Pay Transparency and Gender Equality*

Jack Blundell[†] Emma Duchini[‡] Ştefania Simion[§] Arthur Turrell[¶]

June 2024

Abstract

Since 2018, UK firms with at least 250 employees have been mandated to publicly disclose gender equality indicators. Exploiting variations in this mandate across firm size and time, we show that pay transparency closes 19 percent of the gender pay gap by reducing men's pay growth. By combining different sources of data, we also provide suggestive evidence that the public availability of the equality indicators enhances public scrutiny. In turn, employers more exposed to public scrutiny seem to reduce their gender pay gap the most.

JEL codes: J08, J16, J24.
Keywords: pay transparency; gender pay gap; public disclosure.

*This paper combines two studies that previously circulated separately: "Pay Transparency and Gender Equality" by E. Duchini, S. Simion, and A. Turrell, and "Wage responses to gender pay gap reporting requirements" by J. Blundell. We are grateful to the editor Erzo Luttmer and three anonymous referees for their insightful suggestions. We further thank Ghazala Azmat, Manuel Bagues, Nick Bloom, Luis Candelaria, Thomas Cornelissen, Mirko Draca, Gabrielle Fack, James Fenske, Caroline Hoxby, Libertad Gonzalez, Victor Lavy, Luigi Pistaferri, Giuseppe Pratobevera, and Giulia Vattuone for their valuable advice. We further thank Leonardo Bursztyn, Rebecca Diamond, Mitch Hoffman, Matt Gentzkow, Maia Güell, Guido Imbens, Ed Lazear, Joseph Mullins, Muriel Niederle, Bobby Pakzad-Hurson, Paul Oyer, Ricardo Perez-Truglia, Barbara Petrongolo, Joahnna Rickne, Paul Robinson, Daniel Schaefer, Sebastian Seitz, Kathryn Shaw, Hans H. Sievertsen, Carl Singleton, Isaac Sorkin, and Simone Traini for their useful comments. A special thanks goes to Mihir Chandraker, Sam Pickering, George Taylor, Marios Tsoukis, and Giulia Vattuone for excellent research assistance. We also acknowledge participants in internal seminars, conferences, and workshops for their constructive suggestions. We finally thank Miranda Kyte for facilitating our access to YouGov data. Blundell thanks Stanford Institute for Economic Policy Research for support, in particular the Dixon and Carol Doll Graduate Fellowship and the George P. Shultz Graduate Fellowship. 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 Bank of England and the wider UK government. All mistakes are our own.

[†] Centre for Economic Performance, London School of Economics, United Kingdom. *Email:* j.r.blundell@lse.ac.uk.
[‡] Corresponding Author. University of Essex, Department of Economics, United Kingdom. *Email:* e.duchini@essex.ac.uk.
[§] University of Bristol, School of Economics, United Kingdom. *Email:* stefania.simion@bristol.ac.uk.
[¶] Bank of England, United Kingdom. *Email:* arthur.turrell@bankofengland.co.uk.

1 Introduction

In recent years, over half of OECD countries have adopted pay transparency policies with the aim of improving gender equality (OECD 2023). By increasing the salience of gender gaps in the labor market, transparency measures are meant to act as an information shock that alters the relative bargaining power of male and female employees vis-à-vis the firm, and, in turn, improve women's relative pay and career outcomes. Importantly, the magnitude of these effects is also likely to depend on how strong and salient the information shock is. Notably, Baker et al. (2022) show that pay transparency reduces the gender pay gap in the context of Canadian universities, where the information on faculty salaries is publicly available, whereas the impact of pay transparency on gender equality is mixed in contexts where employers only have to provide this information internally, to workers' representatives (Bennedsen et al. 2022, Gulyas et al. 2023).

This paper studies the causal impact of pay transparency on the gender pay gap in a context where all large private sector firms have to publicly disclose their gender equality performance, and provides suggestive evidence that the public availability of this information strengthens the disciplinary effects of the policy. 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 websites 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 firm's wage distribution.

We begin our analysis by studying the impact of this policy on the gender gap in hourly pay using the Annual Survey of Hours and Earnings (ASHE), the UK matched employer-employee data set, from 2013 to 2021. To identify causal effects, we exploit 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. To enhance comparability between the treatment and control groups, we restrict the sample to firms with +/-50 employees from the 250-employee threshold in our main specification.

Our results show that the UK pay transparency policy leads to a 19 percent reduction in the gender pay gap, off a base of £2.60 per hour. Importantly, and consistent with the hypothesis that pay transparency reduces the relative bargaining power of better-paid employees (Cullen and Pakzad-Hurson 2023), we find that this effect is driven by a significant 2.9 percent slowdown of men's pay growth in treated firms relative to control ones. Further evidence suggests that this effect comes from a combination of lower growth in bonus payments and fewer promotions for male employees in treated firms compared to control ones.

Event-study exercises show that these results are not driven by different pre-policy trends in the outcomes of interest between treatment and control groups. A battery of placebo regressions further mitigates the concern 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 around the 250-employee cutoff.

As for the mechanisms behind these results, we provide two pieces of suggestive evidence that point to the importance of the public availability of the equality indicators to increase firms' accountability. First, we use two YouGov surveys that, in 2018 and 2019, measured firms' reputation using representative samples of, respectively, British women and British employees, to show that, each year, firms publishing a larger gender pay gap obtain worse placements in both the Women's Rankings and the Workforce Rankings. In other words, the public availability of firms' gender equality performance seems to increase public scrutiny. Second, we provide suggestive evidence that firms that are potentially more exposed to public scrutiny, as measured by their pre-policy investment in advertising, exhibit a larger response to the pay transparency policy. These results do not provide causal evidence for the role of the public disclosure of the equality indicators. However, taken together, they are consistent with the hypothesis that, by enhancing public scrutiny, the public disclosure of the equality indicators may have magnified the disciplinary effects of the policy (Perez-Truglia and Troiano 2018, Luca 2018, Johnson 2020).

Our paper makes 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 2023, Burn and Kettler 2019, Dube et al.

2019, Bennedsen et al. 2022, Cullen and Pakzad-Hurson 2023, Cullen and Perez-Truglia 2022, Baker et al. 2022, Gulyas et al. 2023).[1] The closest studies to ours are Bennedsen et al. (2022), Baker et al. (2022), and Gulyas et al. (2023). Baker et al. (2022) 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. (2022) and Gulyas et al. (2023) analyze the effect on the gender pay gap of, respectively, a 2006 Danish law and a 2011 Austrian law, both of which mandate that private firms provide their employees with pay data by gender and occupation. Both Bennedsen et al. (2022) and Baker et al. (2022) find that transparency leads to pay compression by slowing down men's wage growth. In contrast, Gulyas et al. (2023) find no impact on individuals' wages or the gender pay gap, and suggest that the fact that, in Austria, the pay information is not disclosed publicly may contribute to explain these null results. 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, and more importantly, in the UK setting, firms' gender equality performance is disclosed publicly rather than provided exclusively to employees' representatives. This allows us to study if the public availability of the information enhances public scrutiny, and if, in turn enhanced public scrutiny influences firms' response to the policy. Finally, although the small sample size of the UK matched-employer-employee data set limits our ability to study the compositional effects of the policy, this data set provides rich information on employees' pay, which allows us to unpack the impact of the policy on different pay components, such as bonuses and promotions, and shed light on how firms restructure their rewarding schemes to tackle the gender pay gap.

The paper proceeds as follows. Section 2 describes the institutional setting and the UK transparency policy. Section 3 presents the identification strategy, data, impact of the policy on the gender pay gap, and its compositional effects. Section 4 illustrates the robustness checks. Section

[1] See Bennedsen et al. (2023), Cullen (2023), and Duchini et al. (2023) for recent reviews of the pay transparency literature.

5 studies the role of the publicly availability of the gender equality indicators in influencing firms' response. Section 6 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 (hereafter we will refer to this website as the Gender Pay Gap Reporting website).[2] Also, organizations that are part of a group must report individually. In sum, around 10,500 firms are subject to this mandate 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).[3] To the best of our knowledge, no other substantial law exclusively targeted firms in this size band when the transparency mandate was introduced.[4]

[2]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.

[3]The Business Structure Database (BSD) provides information on firm output, employment, and turnover for almost 99 percent of business organizations registered in the UK. The data come from the Inter-Departmental Business Register (IDBR), a live register of firms collected by the tax authorities via VAT and employee tax records. Office for National Statistics. (2021). Business Structure Database, 1997-2021: Secure Access. [data collection]. 14th Edition. UK Data Service. SN: 6697, DOI: 10.5255/UKDA-SN-6697-14.

[4]Since 2010, employees working in firms with at least 250 employees have the right to request time off for training.

The timing of the publication of the equality indicators works as follows. Each year, if a firm has at least 250 employees on April 5th (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 April 5th of the following year. Firms themselves must calculate their number of employees using guidelines provided by the government. Importantly, they have to adopt an extended definition of an employee that includes agency workers. Partners of firms are also included in the definition of an 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 percentage of female employees in each quartile of a company's pay distribution. Table 1 provides sample means of these indicators. The first thing to note is that the sample size substantially drops in 2019/20. This is because in mid-March 2020 the government temporarily lifted the transparency mandate due to the Coronavirus outbreak, and firms were only asked to start publishing the equality indicators again in October 2021. By the time the mandate was paused, just over half of the firms that were deemed to publish the equality indicators in 2020 had done so.[5]

The first row of Table 1 shows that the gender gap in median pay is 12 percent in 2017/18 and remains between 12 and 13 percent in the following years. The gap in mean pay is around 14 percent in 2017/18, and only decreases to 13 percent by 2022/23. Both gaps in median and mean bonuses tend to be smaller than pay gaps but it is also worth noting that in the first two years some firms mistakenly reported their level gap rather than a percentage, making it difficult to interpret

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.

