real pay (right-hand side graph of Panel B).

Why. Until now we have shown that firms react to the transparency policy. What we also want to understand is why they have reacted. Reputation concerns are a plausible reason if the policy increases public scrutiny. We bring descriptive evidence supporting this hypothesis, by using YouGov's Women's Rankings for 2018 and 2019. Every year since 2018, YouGov surveys a representative sample of 50 to 100 women per day between February in one year and January of the following year, asking for their impression of a list of 1,340 (self-selected) brands. YouGov then constructs impression scores based on women's answers to the question: "Overall, of which of the following brands do you have a positive/negative impression?". The resulting score is the percentage difference between all the positive and negative answers, relative to all the answers received.

Our objective is to investigate whether firms' rank is associated with their performance in gender equality indicators. For this, we link YouGov data with firms' gender equality indicators using a name-matching algorithm described in Appendix B supplemented by manual matching. We are able to match 996 companies in 2018 and 1,018 in 2019.[18] Note that firms voluntarily ask YouGov to be included in their surveys. While Appendix Table B1 shows that firms included in the YouGov list do not perform better in terms of the main equality indicators, this group of firms potentially care more about their reputation.

Despite this, Panel A of Table 5 shows that, over the two years of data, a larger median gender pay gap is negatively correlated with the Women's Impression Score, while a higher percentage of women at the top of the wage distribution is positively associated with it. Panels B and C further break down the analysis by year and show that these correlations tend to become larger in magnitude and more significant from the first to the second year of data, when, potentially, more

[17]We find 6.748 firms targeted by the policy in ASHE, or two thirds of all businesses that have to publish gender equality indicators.

[18]Most of the YouGov companies that we cannot link with the GEO data are not registered in the UK.

people became aware of the gender equality indicators. To us, this suggests that firms are under the scrutiny of women, and this could help explain why firms have responded to the transparency policy.[19] In Appendix C, we further show that firms targeted by the mandate experience a negative, though short-lived, drop in their stock market returns right after their publication of the equality indicators, suggesting that firms are also under the scrutiny of investors. Overall, these two pieces of evidence indicate that a reputation motive may help explain why firms react to this policy.

7 Impact on firm-level outcomes

7.1 Labor productivity and profits

To evaluate the effectiveness of transparency policies, it is important to consider all of their implications for workers and firms. Their effect on labor productivity is especially important and a priori ambiguous. On the one hand, in accordance with the "fair wage-effort hypothesis" (Akerlof and Yellen 1990), the information revealed may decrease the job satisfaction of lower-paid employees if they perceived to be treated unfairly by the information they acquire, while the higher-paid may feel threatened by any attempt by the firm to mitigate inequality (Card et al. 2012, Breza et al. 2018, Dube et al. 2019). In addition, as discussed in Section 6, the public disclosure of gender equality indicators induces comparisons across firms. Taken together, pay comparisons among employees within and outside the firm may lower job satisfaction, with negative knock-on effects for labor productivity. On the other hand, if firms respond to the policy by improving gender equality, this could boost the productivity of those workers who care about working in a fair environment (Bennedsen et al. 2021). In turn, any impact on labor productivity could translate into an effect on firms' profits, taking into account that our results imply a reduction of labor costs.

[19] In Appendix Table B2, we also explore the correlation between firms' gender equality indicators and their score in the YouGov's Workforce Ranking. This is obtained by asking both men and women the following questions about a sample of 1342 firms: "Imagine you (or your friend) were applying for the same sort of role at the following brand that you currently have or would apply for?" coupled with "Which of the following brands would you be proud to work for?" and "Which of the following brands would you be embarrassed to work for?". The point estimates in Table B2 suggest that a worse performance on gender equality indicators is associated with a lower rank in this index too, though the correlations are not significant.

We study these implications in Table 6, by estimating a firm-level version of regression 1 on labor productivity, labor costs, and firms' profits. We proxy labor productivity with log output from the Business Structure Database (BSD).[20] The BSD provides information on firm output and employment for almost 99 percent of business organizations in the UK. The data are reported as of April each year, and come from the Inter-Departmental Business Register (IDBR), a live register of firms collected by the tax authorities via VAT and employee tax records.[21]. Note that we do not divide turnover by the number of employees to avoid capturing the complex evolution of firm size in a constrained interval of +/- 50 employees from the 250-employee threshold.[22] Next, we use the Annual Business Survey (ABS) to extract information on firms' wage costs and gross value added.[23] ABS is an annual survey of businesses covering the production, construction, distribution, and service industries, and as such represents about two-thirds of the UK economy in terms of gross value added (GVA). So far, it is only available up to 2018. We measure labor costs using the log of wage costs, and consider the inverse hyperbolic transformation of GVA over assets to account for negative profits. All monetary values are deflated using the ONS 2015 CPI Index.

Results reported in Table 6 point to a negative but insignificant effect on labor productivity, remarkably similar to what Bennedsen et al. (2021) find for Denmark and consistent with the potential increase in employees' separations discussed in Section 4. In line with the negative effect we find on men's pay, the policy also leads to a significant reduction in firms' wage bill.[24] In turn, the negative effects on labor productivity and labor costs seem to compensate each other, as the impact on profits is positive but not significant.

[20]Office for National Statistics. (2019). Business Structure Database, 1997-2019: Secure Access. [data collection]. 10th Edition. UK Data Service. SN: 6697, http://doi.org/10.5255/UKDA-SN-6697-10.

[21]If a business is liable for VAT (turnover exceeds the VAT threshold) and/or has at least one member of staff registered for the Pay-as-You-Earn tax collection system, then it will appear on the IDBR (and hence in the BSD). As a result, only very small businesses do not appear in the IDBR.

[22]Ideally, we would use log ouput over hours worked, but unfortunately, to the best of our knowledge, there are no data on employees' hours worked at the firm level.

[23]Office for National Statistics. (2021). Annual Business Survey, 2005-2018: Secure Access. [data collection]. 14th Edition. UK Data Service. SN: 7451, http://doi.org/10.5255/UKDA-SN-7451-14.

[24]Admittedly, point estimates are larger than what a 2.6 percent reduction in the pay of male employees could imply. As shown in Appendix Figure D1, these estimates may in part capture the effect of pre-policy differential trends between treated and control firms.

7.2 Hiring practices

We conclude our analysis by studying the impact of the transparency policy on firms' hiring practices, and in particular the wage-posting decision. While so far the policy does not seem to affect firms' ability to recruit and retain female employees, increasing efforts to attract women in highly-paid occupations could be key to reducing firms' gender pay gap without changing men's pay.

Many studies document that there exists a gender gap in bargaining skills. In particular, women are less likely to ask for wage increases (Babcock et al. 2003, Bowles et al. 2007), and tend to avoid bargaining when they apply for jobs that leave wage negotiation ambiguous (Leibbrandt and List 2015). Providing information on wages posted upfront in job adverts could therefore help firms attract female job applicants. We study this hypothesis by combining the difference-in-differences strategy with a text-analysis of online job listings from Burning Glass Technologies (BGT). BGT scrapes online job ads from company websites and job boards. UK data are available since the 2013 financial year and cover more than 50 million vacancies. They provide information on the full job-ad text, occupational code (95 percent of vacancies), job location (90 percent of vacancies) and employer name, whenever this is posted, which occurs in 30 percent of the vacancies over the entire period. We focus on the restricted sample providing employers' names and merge it with FAME, the UK version of Amadeus, using the name-matching strategy introduced in Section 6, to retrieve information on firms' number of employees. Appendix E explains in detail how we construct the estimation sample and the outcome variables, and addresses potential selection concerns regarding the restricted sample of vacancies providing the employer's name. In a nutshell, we define a vacancy as posting wage information if it mentions either a salary interval or a single salary amount. A series of validation exercises conducted by research assistants show that we correctly classify the presence of wages in 98 percent of cases. The residual 2 percent are false negatives, meaning that our code indicates that there is no wage posted when there is actually one.

Before embarking in the difference-in-differences analysis, in Appendix Figure E4, we first investigate how firms' wage-posting decision correlates with gender equality performance. Specifically, this bar graph reports the correlation between GEO firms' average percentage of vacancies posting wage information between 2013 and 2019 and, respectively, the average percentage of women in top-paid occupations (blue bar) and the average gender pay gap (red bar) between 2018 and 2019, conditional on firms' sector and the occupational composition of vacancies. This analysis shows that firms that are more likely to post wage information also tend to have a larger percentage of women in top-paid occupations and a lower gender pay gap, which strongly motivates us to study the impact of the transparency policy on the use of this hiring practice.

For this, we run a version of regression 1 at the vacancy level controlling for employer, quarter, and region-specific time shocks. According to the results presented in Table 7, the policy significantly increases the probability that employers report wage information on their job ads by 4 percentage points, or 10 percent compared to the pre-policy sample mean. Appendix Figure E5 reports the corresponding event-study. While some leads of the policy are significant at the beginning of the sample, overall, there are little signs of differential pre-trends before the introduction of the policy, while its effect becomes visible in the second year of the transparency regime.[25] A comparison between Columns 2 and 3 of Table 7 also suggests that the effect is larger in vacancies for top-paid professions, though the effect is not statistically different across the two subgroups.

In Appendix Table E5, we complement this analysis by studying how the policy affects the value of posted wages, converting all salaries to their annual equivalents. A priori, the effect of the transparency policy on this outcome is ambiguous. On the one hand, as the share of firms posting wages increases, we may expect the marginal firm to be more cautious and post lower wages. On the other hand, treated firms may be willing to post higher entry wages to mitigate reputation and recruitment concerns. Results in Column 2 show that the policy leads to a significant increase in offered entry wages. Columns 3 and 4 suggest that this may be in part due to the fact that, after the introduction of the transparency policy, treated firms are less likely to post vacancies for part-time positions, potentially in favor of full-time flexible work arrangements, such as flexi-work

[25]To further ensure that we are not capturing the impact of pre-policy differential trends and also take into account that BGT data have been shown to be less representative of the UK labor market in the fiscal years 2013 and 2014, in Appendix Table E4, we show that this result does not change when estimating the regression from the fiscal year 2015 onward. Additionally, the estimate is also barely affected when restricting the sample to firms that are perfectly matched (match-score=1) in the name-matching of BGT and FAME.

or remote work. Even more importantly, the increase in posted wages is consistent with another piece of suggestive evidence drawn from ASHE; when comparing the impact of the policy on the wages of recently hired workers (at most two years of tenure) and the wages of incumbent employees in Appendix Table A7, women who have recently joined the firm seem to experience up to a 5 percent increase in their wages. Moreover, the slowdown of men's wage growth seems to be driven by incumbent workers. Taken together, these findings suggest that targeted firms react to the pay transparency policy by decreasing the wage growth of highly-paid and more experienced men, and try to attract women by posting wage information upfront in their job vacancies and offering higher entry wages.[26]

8 Conclusion

To tackle the persistence of gender inequality in the labor market, many governments are promoting pay transparency policies. Exploiting the variation across firm size and over time in the application of the UK's transparency policy, this paper shows that increased transparency significantly reduces the gender pay gap by 14 percent. Importantly, this effect is the result of a slowdown in men's wage growth, while the policy has no significant effect on women's pay. In other words, pay transparency leads to pay compression from the upper part of the wage distribution.

To give a comprehensive assessment of the policy, we also show that it has no detrimental effect on firms' profits, though we find suggestive evidence that labor productivity and employees' retention may have been negatively affected. Overall, our results suggest that transparency policies can reduce the gender pay gap with limited costs for firms, but may not be suited to achieve the objective of improving outcomes of lower-paid employees. As an increasing number of studies confirm that transparency policies mainly lead to pay compression from the top (Mas 2017, Baker et al. 2019, Cullen and Pakzad-Hurson 2021, Bennedsen et al. 2021), policy makers should ask

[26] Of course, these hiring practices may also attract male candidates, but what is important for firms is to avoid experiencing a drop in female job applicants to be able to choose between the two genders, especially in highly paid occupations.

whether this is a desirable way to tackle wage inequality.

To conclude, it is important to stress that our analysis only identifies short-term effects. Potentially, pay transparency may only have the effect of stimulating responses from firms in the short run, when it acts as an information shock that attracts strong attention from the public and leads firms to fear for their reputation. If the policy does not produce an actual change in the culture within firms, its effect may fade away over time as the strength of the information shock weakens (Giuliano 2021). In this respect, our results also show that the UK policy increases the transparency of the hiring process. If, over time, this helps firms attract more women into in highly-paid positions, pay transparency may result in positive long-term effects on gender equality. Considering these ambiguous long-run effects, it is necessary to keep monitoring the impact of these types of policies in order to fully assess their effectiveness.

References

ACAS, "Flexible Working and Work-Life Balance," Technical Report, Advisory, Conciliation, and Arbitration Service 2015.

Adams, Abigail, Maria Balgova, and Qian Matthias, "Flexible Work Arrangements in Low Wage Jobs: Evidence from Job Vacancy Data," CEPR Discussion Paper No. 15263, Centre for Economic Policy Research 2020.

