Who works for whom and the UK gender pay gap?

Jewell, Sarah and Razzu, Giovanni and Singleton, Carl

University of Reading, University of Reading, University of Edinburgh

1 April 2018

Who works for whom and the UK gender pay gap?[*]

Sarah Louise Jewell[†]   Giovanni Razzu[‡]   Carl Singleton[§]

April 2018
Link to latest version

Abstract

This study reports novel facts about the UK gender pay gap. We use a large, longitudinal, representative and employer-employee linked dataset for the years 2002-16. Men's average log hourly wage was 22 points higher than women's in this period. We ask how much of this gap is accounted for by the differences in whom men and women worked for; how much is explained by the relative wage premiums that firms paid their employees, after adjusting for the influence of other factors, such as occupations and tenure? The answer is less than 1 percentage point, or about one eighteenth of the adjusted hourly gender pay gap. We also find that the allocation of men and women to occupations was as unimportant as how workers were allocated to firms. These results show that in the United Kingdom what happens within firms and occupations is far more important than what jobs men and women have. Therefore, attention should focus on why men and women within UK firms tend to receive different rates of pay.

Keywords: gender wage gap; firm-specific wages; occupation premiums
JEL codes: J16; J31; J70

1 Introduction

The gap between the average hourly pay of men and women in the United Kingdom stood at 17.4% in April 2017.[1] Since April 2018, all British employers with 250 employees or more must annually publish their own gender pay gaps.[2] This recent change in UK equalities legislation has placed firms right at the centre of public and policy-making attention on the issue. More generally, there is a renewed focus on identifying how firms in several countries have shaped and driven wage inequality patterns, of which the gender pay gap is one aspect.[3] Specifically for the United Kingdom, Faggio et al. (2010) found that large increases in earnings inequality throughout the 1980s and early 1990s were driven by increasing dispersion in the average wages paid by firms. Although Mueller et al. (2017) and Schaefer and Singleton (2017) have since described how the majority of British wage inequality in the past two decades still lies within firms, nonetheless, a sizeable fraction is accounted for by between-firm differences. In this paper, we ask what role firm-specific wage premiums have in accounting for the UK gender pay gap.

To answer this question, we use the UK-representative Annual Survey of Hours and Earnings (ASHE) datasets for the years 2002-16. We estimate a wages model which allows for worker- and firm-specific wage effects, following the contributions of Abowd et al. (1999) and Abowd et al. (2002). Then, using an application of the Gelbach (2016) decomposition, we account for how much of an observable-covariate-adjusted difference in the gender averages of log real hourly pay was due to the distribution of workers over firms. In other words, we estimate the role of which workers were employed by which firms between 2002 and 2016 for the gender wage gap. The observed mean hourly pay gap (male minus female) among all employees, aged 25-64, was 22.3 log points in our analysis sample. Out of an adjusted gap of 14.5 log points, an estimated share of just 6% was contributed by the allocation of employees over firm wage premiums. These 0.8 log points tell us how far the UK hourly pay gap would close if men and women were equally distributed across firms, conditional on other relevant observable factors for wage determination, such as occupations, tenure, full-time work status and industry sectors. We also find that the allocation of men and women to occupations contributed a similarly small amount to gender wage inequality. Retail and hospitality services was

---

[1]Excluding overtime, where the gap is measured as the ratio of the difference between male and female pay over male pay. The preferred measure of the UK Government is median wages rather than the mean and excludes part-time workers. The median gap for full-time employees was 9.1% in 2017. The median gap for all employees, including part-time, was 18.4% in 2017. Source: ONS, ASHE Total Table 1.6a, Hourly Pay excluding overtime; https://www.ons.gov.uk/...

[2]Legislation titled "The Equality Act 2010 (Gender Pay Gap Information) Regulations 2017"; https://www.legislation.gov.uk/... Public sector employers throughout the UK already had duties to report pay gaps, following from the Equality Act 2010.

[3]See for examples, using new, large-scale, longitudinal and employer-employee linked datasets: Card et al. (2013) for Germany; Song et al. (2016) and Barth et al. (2016) for the United States; and Card et al. (2018) for Portugal and a summary of this literature.

the only broad industry group where women were to any great extent disproportionately employed by low-wage firms, compared with the men working in the same industry. Men and women aged 25-34 were similarly distributed across low- and high-wage firms. But for those aged 35-44, gaps in representation open up, with men then being relatively more likely to work for high-wage firms. But overall, in order to identify the causes of the UK hourly gender pay gap, the focus should be firmly placed on what determines the wage differences within firms.

A vast literature has studied the determinants of gender pay gaps (see for reviews Altonji and Blank, 1999; Weichselbaumer and Winter-Ebmer, 2005; Blau and Kahn, 2017). Explanations for the labour market differences between men and women are typically grouped into three broad categories: productivity, preferences and discrimination, which are all interrelated (Altonji and Blank, 1999). However, with the diminishing of gender gaps in the majority of developed countries, the importance and focus on explanations from the first category, especially human capital-based ones, has lessened. Nonetheless, pay gaps persist and are pervasive. More recent work has looked to gender differences in preferences and psychological attributes, and how these impact on productivity, choices and beliefs (see for reviews Croson and Gneezy, 2009; Bertrand, 2011; Azmat and Petrongolo, 2014). The role of firms, in particular where men and women work, cuts across across these sets of explanations. Early work found that US women were more likely to work for lower wage firms than men, and vice versa regarding higher wage firms (Blau, 1977; Groshen, 1991; Bayard et al., 2003). More recent studies have found that low wage growth within an establishment for women plays a bigger role in the US gender pay gap than how women are (not) sorted into higher wage firms (Goldin et al., 2017; Barth et al., 2017).