[5]Note also that the number of observations in each year varies slightly depending on the date the data are downloaded from the Gender Pay Gap Reporting website. This happens, for instance, because some firms report data retrospectively.

these bonus gaps.[6] The proportion of women receiving bonus pay is smaller than for men in each year, but the former exhibits a larger percentage increase over time. The gender ratio along the pay distribution is in favor of women at the bottom, but at most 41 percent of employees in the upper part of the wage distribution are women by 2022/23. Lastly, the proportion of women in each quartile of the pay distribution slightly increases over the years.

Clearly, the magnitude of these raw firm-level indicators depends both on compositional and observable factors, such as gender differences in educational choices, occupation held and experience, and unobservable factors such as employers' unconscious biases and subtle discrimination in the workplace (Azmat et al. 2020, Bertrand 2020). As they are, these aggregate measures do not allow one to distinguish the importance of each underlying factor, and statistics broken down by occupation, or even better hierarchy position, would be more informative in this respect. Yet, the firm-level indicators may reflect a compromise between the government's will to disclose these statistics publicly, and firms' privacy concerns, and it is a matter of empirical analysis to understand how effective they are at pushing firms to tackle the underlying causes of gender inequality. From now on, we will refer to these data as the GPG data, or data published by GPG firms.

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 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 pooling all the publication years together. While some firms do not meet the deadline, the majority publish their data in the last month before it. Note also that less than 600 firms with fewer than 250 employees publish gender equality indicators each year. 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

[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.

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 the Government Equalities Office 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 both a dedicated government website and companies' own 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, *The Guardian* 2021, *Financial Times* 2023), and firms are not spared from "naming and shaming" articles.[7] Importantly, Panel B of Figure 1 shows that Google searches for the term "gender pay gap" spiked around the first deadline, indicating that this policy attracted significant public interest. Moreover, although searches for this topic have diminished since then, especially at the peak of the pandemic, at each reporting deadline public attention re-surges. And while there is no direct evidence that employees of targeted firms consult the gender equality indicators, the law requires that firms publish this information on their website "in a manner that is accessible to all its employees and to the public; and for a period of at least three years beginning with the date of publication",[8] which makes it unlikely that employees are completely unaware of it.

[7]For example, the Independent, a prominent daily newspaper ran a story titled "Gender pay gap: worst offenders in each sector revealed as reporting deadline passes"(*Independent* 2018). In this article a championship football club and an airline were revealed as having among the greatest gender pay gaps in the country.

[8]The full text of the law is available at https://www.legislation.gov.uk/ukdsi/2017/9780111152010.

3 Impact on the gender pay gap

3.1 Identification strategy

Our primary goal is to identify the impact of the UK pay transparency policy on the gender gap in hourly pay and unpack this into the effect on women's and men's pay. For this, we exploit the variation in the implementation of the policy across firm size and over time, and compare the evolution of the outcomes of interest in firms whose size is slightly above (treatment 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.[9] Moreover, to enhance comparability between treatment and control group, we consider firms with $+/-50$ employees from the 250 threshold in the main specification. 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. Based on these choices, we estimate the following triple-differences regression model that aims to estimate the relative impact of the policy on men's and women's outcomes:

$$\begin{aligned} Y_{ijt} &= \alpha_{ij} + \theta_{rt}^{M} + \theta_{rt}^{F} + \beta\left(TreatedFirm_j * Post_t\right) \\ &\quad + \gamma\left(TreatedFirm_j * Post_t * Fem_i\right) \\ &\quad + X_{ijt}'\pi + u_{ijt}, \end{aligned} \tag{1}$$

where i is an employee working in firm j, which has between 200 and 300 employees, in year t, running between 2013 and 2021,[10] M and F stand respectively for men and women, and r

[9] In the estimation sample, which we describe in Section 3.2, we find that 69 percent of firms still fall on the same side of the cutoff in the post period. Appendix Figure A1 further shows the distribution of firms around the 250-employee cutoff in each year since the announcement of this threshold. Data are drawn from the Business Structure Database, covering 99 percent of UK firms. 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.

[10] As explained in Section 3.2, we choose this time window because it is the maximum number of years over which we observe all outcomes of interest.

stands for one of the 11 UK macro-regions. The outcome Y_{ijt} is either a pay or career outcome, as defined in the next section. As for the regressors, α_{ij} are individual-firm fixed effects that capture the impact of individual-firm-specific time-invariant characteristics such as the quality of the match between the employee and the employer; θ^M_{rt} and θ^F_{rt} are gender-region-year fixed effects that control for local time shocks common to all firms operating in a region but gender-specific such as the local expansion of public child care; Fem_i is a dummy variable that is equal to one if i is a woman; $TreatedFirm_j$ is a dummy variable equal to one if a firm has at least 250 employees in 2015; as for $Post_t$, in our main specification we constructed it as a dummy variable equal to one from 2018 onward. We explain this choice in the next section, after describing the timing of the UK matched-employer-employee data set. In our main specification, we also do not include any further controls, but in Section 4, we present robustness checks where the vector X_{ijt} includes workers' age controls, as well as an alternative specification with gender-1-digit-SIC specific time shocks, in place of gender-region-year fixed effects. Standard errors are clustered at the firm level.

Our main coefficient of interest is γ which, conditional on the validity of this identification strategy, should capture any deviation from a parallel evolution in the outcome's gender gap between the treatment and the control group due to the introduction of the mandate. Put differently, γ should identify the differential effect of the policy on women compared to men. Equally important are β and $\beta + \gamma$, which identify, respectively, the effect of the policy on male and female employees. Thus, at the bottom of results' tables we also report the p-value of the t-test on women's effect. If the policy was effective at reducing the gender pay gap, we would observe a $\gamma > 0$. A negative β would tell us that, in order to narrow the gender pay gap, treated employers slowdown men's pay growth relative to the control group. Finally, if $\beta + \gamma > 0$, this would mean that women's pay increases faster in treated firms compared to control firms after the implementation of the policy.

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

$$
\begin{aligned}
Y_{ijt} &= \alpha_{ij} + \theta^M_{rt} + \theta^F_{rt} + X'_{ijt}\pi \\
&+ \sum_{k=2013}^{2021} \beta_k (TreatedFirm_j * \mathbf{1}[t=k]) \\
&+ \sum_{k=2013}^{2021} \gamma_k (Fem_i * TreatedFirm_j * \mathbf{1}[t=k]) + \nu_{ijt},
\end{aligned} \tag{2}
$$

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 first reporting deadline, as the reference year.

Next, in Section 4, we will provide evidence that our results do not capture the effect of other time shocks that coincide with the introduction of the pay transparency policy and affect firms on either side of the 250-employee cutoff differently. And we will show that the results do not depend on the size of the bandwidth considered around the policy cutoff, nor on the year chosen to define the treatment status.

3.2 Data

To study the impact of the policy on the gender pay gap, we use the Annual Survey of Hours and Earnings (ASHE).[11] 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. 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 date in April each year. Information on age, occupation (SOC), and firm's sector (SIC) 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.

[11] Office for National Statistics (2022). Annual Survey of Hours and Earnings, 1997-2021: Secure Access. [data collection]. 20th Edition. UK Data Service. SN: 6689, DOI: 10.5255/UKDA-SN-6689-19.

The main limitation of ASHE is its small sample size and the fact that we do not observe all the employees of a firm, which does not allow us to compute a firm-level measure of the gender pay gap. However, this is the only data set available in the UK that provides both a large range of employees' outcomes, including salary components, and information on the total number of employees in a firm and year, which allows us to define the treatment status in our identification strategy.[12]

Treatment timing. As explained in the previous section, in our main specification, we assume that the treatment period starts in 2018. In principle, employers could start reacting in 2017, considering that by the first deadline (April 5 2018) they had to report equality indicators calculated as of April 5 2017. Such a reaction would be visible in ASHE 2017, as this wave provides pay variables measured in April 2017. However, given that the law was approved in February 2017, employers had very little time to adjust employees' pay by April 5th. Moreover, most employers had never computed their gender pay gap before then, nor they had information on the gender pay gap of their competitors. On top of this, in 2017 it was not yet clear what the media coverage or the public audience's interest in these statistics would have been by the time of the first deadline, April 2018. For all these reasons, in our main specification we assume that the treatment period starts in 2018. However, ultimately, whether employers started reacting in 2017 or 2018 is a matter of empirical analysis. The event-study specification discussed in the next section will be informative in this respect. And to explore this further, in Section 4 we will estimate a specification where $Post_t$ is equal to 1 from 2017 onward. Importantly, note that if there had been any employers' reaction prior to 2018, our main estimates would be downward-biased.