Akerlof, George A and Janet L Yellen, "The Fair Wage-Effort Hypothesis and Unemployment," *The Quarterly Journal of Economics*, 1990, *105* (2), pp. 255–283.

Allcott, Hunt and Judd B Kessler, "The Welfare Effects of Nudges: A Case Study of Energy Use Social Comparisons," *American Economic Journal: Applied Economics*, 2019, *11* (1), pp. 236–276.

Aumayr-Pintar, Christine, Anna-Karin Gustafsson, Anna Savolaine, and Carsten Jørgensen, "Pay Transparency in Europe: First Experiences with Gender Pay Reports and Audits in Four Member States," *Publications Office of the European Union, Luxembourg*, 2018.

Azar, José, Ioana Marinescu, and Marshall Steinbaum, "Labor Market Concentration," *Journal of Human Resources*, 2020, pp. 1218–9914R1.

_ , _ , _ , and Bledi Taska, "Concentration in US Labor Markets: Evidence from Online Vacancy Data," *Labour Economics*, 2020, *66*, pp. 101886.

Azmat, Ghazala and Rosa Ferrer, "Gender Gaps in Performance: Evidence from Young Lawyers," *Journal of Political Economy*, 2017, *125* (5), pp. 1306–1355.

Babcock, Linda, Sara Laschever, Michele Gelfand, and Deborah Small, "Nice Girls Don't Ask," *Harvard Business Review*, 2003, *81* (10), pp. 14–16.

Baker, Michael, Yosh Halberstam, Kory Kroft, Alexandre Mas, and Derek Messacar, "Pay Transparency and the Gender Gap," NBER Working Paper No. 25834, National Bureau of Economic Research 2019.

Bell, Brian and Stephen Machin, "Minimum Wages and Firm Value," *Journal of Labor Economics*, 2018, *36* (1), pp. 159–195.

Bennedsen, Morten, Elena Simintzi, Margarita Tsoutsoura, and Daniel Wolfenzon, "Do Firms Respond to Gender Pay Gap Transparency?," *Journal of Finance, forthcoming*, 2021.

Bertrand, Marianne, Claudia Goldin, and Lawrence F Katz, "Dynamics of the Gender Gap for Young Professionals in the Financial and Corporate Sectors," *American Economic journal: Applied Economics*, 2010, *2* (3), pp. 228–255.

Blundell, Jack, "Wage Responses to Gender Pay Gap Reporting Requirements," CEPR Discussion Paper No. 1750, Centre for Economic Policy Research 2021.

Bowles, Hannah Riley, Linda Babcock, and Lei Lai, "Social Incentives for Gender Differences in the Propensity to Initiate Negotiations: Sometimes it Does Hurt to Ask," *Organizational Behavior and human decision Processes*, 2007, *103* (1), pp. 84–103.

Breza, Emily, Supreet Kaur, and Yogita Shamdasani, "The Morale Effects of Pay Inequality," *The Quarterly Journal of Economics*, 2018, *133* (2), pp. 611–663.

Burn, Ian and Kyle Kettler, "The More You Know, the Better You're Paid? Evidence from Pay Secrecy Bans for Managers," *Labour Economics*, 2019, *59*, pp. 92–109.

Card, David, Alexandre Mas, Enrico Moretti, and Emmanuel Saez, "Inequality at Work: the Effect of Peer Salaries on Job Satisfaction," *American Economic Review*, 2012, *102* (6), pp. 2981–3003.

Cortes, Patricia and Jessica Pan, "When Time Binds: Substitutes for Household Production, Returns to Working Long Hours, and the Skilled Gender Wage Gap," *Journal of Labor Economics*, 2019, *37* (2), pp. 351–398.

Cullen, Zoë and Bobak Pakzad-Hurson, "Equilibrium Effects of Pay Transparency," NBER Working Paper No. 28903, National Bureau of Economic Research 2021.

_ and Ricardo Perez-Truglia, "The Salary Taboo: Privacy Norms and the Diffusion of Information," NBER Working Paper No. 25145, National Bureau of Economic Research 2018.

_ and _ , "How Much Does Your Boss Make? The Effects of Salary Comparisons," *Journal of Political Economy, forthcoming*, 2021.

Dahlgreen, Will, Ransome Mpini, Daniele Palumbo, and Clara Guibourg, "What is the Gender Pay Gap at Your Company?," *BBC News*, 2018.

Deming, David and Lisa B Kahn, "Skill Requirements across Firms and Labor Markets: Evidence from Job Postings for Professionals," *Journal of Labor Economics*, 2018, *36* (S1), pp. S337–S369.

Downing, Christabel, Erica Garnett, Katie Spreadbury, and Mark Winterbotham, "Company Reporting: Gender Pay Data," Technical Report, IFF Research 2015.

Dube, Arindrajit, Laura Giuliano, and Jonathan Leonard, "Fairness and Frictions: the Impact of Unequal Raises on Quit Behavior," *American Economic Review*, 2019, *109* (2), pp. 620–63.

Giuliano, Paola, "Gender and culture," *Oxford Review of Economic Policy*, 2021, *36* (4), pp. 944–961.

Goldin, Claudia, "A Grand Gender Convergence: Its Last Chapter," *American Economic Review*, 2014, *104* (4), pp. 1091–1119.

Government Equalities Office, GEO, "Closing the Gender Pay Gap. Government Consultations," Technical Report 2015.

Grembi, Veronica, Tommaso Nannicini, and Ugo Troiano, “Do Fiscal Rules Matter?,” *American Economic Journal: Applied Economics*, 2016, pp. 1–30.

Gulyas, Andreas, Sebastian Seitz, and Sourav Sinha, “Does Pay Transparency Affect the Gender Wage Gap? Evidence from Austria,” CRC TR 224 Discussion Paper Series No. 194, University of Bonn and University of Mannheim, Germany 2020.

Johnson, Matthew S, “Regulation by Shaming: Deterrence Effects of Publicizing Violations of Workplace Safety and Health Laws,” *American Economic Review*, 2020, *110* (6), pp. 1866–1904.

Kommenda, Niko, Caelainn Barr, and Josh Holder, “Gender Pay Gap: What We Learned and How to Fix It,” *The Guardian*, 2018.

Lee, David S and Alexandre Mas, “Long-Run Impacts of Unions on Firms: New Evidence from Financial Markets, 1961–1999,” *The Quarterly Journal of Economics*, 2012, *127* (1), pp. 333–378.

Leibbrandt, Andreas and John A List, “Do Women Avoid Salary Negotiations? Evidence from a Large-Scale Natural Field Experiment,” *Management Science*, 2015, *61* (9), pp. 2016–2024.

Luca, Dara Lee, “The Digital Scarlet Letter: The Effect of Online Criminal Records on Crime,” *Available at SSRN 1939589*, 2018.

Mas, Alexandre, “Does Transparency Lead to Pay Compression?,” *Journal of Political Economy*, 2017, *125* (5), pp. 1683–1721.

Perez-Truglia, Ricardo and Ugo Troiano, “Shaming Tax Delinquents: Theory and Evidence from a Field Experiment in the United States,” *Citeseer*, 2015.

Siniscalco, Gary R, Erin M Connell, and Chad Smith, “State Pay Equity Laws: Where a Few Go, Many May Follow,” Technical Report, Mimeo 2017.

Sran, Gurpal, Felix Vetter, and Matthew Walsh, “Employer Responses to Pay History Inquiry Bans,” *Available at SSRN 3587736*, 2020.

Strauss, Delphine, “Gender Pay Reporting: Does it Make a Difference?,” *Financial Times*, 2019.

Wisniewska, Aleksandra, Billy Ehrenberg-Shannon, and Sarah Gordon, “Gender Pay Gap: How Women are Short-Changed in the UK,” *Financial Times*, 2018.

Wiswall, Matthew and Basit Zafar, “Preference for the Workplace, Investment in Human Capital, and Gender,” *The Quarterly Journal of Economics*, 2018, *133* (1), pp. 457–507.

Figures and Tables

Figure 1: Institutional Setting

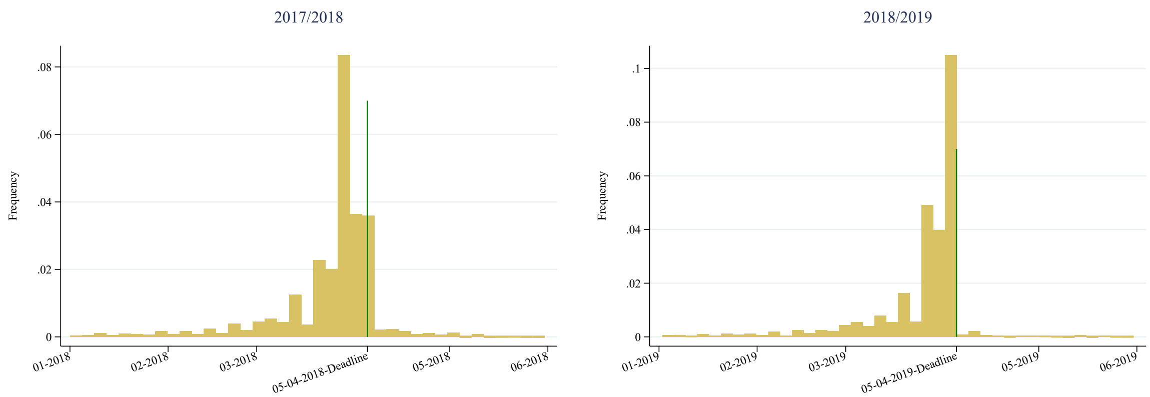

(A) Distribution of submission dates

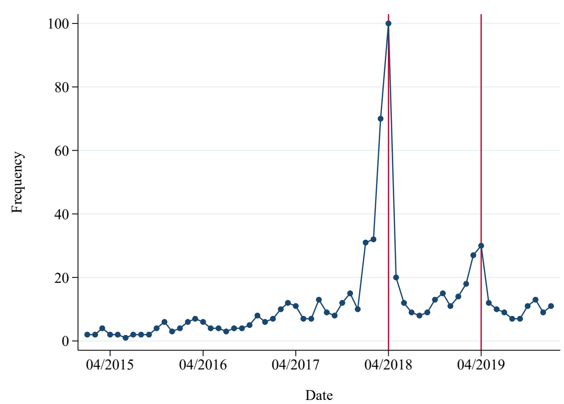

(B) Google searches for “gender pay gap”

Source: UK Government Equalities Office (GEO); Google, 2015-2019.
Notes: The figures in Panel A show the distribution of days when firms published their gender equality indicators. The graph on the left refers to the 2017/18 data (10,557 observations), while the one on the right refers to 2018/19 (10,812 observations). Around 5 percent of firms publish before January of the deadline year. The graph in Panel B reports the UK relative search volume for the term “gender pay gap” between April 2015 and June 2019 using Google’s search services. The frequency is indexed to the peak, which occurred in the week commencing 1st April 2018, when firms faced the first deadline to publish gender equality indicators.

Figure 2: Event studies - log hourly pay

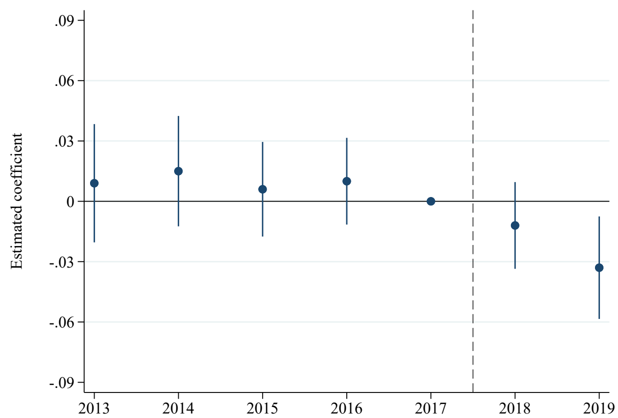

(A) Men

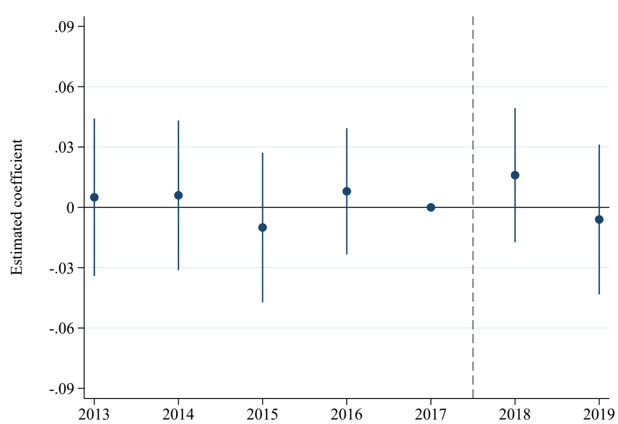

(B) Women

Source: ASHE, 2013–2019.
Notes: These graphs present the estimates of the leads and lags of the policy on log real hourly pay. These results are obtained from the estimation of regression 2. The graph in Panel A refers to men, the one in Panel B to women. In each graph, the estimation sample includes men (women) employed in firms with 200-300 employees. Both regressions are estimated using LFS weights. 95 percent confidence intervals associated with firm-level clustered standard errors are also reported. The dash vertical line indicates the month when the mandate is approved, i.e., February 2017.