The only papers to have previously looked at the role of where UK men and women work, and whether this could explain part of the overall pay gap, are Mumford and Smith (2007, 2009). These authors used cross-sectional data from the 2004 Workplace Employment Relations Survey, to disentangle the influence of observable employer and employee characteristics. They showed that the proportions of women relative to men in occupations and workplaces did account for part of the UK gender pay gap. However, these studies, as well as those for other countries that have used a similar method, were hampered by not being able to account for unobservable heterogeneity over workers and firms in the determination of wages. Therefore, their results cannot be directly compared with what we find here.

Similarly, our main findings are not directly comparable with the majority of recent UK gender pay gap studies, which have used sources of longitudinal household survey data. One of these studies, which made a significant contribution to our understanding of gender pay inequality more generally, was by Manning and Swaffield (2008). They looked at what determined the differences between men and women in early career

wage growth, asking why the gender pay gap started out small, but then grew over the life-cycle as workers accumulated experience. They found that over 50% of the pay gaps that emerged in this way in the United Kingdom could not be explained, and that the occupational allocation and changes of workers explained only a small fraction overall. In an up-to-date study, using one of these longitudinal household data sources, Costa Dias et al. (2018) demonstrate the importance of accumulated work experience in determining pay gaps. They find that among UK college graduates, the majority of the gender pay gap twenty years after the first childbirth can be explained by differences in work experience, mostly through accumulated working hours. Another recent contribution to our understanding of UK gender pay inequality comes from Fortin et al. (2017). They used the ASHE, like ourselves, but did so to document the differences in male and female representation throughout the overall employee annual earnings distribution, with particular interest in the relative lack of women at the top. They then used the household-based UK Labour Force Survey to estimate how much of the average hourly gender pay gap these differences in representation could account for. They found that the under-representation of women in the top centiles of employee pay could explain as much as a half of the pay gap in 2015.

However, none of these aforementioned UK-based studies could control robustly for the potential influence of how male and female employees were allocated across firms. It is plausible that omitting this factor, which generally explains a large and significant fraction of overall UK wage variation, could have affected their results. For example, it is well-known that measured returns to tenure and work experience are upwards biased unless the unobserved worker-firm-match quality is controlled for (Topel, 1991; Snell et al., 2018). However, our main findings show that the way in which employees are assigned to firms can explain only a small fraction of the adjusted UK hourly pay gap. Therefore, we add support to some of these previous conclusions about the determinants of gender pay gaps, which were based on UK household survey data from a similar period, but which could not have addressed this potential role of firms to begin with.

More direct comparisons of our main findings are possible with a few recent studies of pay gaps in other countries.[4] In particular, Cardoso et al. (2016) looked at how much of the Portuguese hourly pay gap could be accounted for by the allocation of men and women over establishments and job-titles. They found that these two factors could each explain around a fifth of an adjusted measure of the Portuguese pay gap over three decades. Our methodological approach is close to that of Cardoso et al., though we expand on it. We use the Gelbach decomposition to identify the role of unobservable worker and firm fixed factors in the pay gap, after also accounting for the influence of some observable fixed characteristics on wages, such as employee birth cohorts or employer industry sectors. If we exclude some relevant time-varying covariates from our UK wages model, then we substantially overestimate the contribution to the adjusted pay

---

[4]see Section 4.3 & Table 6 for a full comparison with these relevant studies.

gap from firm fixed effects. We suggest this as a plausible explanation for the order of magnitude difference between our own results and those from Cardoso et al. A second potential explanation is that our estimated firm-specific wage premiums explain less of the overall log hourly wage variance in the United Kingdom (approx. 12-13%) than similarly obtained estimates do in Portugal (approx. 20%: Card et al., 2016, 2018).

As the April 2018 deadline for large UK firms to publish their pay gaps approached and passed, there was significant interest and scrutiny of the reported figures among the British public and media.[5] The UK government stated that the new legislation's motivation was to highlight gender pay inequality and discrimination, and to encourage firms to tackle these issues within the workplace.[6] But the policy change could also lead to greater regulatory burden and stress on employers, as emphasised in the government's many regulatory impact assessments on the policy proposal.[7] Despite these potential economic costs of the legislation, we could find no robust empirical evidence, representative of the whole UK labour market, that showed the gender pay gap was not an issue of differences in pay between firms, rather than within them. It will be impossible to address this evidence gap using the pay gap data actually collected from firms under the new legislation. Quite simply, there is no robust way to address how much of these reported pay gaps, and the differences therein between firms, are explained by workforce composition. Therefore with this paper, we address an important gap in the current UK equalities policy evidence base, on whether requiring firms to publish their own pay gaps is relevant to what explains the overall UK gender pay gap.