Next, from ASHE, we create the following variables:

Pay measures. Our main variable of interest is employees' hourly pay, including additional

[12]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 (see Footnote 5 and Appendix Section A.1 for more information on this data set). This concerns 28 percent of firms in our estimation sample. Note that the correlation between the firm size variables for firms present in both data sets is 0.997. More importantly, the values coincide for 74 percent of firms, and the average difference is 1 employees for firms with different values in the two data sets. More information on the matching between ASHE and BSD is provided in Appendix Section A.1. In Section 4 we further show that our results are not affected if firms with missing firm size in ASHE are excluded from the estimation sample.

payments, i.e., allowances, bonuses, and shift pay, but excluding overtime pay; we also separately consider the basic hourly pay and additional payments, as well as weekly pay and hours worked. To study the impact of the policy on pay variables, we take log transformations. As for the additional payments, to take into account that 80 percent of workers do not receive any of them, we consider the ratio between these payments per hour and the employees' hourly basic pay.[13] When studying the impact of the policy on this variable, we exclude workers with a ratio of additional payments to basic pay that is greater or equal than one, that is workers who are mostly payed in the form of allowances or bonuses (0.4 percent of the sample). All monetary values are deflated using the ONS' 2015 CPI Index. To complement the analysis on pay outcomes, we also study the impact of the policy on employees' promotion prospects. For this, we consider the ONS' definition of promotion as an event whereby an employee has experienced at least a 30 percent increase in his/her hourly pay since the previous year or has acquired managerial responsibilities (ONS 2020). As ASHE does not provide information on the acquisition of managerial responsibilities for all the years of the estimation sample, we measure promotions using a dummy variable that is equal to one if, within the same firm, an employee has experienced at least a 30 percent increase in his/her hourly pay since the previous year.

Occupation and job mobility. To get a full understanding of the impact of the policy on employees, in our analysis we also consider its compositional effects. First, to study mobility into the firm, we use a dummy variable that is equal to one if the worker has at most one year of tenure in the firm, though the tenure variable is missing for around 3 percent of workers. Second, to study separations, we construct a dummy variable that is equal to one if the employee has left the firm by $t+1$. By construction, this variable is missing in 2021. Third, we consider employees' occupation. To take into account that we only observe few employees per firm, we consider three groups of 1-digit SOC occupations, the bottom, middle, and top terciles of the pay distribution, based on the ranking of the 1-digit SOC median hourly pay pre-policy.

[13]While it would be interesting to explore separately the impact of the policy on the probability of receiving additional payments, and on the amount received, the fact that the latter outcome is only observed for 20 percent of the sample strongly limits our ability to do this.

Estimation period and sample restrictions. As for the estimation period, we use data over the years 2013–2021. We start from 2013, as we can observe all outcomes since then, and stop in 2021, as this is the last available year of data at the time of writing. However, note that we use information from 2012 to construct the promotion dummy. In terms of sample restrictions, we drop individuals with missing id or missing firm id (0.4 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). Finally, we drop individuals with a basic pay equal to 0, representing 0.7 percent of observations. Our resulting sample in the main specification is formed of 13,063 men and 11,995 women, for a total of 27,051 individual-year observations for men and 24,226 for women. We observe men across 5,981 firms and women across 5,558 firms.

Summary statistics. Table 2 provides summary statistics for the main outcomes measured in the pre-policy period, 2013–2017, separate for men and women and treated and control firms. 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 gap in hourly pay amounts to £2.58, or 16 percent of men's pay. This gap reaches 29 percent when looking at weekly pay, as there is also a 16 percent gap in hours worked. As for additional payments, there is a large gender gap in the probability of receiving them (34 percent), and an even larger gap in the amount received per week (61 percent). In turn, additional payments per hour constitute a larger share of men's hourly base pay than women's hourly base pay (4 vs. 2 percent). And while there is no gender promotion gap, there is a 5 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, 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.

3.3 Results

This section presents the impact of the pay transparency policy on employees' pay. Figure 2 shows the event studies for the gender gap in hourly pay, in Panel A, and the log hourly pay, separately for men and women, in Panels B and C.[14] From these figures, we observe, first, that the evolution of the outcomes in the pre-policy period is comparable across treatment and control groups, both for what concerns the gender pay gap, and separately for male and female employees' pay. Note that the absence of significant pre-trends also suggests that employers did not react to the policy before the first publication deadline. However, Panel A shows that women's pay increases relatively to men's pay from 2018 onwards, with this effect being significant at 10 percent in 2019. In 2020 and 2021, the estimates become noisier, most likely because the number of observations in the sample falls by 25 percent in these two years relative to 2019, as the pandemic substantially reduced labor force participation (Barrero et al. 2022, Li and Granados 2023).[15] Interestingly, in (April) 2020, when the government pauses the mandate to publish gender equality indicators and the attention of the public audience shifts towards the Coronavirus outbreak, the magnitude of the policy effect slightly decreases compared to 2019, which suggests that the policy has indeed disciplinary effects on firms' behavior.

The third point that emerges from Panel B of Figure 2 is that the effect on the gender pay gap is driven by a slowdown of men's pay growth in treated firms relative to control firms after the introduction of the mandate, with this effect being significant at the 5 percent level in 2019, and at 10 percent in 2021. In contrast, the policy does not appear to have any visible impact on women's pay (Panel C).

Table 3 reports the estimates of the corresponding average effects of the policy. Each column shows a different outcome. At the bottom of the table, we report the p-value of the t-test on the effect on women and the pre-policy mean for the treatment group calculated over the period 2013–2017. Consistent with the dynamics seen in the event studies, Column 1 shows that the

[14] See Table A1 for the detailed regression results.

[15] This is true for both men and women, and treated and control firms, and it is not specific to the estimation sample, but is visible in the overall ASHE sample.

policy leads to a significant 3 percentage-point increase in women's hourly pay compared to men's pay. Relative to the pre-policy value of 16 percent, this corresponds to a 19 percent decrease in the gender gap in hourly pay. Importantly, the coefficient on *TreatedFirm*Post* confirms that this effect is driven by a 2.9 percent significant decrease in men's real pay, while on average, the policy has no impact on women's pay. These results are remarkably consistent with the estimated effects of pay transparency in other settings (Bennedsen et al. 2022, Baker et al. 2022). Here we exploit the richness of information provided by ASHE to further decompose our results into the effect on additional payments and the impact on promotions.

To open this discussion, note that our specification includes worker times firm fixed effects, which implies that these results are not driven by compositional effects, such as high-paying men leaving treated firms or inexperienced women joining them after the introduction of the policy. Instead, they are driven by differential changes between treated and control firms in wages of employees that were already employed at these firms before the implementation of the policy. Nonetheless, as compositional effects are interesting per se, we analyze them separately in Section 3.4.

Columns 2 to 4 of Table 3 then unpack the wage effects into the impact on the different pay components. The slowdown of men's pay growth is clearly visible when considering the contractual pay in Column 2. Column 3 adds to this that the policy has no significant impact on the ratio of additional payments to base pay, for either men or women. Note that the null effect on this variable for men implies a negative effect on the actual amount of additional payments, given that the denominator of the ratio is negatively affected by the policy. Finally, Column 4 shows that, on average the policy has no significant effect on either men's or women's probability of promotion. However, the dynamic specification depicted in Figure 3, Panel C, shows a significant 10 percent drop in men's promotion prospects in 2019, which persists but becomes insignificant in 2020, and turns positive in 2021 (though with a very large confidence interval). Overall, these results show that the slowdown of men's pay growth in treated firms compared to control firms is driven by a combination of lower growth in additional payments and fewer promotions.[16]

As for the null effect on women's pay, we cannot rule out that both treated and control firms have raised women's pay as they compete for the same workers. To investigate this further, in the online appendix we plot men and women time effects for the control group from regression 1 and show that the hourly pay evolves similarly across genders, which speaks against general equilibrium effects (See Appendix Figure A2). However, any conclusion from this exercise has to be taken with a grain of salt given that it does not provide causal evidence that the control group is not raising women's pay.[17]

The next section discusses the compositional effects of the policy. Next, Section 4 is dedicated to showing that our 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 5 will explore the contribution of different channels in explaining these results.

3.4 Compositional effects

So far, we have shown that the UK pay transparency policy leads to a reduction of the gender pay gap, driven by a slowdown of men's pay growth. Equally relevant from a policy point view are its compositional effects. For instance, treated firms may try to reduce their gender pay gap by hiring more women in better-paid occupations. At the same time, firms disclosing their gender pay gap may find it more difficult to recruit talented women. And more productive women may decide to quit their firm if they are disappointed by their employer's performance in terms of gender equality or unsatisfied with the firm's response to the policy. Note that some of these effects may counteract each other, and one challenge that we face in disentangling them is that we do not have

[16] In the online appendix we further show that the effect on men's hourly pay comes from a negative impact on weekly pay rather than an increase in hours worked (See Appendix Table A2). Further results suggest that men in treated firms do not experience any nominal or real pay cut, which supports the interpretation of the main effect as a slowdown of pay growth in treated vs. control firms (See Appendix Table A3).

[17] While the small sample size in ASHE limits our ability of conducting subgroup analysis, we also explore whether the policy has a differential impact across occupations, but do not find strong evidence for heterogeneous effects along this dimension (See Appendix Table A4).

worker-level measures of productivity. Bearing this in mind, in Table 4, we provide evidence on net compositional effects. To this aim, the regressions in this table include firm fixed effects instead of firm times individual fixed effects.