Figure 3: Robustness checks

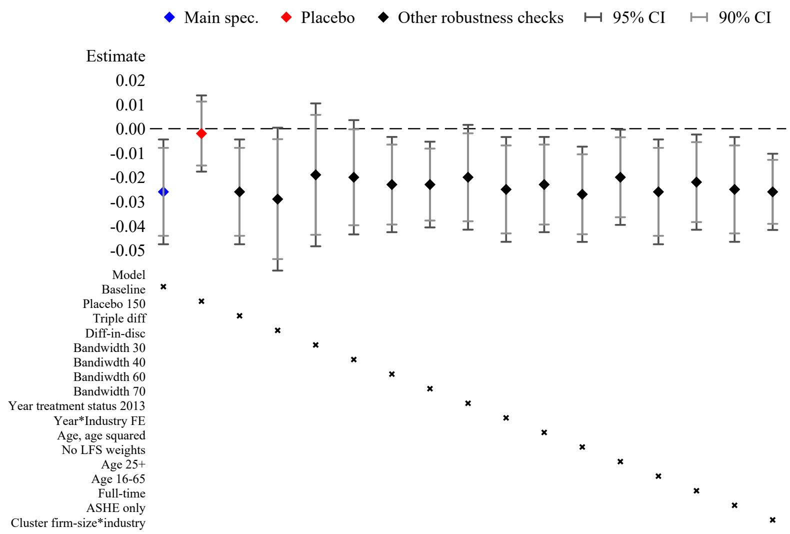

(A) Men

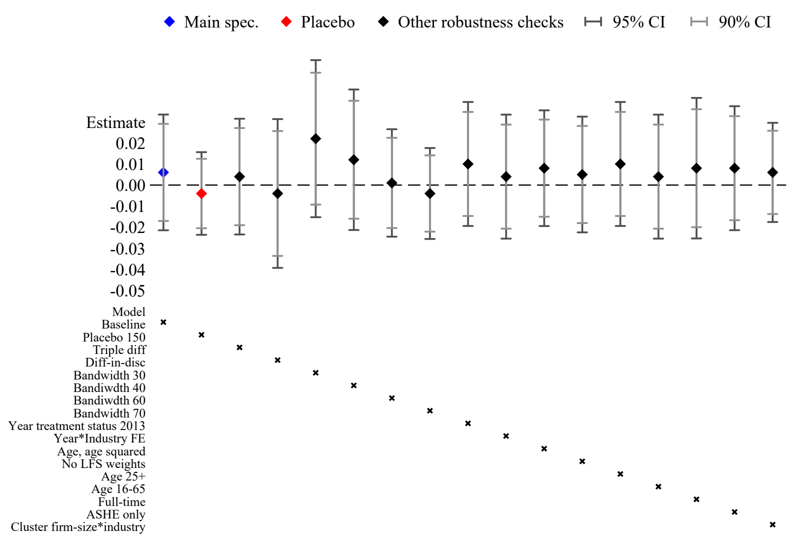

(B) Women

Source: ASHE, 2013–2019.
Note: These graphs present a series of robustness checks on the impact of the policy on employees' pay. The graph in Panel A refers to men, the one in Panel B refers to women. Detailed results are presented in Appendix Tables A1-A5.

Figure 4: Behavioral response - changes compared to baseline - 2019 vs 2018

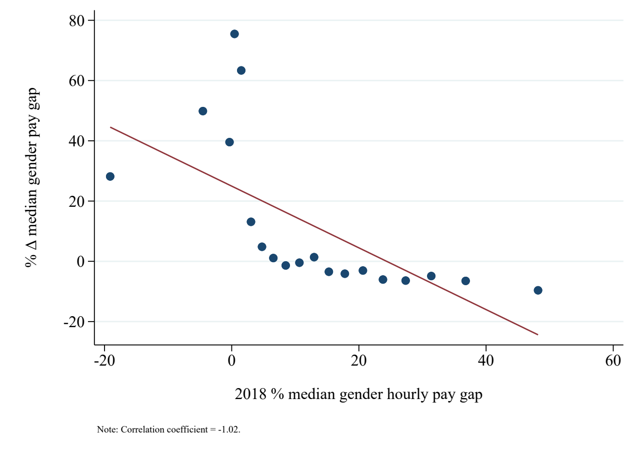

(A) Gender pay gap

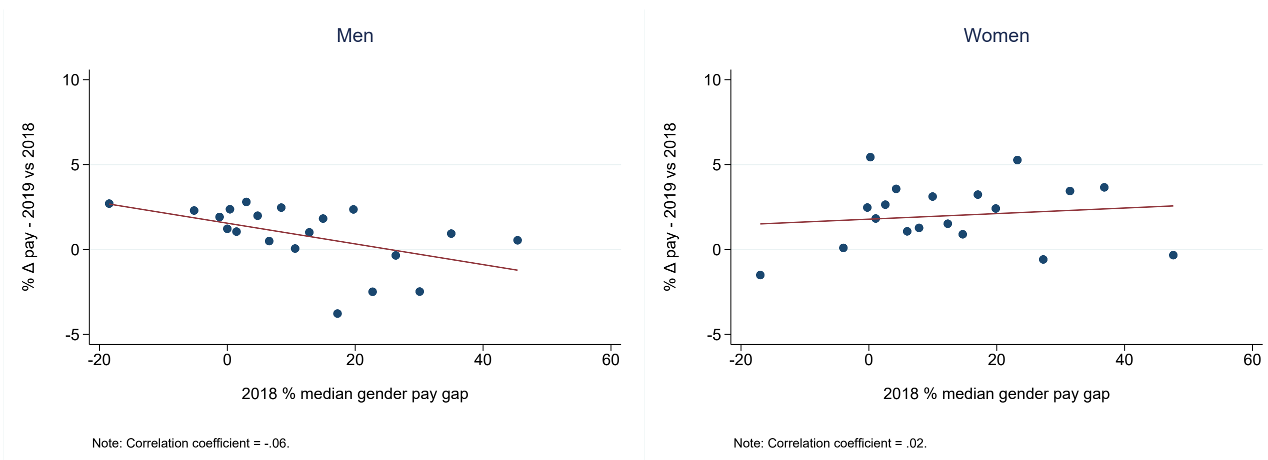

(B) Pay by gender

Source: UK Government Equality Office (GEO), ASHE, 2018–2019.
Notes: These graphs show how changes in firms' gender pay gap (Panel A), men's pay (Left graph, Panel B), and women's pay (Right graph, Panel B) between 2018 and 2019 correlate with firms' gender pay gap in 2018. The sample in Panel A includes all 8,626 firms that publish gender equality indicators in both 2018 and 2019. The graphs in Panel B include the subgroup of 6,748 firms that are also interviewed in ASHE in 2018 and 2019. In all graphs, we exclude outliers (the bottom and top 1 percent) in the distribution of the y variable.

Table 1: Public indicators of gender equality

	2017/18 (1)	2018/19 (2)
Median gender hourly pay gap (%)	11.79 (15.84)	11.88 (15.51)
Mean gender hourly pay gap (%)	14.33 (14.91)	14.19 (14.21)
Median gender bonus gap (%)	-21.72 (1,399.04)	-0.86 (270.51)
Mean gender bonus gap (%)	7.66 (833.06)	15.44 (200.70)
% men receiving bonus	35.39 (36.33)	35.72 (36.68)
% women receiving bonus	33.92 (36.01)	34.40 (36.38)
% women lower quartile	53.67 (24.13)	53.88 (24.11)
% women lower-middle quartile	49.49 (26.09)	49.82 (26.19)
% women upper-middle quartile	45.15 (26.22)	45.62 (26.32)
% women top quartile	39.20 (24.41)	39.75 (24.48)
Observations	10,557	10,812

Source: UK Government Equalities Office (GEO).
Notes: This table reports mean and standard deviation of gender equality indicators published by targeted firms, separately by year of publication.

Table 2: ASHE Summary statistics - pre-policy period

	Treated men (1)	Control men (2)	Treated women (3)	Control women (4)
Hourly pay (£)	17.07	16.70	13.88	13.94
	(15.08)	(12.35)	(9.20)	(10.84)
Weekly pay (£)	623.50	609.45	434.14	431.28
	(563.33)	(454.00)	(320.30)	(330.67)
Weekly hours	36.53	36.67	30.92	30.71
	(8.11)	(8.07)	(10.34)	(10.55)
Receiving bonuses	0.29	0.29	0.20	0.18
	(0.45)	(0.46)	(0.40)	(0.38)
Bonus amount (£)	27.52	27.88	10.54	10.10
	(108.49)	(125.81)	(40.11)	(44.77)
Promotion	0.07	0.06	0.07	0.07
	(0.25)	(0.24)	(0.25)	(0.25)
Top-paid occupations	0.50	0.50	0.46	0.44
	(0.50)	(0.50)	(0.50)	(0.50)
Tenure in months	88.31	87.24	74.23	71.84
	(97.66)	(96.90)	(80.52)	(80.15)
Leaving firm in t+1	0.35	0.34	0.37	0.35
	(0.48)	(0.48)	(0.48)	(0.48)
Private sector	0.90	0.91	0.79	0.78
	(0.29)	(0.28)	(0.40)	(0.42)
Covered by collective agreement	0.28	0.27	0.32	0.34
	(0.45)	(0.44)	(0.47)	(0.47)
Observations	6,911	8,660	5,867	7,682

Source: ASHE, 2013–2017.

Notes: This table reports mean and standard deviation of the main variables used in the analysis, separately for men and women, and treatment and control group, before the implementation of the mandate.

Table 3: Impact on pay outcomes

	Log hourly pay (1)	Log weekly pay (2)	Weekly hours worked (3)	Log hourly basic pay (4)	Bonuses (per hour) (5)
Panel A: Men					
Treated firm*post	-0.026**	-0.020	0.087	-0.025**	-0.028
	(0.011)	(0.012)	(0.220)	(0.011)	(0.031)
Observations	15,899	15,899	15,899	15,899	15,899
Pre-policy Mean	17.07	623.5	36.53	16.34	0.72
Panel B: Women					
Treated firm*post	0.006	-0.012	-0.477	0.009	-0.024
	(0.014)	(0.019)	(0.391)	(0.014)	(0.031)
Observations	13,568	13,568	13,568	13,568	13,568
Pre-policy Mean	13.88	434.14	30.92	13.54	0.35
Year, Firm, Individual FE	✓	✓	✓	✓	✓
Year*Region FE	✓	✓	✓	✓	✓
P-value Men Vs Women	0.064	0.709	0.196	0.043	0.931

Source: ASHE, 2013–2019.

Notes: This table reports the impact of pay transparency on the main outcomes of interest, obtained from the estimation of regression 1. Panel A presents results for men, Panel B for women. Each column refers to a different outcome, as specified at the top of it. The estimation sample comprises men (women) working in firms that have between 200 and 300 employees. All regressions include firm, year, individual fixed effects, and region-specific time shocks. A treated firm is defined as having at least 250 employees in 2015. The post dummy is equal to one from 2018 onward. All regressions are estimated with LFS weights. Heteroskedasticity-robust standard errors clustered at firm level in parentheses. The pre-policy mean represents the mean of the outcome variable for the treated group between 2013 and 2017. The p-value at the bottom of the table refers to the t-test on the equality of coefficients for men and women (reported in Panels A and B).

*** $p<0.01$, ** $p<0.05$, * $p<0.1$.

Table 4: Impact on career outcomes

	Promotion (1)	Top-paid occupations (2)	Tenure in months (3)	Leaving the firm in t+1 (4)
Panel A: Men				
Treated firm*post	-0.005	-0.002	0.122	0.013
	(0.012)	(0.011)	(0.707)	(0.030)
Observations	15,899	15,899	15,637	15,899
Pre-policy Mean	0.07	0.50	88.31	0.35
Panel B: Women				
Treated firm*post	0.009	-0.015	0.025	0.025
	(0.015)	(0.012)	(0.210)	(0.035)
Observations	13,568	13,568	13,278	13,568
Pre-policy Mean	0.07	0.46	74.23	0.37
Year, Firm, Individual FE	✓	✓	✓	✓
Year*Region FE	✓	✓	✓	✓
P-value Men Vs Women	0.471	0.442	0.896	0.792

Source: ASHE, 2013–2019.

Notes: This table reports the impact of pay transparency on career outcomes, obtained from the estimation of regression 1. Panel A presents results for men, Panel B for women. Each column refers to a different outcome, as specified at the top of it. The estimation sample comprises men (women) working in firms that have between 200 and 300 employees. All regressions include firm, year, individual fixed effects, and region-specific time shocks. A treated firm is defined as having at least 250 employees in 2015. The post dummy is equal to one from 2018 onward. All regressions are estimated with LFS weights. Heteroskedasticity-robust standard errors clustered at firm level in parentheses. The pre-policy mean represents the mean of the outcome variable for the treated group between 2013 and 2017. The p-value at the bottom of the table refers to the t-test on the equality of coefficients for men and women (reported in Panels A and B).