The remainder of the paper proceeds as follows: Section 2 describes the method and our employer-employee panel sample construction; Section 3 presents the main results; Section 4 discusses the robustness of our findings, expands on the possible role of occupations and compares these UK results with those from elsewhere; and Section 5 concludes.

---

[5]See for example these articles, which appeared high on the BBC news website and featured among the most read articles among visitors at that time: "Gender pay gap: More than 500 firms reveal their figures" (6 January, 2018; http://www.bbc.co.uk/news/...); "Gender pay gap: Men still earn more than women at most firms" (21 February, 2018; http://www.bbc.co.uk/news/...).

[6]See for example the joint Government Equalities Office and Chartered Institute of Personnel and Development guidance document published for employers in December 2017: "GENDER PAY GAP: Closing it together - Actions for employers"; https://www.gov.uk/...

[7]There have been several, since before the enactment of the Equality Act 2010. For example, the Regulatory Impact Assessment of April 2016, relating to this new legislation, suggested estimates of the net economic costs to private sector employers of £3.8million per annum (which figure was based on seemingly conservative assumptions compared to the actual consultation responses published by and collected from businesses); https://www.gov.uk/...

2 Method and Data

We estimate a so-called AKM-type wages model, which features both worker and firm fixed effects (Abowd et al., 1999). We then apply a decomposition method suggested by Gelbach (2016), to account for how much of the UK gender pay gap is accounted for by which of these firm fixed wage effects men and women receive. This method can address the role in the pay gap of whom employees work for, after conditioning on the influence of some observable worker, firm and job characteristics. Cardoso et al. (2016) have identified and described thoroughly the usefulness of this method for studying the role of firm fixed factors in the conditional gender pay gap. Therefore, our treatment here is concise. However we suggest one expansion on Cardoso et al. We emphasise that the Gelbach decomposition allows us to address the observable fixed characteristics of workers and firms, such as birth cohort or industry sector, even in an estimated wages model which includes fixed effects, before attributing the remaining pay gap to unobservable fixed factors, such as worker preferences or firm productivities.

The Gelbach decomposition quantifies how much of the change in some coefficient of interest, in an estimated linear regression model, is due to adding further covariates to the model, without concern for the order in which those covariates are added. In the same terminology used by Gelbach, the 'Full' model includes the coefficient of interest, which in our application is the marginal effect on wages of gender, and a full set of covariates. Our AKM-type Full wages model is given by

$$w_{it} = \mathbf{x}'_{it}\boldsymbol{\beta} + \alpha_i + \phi_{J(it)} + \varepsilon_{it} \, ,$$

or in stacked matrix notation as

$$\mathbf{y} = \mathbf{X}\boldsymbol{\beta} + \mathbf{A}\boldsymbol{\alpha} + \mathbf{F}\boldsymbol{\phi} + \boldsymbol{\varepsilon} \, . \tag{1}$$

In this equation, $\mathbf{y}$ gives an $(Nx1)$ vector containing the natural logarithm of the real hourly wage, w_{it}, for individual i in period t. For the remainder of the paper, any reference to 'log' wages concerns w_{it}. Similarly, any reference to a 'pay (wage) gap' concerns the difference in log real hourly wages between men and women. The (Nxk) matrix $\mathbf{X}$ and the vector $\mathbf{x}_{it}$ contain k time-varying covariates, with associated $(kx1)$ coefficient vector $\boldsymbol{\beta}$. The (NxP) and (NxL) matrices $\mathbf{A}$ and $\mathbf{F}$ are designs for the P workers and L firms covered by the model, respectively. The worker fixed effects, α_i, are contained within the $(Px1)$ vector $\boldsymbol{\alpha}$. Similarly, the firm fixed effects $\phi_{J(it)}$ are contained in the $Lx1$ vector $\boldsymbol{\phi}$, where J is a function denoting whether worker i in period t is employed by firm j. Finally, the vector $\boldsymbol{\varepsilon}$ contains the N error terms, ε_{it}, which are assumed to have the standard properties $E[\boldsymbol{\varepsilon} \mid \mathbf{X}, \mathbf{A}, \mathbf{F}] = 0$.

AKM-type models intrinsically concern firm switching. The worker fixed effects, α_i, are transferable, affecting an employee's wages to the same degree wherever and

whenever he works, and in whatever job. The firm fixed effects, $\phi_{J(it)}$, measure relative wage premiums, which employees receive upon switching firms. The Full wages model then estimates the systematically higher or lower wages that firms pay relative to other firms, for whatever reason. If such relative firm wage premiums do exist in the labour market, we can ask whether or not men or women are disproportionately benefiting from their existence, and if so, quantify their importance in terms of the gender pay gap.