Column 1 tells us that the policy has no overall impact on hiring or the gender composition of new hires. Yet, note that there could still be heterogeneous effects across occupations, as results in Columns 3-5 suggest. Column 2 shows that the policy seems to increase the probability of women's separations in treated firms compared to control firms, both relative to men and in absolute terms. However, this result does not pass all the robustness checks, and hence we do not want to over-interpret it (see Appendix Figures A3 and A4). It is however consistent with the findings of studies outside of the gender literature, which, taken together, show that pay transparency may increase absences and the intention of quitting of lower-paid employees, when pay differences are perceived to be unfair (Card et al. 2012, Breza et al. 2018, Dube et al. 2019, Cullen and Perez-Truglia 2022). Next, although we do not find significant effects on firms' occupational composition, point estimates in Columns 3 to 5 suggest that treated firms have indeed tried to hire more women in better paid occupations.[18] Finally, in Column 6, we re-estimate our main regression on log hourly pay to measure the total effect of the policy on employees' pay, including any compositional effect. Interestingly, the effects on men's pay and the impact on the gender pay gap become smaller in magnitude and insignificant, while the total effect on women is larger than in the main specification (Table 3, Column 1), although still insignificant. This last result suggests that treated firms' efforts to hire women in better-paid occupations may have counteracted any difficulty in recruiting or replacing talented women. As for men, while finding a smaller (in absolute terms) and insignificant estimate in this specification may seem surprising, we can only speculate that the policy may have induced more low-productive men to leave treated firms for fear of the employer's response to the

[18]In particular, the policy has a significant negative impact on the probability of observing women employed in the bottom tercile of the wage distribution in treated firms compared to control ones. To further explore whether firms have increased their efforts to reduce their gender pay gap at entry level, we have studied the impact of the policy on firms' hiring practices, and specifically on firms' wage posting decision, using vacancy data from Lightcast (previously known as Burning Glass Technologies). Although we do not find robust evidence that the policy affects this margin of decision, interestingly, we find that firms that are more likely to post wage information in their vacancies tend to have a lower gender pay gap and a larger share of women in the top quartile of the wage distribution (see Appendix Section C).

policy, rather than pushing high-productive men to leave treated firms for their actual response to the policy.

4 Robustness checks

This section presents two sets of robustness checks on the estimated effect of the policy on the gender pay gap. First, we show that our results are unlikely to be driven by time shocks that happen at the same time of the policy and 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 4, and report detailed regression tables in the online appendix (see Appendix Section A).

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 run a series of placebo tests pretending that the mandate binds at different firm size thresholds. The estimated effects of these placebo reforms on the gender pay gap, together with 95% and 90% confidence intervals, are displayed in Panel A of Figure 4. In each regression, the estimation sample includes firms with $+/-50$ employees from the threshold indicated on the vertical axis. Reassuringly, none of the placebo mandates has a significant impact on the gender pay gap. This exercise helps 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.[19]

Specification. Our second set of robustness checks aims to verify that our results are robust to the choice of the bandwidth around the 250-employee cutoff, do not depend on the year we

[19]Note that the regressions corresponding to placebo cutoff values ”300” to ”450” include all treated firms. The fact that the point estimates are positive 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.

use to define the treatment status, and are not sensitive to the other choices made in the main specification. Panel B of Figure 4 shows that the estimates of γ from equation 1 change very little when restricting or enlarging the bandwidth around the 250-employee cutoff. Specifically, the estimated effect of the policy on the gender pay gap are only marginally insignificant when using a bandwidth of $+/-60$ (p-value 0.102) and $+/-90$ (p-value 0.135), while become smaller and insignificant with a bandwidth $+/-100$, where treated and control firms may start to be less comparable. Importantly, in the online appendix we show that the estimated negative effect on men's pay is always significant and comparable in magnitude across the different regressions (see Appendix Table A6).

Next, Panel C of Figure 4 compares the impact of the policy on the gender pay gap when changing the year used to define the treatment status. Note that, to avoid capturing any impact of the policy on firm size, we only consider years before the announcement of the employee-cutoff, which took place in the fall of 2015. While the estimates on the gender pay gap become insignificant when using the firm size in 2014 to define the treatment status, the other specifications give significant and comparable effects to the main specification. In the online appendix, we also show that the negative impact on men's pay is always significant and comparable in magnitude across the different regressions (see Appendix Table A7).

Finally, Panel D of Figure 4 shows that the estimated impact on the gender pay gap changes little when: including gender-industry specific time shocks in place of gender-region specific time shocks; adding age controls; restricting the sample to either workers aged 16-65 or those aged 25+; considering only full-time employees or those working in the private sector; using Labour Force Survey weights; or restricting the sample to firms for which we can use only ASHE-based information on the number of employees to define the treatment status. Finally, note that the effect on the gender pay gap remains significant when including 2017 in the treatment period, but point estimates are slightly smaller, again pointing to little employers' response before 2018.

5 Mechanisms

Our results show that the UK pay transparency policy reduces the gender pay gap through a slowdown of men's pay growth. This finding is remarkably consistent with the evidence produced by contemporaneous studies (Bennedsen et al. 2022, Baker et al. 2022). As an increasing number of countries introduce pay transparency policies, it is especially important to understand in what circumstances these laws are effective at reducing gender inequality.[20]

One of the most innovative features of the UK transparency policy as compared to the mandates introduced in other countries is that firms must make their equality indicators publicly available. The public disclosure of this information has the potential to magnify the disciplinary effects of transparency policies (Perez-Truglia and Troiano 2018, Luca 2018, Johnson 2020) through two channels. First, it enables performance comparisons across firms, which may induce worse performing employers - that is those with a larger gender pay gap - to improve gender equality the most. Second, the public availability of the equality indicators may enhance public scrutiny, which, in turn, may induce firms that are more exposed to it to reduce the gender pay gap the most.

Notably, the evidence on the impact of pay transparency is mixed in contexts where the information on firms' gender equality performance is only revealed internally, to employees' representatives (Bennedsen et al. 2022, Gulyas et al. 2023). In contrast, our findings, and those of Baker et al. (2022), show that pay transparency enhances gender equality in contexts where this information is publicly available. While there could be different reasons why a policy works in a context and not in others, in light of these results, it seems important from a policy perspective to explore the role of public disclosure.

In this section, we thus investigate whether the public availability of the gender equality indicators may have been important to increase firms' accountability. While we find little evidence that performance comparisons across firms influence their response to the policy, we provide two pieces of evidence pointing to the importance of public scrutiny. First, we use two YouGov surveys that, in 2018 and 2019, measured firms' reputation using representative samples of, respectively,

[20]In the conclusion, we will return to the failure of these policies to increase the salaries of low-paid workers.

British women and British employees, to show that, each year, firms publishing a larger gender pay gap obtain worse placements in both the Women's Rankings and the Workforce Rankings. In other words, the public availability of firms' gender equality performance seems to increase public scrutiny. Second, we provide suggestive evidence that firms that are potentially more exposed to public scrutiny, as measured by their pre-policy investment in advertising, exhibit a larger response to the pay transparency policy.

Performance comparisons. The behavioural economics literature provides evidence that when individuals receive information on their relative performance, those performing worst improve the most afterwards (Allcott and Kessler 2019). The same may be true of firms comparing their relative performance in terms of gender equality. Unfortunately, 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 in ASHE.[21] However, we explore this mechanism descriptively using the publicly available data on the gender gap in the median hourly pay that firms subjected to the mandate publish each year. In particular, in Figure 5, we compare the yearly distribution of the gender pay gap. If worse performing firms had reduced the gender pay gap the most over time, we should see the distribution shrinking from its right-tail. Although large values seem to become progressively less frequent, the comparison in Figure 5 does not offer strong evidence in favor of such a behavioral response.[22] In other words, performance comparisons across firms publishing the equality indicators do not seem the key driver of businesses' response to the policy. We now turn to study the role of public scrutiny.

Firms' reputation. In 2018 and 2019, the polling organization YouGov compiled two distinct rankings of 1,342 firms operating in the UK, called, respectively, YouGov Women's Rankings and YouGov Workforce Rankings. These rankings are parts of a larger initiative that YouGov launched

[21]ASHE does not provide information on all employees in a firm.

[22]In the online appendix, we further perform a Kolmogorov-Smirnov test to formally compare year-on-year distributions, and only in the comparison of the 2021/22 and 2022/23 distributions we reject the null that the two distributions are equal in favor of the hypothesis that the 2022/23 has smaller values than the 2021/22 distribution (see Appendix Figure B1).

at the end of the 2000s to keep track of “brands’ health”.[23] Between January and December of each year, YouGov interviews a representative sample of 50 to 100 UK individuals per day, and, among other queries, it poses a series of questions to measure what individuals think of specific brands. Women’s Rankings are based on women’s answers to the question: “Overall, of which of the following brands do you have a positive/negative impression?”. The Workforce Rankings are instead obtained by asking both employed men and women: “Which of the following brands would you be either proud or embarrassed to work for?”. The resulting “impression score” in the case of Women’s Rankings, and “reputation score” in the case of the Workforce Rankings, are constructed as the percentage difference between all the positive and negative answers relative to all the answers received in the survey; the higher the score that a firm receives in a survey, the better its placement in the corresponding ranking. The YouGov Workforce Ranking was discontinued after 2019, while the YouGov Women’s Ranking was compiled for an additional year and was then also discontinued after 2020. YouGov has kindly shared with us the 2018 and 2019 data for the two rankings.[24]

Ideally, we would like to compare the evolution of firms’ placement in the two rankings before and after the introduction of the policy, but unfortunately YouGov only started compiling these rankings in 2018. However, we can study descriptively how a firm’s placement correlates with its gender equality performance. For this, we manually matched YouGov data with firms’ gender equality indicators. Taking into consideration that more than one YouGov firm is associated with the same GPG parent company, we match 943 (924) YouGov companies, or 70 (69) percent of the YouGov sample, to 540 (527) companies disclosing their equality indicators in 2018 (2019), or around 5 percent of GPG companies each year.[25] In terms of sample selection, GPG firms included

[23] YouGov could not tell us what criteria it adopted to build its initial list of brands, but it told us that over time this list has been updated according to both employers’ demand and market conditions.