*** $p<0.01$, ** $p<0.05$, * $p<0.1$.

Table 5: Gender equality indicators and firms' reputation

	Women's impression score	
	(1)	(2)
Panel A: Two years pooled		
Median gender pay gap	-0.039**	
	(0.019)	
% women at the top		0.035***
		(0.013)
Observations	2,014	2,014
Panel B: 2017/2018		
Median gender pay gap	-0.032	
	(0.025)	
% women at the top		0.032*
		(0.018)
Observations	996	996
Panel C: 2018/2019		
Median gender pay gap	-0.049*	
	(0.029)	
% women at the top		0.038**
		(0.018)
Observations	1,018	1,018

Source: GEO, YouGov, 2018–2019.

Notes: This table shows the raw correlation between firms' gender equality indicators and their score in YouGov's Women's Rankings. Panel A refers to both years, Panel B refers to 2017/2018, while Panel C refers to 2018/2019. In each panel, the sample includes the GEO firms that have been perfectly matched with YouGov entries.

*** $p<0.01$, ** $p<0.05$, * $p<0.1$.

Table 6: Impact on firm-level outcomes

	Log output (1)	Log wage costs (2)	Profits over assets (3)
Treated firm*post	-0.049	-0.040**	0.008
	(0.032)	(0.018)	(0.063)
Observations	24,117	6,149	6,149
Pre-policy Mean	64,611	7,662	3
	(£000's)	(£000's)	(%)
Year, Firm FE	✓	✓	✓
Year*Region FE	✓	✓	✓
Data set	BSD	ABS	ABS

Source: BSD, 2013–2019, ABS 2013–2018.

Notes: This table reports the impact of pay transparency on various firm-level outcomes, obtained from the estimation of regression D.1. Each column refers to a different outcome, as specified at the top of it. The estimation sample comprises firms that have between 200 and 300 employees. All regressions include year, firm and year interacted with region fixed effects. A treated firm is defined as having at least 250 employees in 2015. The post dummy is equal to one from 2018 onward. Heteroskedasticity-robust standard errors clustered at firm level in parentheses. The pre-policy mean represents the mean of the outcome variable for the treated group between 2013 and 2017.

*** $p<0.01$, ** $p<0.05$, * $p<0.1$.

Table 7: Impact on wage posting

	Entire sample (1)	Top-paid Occupations (2)	Other Occupations (3)
Treated Firm*Post	0.039*	0.047**	0.032
	(0.021)	(0.023)	(0.026)
Observations	180,679	94,527	85,409
Pre-policy Mean	0.39	0.33	0.48
P-value difference		0.594	

Source: BGT, FAME, 2013–2019.

Notes: This table reports the impact of pay transparency on firms' wage-posting decision, obtained from the estimation of regression E.1. The estimation sample comprises firms that have between 200 and 300 employees. All regressions include firm and quarter fixed effects, and region-specific time-shocks. The post dummy is equal to one from 2018 onward. A treated firm is defined as having at least 250 employees in 2015. Top-paid occupations include managerial, professional, and technical occupations. Heteroskedasticity-robust standard errors clustered at firm level in parentheses. The pre-policy mean represents the mean of the outcome variable for the treated group between 2013 and 2017. The p-value reported at the bottom of each panel refers to the test of equality of coefficients on top-paid and other occupations.

*** $p<0.01$, ** $p<0.05$, * $p<0.1$.

Appendix

A Further results and robustness checks

Figure A1: Firm size distribution

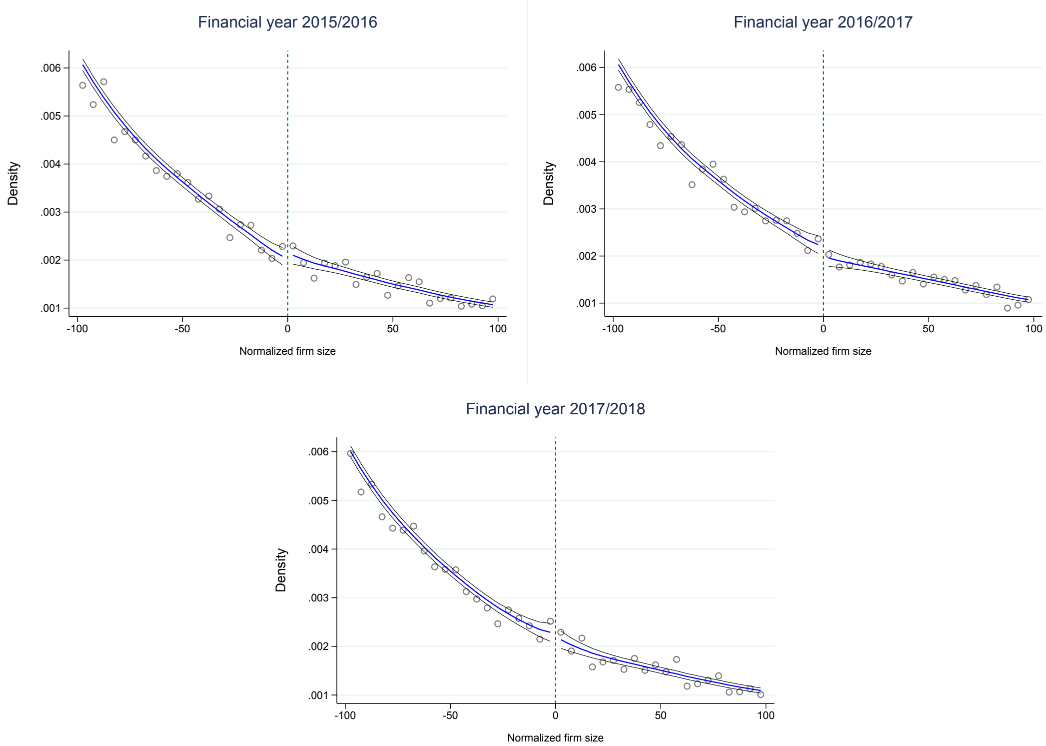

Source: BSD, 2016–2018.
Note: These graphs show the distribution of firms around the 250-employee cutoff in each year since the announcement of the policy. In each figure, the sample includes firms with +/100 employees from the threshold, grouped in 20 bins. Each dot represents the share of firms with a number of employees comprised in the corresponding bin.

Figure A2: Event studies - other pay outcomes

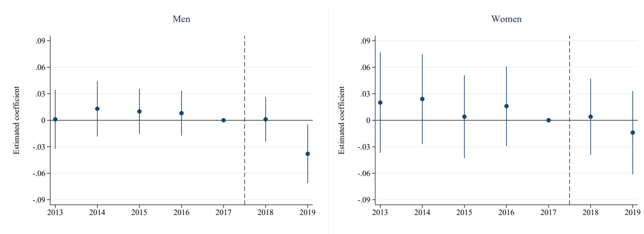

(A) Log weekly pay

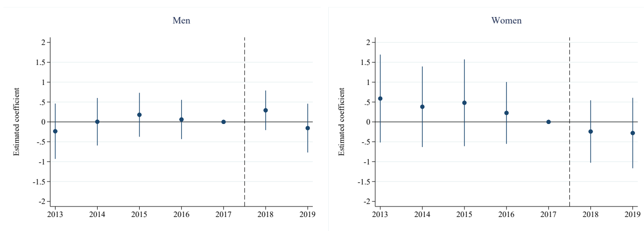

(B) Weekly hours worked

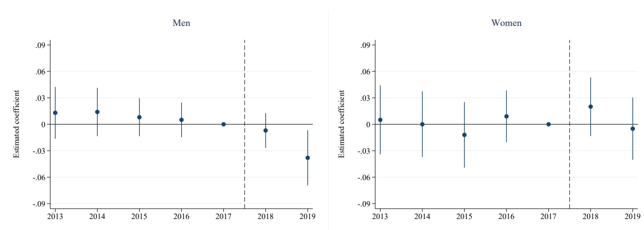

(C) Log hourly basic pay

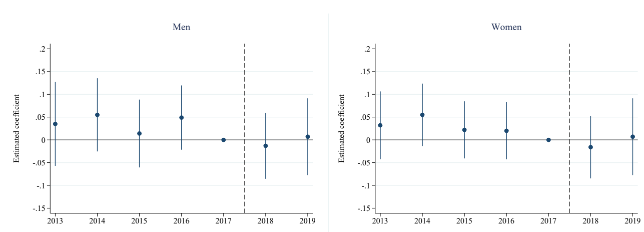

(D) Bonuses

Source: ASHE, 2013–2019.
Notes: These graphs present the estimates of the leads and lags of the policy on different outcomes. These results are obtained from the estimation of regression 2. In each Panel, the graph on the left refers to men, the one on the right to women. In each graph, the estimation sample includes men (women) employed in firms with 200-300 employees. All regressions are estimated using LFS weights. 95 percent confidence intervals associated with firm-level clustered standard errors are also reported. The dash vertical line indicates the month when the mandate is approved, i.e., February 2017.

Figure A3: Event studies - career outcomes

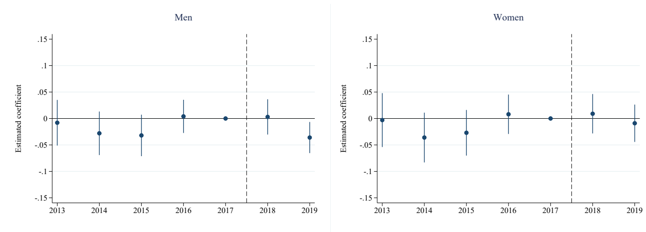

(A) Promotions

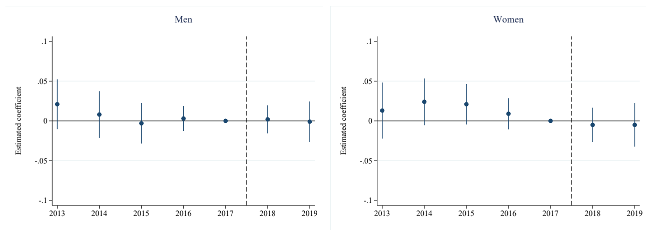

(B) Working in top-paid occupations

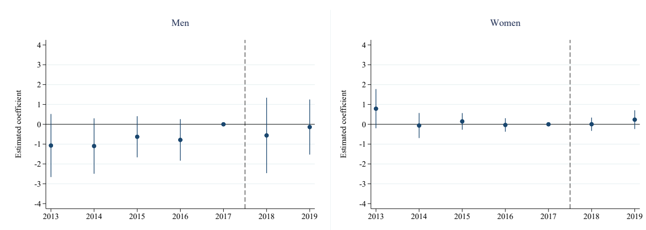

(C) Months of tenure in the firm

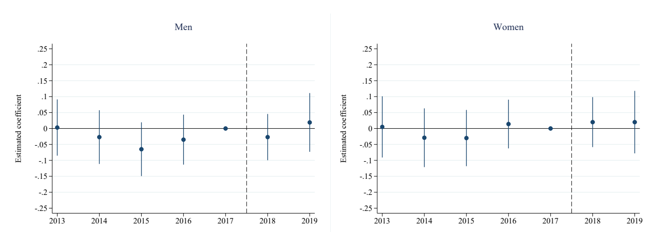

(D) Leaving the firm in t+1

Source: ASHE, 2013–2019.
Notes: These graphs present the estimates of the leads and lags of the policy on career outcomes. These results are obtained from the estimation of regression 2. In each Panel, the graph on the left refers to men, the one on the right to women. In each graph, the estimation sample includes men (women) employed in firms with 200-300 employees. All regressions are estimated using LFS weights. 95 percent confidence intervals associated with firm-level clustered standard errors are also reported. The dash vertical line indicates the month when the mandate is approved, i.e., February 2017.

Figure A4: Hours worked and bonuses - ASHE

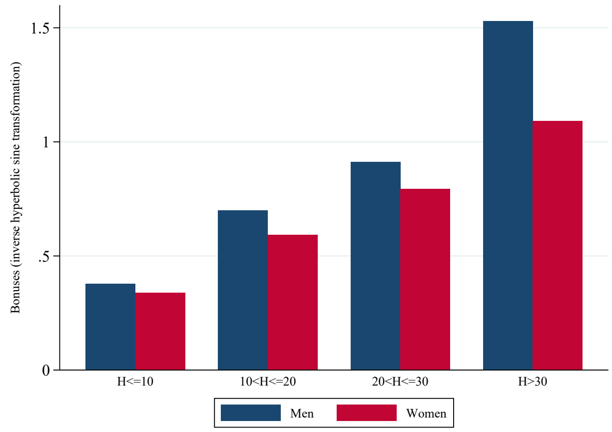

Source: ASHE, 2013–2019.
Notes: This figure shows the correlation between employees' hours worked and bonuses in ASHE, separately by gender. We use the inverse hyperbolic sine transformation of bonuses to take into account that many workers do not receive bonuses.