To demonstrate our application of the Gelbach decomposition, it is useful if we re-write the Full model as follows:

$$\mathbf{y} = \mathbf{X}\boldsymbol{\beta} + \underbrace{\mathbf{g}\widetilde{\lambda} + \mathbf{W}\boldsymbol{\beta_w} + \mathbf{A}\widetilde{\boldsymbol{\alpha}}}_{=\mathbf{A}\boldsymbol{\alpha}} + \underbrace{\mathbf{Z}\boldsymbol{\beta_z} + \mathbf{F}\widetilde{\boldsymbol{\phi}}}_{=\mathbf{F}\boldsymbol{\phi}} + \boldsymbol{\varepsilon} \, , \tag{2}$$

where $\mathbf{g}$ contains a dummy variable for whether or not an individual is male, and $\widetilde{\lambda}$ measures the wage gap conditional on the other factors in the model. The (Nxp) and (Nxl) matrices $\mathbf{W}$ and $\mathbf{Z}$ contain p and l time-invariant observable worker and firm characteristics, respectively. Therefore, $\widetilde{\boldsymbol{\alpha}}$ and $\widetilde{\boldsymbol{\phi}}$ measure the effects on employee wages of all unobserved (or residual) fixed worker and firm factors. But the coefficients contained in $\{\widetilde{\lambda}, \boldsymbol{\beta_w}, \widetilde{\boldsymbol{\alpha}}, \boldsymbol{\beta_z}, \widetilde{\boldsymbol{\phi}}\}$ cannot all be separately identified using least squares estimation of this single-equation model.

Still keeping to the terminology of Gelbach, the 'Basic' model omits the particular covariates whose effects in the Full model on the estimated coefficient of interest we wish to account for, which here is the observable-covariate-adjusted gender pay gap (henceforth referred to as the 'adjusted gender pay gap'). Our Basic wages model is given by

$$w_{it} = g_i \lambda + \widetilde{\mathbf{x}}'_{it}\widetilde{\boldsymbol{\beta}} + e_{it} \, ,$$

or in stacked matrix notation as

$$\mathbf{y} = \mathbf{g}\lambda + \widetilde{\mathbf{X}}\widetilde{\boldsymbol{\beta}} + \mathbf{e} \, . \tag{3}$$

The adjusted gender pay gap in (3) is given by the coefficient λ. The covariates are contained in the $(Nx[k + p + l])$ matrix $\widetilde{\mathbf{X}} = [\mathbf{X}, \mathbf{W}, \mathbf{Z}]$. Although the effects of each factor in $\mathbf{W}$ and $\mathbf{Z}$ cannot be identified in the single-equation Full model, this is not a good reason to exclude them from our Basic wages model. In a standard exploration of the gender pay gap and estimation of a wages model, such as those which do not use longitudinal employer-employee data, we would anticipate that time-invariant factors could be significant, such as an individual's birth cohort, or whether a firm is in the private sector. This is the main way in which we differ from the application of the Gelbach decomposition by Cardoso et al., who only included a set of k time-varying

covariates in their Basic model. The vector **e** contains the N error terms, e_{it}, which are again assumed to have standard properties.

The Gelbach decomposition is then applied by pre-multiplying components of the estimated version of the Full wages model, equation (1), with the $(1xN)$ row vector

$$
\begin{aligned}
\boldsymbol{\gamma}' &= \left(\mathbf{g}'\left[\mathbf{I} - \widetilde{\mathbf{X}}(\widetilde{\mathbf{X}}'\widetilde{\mathbf{X}})^{-1}\widetilde{\mathbf{X}}'\right]'\left[\mathbf{I} - \widetilde{\mathbf{X}}(\widetilde{\mathbf{X}}'\widetilde{\mathbf{X}})^{-1}\widetilde{\mathbf{X}}'\right]\mathbf{g}\right)^{-1}\mathbf{g}'\left[\mathbf{I} - \widetilde{\mathbf{X}}(\widetilde{\mathbf{X}}'\widetilde{\mathbf{X}})^{-1}\widetilde{\mathbf{X}}'\right]' \\
&= (\boldsymbol{\Gamma}'\boldsymbol{T})^{-1}\boldsymbol{\Gamma}' \,,
\end{aligned}
\tag{4}
$$

where **I** is the (NxN) identity matrix. Pre-multiplying some variable, contained within an $(Nx1)$ vector, by $\boldsymbol{\gamma}'$ computes the gender gap in the average values of that variable, conditional on the estimated linear additive effects of any other variables contained within $\widetilde{\mathbf{X}}$: i.e. it 'adjusts' the observed gender gap. If we replaced $\widetilde{\mathbf{X}}$ in (4) with a vector of 1's, then pre-multiplying **y** by $\boldsymbol{\gamma}'$ would give the actual mean gender pay gap observed in the data. Using (4), we can write the following decomposition equation for the components of the adjusted pay gap:

$$
\underbrace{\boldsymbol{\gamma}'\mathbf{y}}_{\widehat{\delta_y}=\widehat{\lambda}\text{ - Adjusted gender pay gap}} = \underbrace{\boldsymbol{\gamma}'\mathbf{A}\widehat{\boldsymbol{\alpha}}}_{\widehat{\delta_a}\text{ - dist. worker effs}} + \underbrace{\boldsymbol{\gamma}'\mathbf{F}\widehat{\boldsymbol{\phi}}}_{\widehat{\delta_f}\text{ - gender-firm sorting}} + \underbrace{\boldsymbol{\gamma}'\widehat{\boldsymbol{\varepsilon}}}_{=0} \,.
\tag{5}
$$