[24] YouGov later informed us that from 2018 onward they have also been compiling the so-called BrandIndex Buzz Rankings, based on answers of all respondents to the question ‘Over the past two weeks, which of the following brands have heard something positive/negative about?” While these data would have also been useful for our analysis, unfortunately, we have not been able to obtain either these data, or data on other questions concerning brands’ health.

[25] The YouGov companies that we cannot link with the GPG data are mostly below the 250-employee cutoff or not registered in the UK. Note also that the impression scores of women are not available for three companies. Finally, while the list of firms included in YouGov surveys does not change over time, the pool of GPG employers varies from one year to another as it only includes firms with at least 250 employees as of that year.

in the YouGov list have a slightly larger gender pay gap than the other GPG companies, especially in 2018, and they tend to be among the largest employers that publish gender equality indicators (see Appendix Table B1 and Appendix Figure B2).

Before exploring the patterns of correlation between firms' placement in YouGov Rankings and their gender equality performance, it is important to consider the timing of the two data sets. Most GPG firms publish their gender equality indicators by April each year, and YouGov surveys are run from January to December. This implies that, each year, at least two thirds of people interviewed by YouGov have access to the information on firms' gender equality performance for the year when the interview takes place. Given this timing, we explore within-year correlations between firms' equality indicators and their placements in YouGov Rankings.[26] Importantly, the availability of two years of data allows us to compute these correlations conditional on firm and year fixed effects. We also cluster standard errors at the level of the GPG company. Lastly, because a larger number in the ranking means a worse placement, we invert the ranking for ease of interpretation. Table 5 shows that a one standard-deviation increase in a firm's gender pay gap is associated with a loss of 9 and 10 positions (out of 1,342) in, respectively, YouGov Women's Rankings and YouGov Workforce Rankings. While these dynamics could be influenced by other factors in addition to year and firm fixed effects, they are consistent with the hypothesis that the public availability of the equality indicators increases the attention of the public audience. This further motivates us to study firms' response to increased public scrutiny.

Firms' response to public scrutiny. The negative correlation between firms' gender pay gap and YouGov's Rankings suggest that increased public scrutiny plays an important role in shaping firms' response to the pay transparency policy. To explore this hypothesis further, we would like to compare the response to the policy across firms that may be more or less exposed to public

[26]It would also be interesting to study how firms' placement changes from one year to the next depending on their gender equality performance the first year. However, the fact that, each year, the majority of YouGov interviews take place after the equality indicators for that year become available makes it difficult to isolate the influence that their disclosure has on the evolution of firms' reputation. Similarly, it would be interesting to study whether firms that perform worse in the 2018 YouGov Rankings reduce their gender pay gap the most by the following year, but, unfortunately, it is unlikely that firms publishing gender equality indicators by April 2019 already have the information on their performance in 2018 YouGov Surveys, as these are run until December 2018.

scrutiny, but face the challenge that there is no official definition of what it means to be exposed to public scrutiny. To overcome this challenge, we proceed as follows: firms that are more exposed to public scrutiny are likely to be firms that the public audience is more familiar with. In turn, firms that spend a larger share of their budget on advertising are likely to be more renowned among the public audience. We therefore ask whether firms that have traditionally spent more on advertising exhibit a larger response to the pay transparency policy.

To answer this question, we exploit data on firms' annual advertising costs provided by the Annual Business Survey.[27] The Annual Business Survey (ABS hereafter) is an annual survey of businesses covering the production, construction, distribution, and service industries, which represent about two-thirds of the UK economy in terms of gross value added. From ABS, we constructed an advertising-to-sales ratio, as the ratio between advertising costs and sales, and computed the average ratio for each firm between 2013 and 2017. We then matched ABS and ASHE, using the common anonymized firm identifier, and found 78 percent of firms included in the ASHE estimation sample. Finally, we ranked ASHE firms based on their average pre-policy advertising-to-sales ratio and grouped employers with below- and above-median advertising-to-sales ratios.[28]

Table 6 compares the impact of the policy on employees' pay across these two groups. While the coefficients are not statistically different across subgroups for either men or women, point estimates suggest that firms that care more for their public image, as proxied by their pre-policy advertising-to-sales ratio, have a larger response to the policy.[29] Although these results are merely suggestive, and advertising costs are only an imperfect measure of firms' exposure to public scrutiny, these estimates are consistent with the hypothesis that reputation concerns play an important role in influencing firms' response to the policy.

Overall, in this section, we have provided suggestive evidence that the public availability of

[27] Office for National Statistics. (2021). Annual Business Survey, 2005-2019: Secure Access. [data collection]. 15th Edition. UK Data Service. SN: 7451, DOI: 10.5255/UKDA-SN-7451-15.

[28] More information on the construction of the advertising-to-sales ratio in ABS and the matching between ASHE and ABS is provided in Appendix Section B.3.

[29] Focusing on the effect on the gender pay gap, point estimates imply that high-advertising costs firms close up to 34 percent of their pre-policy pay gap, while low-advertising costs firms do not reduce their gender pay gap in response to the pay transparency policy.

firms' gender equality indicators enhances public scrutiny, which, in turn, pushes firms that are potentially more exposed to public scrutiny to react the most to the policy.

6 Conclusion

To tackle the persistence of gender inequality in the labor market, many governments are introducing pay transparency policies. Exploiting the variation across firm size and over time in the application of the UK transparency policy, we provide causal evidence that this policy leads to an 19 percent significant reduction in the gender pay gap. Moreover, we provide suggestive evidence that the public disclosure of firms' gender equality performance may have contributed to increase firms' accountability.

To conclude, we discuss two points to reflect on the policy implications of our analysis. First, to evaluate the overall effectiveness of this policy, one would need to consider all of its implications for workers and firms. In particular, the slowdown of men's pay growth in treated firms compared to control ones could translate into higher relative profits. At the same time, the publication of the equality indicators, coupled with employers' response to the policy may decrease workers' job satisfaction and productivity, with negative knock-on effects on profits. An important limitation of our analysis is that we do not have worker-level measures of productivity or comprehensive data on firms' profits to provide robust evidence on these effects.

Second, focusing on the impact of the policy on gender equality, our analysis shows that the reduction in the gender pay gap is the result of a slowdown in men's pay growth, while the policy has no significant effect on women's pay. In other words, these 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 women's outcomes. As an increasing number of studies confirm that transparency policies mainly generate pay compression by slowing down the pay growth of better-paid employees (Mas 2017, Bennedsen et al. 2022, Cullen and Pakzad-Hurson 2023, Baker et al. 2022), policy makers should consider whether this is a desirable way to tackle wage inequality.

References

Adams-Prassl, Abi, Maria Balgova, Matthias Qian, and Tom Waters, "Firm Concentration & Job Design: the Case of Schedule Flexible Work Arrangements," 2023.

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.

Azmat, Ghazala, Vicente Cuñat, and Emeric Henry, "Gender Promotion Gaps: Career Aspirations and Workplace Discrimination," *Available at SSRN 3518420*, 2020.

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," *American Economic Journal: Economic Policy*, 2022, *15* (2), pp. 157–183.

Banfi, Stefano and Benjamin Villena-Roldan, "Do High-wage Jobs Attract More Applicants? Directed Search Evidence from the Online Labor Market," *Journal of Labor Economics*, 2019, *37* (3), pp. 715–746.

Barrero, Jose Maria, Nicholas Bloom, and Steven J Davis, "Long Social Distancing," NBER Working Paper No. 30568, National Bureau of Economic Research 2022.

Belot, Michèle, Philipp Kircher, and Paul Muller, "How Wage Announcements Affect Job Search – A Field Experiment," *American Economic Journal: Macroeconomics*, 2022, *14* (4), pp. 1–67.

Bennedsen, Morten, B. Larsen, and J Wei, "Gender Wage Transparency and the Gender Pay Gap: a Survey," *Journal of Economic Surveys*, 2023, *00*, pp.1–35.

_, Elena Simintzi, Margarita Tsoutsoura, and Daniel Wolfenzon, "Do Firms Respond to Gender Pay Gap Transparency?," *The Journal of Finance*, 2022, *77* (4), 2051–2091.

Bertrand, Marianne, "Gender in the Twenty-first Century," in "AEA Papers and proceedings," Vol. 110 2020, pp. 1–24.

Biasi, Barbara and Heather Sarsons, "Flexible Wages, Bargaining, and the Gender Gap," *The Quarterly Journal of Economics*, 2022, *137* (1), pp. 215–266.

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.

_ **, Ana Rute Cardoso, and Patrick Kline**, "Bargaining, Sorting, and the Gender Wage Gap: Quantifying the Impact of Firms on the Relative Pay of Women," *The Quarterly journal of economics*, 2016, *131* (2), pp. 633–686.

Chu, Ben, "Company Gender Pay Gap Reporting Is Already Helping, and Will Continue to Help Tackle Inequalities," *The Independent*, 2018.

Cullen, Zoë and Bobak Pakzad-Hurson, "Equilibrium Effects of Pay Transparency," *Econometrica*, 2023, *91* (3), 765–802.

_ **and Ricardo Perez-Truglia**, "How Much Does Your Boss Make? The Effects of Salary Comparisons," *Journal of Political Economy*, 2022, *130* (3), 766–822.

_ **and** _ , "The Salary Taboo: Privacy Norms and the Diffusion of Information," *Journal of Public Economics*, 2023, *222*, 104890.

Cullen, Zoe B, "Is Pay Transparency Good?," NBER Working Paper No. w31060, National Bureau of Economic Research 2023.