Figure A5: Raw trends - log hourly pay

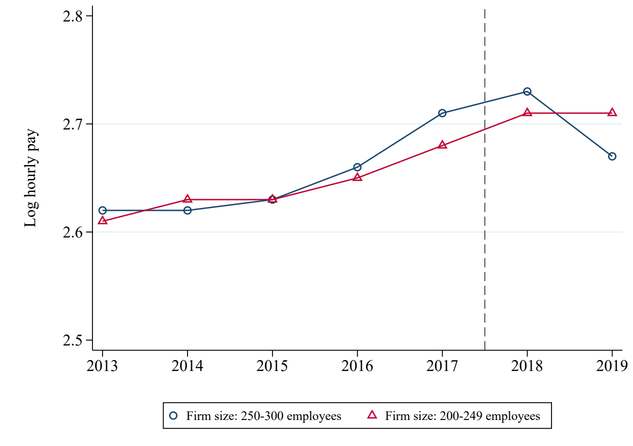

(A) Men

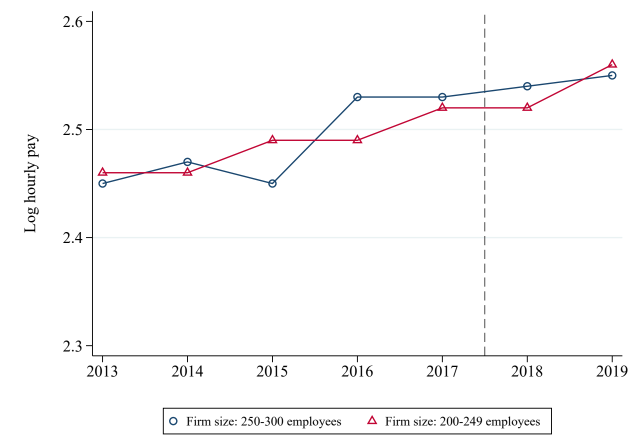

(B) Women

Source: ASHE, 2013–2019.
Notes: This figure presents the trends in log real hourly pay, separately for treated and control groups. The graph in Panel A refers to men, the one in Panel B to women. The blue line represents the treatment group, individuals working in firms with 250-300 employees, and the red line the control group, individuals working in firms with 200-249 employees. The vertical dash line indicates the month when the mandate is approved, i.e., February 2017.

Figure A6: Estimation sample versus entire ASHE

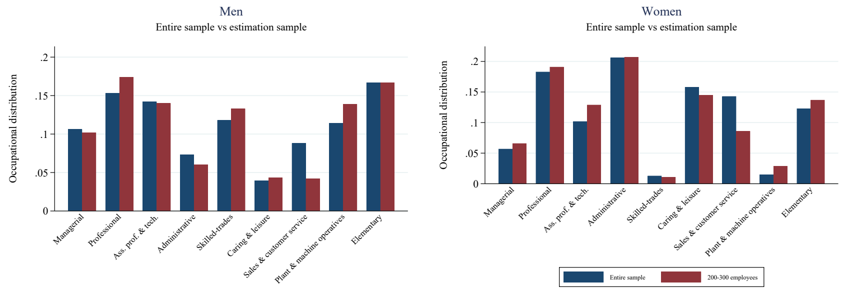

(A) Occupational distribution

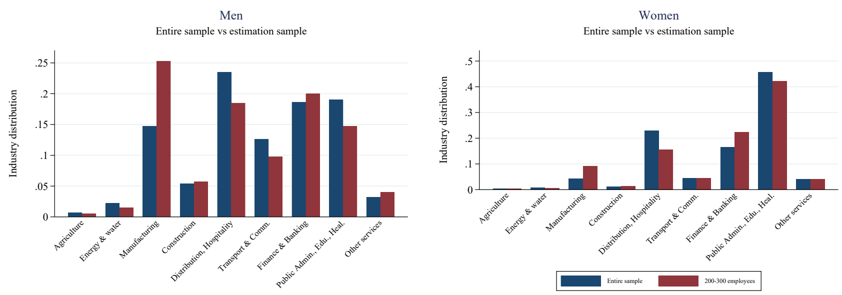

(B) Industry distribution

Source: ASHE, 2013–2019.
Note: These figures compare the occupational and industry distribution of men and women in the estimation sample and in the entire population of ASHE, over the period of analysis.

Table A1: Impact on log hourly pay - placebo regressions

	150	**250**	350	450
	(1)	(2)	(3)	(4)
Panel A: Men				
Treated firm*post	-0.002	-0.026**	-0.013	-0.009
	(0.008)	(0.011)	(0.012)	(0.019)
Observations	28,173	15,899	10,678	7,543
Pre-policy Mean	16.87	17.07	18.43	18.55
Panel B: Women				
Treated firm*post	-0.004	0.006	0.003	0.017
	(0.010)	(0.014)	(0.021)	(0.019)
Observations	24,746	13,568	8,878	6,214
Pre-policy Mean	13.94	13.88	14.40	13.92
P-value Men Vs Women	0.912	0.064	0.646	0.891

Source: ASHE, 2013–2019.
Notes: This table reports the impact of placebo policies on log real hourly pay, obtained from the estimation of regression 1. Panel A presents results for men, Panel B for women. In each regression, the estimation sample comprises men (women) working in firms that have +/- 50 employees from the threshold c specified at the top of each column. All regressions include individual, firm and year fixed effects, and region-specific time shocks. The post dummy is equal to one from 2018 onward. A treated firm is defined as having at least c employees in 2015. Heteroskedasticity-robust standard errors clustered at firm level in parentheses.
*** p<0.01, ** p<0.05, * p<0.1.

Table A2: Diff-in-Diff vs. Triple Diff-in-Diff

	Log real hourly pay		
	Men	Women	Triple Diff
	(1)	(2)	(3)
Treated firm*post	-0.026**	0.006	-0.026**
	(0.011)	(0.014)	(0.011)
Treated firm*post*fem			0.029*
			(0.017)
Post*Female			-0.017
			(0.012)
Treated Firm*Female			-0.398
			(0.269)
Observations	15,899	13,568	29,557
P-value Women			0.781

Source: ASHE, 2013–2019.

Notes: This table compares the results of the difference-in-differences model to those obtained from the estimation of a triple-differences model as specified in equation 3. The estimation sample comprises men (women) working in firms that have between 200 and 300 employees. All regressions control for individual, firm and year fixed effects, and region-specific time shocks. A treated firm is defined as having at least 250 employees in 2015. All regressions are estimated with LFS weights. Heteroskedasticity-robust standard errors clustered at firm level in parentheses. The p-value at the bottom of the table refers to the t-test on the effect for women in the triple-differences model (Treated Firm*Post+Treated Firm*Post*Female).

*** p<0.01, ** p<0.05, * p<0.1.

Table A3: Diff-in-Diff vs. Diff-in-Disc

	Log real hourly pay	
	Diff-in-Diff	Diff-in-Disc
	(1)	(2)
Panel A: Men		
Treated firm*post	-0.026**	-0.029*
	(0.011)	(0.015)
Observations	15,899	15,899
Panel B: Women		
Treated firm*post	0.006	-0.004
	(0.014)	(0.018)
Observations	13,568	13,568
Year FE	✓	
Year*Region FE	✓	
Post		✓
Post*Region FE		✓
Norm. Firm Size*Post		✓
Norm. Firm Size*Treated Firm*Post		✓
P-value Men Vs Women	0.080	0.298

Source: ASHE, 2013–2019.

Notes: This table compares the results of the difference-in-differences model to those obtained from the estimation of a difference-in-discontinuities model as specified in equation 4. Panel A presents results for men, Panel B for women. The estimation sample comprises men (women) working in firms that have between 200 and 300 employees. All regressions include individual and firm fixed effects. A treated firm is defined as having at least 250 employees in 2015. All regressions are estimated with LFS weights. Heteroskedasticity-robust standard errors clustered at firm level in parentheses. The pre-policy mean represents the mean of the outcome variable for the treated group between 2013 and 2017. The p-value at the bottom of the table refers to the t-test on the equality of coefficients for men and women (reported in Panels A and B).

*** p<0.01, ** p<0.05, * p<0.1.

Table A4: Impact on log real hourly pay - different bandwidths

	30	35	40	45	**50**	55	60	65	70
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Panel A: Men									
Treated firm*post	-0.019	-0.020	-0.020	-0.023**	-0.026**	-0.024**	-0.023**	-0.025***	-0.023***
	(0.015)	(0.014)	(0.012)	(0.012)	(0.011)	(0.010)	(0.010)	(0.009)	(0.009)
Observations	8,743	10,402	12,321	13,971	15,899	17,763	19,687	21,590	23,567
Pre-policy Mean	17.22	17.18	17.29	17.21	17.07	16.98	16.90	17.05	17.00
Panel B: Women									
Treated firm*post	0.022	0.007	0.012	0.010	0.006	0.000	0.001	-0.004	-0.004
	(0.019)	(0.018)	(0.017)	(0.016)	(0.014)	(0.014)	(0.013)	(0.012)	(0.011)
Observations	7,191	8,751	10,351	11,867	13,568	15,191	16,869	18,569	20,429
Pre-policy Mean	13.88	13.79	13.89	13.92	13.88	13.93	13.88	13.95	13.96
P-value Men Vs Women	0.075	0.206	0.117	0.079	0.064	0.148	0.119	0.147	0.163

Source: ASHE, 2013–2019.

Notes: This table reports the impact of pay transparency on log real hourly pay, obtained from the estimation of regression 1. Panel A presents results for men, Panel B for women. In each regression, the estimation sample comprises men (women) working in firms that have +/- h employees from the 250 threshold, where h is indicated at the top of each column. All regressions include individual, firm and year fixed effects, and region-specific time shocks. A treated firm is defined as having at least 250 employees in 2015. The post dummy is equal to one from 2018 onward. All regressions are estimated with LFS weights. Heteroskedasticity-robust standard errors clustered at firm level in parentheses.

*** p<0.01, ** p<0.05, * p<0.1.

Table A5: Impact on log real hourly pay - other robustness checks

	Year treatment status 2013	Year* industry FE	Age controls	No LFS weights	Age 25+	Age 16-65	Full-time	ASHE only	Cluster firm size* industry
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Panel A: Men									
Treated firm*post	-0.020*	-0.025**	-0.023**	-0.027***	-0.020**	-0.026**	-0.022**	-0.025**	-0.026***
	(0.011)	(0.011)	(0.010)	(0.010)	(0.010)	(0.011)	(0.010)	(0.011)	(0.008)
Observations	15,727	15,899	15,899	15,899	14,491	15,506	14,442	14,098	15,899
Pre-policy Mean	17.25	17.07	17.07	15.93	18.07	17.08	17.59	17.13	17.07
Panel B: Women									
Treated firm*post	0.010	0.004	0.008	0.005	0.010	0.004	0.008	0.008	0.006
	(0.015)	(0.015)	(0.014)	(0.014)	(0.015)	(0.015)	(0.017)	(0.015)	(0.012)
Observations	13,278	13,568	13,568	13,568	12,396	13,290	8,783	11,952	13,568
Pre-policy Mean	13.94	13.88	13.88	13.36	14.67	13.92	14.56	13.95	13.88
P-value Men Vs Women	0.092	0.093	0.069	0.058	0.075	0.086	0.111	0.066	0.026

Source: ASHE, 2013–2019.

Notes: This table reports a series of robustness checks on the impact of pay transparency on log real hourly pay, obtained from the estimation of regression 1. Panel A presents results for men, Panel B for women. In each regression, the estimation sample comprises men (women) working in firms that have +/- 50 employees from the 250 threshold. All regressions include individual, firm and year fixed effects, and region-specific time shocks. A treated firm is defined as having at least 250 employees in 2015. The post dummy is equal to one from 2018 onward. All regressions, except those in Columns 4, are estimated with LFS weights. Heteroskedasticity-robust standard errors clustered at firm level in parentheses (at firm-size times 1-digit industry code in Column 9: 815 clusters in men's regressions, 757 in women's regressions).

*** $p<0.01$, ** $p<0.05$, * $p<0.1$.

Table A6: Impact on log real hourly pay by occupation

	Entire sample (1)	Top-paid occupations (2)	Other occupations (3)
Panel A: Men			
Treated Firm*Post	-0.026**	-0.029*	-0.018
	(0.011)	(0.017)	(0.013)
Observations	15,899	6,834	8,755
Pre-policy Mean	17.06	23.51	10.72
P-value Low vs High		0.609	
Panel B: Women			
Treated Firm*Post	0.006	0.0144	-0.002
	(0.014)	(0.024)	(0.017)
Observations	13,568	5,311	7,942
Pre-policy Mean	13.88	18.73	9.70
P-value Low vs High		0.573	

Source: ASHE, 2013–2019.