The term on the left-hand-side of (5) is an estimate of the adjusted gender pay gap, and is equivalent to the least squares estimate $\widehat{\lambda}$ of the Basic model (3). The first term on the right-hand-side of (5), $\widehat{\delta_a}$, gives an estimate of how much of the adjusted gender pay gap is accounted for by the gendered distributions of the unobservable worker-fixed effects (worker-specific and time-invariant heterogeneity), after partialling out the observable worker, job or firm characteristics in $\widetilde{\mathbf{X}}$. Similarly, $\widehat{\delta_f}$ estimates how much of the adjusted pay gap is accounted for by who works for whom, i.e. by the gendered conditional distributions of the estimated unobservable firm-fixed effects. Henceforth, we refer to this as the 'gender-firm sorting effect'.[8] We can also write this effect as follows:

$$
\widehat{\delta_f} = E_{it}\left[\widetilde{\phi}_{J(it)} \mid i \in M, \widetilde{\mathbf{X}}\widetilde{\boldsymbol{\beta}}\right] - E_{it}\left[\widetilde{\phi}_{J(it)} \mid i \in F, \widetilde{\mathbf{X}}\widetilde{\boldsymbol{\beta}}\right] \,,
\tag{6}
$$

where M and F denote the sets of men and women in N, respectively. The final term in (5) is equal to zero by the standard assumptions of the AKM-type Full wages model: $E\left[\widehat{\boldsymbol{\varepsilon}} \mid \widetilde{\mathbf{X}}\widetilde{\boldsymbol{\beta}}\right] = 0$. If in the expression for $\boldsymbol{\gamma}'$ we replace all occurrences of $\widetilde{\mathbf{X}}$ with **X** (i.e. if

---

[8]We use the term sorting here loosely, and only to make our discussion of the results more concise. We caveat this label because of recent theoretical (e.g. Eeckhout and Kircher, 2011) and quantitative (e.g. Lopes de Melo, 2018) contributions, which have emphasised the difficulties in identifying general sorting patterns in the labour market using wages data alone; i.e. there is a danger in inferring too much about the extent of sorting from correlations between estimates of worker and firm fixed effects.

we assume that the Basic wages model does not include any time-invariant factors: $\mathbf{W}$ and $\mathbf{Z}$ are equal to $\mathbf{0}$), then equation (5) accounts for the role of all estimated worker or firm time-invariant heterogeneity, as per the application by Cardoso et al.: i.e. the gender-firm sorting effect becomes

$$= E_{it}\left[\widehat{\phi}_{J(it)} \mid i \in M, \mathbf{X}\boldsymbol{\beta}\right] - E_{it}\left[\widehat{\phi}_{J(it)} \mid i \in F, \mathbf{X}\boldsymbol{\beta}\right].$$

Consistent with the AKM-type models being characterised by firm switching, estimates of the fixed effects obtained from our Full model are only comparable within mobility groups, i.e. connected sets of workers and firms (Abowd et al., 2002).[9] Therefore, in all our following analysis, we restrict attention only to the largest identified mobility group in the panel dataset. Nonetheless, this still means estimating a large number of coefficients, $(k + P + L - 2)$, for the Full model. Fortunately, algorithms have been developed, such as by Abowd et al. (2002) and Guimaraes and Portugal (2010), which identify mobility groups in large datasets and models with high-dimensional fixed effects, and also allow an exact solution of the least squares estimation of these models.[10]

It is important to acknowledge that the key assumptions of the AKM model have been criticised as being unrealistic (see for a summary Card et al., 2018). Most significantly, the interpretation of the firm fixed effects in the model, as allowing for consistent estimates of firm-specific wage premiums, which workers gain or lose symmetrically upon switching firms, relies on a strong assumption that the mobility of workers is exogenous, conditional on all observable and unobservable factors. This is equivalent to the assumption that the vector of model errors $\boldsymbol{\varepsilon}$ is orthogonal to $\mathbf{X}$, $\mathbf{A}$ and $\mathbf{F}$. For example, OLS estimates of the fixed effects would be biased if employees switch firms when they experience shocks to the match-specific component of their wages, which in our Full model is assumed away within the residual ε_{it}. We discuss the validity of this assumption in the context of our data and application in Appendix B.

2.1 Data and sample construction

The Annual Survey of Hours and Earnings (ASHE), 2002-16, is based on a 1% random sample of employees, drawn from HM Revenue and Customs Pay As You Earn (PAYE) records, collected and administered by the Office for National Statistics. Questionnaires are sent to employers, who are legally required to complete them with reference to payrolls for a period in April. The dataset can be viewed as a panel of employees without attrition, where each annual cross-section forms an approximate 1% random

---

[9]Abowd et al. (2002) describe the concept of connectedness in this context as: "When a group of person and firms is connected, the group contains all the workers who ever worked for any of the firms in the group and all the firms at which any of the workers were ever employed. In contrast, when a group of person and firms is not connected to a second group, no firm in the first group has ever employed a person in the second group, nor has any person in the first group ever been employed by a firm in the second group."

[10]We use the STATA programme of Guimaraes and Portugal (2010).

and representative sample of UK employees.[11] Particularly valuable for our analysis are the longitudinal identifiers for individuals and enterprises. We use the terms 'firm' and 'enterprise' synonymously. The latter in this case is a specific administrative definition of UK employers, which could contain several local units (or establishments). These identifiers allow us to construct a longitudinal, linked employer-employee dataset, which contains the records of approximately 1% of the UK employee population in every year between 2002 and 2016. We include only the main job observation of an individual, which must not be at a trainee or an apprentice level, and not have incurred a loss of pay in the reference period for whatever reason. To avoid some spurious derived hourly wage rates, we only keep observations with 1-100 basic paid weekly hours, and drop any observations with missing values for gross weekly earnings. For further information on how we constructed this dataset from annual ASHE cross-sections, other adjustments made to the data, and the details of some trimming (or cleaning) of the sample, see Appendix A.