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

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.

Duchini, Emma, Stefania Simion, and Arthur Turrell, "A Review of the Effects of Pay Transparency," *Oxford Research Encyclopedia of Economics and Finance*, 2023, *Forthcoming.*

Duncan, Pamela, Thomas Tobi, and Alexandra Topping, "Gender pay gap at UK's biggest firms is growing, data suggests," *The Guardian*, 2021.

Flinn, Christopher and Joseph Mullins, "Firms' Choices of Wage-Setting Protocols," Technical Report, Discussion paper, New York University 2021.

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

Gulyas, Andreas, Sebastian Seitz, and Sourav Sinha, "Does Pay Transparency Affect the Gender Wage Gap? Evidence from Austria," *American Economic Journal: Economic Policy*, 2023, *15* (2), 236–255.

Hall, Robert E and Alan B Krueger, "Evidence on the Incidence of Wage Posting, Wage Bargaining, and On-the-job Search," *American Economic Journal: Macroeconomics*, 2012, *4* (4), pp. 56–67.

Hawkins, Oliver and Daniel Thomas, "Nearly 80more than women," *Financial Times*, 2023.

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.

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.

Li, Gloria and Carlos Mulas Granados, "The Recent Decline in United Kingdom Labor Force Participation: Causes and Potential Remedies," 2023.

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

Marinescu, Ioana and Ronald Wolthoff, "Opening the Black Box of the Matching Function: the Power of Words," *Journal of Labor Economics*, 2020, *38* (2), 535–568.

Martin, Josh, "Homeworking Hours, Rewards and Opportunities in the UK: 2011 to 2020," *ONS*, 2020.

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

OECD, *Reporting Gender Pay Gaps in OECD Countries* 2023.

Pedregosa, F., G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay, "Scikit-learn: Machine Learning in Python," *Journal of Machine Learning Research*, 2011, *12*, 2825–2830.

Perez-Truglia, Ricardo and Ugo Troiano, "Shaming Tax Delinquents: Theory and Evidence from a Field Experiment in the United States," *Journal of Public Economics*, 2018, *167*, 120–137.

Roussille, Nina, "The Central Role of the Ask Gap in Gender Pay Inequality," *Working Paper*, 2020.

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

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.

Wright, Randall, Philipp Kircher, Benoît Julien, and Veronica Guerrieri, "Directed Search and Competitive Search Equilibrium: A Guided Tour," *Journal of Economic Literature*, 2021, *59* (1), pp. 90–148.

7 Figures and Tables

Figure 1: Institutional Setting

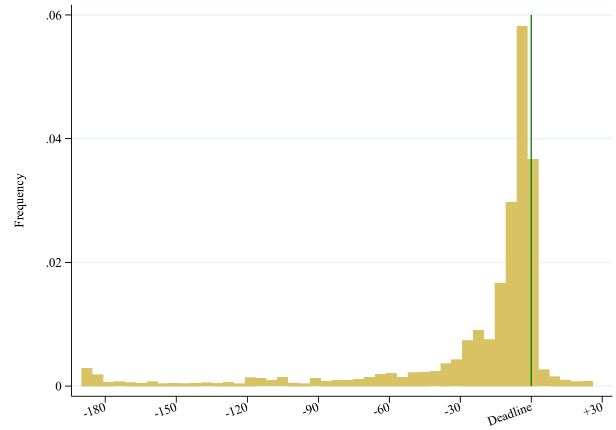

(A) Distribution of submission dates

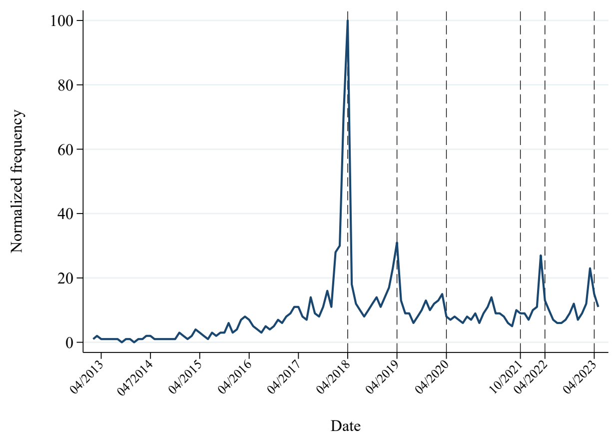

(B) Google searches for "gender pay gap"

Source: UK Government Equalities Office 2018-2023; Google, 2013-2023.
Notes: The figures in Panel A show the distribution of days when firms published their gender equality indicators, relative to the deadline. The bottom and top 5 percent of the distribution are not displayed. The graph in Panel B reports the UK relative search volume for the term "gender pay gap" between April 2013 and April 2023 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

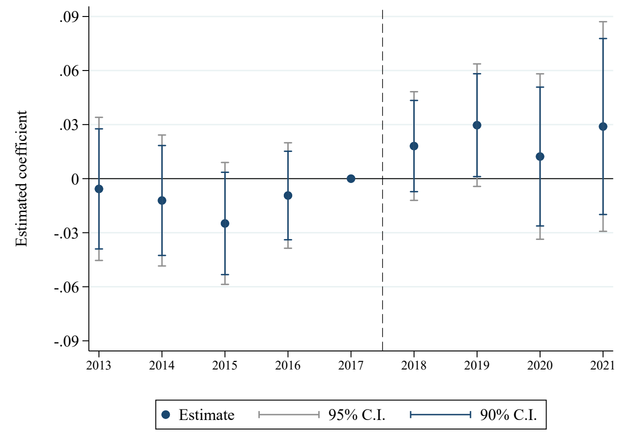

(A) Gender pay gap

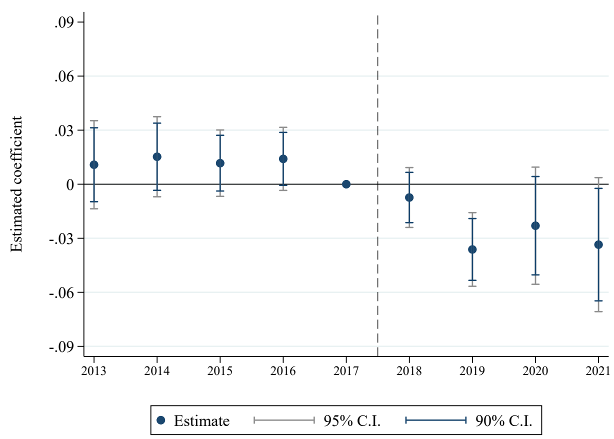

(B) Men

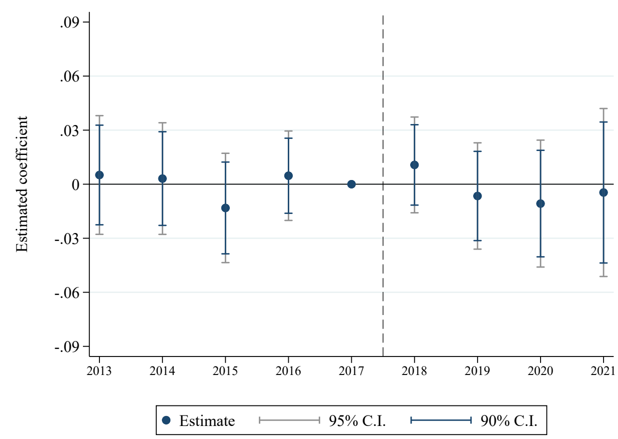

(C) Women

Source: ASHE, 2013–2021.
Notes: These graphs present the estimates of the leads and lags of the policy on the gender pay gap (Panel A), and men's and women's pay (Panel B and C, respectively). These results are obtained from the estimation of regression 2. In each graph, the estimation sample includes workers employed in firms with 200 to 300 employees. The graphs also report 90 and 95 percent confidence intervals associated with firm-level clustered standard errors. The dash vertical line indicates the month when the mandate is approved, i.e., February 2017.