Notes: This table compares the impact of pay transparency on employees' pay in top-paid occupations and other professions. Top-paid occupations include managerial, professional, and technical occupations. Panel A presents results for men, Panel B for women. The estimation sample comprises men (women) working in firms that have between 200 and 300 employees. All estimates are obtained by running regression 1 on the sample indicated on top of each column. All regressions include firm, year, and individual fixed effects. Regressions in Column 1 also include region-specific time shocks. A treated firm is defined as having at least 250 employees in 2015. The post dummy is equal to one from 2018 onward. All regressions are estimated with LFS weights. Heteroskedasticity-robust standard errors clustered at firm level in parentheses. The pre-policy mean represents the mean of the outcome variable for the treated group between 2013 and 2017. The p-value at the bottom of each panel refers to the t-test on the equality of coefficients across the two subgroups.

*** $p<0.01$, ** $p<0.05$, * $p<0.1$.

Table A7: Impact on log real hourly pay by tenure

	Entire sample (1)	Incumbent workers (2)	New hires (3)
Panel A: Men			
Treated Firm*Post	-0.026**	-0.024**	0.0102
	(0.011)	(0.011)	(0.0282)
Observations	15,899	11,629	2,550
Pre-policy Mean	17.06	18.70	13.40
P-value Low vs High		0.252	
Panel B: Women			
Treated Firm*Post	0.006	0.000	0.054
	(0.014)	(0.017)	(0.046)
Observations	13,568	9,627	2,368
Pre-policy Mean	13.88	15.01	11.74
P-value Low vs High		0.260	

Source: ASHE, 2013–2019.

Notes: This table compares the impact of pay transparency on pay of newly-hired employees (at most two years of tenure) versus incumbents. Panel A presents results for men, Panel B for women. The estimation sample comprises men (women) working in firms that have between 200 and 300 employees. All estimates are obtained by running regression 1 on the sample indicated on top of each column. All regressions include firm, year, and individual fixed effects. Regressions in Column 1 also include region-specific time shocks. A treated firm is defined as having at least 250 employees in 2015. The post dummy is equal to one from 2018 onward. All regressions are estimated with LFS weights. Heteroskedasticity-robust standard errors clustered at firm level in parentheses. The pre-policy mean represents the mean of the outcome variable for the treated group between 2013 and 2017. The p-value at the bottom of each panel refers to the t-test on the equality of coefficients across the two subgroups.

*** $p<0.01$, ** $p<0.05$, * $p<0.1$.

B YouGov data

B.1 Name matching algorithm

We used a combination of techniques to merge two different data sets A and B through name matching. We first collapsed all firm names in each data set down to a unique set of firm names using standard text cleaning procedures. We identified any exact matches between firm names in data sets A and B, giving these a match score of unity. We matched the remaining N firm names between the two data sets, with M unique entries, using a combination of techniques provided in the scitkit-learn software package. First, firm names in data set A are expressed as character-level 2- and 3-grams with a maximum of 8,000 features, creating a matrix T with dimensions (number of entries in A) X (number of features). The 8,000 features define a vector space that we used to express the official firm names into, with a matrix G. Matching directly with these matrices would require NXM inner products of 8,000 dimensional vectors. Instead, we created a reduced vector space of just 10 dimensions using truncated singular value decomposition on T, creating a reduced dimension matrix $\hat{T}$ and expressing G as $\hat{G}$ in the reduced space. The vectors representing $\hat{G}$ and $\hat{T}$ were then sorted into 500 clusters using k-means, providing an associated cluster for each firm name on both sides of the matching problem. For each cluster c_i with $i \in \{1, 500\}$ the problem was reduced to finding matches between $c_i(N) \leq N$ and $c_i(M) \leq M$ entries - where the equality holds for at most one of the clusters respectively (and rarely holds in practice). Within each cluster, we computed all of the pair-wise cosine similarities between $c_i(T)$ and $c_i(G)$; i.e., within a cluster, and with features indexed by f, the matches for T are found by solving

$$\arg\max_m \{T_{nf} \cdot G_{fm}\}$$

The score is the cosine similarity of the matched vectors scaled by 0.99 (to distinguish exact matches from exact-in-the-vector-space matches).

Table B1: Gender equality performance and presence in YouGov

	Entire sample	Merged with YouGov No	Merged with YouGov Yes	P-value difference	Entire sample	Merged with YouGov No	Merged with YouGov Yes	P-value difference
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Median gender hourly pay gap	11.79	11.82	11.46	0.49	11.88	11.92	11.50	0.42
	(15.84)	(15.80)	(16.23)		(15.51)	(15.65)	(14.06)	
% women top quartile	39.20	39.58	35.56	0.00	39.75	40.11	36.28	0.00
	(24.41)	(24.57)	(22.45)		(24.48)	(24.67)	(22.38)	
Observations	10,557	9,561	996		10,812	9,794	1,018	

Source: GEO, FAME, 2018–2019.

Notes: This table explores potential selection patterns of GEO firms that we merge with YouGov. Columns 1 and 5 report the main gender equality indicators for all GEO firms in 2017/2018 and 2018/2019, Column 2 and 6 refer to firms that we do not find in YouGov, Columns 3 and 7 refer to firms that we merge with in YouGov, and Columns 4 and 8 report the p-value of the difference in the sample means of these two groups in each year.

*** $p<0.01$, ** $p<0.05$, * $p<0.1$.

Table B2: Gender equality indicators and firms' reputation in the workforce

	Workforce's reputation score	
	(1)	(2)
Panel A: Two years pooled		
Median gender pay gap	-0.006	
	(0.014)	
% women at the top		0.011
		(0.009)
Observations	2,018	2,018
Panel B: 2017/2018		
Median gender pay gap	-0.004	
	(0.018)	
% women at the top		0.011
		(0.013)
Observations	998	998
Panel C: 2018/2019		
Median gender pay gap	-0.009	
	(0.021)	
% women at the top		0.011
		(0.013)
Observations	1,020	1,020

Source: GEO, YouGov, 2018–2019.

Notes: This table shows the raw correlation between firms' gender equality indicators and their score in YouGov's Workforce Rankings. Panel A refers to both years, Panel B refers to 2017/2018, while Panel C refers to 2018/2019. In each panel, the sample includes the GEO firms that have been perfectly matched with YouGov entries.

*** p<0.01, ** p<0.05, * p<0.1.

C Stock Market Response

Section 6 supports the hypothesis that the public disclosure of firms' gender equality indicators may induce businesses to tackle gender pay differentials to preserve their reputation among potential customers and employees. But firms may also be concerned about what investors think. A negative reaction of the stock market to the publication of the gender equality indicators may constitute a strong incentive for a firm to improve its performance on gender equality. Importantly, a priori it is not clear how the stock market may react. On the one hand, investors could punish firms, and especially those with a high gender pay gap, assuming that these will have to increase women's wages, with a resulting increase in the wage bill and lower profits. Also, the stock market may fear potential negative repercussions of the policy on labor productivity. On the other hand, investors may reward firms that publish gender equality indicators, if this may stimulate improvements in management practices, with positive knock-on effects on workers' productivity and firms' profits.

To investigate these dynamics, we adopt the traditional event-study methodology (Lee and Mas 2012, Bell and Machin 2018). In particular, we focus on the first year of publication as this is when gender equality indicators are more likely to represent an information shock for the market. We first combine the list of firms publishing gender equality indicators in the financial year 2017/18 with FAME to identify both firms that are directly publicly listed on the London Stock Exchange (LSE), and those that have a parent company that is publicly listed. This leads us to identify 926 firms, or around 10 percent of firms publishing gender equality indicators. Of this group, 101 are directly publicly listed, while the others have a publicly listed parent company. Importantly, different firms can have the same parent company. As a result, we follow 405 distinct publicly listed firms, or 35 percent of all firms listed on the main market of the London Stock Exchange in 2018. Also note that 80 percent of firms belonging to the same group publish gender equality indicators on the same date. Hence, in what follows, we consider the publication date of the first one to publish. Extracting daily stock prices from Datastream, we then construct firms' abnormal returns, or AR, as the difference between a stock's actual return and the expected return, where this is estimated using a simple market model over the previous year of data:[A.1]

$$AR_{jt} = r_{jt} - (\widehat{\alpha}_j + \widehat{\beta}_j r_{mt}), \tag{C.1}$$

where r_{jt} is firm j stock market return on day t, and r_{mt} is the return of the LSE-all-shares index on day t.[A.2] As it is standard when employing this methodology (Lee and Mas 2012, Bell and Machin 2018), we then look at the evolution of 3-day cumulative abnormal returns, or $CARS(-1, 1) = \sum_{k=-1}^{+1} AR_{jk}$. This allows us to take into account both potential leaks of information in the day prior to the event of interest, as well as lagged responses in the day following this event. Figure C1 plots the $CARS(-1, 1)$, in the five days before and after the publication date. While these are not statistically different from zero in the days prior to the publication date, they start to become negative from the publication date up to four days afterwards, with an average loss per day of around 35 basis points.[A.3] Table C1 further investigates how this drop relates to the performance on

[A.1] The 15 days before t are excluded from the estimation of predicted returns to avoid capturing any potential anticipation effect of events happening at time t.

[A.2] Alternatively, one could use a CAPM model or Four-Factor model to predict firms' returns, but this is beyond the scope of this paper.

[A.3] As a comparison, note that Bell and Machin (2018) find that the sudden increase in the minimum wage, an-the gender equality indicators. Column one regresses the 3-day CARs on the day of the publication on a constant, the average median gender pay gap reported by firms related to the same publicly listed firm, called “Group-avg GPG” in the table, a dummy equal to one if the gender pay gap is in favor of men, called “Group-avg GPG negative”, and an interaction term between these two. Column 2 adds the following controls: a categorical variable for whether the listed firm directly publishes the gender equality indicators, or has a subsidiary publishing them; and the number of firms in the group publishing the gender equality indicators. Column 3 adds industry fixed effects, and column 4 also controls for the log of market capitalization at t-1, the book-to-market value at t-1 and the return on assets at t-1. While it does not seem that firms publishing a gender pay gap in favor of men are penalized more than others, the main message of this analysis is that firms publishing gender equality indicators are under the scrutiny of investors. In turn, this supports the hypothesis that the reputation motive may have played an important role in explaining the reaction of treated firms.

Figure C1: 3-day cumulative abnormal returns around 2017/18 publication date

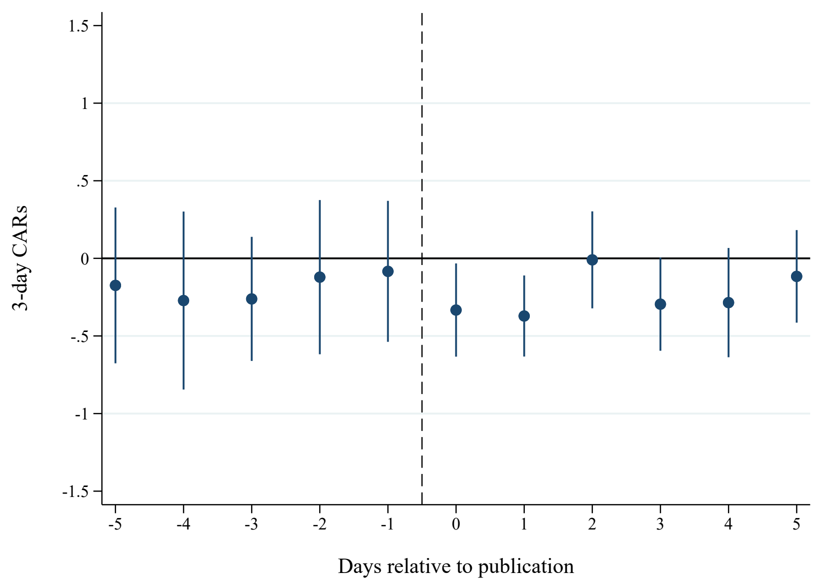

Source: Datastream, FAME, GEO, 2017–2018.
Note: This figure plots 3-day cumulative abnormal returns around the publication date of gender equality indicators in 2017/2018. In particular, it shows CARS(-1, 1) around the day reported on the graph. 95 percent confidence intervals associated with standard errors clustered at the level of publication date are also displayed. The sample includes firms that had to publish gender equality indicators by April 5th 2018, or that have a subsidiary that had to publish these data.

nounced by the UK government in May 2015, leads to a 70 basis points immediate decrease in abnormal returns of low-wage firms.

Table C1: CARs(-1,1) around publication

	(1)	(2)	(3)	(4)
Group-avg GPG performance negative	0.618	0.669	0.651	0.392
	(0.868)	(0.932)	(0.956)	(1.020)
Group-avg GPG performance	-0.027	-0.027	-0.028	-0.043
	(0.056)	(0.055)	(0.056)	(0.058)
Group-avg perf.*group-avg perf. negative	0.037	0.036	0.036	0.048
	(0.056)	(0.056)	(0.057)	(0.059)
Constant	-1.057	2.163**	2.823***	0.560
	(0.742)	(0.893)	(0.952)	(1.552)
Observations	405	405	405	383
Ownership structure		✓	✓	✓
N. Firms in the group		✓	✓	✓
Industry FE			✓	✓
Other controls				✓

Source: Datastream, FAME, GEO, 2017-2018.
Notes: This table shows the estimates of the cumulative abnormal returns around the publication of gender equality indicators. The dependent variable is the sum of abnormal returns in the 3-day window around the publication date. The sample includes firms which had to publish gender equality indicators by April 5th 2018, or which have a subsidiary that had to publish these data. From Column 2 onward we include a variable measuring the number of firms in the group publishing the gender equality indicators, and dummies for whether it is the listed firm or the immediate, domestic or global owner of a firm that has to publish the gender equality indicators. Other controls in Column 4 include the lagged values of log of market capitalization, price to book value ratio, and the return on assets.
*** p<0.01, ** p<0.05, * p<0.1.