Our analysis focuses on the hourly wage rate, which equals the ratio of employee gross weekly earnings to the corresponding record of basic weekly paid hours, all excluding overtime. We refer to this simply as wages. We deflate the wages using the corresponding April values of the Consumer Price Index (CPI), and all values are then presented in April 2002 prices.[12] We consider prime-working-age employees, aged 25-64, who have non-missing records of earnings and hours. The ASHE does not contain any information on an individual's education. We drop observations under the age of 25, such that we only study the clear majority who would have completed full-time education by this age in the United Kingdom. Therefore, the worker fixed effects in our Full model should account for most of the employee heterogeneity associated with human capital accumulation before individuals have begun to fully engage with the labour market. The following analysis uses these observable characteristics of employees and firms: gender, age (years), tenure (completed consecutive years with a firm), whether a job is full- or part-time (thirty hours or less), occupation, industry sector of the firm, whether or not the firm is in the private sector, employee birth cohort (year) and the number of employees working for the firm. Details of these variables, their discrete categories, and how some of these were derived, are all described in Appendix A. Before describing the dataset, we first drop the small number of employee-year observations which have missing values for any of these variables.

From what we then call the 'Whole' ASHE sample, the mean and median values of the raw (observed/unadjusted) gender pay gap, for all employees and each year in 2002-16, are shown in Figure 1. The mean value declined from around 28 log

---

[11]The two main reasons why an individual might not be observed in some year are: either being truly non-employed, or having changed employer between January and April. Since the survey questionnaires are in most cases sent in April to the employer's registered address from January PAYE records, workers who switch employers during these months are under sampled.

[12]Obtained from UK National Statistics, accessed 24/4/2017.

points in 2002 to 16 points in 2016. The median gap similarly declined throughout the period.[13] Figure 1 also shows comparable pay gap statistics for what we call the 'Analysis' sample. This is the 87% sub-sample of the Whole ASHE sample constituting the largest mobility group of workers and firms. All our main results in this paper use the Analysis sample. It contains 1.71 million employee-year observations, with 131,903 men and 124,501 women. The median and mean number of years that men and women appear in the sample are for both genders 10 and 9, respectively. The sample covers 86,779 different firms, where 18,404 are present in 2002, and thereafter, approximately 4-6,000 thousand new firms enter in each subsequent year. The median, mean and standard deviation of the number of employee-year observations per firm are 300-400, 1,815 and 3,910 respectively.[14] Appendix Table A1 compares some descriptive statistics over employee-year observations between the Analysis and Whole samples of the ASHE. The 87% sub-sample is similar in most regards, though as we would anticipate in the largest connected set, the average firm size (actual number of employees) and its variance is higher. This implies that large firms are marginally over-represented in our analysis compared with the true UK employee population.

FIGURE 1: Gaps between the mean and median log real hourly wage of UK men and women, 2002-2016

A. Mean

B. Median
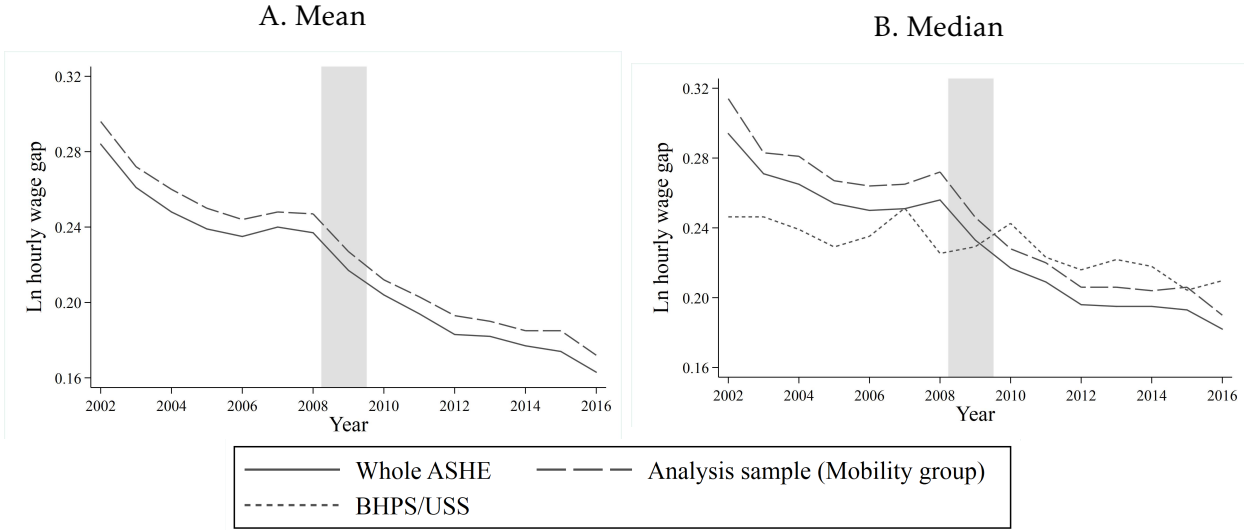
Notes.- author calculations using the ASHE 2002-16 and BHPS/USS 2002-2016, all employees aged 25-64. Pay excludes overtime. Gap is male minus female. "Analysis sample" is an 87% sub-sample of the "Whole ASHE" and represents statistics using only jobs in the largest mobility group, i.e those used to generate the paper's main results. See text and Appendix A for further details of the sample construction. Shaded area represents an official UK recession.