Figure 3: Event studies - other pay outcomes

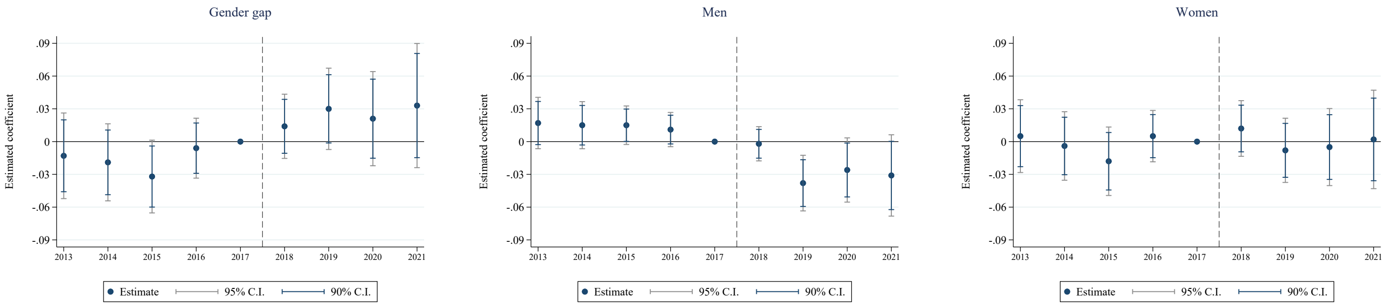

(A) Log hourly basic pay

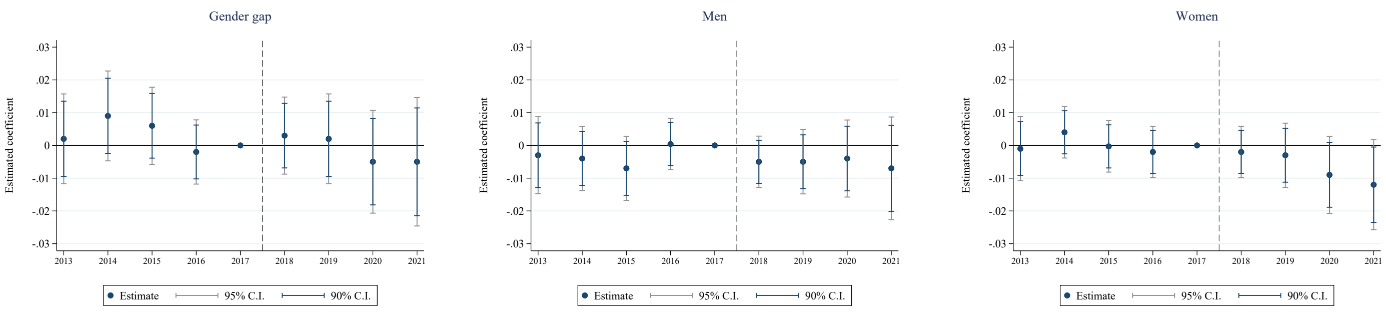

(B) Additional payments / base pay

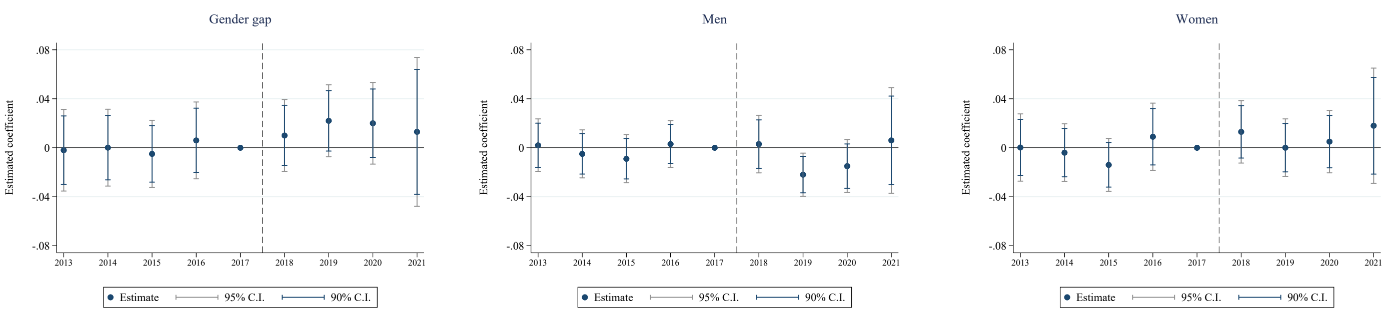

(C) Promotion

Source: ASHE, 2013–2021.
Notes: These graphs present the estimates of the leads and lags of the policy on different pay outcomes. These results are obtained from the estimation of regression 2. In each graph, the estimation sample includes workers employed in firms with 200 to 300 employees. The graphs also report 90 and 95 percent confidence intervals associated with firm-level clustered standard errors. The dash vertical line indicates the month when the mandate is approved, i.e., February 2017.

Figure 4: Robustness checks - gender pay gap

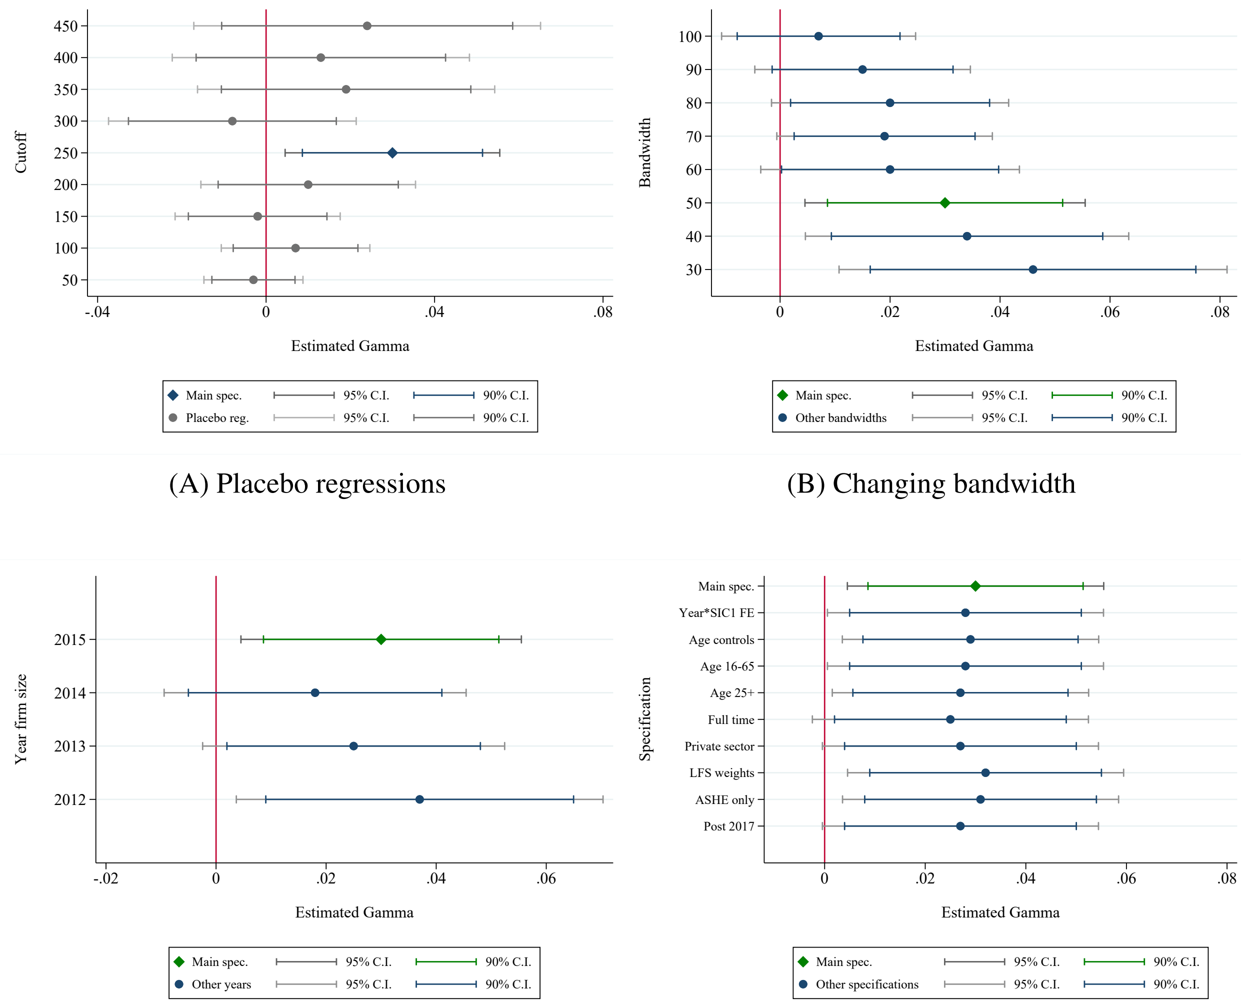

(A) Placebo regressions

(B) Changing bandwidth

(C) Changing year treatment status

(D) Other robustness checks

Source: ASHE, 2013–2021.
Notes: These graphs present a series of robustness checks on the impact of the policy on the gender pay gap. Detailed results are presented in Appendix Tables A5-A8.

Figure 5: Gender pay gap distribution over time

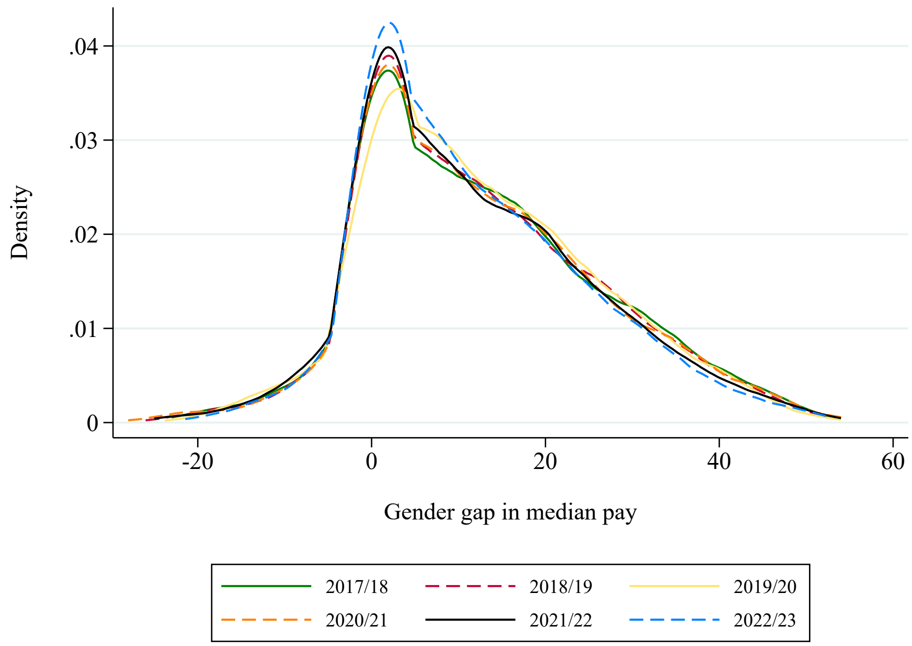

Source: UK Government Equalities Office, 2018-2023.
Notes: This figure plots the gender pay distribution over time. The data are drawn from the Gender Pay Gap Reporting website. The sample is restricted to a balanced sample of firms that publish equality indicators in all of the following years: 2018, 2019, 2021, 2022, and 2023. Note that this excludes 2020, as the disclosure mandate was lifted in March 2020 due to the COVID-19 pandemic, and by then only half of targeted firms had published their data for that year. Outliers (bottom and top 1 percent) are excluded from the graphs.