D Firm-level outcomes

To study the effect of the pay transparency policy on firm-level outcomes, we estimate the difference-in-differences model at the firm level:

$$Y_{jt} = \alpha_j + \theta_t + \beta(TreatedFirm_j * Post_t) + Z'_{jt}\delta + u_{jt}, \tag{D.1}$$

where Y_{jt} is either labor productivity, labor costs, or profits of firm j in year t; α_j and θ_t are firm and year fixed effects respectively, and Z_{jt} includes region-specific time shocks. Standard errors are clustered at the firm level.

Figure D1: Event studies - firm-level outcomes

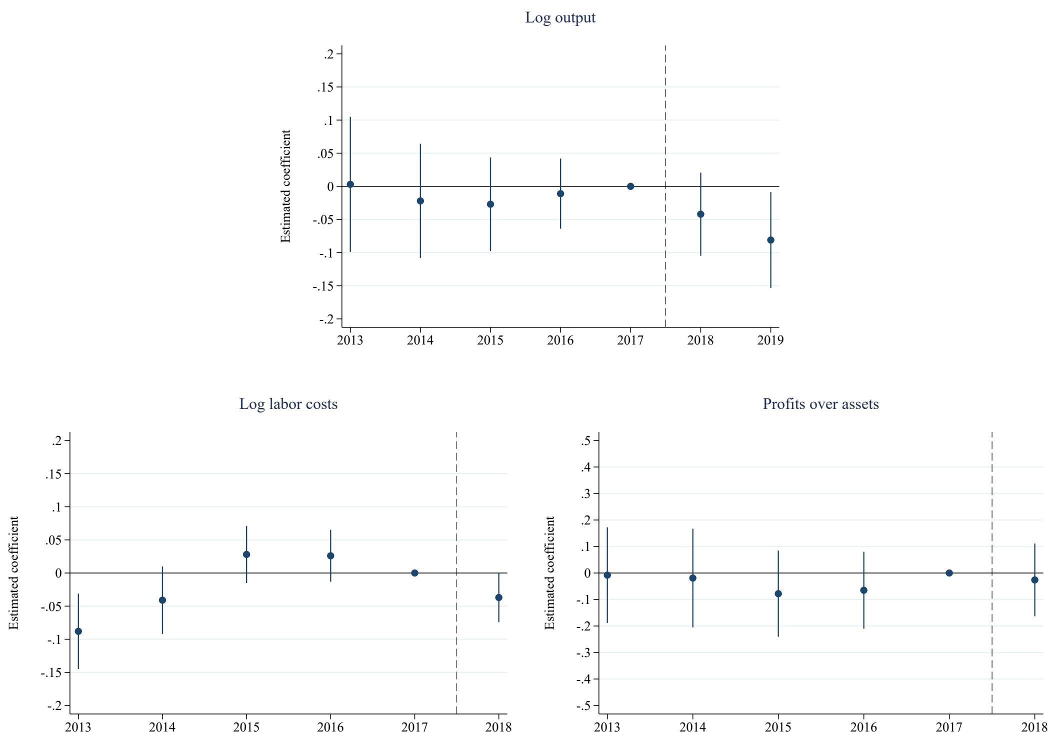

Source: ASHE, 2013–2019.
Notes: These graphs present the estimates of the leads and lags of the policy on firm-level outcomes. These results are obtained from the estimation of a dynamic version of regression D.1. The estimation sample includes firms with 200-300 employees, and present in BSD (ABS) between the financial years 2013 and 2019 (2018). 95 percent confidence intervals associated with firm-level clustered standard errors are also reported. The dash vertical line indicates the month when the mandate is approved, i.e., February 2017.

E Burning Glass Technologies data

Burning Glass Technologies collects data on UK online job vacancies since the fiscal year 2013. Over the fiscal years 2013–2019 these comprise around 50 million (de-duplicated) individual job vacancies, collected from a wide range of online job listing sites. While the data set only includes online advertisements, and hence misses vacancies not posted online (e.g. those advertized informally and internal vacancies), it includes a rich set of information that is especially useful for our analysis. First, each observation includes the text of the job advertisement. Second, more than 95 percent of vacancies have an occupational SOC identifier and 90 percent a county identifier. Finally, around one third of vacancies, or 17 million observations, include the name of the employer. As this is the only variable that can facilitate the merging of BGT data with other firm-level data, we focus on a restricted sample with non-missing employer names. To exclude potential selection issues related to the presence of the firm name, in Figure E1 we show that the industry distribution of the stock of vacancies in BGT and that of vacancies in the ONS Vacancy Survey match well, mitigating concerns regarding the representativity of BGT.

E.1 Difference-in-differences sample

To implement the difference-in-differences strategy, we need to identify a treated and a control group in BGT, and for this we need information on firms' number of employees. We now describe in detail how we create this sample, putting particular care in explaining why we need to exclude some firms at each stage, and what impact this could have in terms of sample selection.

FAME. We start from FAME, the UK version of Amadeus, covering all UK-registered firms. For around 30 percent of them we have information on the number of employees for at least one year in the pre-treatment period. To address selectivity concerns at this stage, in Figure E2 we compare the industry distribution for firms with and without information on the number of employees in 2015, the year used to define the treatment status. Reassuringly, while firms with missing information on employees' numbers also tend to have missing information on industry, the rest of the distribution appears similar.

FAME and GEO. Next, we merge FAME with GEO firms using the company registration number. At this stage, 25 percent of GEO firms do not merge with FAME, either because they do not have a registration number or because they are not present in FAME. Most of these instances are schools or other public sector firms. Table E1 compares the main gender equality indicators of firms that merge or not with FAME. While the median gender pay gap is not statistically different across the two groups in either year, the group of firms that do not merge with FAME tend to have a much higher share of women among top-paid employees. We then retain all FAME firms with a number of employees between 200 and 300 in at least one year, with a resulting sample of 11,250 firms. Moreover, we retain all GEO firms found in FAME independently of their number of employees, in order to explore correlations between firms' hiring practices and gender equality indicators.

FAME, GEO, and BGT. We then merge this sample of firms with BGT using the name-matching algorithm described above. We find 8,607 FAME firms in BGT, and among these we only retain those with a match score of at least 0.85. Table E2 shows that these represent the majority of firms and on average tend to have a lower number of employees in 2015 than those firms matching with BGT with a lower score. By keeping only vacancies with non-missing occupation codes and county information, our final sample includes 180,679 vacancies of 3,450 firms, over the fiscal years 2013–2019.

Figure E1: Industry distribution in BGT and ONS Vacancy Survey

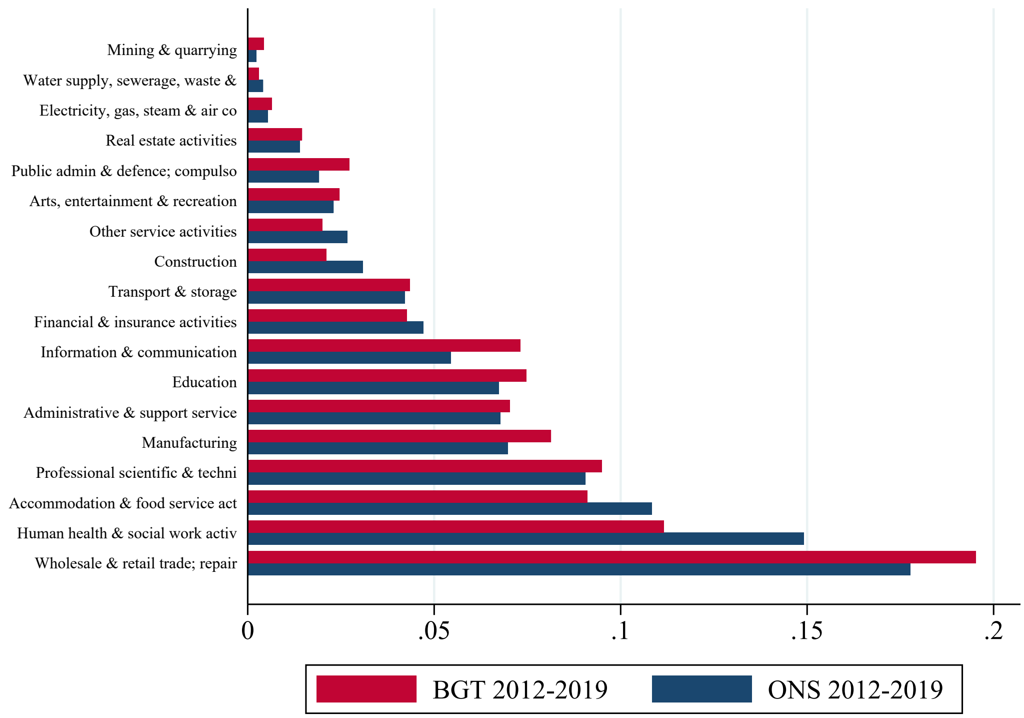

Source: BGT, ONS Vacancy Survey, 2013-2019.
Note: This figure compares the industry distribution in the stock of BGT vacancies with non-missing employer name and in the ONS Vacancy Survey.

Figure E2: Representativity of FAME sample

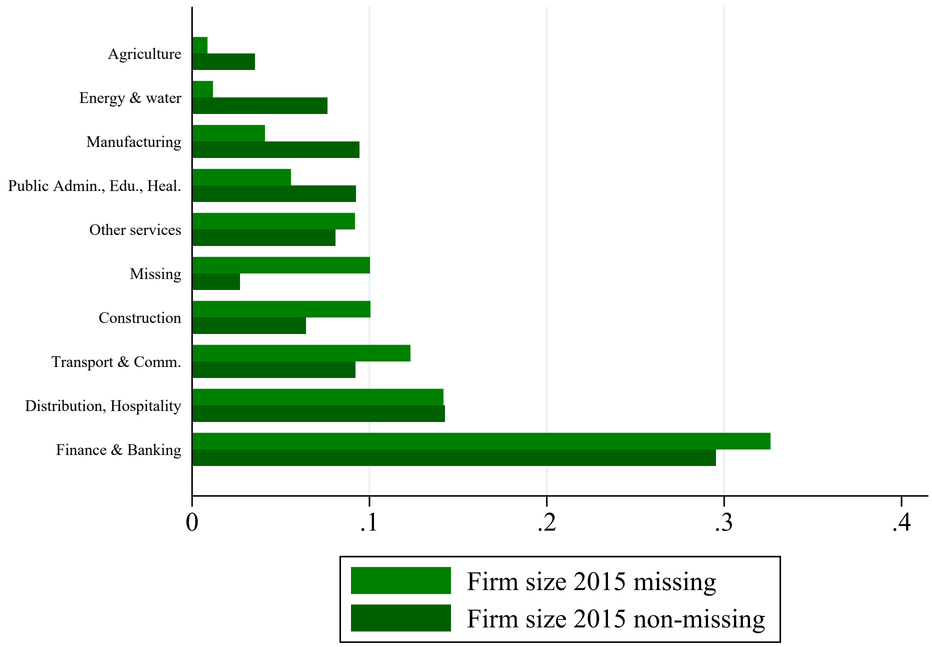

Source: FAME, 2015.
Note: This figure compares the industry distribution of FAME firms with missing and non-missing size information in 2015, the year used to to define firms' treatment status.

Table E1: Gender equality performance and presence in FAME

	Entire sample	Merged with FAME No	Merged with FAME Yes	P-value difference	Entire sample	Merged with FAME No	Merged with FAME Yes	P-value difference
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Median gender hourly pay gap	11.79	11.61	11.81	0.65	11.88	11.24	11.97	0.10
	(15.84)	(13.36)	(16.17)		(15.51)	(13.19)	(15.81)	
% women top quartile	39.20	52.77	37.17	0.00	39.75	53.51	37.80	0.00
	(24.41)	(16.68)	(24.73)		(24.48)	(16.13)	(24.84)	
Observations	10,557	1,373	9,184		10,812	1,346	9,466	

Source: GEO 2018–2019, FAME, 2013–2019.

Notes: This table explores potential selection patterns of GEO firms that we merge with FAME. Columns 1 and 5 report the main gender equality indicators for all GEO firms in 2017/2018 and 2018/2019, Column 2 and 6 refer to firms that we do find in FAME, Columns 3 and 7 refer to firms that we merge with FAME, and Columns 4 and 8 report the p-value of the difference in the sample means of these two groups in each year.