Returning to Figure 1, both the mean and median pay gap in the Analysis sample are consistently higher across all years than in the Whole ASHE sample, by approximately

---

[13]These levels and trends are similar to those one could obtain from UK official National Statistics aggregates, contained historically within the ONS "Patterns of Pay" series; https://data.gov.uk/.... This is not surprising, since the basis of those statistics is also the ASHE.

[14]Throughout the study, all reported median or percentile statistics are rounded to be consistent with the statistical disclosure control policy of the UK Data Service, i.e. these reported statistics should not be viewed as precise values.

one log point. Appendix Figure A1 shows that both male and female average pay are higher in the Analysis sample, though more so for male pay. Figure 1b and Appendix Figure A1b also show pay gap series derived using the British Household Panel Survey (BHPS, 2002-08) and Understanding Society Survey (USS, 2009-16), with as comparable as possible method and sample selection to our Whole ASHE measures (see Appendix A for details). These other datasets are top-coded, so we only compare median values. The household survey-based earnings data overall provide gender pay statistics similar to the ASHE since 2002. But patterns over time within this period do not closely match what we derive from the ASHE, nor what the ONS present as National Statistics aggregates. Specifically, the median male real hourly wage from the BHPS increased more before the 2007-08 financial crisis than measured within the ASHE; and both male and female median real hourly wages from the USS did not decline as far as they did in the ASHE following 2007-08, especially among men.

We estimate the kernel density of log real hourly wages separately for the male and female observations in the Analysis sample, shown by Figure 2. The female wage distribution is more positively skewed towards values of the UK's real National Minimum Wage, whereas observations with very high hourly rates of pay are dominated by men. Appendix Figure 2 shows how the gender-specific wage distributions are practically unchanged when we move between the Whole ASHE sample and the largest mobility group therein.

FIGURE 2: Male and female distributions of employee real hourly wages
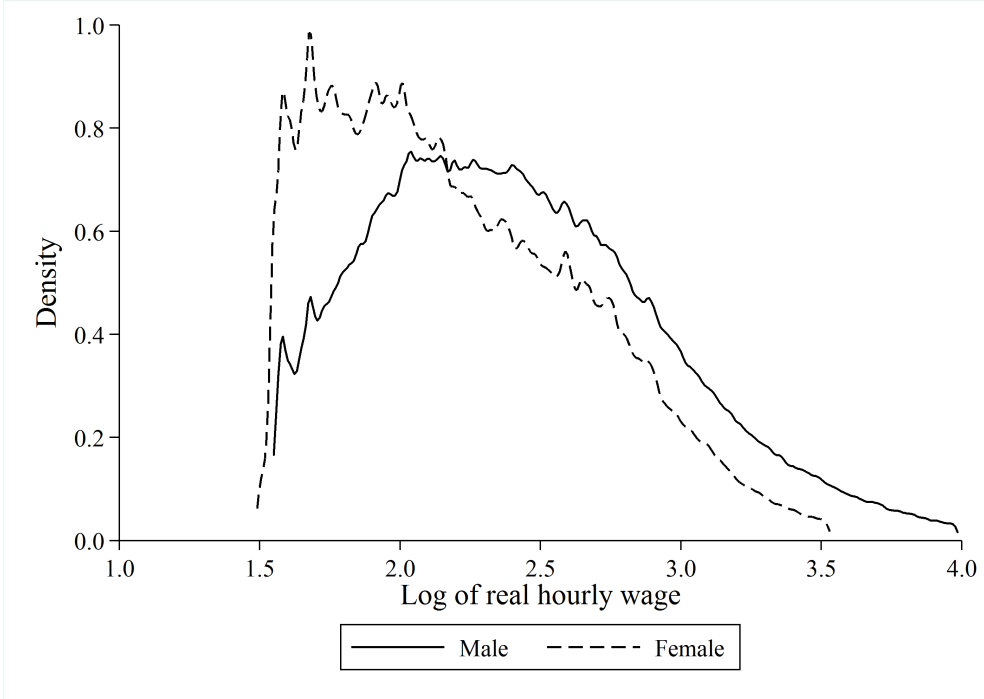
Notes.- see Figure 1. £2002. Uses the "Analysis sample". Both male and female kernel densities were estimated with a bandwidth of one log point. The top and bottom one percent of male and female hourly earners are not displayed.