Table 1: Public gender equality indicators

	2017-18 (1)	2018-19 (2)	2019-2020 (3)	2020-21 (4)	2021-2022 (5)	2022-23 (6)
Median gender hourly pay gap	12 (16)	12 (16)	13 (15)	13 (17)	12 (17)	12 (15)
Mean gender hourly pay gap	14 (15)	14 (14)	14 (15)	14 (15)	14 (15)	13 (14)
Median gender bonus gap	-22 (1,400)	-1 (295)	10 (112)	3 (271)	1 (289)	7 (167)
Mean gender bonus gap	8 (834)	18 (219)	27 (81)	20 (349)	19 (151)	21 (164)
% men receiving bonus	35 (36)	36 (37)	37 (37)	36 (37)	36 (38)	39 (38)
% women receiving bonus	34 (36)	34 (36)	36 (37)	35 (37)	35 (38)	38 (37)
% women lower quartile	54 (24)	54 (24)	55 (24)	55 (24)	55 (24)	55 (24)
% women lower-middle quartile	49 (26)	50 (26)	50 (26)	50 (26)	50 (26)	51 (26)
% women upper-middle quartile	45 (26)	46 (26)	46 (26)	46 (26)	46 (26)	46 (26)
% women top quartile	39 (24)	40 (24)	40 (24)	40 (24)	40 (24)	41 (24)
Observations	10,557	10,812	6,978	10,152	10,529	10,408

Source: UK Government Equalities Office, 2018-2023.

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 (£)	15.94	15.59	13.36	13.39
	(14.24)	(11.68)	(8.87)	(10.70)
Weekly pay (£)	581.73	569.38	414.52	411.71
	(533.46)	(429.76)	(307.33)	(316.99)
Weekly hours	36.41	36.67	30.69	30.49
	(8.54)	(8.50)	(10.53)	(10.69)
Receiving additional payments	0.29	0.29	0.19	0.18
	(0.45)	(0.46)	(0.39)	(0.38)
Additional payments per week (£)	26.05	26.08	10.07	9.47
	(102.38)	(114.65)	(38.72)	(42.15)
Additional payments ph/Hourly base pay	0.04	0.04	0.02	0.02
	(0.10)	(0.10)	(0.08)	(0.07)
Promotion	0.02	0.02	0.02	0.02
	(0.14)	(0.13)	(0.14)	(0.14)
Bottom tercile	0.26	0.25	0.37	0.37
	(0.44)	(0.43)	(0.48)	(0.48)
Middle tercile	0.33	0.33	0.23	0.26
	(0.47)	(0.47)	(0.42)	(0.44)
Top tercile	0.42	0.42	0.40	0.37
	(0.49)	(0.49)	(0.49)	(0.48)
Tenure in months	86.22	84.94	73.13	70.73
	(97.05)	(96.06)	(80.05)	(79.61)
Leaving firm in t+1	0.28	0.28	0.29	0.28
	(0.45)	(0.45)	(0.45)	(0.45)
Private sector	0.91	0.92	0.80	0.78
	(0.29)	(0.27)	(0.40)	(0.41)
Covered by collective agreement	0.28	0.27	0.32	0.34
	(0.45)	(0.44)	(0.47)	(0.47)
Observations	6,910	8,677	5,868	7,710

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 groups, before the implementation of the policy. The variables bottom, middle, and top tercile are three dummies variables that are equal to 1 if a worker is employed, respectively, in the bottom, middle, or top tercile of the occupational distribution, based on the ranking of pre-policy 1-digit SOC-specific median wages.

Table 3: Impact on pay outcomes

	Log hourly pay	Log hourly basic pay	Additional payments / base pay	Promotion
	(1)	(2)	(3)	(4)
Treated firm*post	-0.029***	-0.028***	-0.003	-0.006
	(0.009)	(0.009)	(0.004)	(0.007)
Treated firm*post*fem	0.030**	0.032**	-0.001	0.015
	(0.013)	(0.013)	(0.005)	(0.010)
Observations	35,092	35,092	34,930	35,092
Adjusted R^2	0.894	0.897	0.529	0.005
P-value Women Coeff	0.909	0.656	0.215	0.231
Men's pre-policy mean	15.94	15.26	0.04	0.02
Women's pre-policy mean	13.36	13.02	0.02	0.02

Source: ASHE, 2013–2021.

Notes: This table reports the impact of pay transparency on pay outcomes, obtained from the estimation of regression 1. Each column refers to a different outcome, as specified at the top of it. The estimation sample comprises men and women working in firms that have between 200 and 300 employees. All regressions include firm*individual fixed effects and gender-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. 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 sum of the two reported coefficients, corresponding to the effect of the policy on female employees. 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 4: Compositional effects

	New hire (1)	Leaving firm in t+1 (2)	Bottom tercile (3)	Middle tercile (4)	Top tercile (5)	Log hourly pay (6)
Treated firm*post	0.015	-0.010	-0.010	0.015	-0.006	-0.012
	(0.012)	(0.017)	(0.013)	(0.014)	(0.015)	(0.014)
Treated firm*post*fem	-0.003	0.044**	-0.016	-0.006	0.022	0.021
	(0.014)	(0.019)	(0.020)	(0.021)	(0.022)	(0.021)
Observations	46,098	44,367	48,589	48,589	48,589	48,589
Adjusted R^2	0.126	0.236	0.440	0.372	0.415	0.490
P-value Women Coeff	0.348	0.053	0.075	0.557	0.302	0.534
Men's pre-policy mean	0.19	0.28	0.26	0.33	0.42	15.94
Women's pre-policy mean	0.22	0.29	0.37	0.23	0.40	13.36

Source: ASHE, 2013–2021.

Notes: This table reports the compositional effects of the pay transparency policy, obtained from the estimation of regression 1 with firm fixed effects in place of firm times individual fixed effects. Each column refers to a different outcome, as specified at the top of it. The outcomes in Columns 3-5 are dummy variables that are equal to one if an employee works, respectively, in the bottom, middle, or top tercile of the occupational distribution, based on the ranking of pre-policy 1-digit SOC-specific median wages. The estimation sample comprises men and women working in firms that have between 200 and 300 employees. All regressions include gender-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. 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 sum of the two reported coefficients, corresponding to the effect of the policy on female employees. 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 5: Gender pay gap and placement in YouGov Rankings

	Women's Rankings (1)	Workforce Rankings (2)
Gender pay gap	-0.681*	-0.789**
	(0.380)	(0.375)
Observations	1,807	1,813
Adjusted R^2	0.659	0.707

Source: UK Government Equalities Office, YouGov, 2018–2019.

Notes: This table shows the correlation between firms' gender gap in the median hourly pay and, respectively, firms' placement in YouGov Women's Rankings (Column 1), and YouGov Workforce Rankings (Column 2). The gender pay gap is expressed relative to men's pay. In the regression sample, it is comprised between -41 and 76, with a mean of 11 and a standard deviation of 13 percentage points. Firms' placement in YouGov Rankings is an integer that runs from 1 for the lowest-ranked firm to 1342 for the highest-ranked firm. Both regressions include year and GPG firm fixed effects. Standard errors are clustered at the level of the GPG company. In each column, the sample includes YouGov Rankings' firms that either publish directly or have a parent company that publishes gender equality indicators in at least two consecutive years. Each year, data for YouGov Women's rankings are missing for 3 firms, compared to the list of employers included in YouGov Workforce Rankings. Both regressions exclude GPG firms that publish gender equality indicators after the end of the calendar year when these were due.

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

Table 6: Impact on log hourly pay by advertising costs

	Entire sample	Below-median advertising-to-sales ratio	Above-median	P-value T-test
	(1)	(2)	(3)	(4)
Treated firm*post	-0.028***	-0.016	-0.039**	0.237
	(0.010)	(0.011)	(0.017)	
Treated firm*post*fem	0.030**	-0.001	0.052**	0.232
	(0.015)	(0.016)	(0.023)	
Observations	27,359	13,793	13,566	
Adjusted R^2	0.890	0.900	0.882	
P-value Women Coeff	0.814	0.342	0.461	
Men's pre-policy mean	15.36	15.02	15.78	
Women's pre-policy mean	12.84	12.16	13.34	

Source: ASHE, 2013–2021.

Notes: This table compares the impact of pay transparency on employees' hourly pay across firms with different advertising costs, by estimating regression 1 by subgroup. Specifically, the first column reports the estimate for the entire sample, employees working in firms that have between 200 and 300 employees and have non-missing pre-policy advertising costs in ABS. Columns 2 and 3 compare the impact across firms with below- and above-median advertising-to-sales ratios. Column 4 reports the p-value of the t-test on the equality of estimates in Columns 2 and 3. All regressions include firm*individual fixed effects and gender-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. 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 sum of the two reported coefficients, corresponding to the effect of the policy on female employees. The pre-policy mean represents the mean of the outcome variable for the treated group and subgroup considered between 2013 and 2017.

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