*** p<0.01, ** p<0.05, * p<0.1.

Table E2: Firm size and presence in BGT

	Entire sample (1)	Match score < 0.85 (2)	≥ 0.85 (3)	P-value Difference
Number of employees	233.19	244.38	228.19	0.00
	(235.13)	(350.63)	(158.10)	
Observations	8,607	2,658	5,949	

Source: BGT, FAME, 2013–2019.
Notes: This table explores potential selection patterns of FAME firms that we merge with BGT. Column 1 reports the 2015 number of employees of all FAME firms found in BGT, Columns 2 and 3 reports this number for firms with a match score that is, respectively, lower or greater or equal than 0.85. Finally, Column 4 reports the p-value of the difference in the sample means of these two groups.

E.2 BGT outcomes

Wages. To study firms' wage-posting decision, we extract wages offered from the job ad text using natural language processing. To identify wages in the text, we use a series of targeted regular expressions, such as "30-35k per annum", or "20,000/year". The frequency of the wage offer (annual, weekly, hourly) is similarly inferred from the text, and all values are transformed into annual wages and deflated using the ONS 2015 CPI Index. When a vacancy posts a wage interval, we consider the mid-point of the interval. Using this procedure, we find that less than 40 percent of BGT job listings contain information on wages that can be automatically identified, leaving room for firms to respond to the policy along this margin. As for the value of posted wages, their mean is £27,175.77 in the pre-policy period (2013-2017), while their median is £22,107.5. A series of validation exercises from research assistants indicate that we correctly identify the value of posted wages in 97 percent of cases. The remaining 3 percent are either false negatives, meaning that our code does not detect a wage when the wage is posted, or a wrong value for the posted wage. To further validate our procedure, in Figure E3, we compare the distribution of log annual full-time wages in BGT with those of newly hired workers from the LFS. Note that posted wages do not necessarily need to match with realized wages, as there is a hiring process in between the two, where some vacancies remain unfilled and wages are furhter renegotiated. Yet, while the distribution of LFS wages presents a fatter left tail, reassuringly, the two distributions are similar.

Flexible working arrangements. Gender differences in preferences for temporal flexibility have been shown to play a key role in explaining gender segregation across occupations, which in turn contributes to the persistence of the gender pay gap (Bertrand et al. 2010, Goldin 2014, Wiswall and Zafar 2018, Cortes and Pan 2019). In order to attract more women, firms may have responded to the mandate by expanding the offer of flexible working arrangements (FWA hereafter). Importantly, this can have ambiguous effects on gender pay differentials. On the one hand, offering FWA may help reduce gender occupational segregation and the gender pay gap. On the other hand, wherever flexibility entails a wage penalty because it generates a productivity cost for the firm, the offer of FWA may increase gender pay differentials (Goldin 2014). To investigate this dimension of response, we constructed a vocabulary of flexible work terms using job listings from Timewise, a website specialized in flexible working, the LFS definition of FWA, and that provided by ACAS, the Advisory, Conciliation and Arbitration Service (ACAS 2015). Table E3 presents the full list of expressions included. Note that we do not consider FWA that give the employer discretion over scheduling, such as shift work or on-call work (Adams et al. 2020), but only those arrangements that can give the employee more control over their work-life balance. Also, precisely to investigate the importance of this trade-off posed by FWA, we consider "part-time" and "job sharing" separately from full-time FWA, under the hypothesis that the latter may be less costly for the firm in terms of productivity. Based on this list, we created two categorical variables equal to 1 if a job vacancy mentions, respectively, the possibility of working part-time, or one of the full-time FWA listed in Table E3. In the pre-policy period, 14 percent of vacancies mentioned a part-time FWA, and 3 percent mentioned at least one full-time FWA. A series of validation exercises from research assistants indicate that we correctly identify full-time FWA in 95 percent of cases, with 3 percent of false positives, and 2 percent of false negatives. As for part-time, we correctly identify this working option 99 percent of times, with the remaining 1 percent being false positives.

Figure E3: Log annual FT wages - BGT vs. LFS

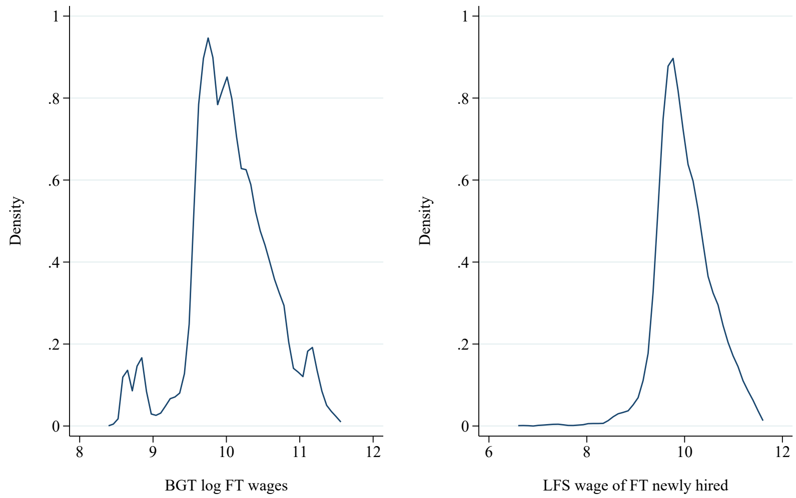

Source: BGT, LFS 2013–2019.
Note: This figure compares the distribution of log posted annual FT wages extracted from BGT, with that of annual wages of newly hired workers (at most 3 months of tenure) computed from the LFS. Both measures are expressed in logs and in real terms.

Table E3: Flexible work arrangements

Flexible work (1)	Remote work (2)	Part-time work (3)
ability to split hours	conducted remotely	job share
annualised hours	home-based	job sharing
compressed hours	home work	parttime
flexible ad hoc hours	home working	part time
flexible days and hours	location remote	
flexible working	position is remote	
flexible start and finish	telecommuting	
flexible start time	telework	
flexible finish time	teleworking	
flexible on days and hours	remote work	
flexible schedule	remote working	
flexible scheduling	virtual job	
flexible staggered hours	work from home	
flextime		
four-and-half-day working		
nine-day fortnight		
returner		

Source: ACAS, LFS, Timewise. and manual inspection.

Notes: This table presents the expressions used to measure the offer of flexible work arrangements in BGT job vacancies.

E.3 BGT results

To investigate the effect of the pay transparency policy on firms' hiring practices, we estimate the difference-in-differences model at the vacancy level:

$$Y_{ijt} = \alpha_j + \theta_t + \beta(TreatedFirm_j * Post_t) + Z'_{jt}\delta + u_{ijt}, \tag{E.1}$$

where Y_{ijt} is either a dummy equal to one if vacancy i of firm j in quarter t offers wage information, full-time FWA or part-time options. Alternatively, it represents the log value of posted wages; α_j and θ_t are firm and quarter fixed effects respectively, and Z_{jt} includes region-specific time shocks. Finally, standard errors are clustered at the firm level.

Figure E4: Wage posting and equality indicators - conditional correlations

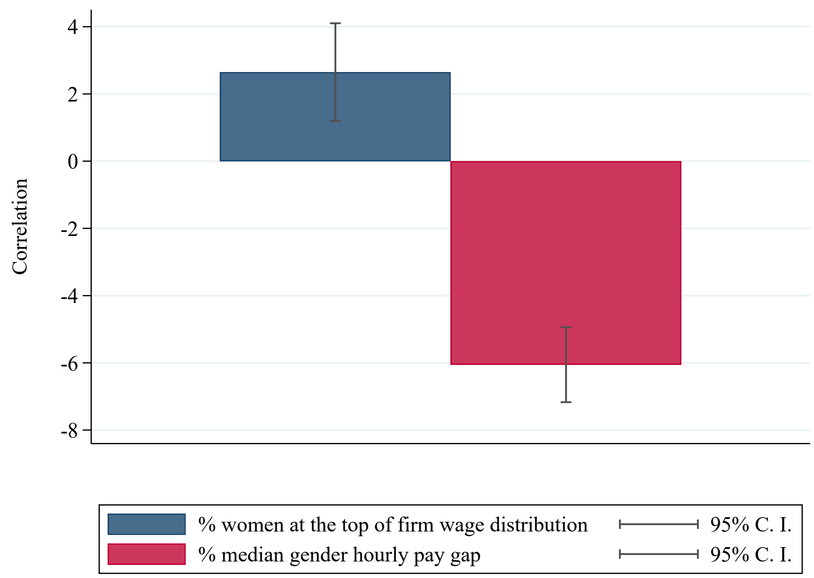

Source: BGT 2013–2019. GEO 2018–2019.
Note: The bar graph reports estimated coefficients from regressions of gender equality indicators (averaged across 2017/18 and 2018/19) on the average percentage of vacancies posting wage information over the period 2013–2019, the occupational composition of firms' vacancies and their sector. The graph also displays 95 percent confidence intervals associated with heteroskedasticity-robust standard errors. The sample includes firms publishing gender equality indicators both in 2017/18 and in 2018/19, with non-missing registration numbers, and matched with BGT with a match score of at least 0.85 (See Appendix B.1 for a description of the name-matching procedure). N. observations = 5,482.

Figure E5: Event study - wage posting decision

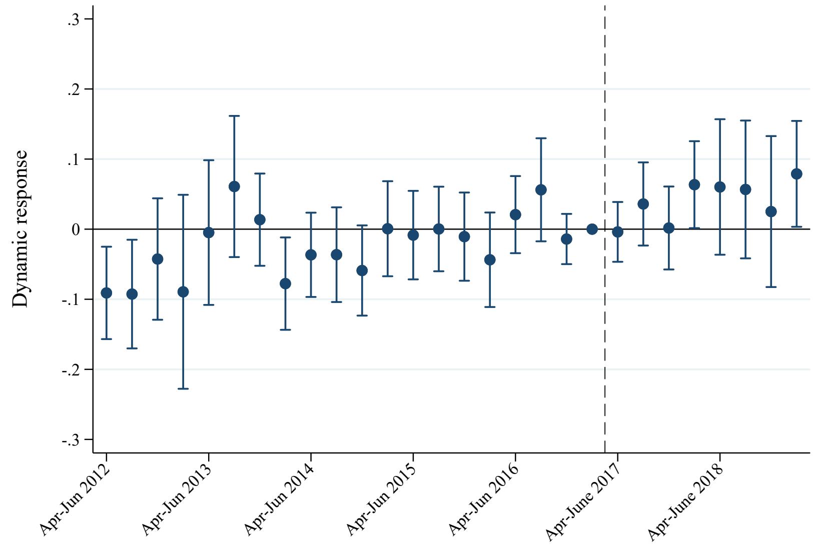

Source: BGT, FAME 2013–2019.
Note: This graph presents the estimates of the leads and lags of the policy on firms' wage-posting decision. These results are obtained from the estimation of the dynamic version of regression E.1. The estimation sample includes firms with 200-300 employees, between the financial years 2013 and 2019. 95 percent confidence intervals associated with firm-level clustered standard errors are also reported. The dash vertical line indicates the month when the mandate is approved, i.e., February 2017.

Table E4: Impact on wage posting - robustness checks

	Main specification (1)	Match score =1 (2)	From 2014/15 (3)
Treated Firm*Post	0.039*	0.047**	0.038*
	(0.021)	(0.023)	(0.023)
Observations	180,679	128,126	96,739
Pre-policy Mean	0.39	0.49	0.48

Source: BGT, FAME, 2013–2019.
Notes: This table presents robustness checks on the impact of pay transparency on wage posting, obtained from the estimation of regression E.1. The estimation sample comprises firms that have between 200 and 300 employees. All regressions include firm and quarter fixed effects, and region-specific time-shocks. The post dummy is equal to one from 2018 onward. A treated firm is defined as having at least 250 employees in 2015. In Column 2, the sample is restricted to firms that have a match score of 1 when merging BGT and FAME. In Column 3, it is further restricted to the fiscal years 2015-2019. The pre-policy mean represents the mean of the outcome variable for the treated group between 2013 and 2017 (2015-2017 in the last column).
*** $p<0.01$, ** $p<0.05$, * $p<0.1$.

Table E5: Impact on other hiring practices

	Wage posted (1)	Log annual pay (2)	Offer Part-time (3)	Offer FT FWA (4)
Treated Firm*Post	0.039*	0.089**	-0.036**	0.008
	(0.021)	(0.036)	(0.015)	(0.012)
Observations	180,679	71,393	180,679	180,679
Pre-policy Mean	0.39	10.05	0.14	0.04

Source: BGT, FAME, 2013–2019.

Notes: This table reports the impact of pay transparency on various firms' hiring practices, obtained from the estimation of regression E.1. The estimation sample comprises firms that have between 200 and 300 employees. All regressions include firm and quarter fixed effects, and region-specific time-shocks. The post dummy is equal to one from 2018 onward. A treated firm is defined as having at least 250 employees in 2015. Heteroskedasticity-robust standard errors clustered at firm level in parentheses. The pre-policy mean represents the mean of the outcome variable for the treated group between 2013 and 2017.

*** $p<0.01$, ** $p<0.05$, * $p<0.1$.