3 Main Results

The raw or observed mean gender pay gap among all employees in the Analysis sample is 22.3 log points. The adjusted pay gap is 14.5 log points, obtained as the least squares estimate $\widehat{\lambda}$ from the Basic wages model, described by equation (3), with the following controls in $\widetilde{\mathbf{X}}$: year effects, squared and cubed employee age, cubic polynomials for employee tenure and firm size, and dummy variables for full-time, private sector, industry groups, occupations (2-digit classification) and employee birth cohorts (years). The estimate of this Basic model is summarised by column (2) of Appendix Table B1. To apply the decomposition of $\widehat{\lambda}$ into contributions from unobserved worker and firm fixed factors, we first estimate the Full wages model, described by equation (1), where the controls in $\mathbf{X}$ are the subset of those in $\widetilde{\mathbf{X}}$ that excludes the dummy variables for private sector, industry groups and employee birth cohorts. The estimate of the Full model is summarised in column (3) of Appendix Table B1. Collecting the least squares estimates of the worker and firm fixed effects, $\widehat{\alpha}_i$ and $\widehat{\phi}_{J(it)}$, we can then apply the Gelbach decomposition described above. Column (1) of Table 1 shows the results of doing so. Of the 14.5 log points adjusted pay gap, 0.8 log points, or a 5.7% share, is accounted for by the estimated gender-firm sorting effect, $\widehat{\delta}_f$. In other words, in a counterfactual where men and women were identically distributed across firms and their associated wage premiums, conditional on their different observable characteristics, the UK gender pay gap would narrow by less than a percentage point: the remaining 94% of the adjusted gap would still persist from within firms, accounted for here by the estimates of unobserved fixed worker factors, which by definition are transferable across firms.

Column (2) of Table 1 also shows the results of the adjusted pay gap decomposition when we exclude the time-invariant worker and firm observable factors from our Basic model. If we don't control for employee birth cohorts and pay premiums associated with very broad UK industry sectors, we find that the amount of the adjusted pay gap contributed by gender-firm sorting increases slightly to 1.1 log points, or a share of 7.4%. If we don't address the greater tendency of women to work in the public sector, for example, then our conditional estimates of the contribution from gender-firm sorting are overestimated by a small amount.

Our main result would not be surprising if overall, irrespective of gender, differences in hourly pay between firms did not account for a sizeable fraction of UK wage inequality. The AKM-type models allow a tractable decomposition of overall log wage variance into the share accounted for by firm-specific wage factors. For the estimate of our Full model, this share is given by $var(\widehat{\phi}_{J(it)})/var(w_{it}) = 12.2\%$.[15] Although comparisons are

---

[15] Alternatively, including the contribution from the covariance of the fixed effects with other factors in the wages model in the measure, the share is given by $covar(w_{it}, \widehat{\phi}_{J(it)})/var(w_{it}) = 13.2\%$.

TABLE 1: Main decomposition results for the adjusted log gender pay gap

	(1)	(2)
Worker effects - $\widehat{\delta_a}$	0.137	0.138
	(0.943)	(0.929)
Gender-firm sorting - $\widehat{\delta_f}$	0.08	0.11
	(0.057)	(0.074)
Adjusted gap - $\widehat{\delta_y}$ or $\widehat{\lambda}$	0.145	0.149
Observed gap (mean)	0.223	0.223

Notes.- author calculations using the ASHE 2002-16, all employees age 25-64. £2002. Pay excludes overtime. Gap is male minus female.

Column (1) gives results where the assumed Basic regression model (3) *does* include time-invariant employee or firm characteristics; i.e. $\widetilde{\mathbf{X}}$ from (3) is *not* identical to $\mathbf{X}$ in the Full model. These Basic model regression results are given by column (2) of Table B1.

Column (2) gives results where the assumed Basic regression model (3) *does not* include time-invariant employee or firm characteristics; i.e. $\widetilde{\mathbf{X}}$ from (3) is identical to $\mathbf{X}$ in the Full model. These Basic model regression results are given by column (1) of Table B1.

Values in parentheses give the share of the Adjusted gap accounted for either by the estimated Worker or Firm effects.

somewhat suspect to the data, especially because of different sample sizes and wage measures, this value is towards the lower end of the range of estimates obtained for other countries using AKM-type models, such as France, Germany, Portugal and the United States (for a summary see Card et al., 2018). Even so, systematic differences in firm wage premiums do account for a sizeable fraction of overall UK wage dispersion. Figure 3 emphasises this further, by plotting kernel density estimates of $\widehat{\alpha_i}$ and $\widehat{\phi_{J(it)}}$ by gender. The male distribution of the firm-specific wage premiums, received by employees in the UK labour market, is visibly more positively skewed than the female distribution, indicating that men are disproportionately more likely to be employed in high-wage firms when compared with women. However, this difference is small in magnitude, when compared with the differences between the distributions of estimated male and female worker fixed wage effects.

We construct another graphical representation, which depicts the role of the estimated firm wage premiums on the hourly wage gap throughout the employee wage distribution. In Figure 4, first we collect all employees in one percentile bins of the observed real hourly pay distribution in the Analysis sample. By design, within each percentile bin, men and women are approximately paid the same. Each bin contains several thousand male and female employee-year observations. For example, the top percentile contains 13,979 men and 3,102 women. We then look up the estimates of worker and firm fixed effects for the employees within a percentile bin, and compute their respective mean values for that percentile by gender. Within a bin, we subtract

FIGURE 3: Distributions over employees of estimated worker- and firm-specific fixed real hourly wage effects

A. Worker-specific effects - $\widehat{\alpha_i}$

B. Firm-specific effects - $\widehat{\phi_{J(it)}}$
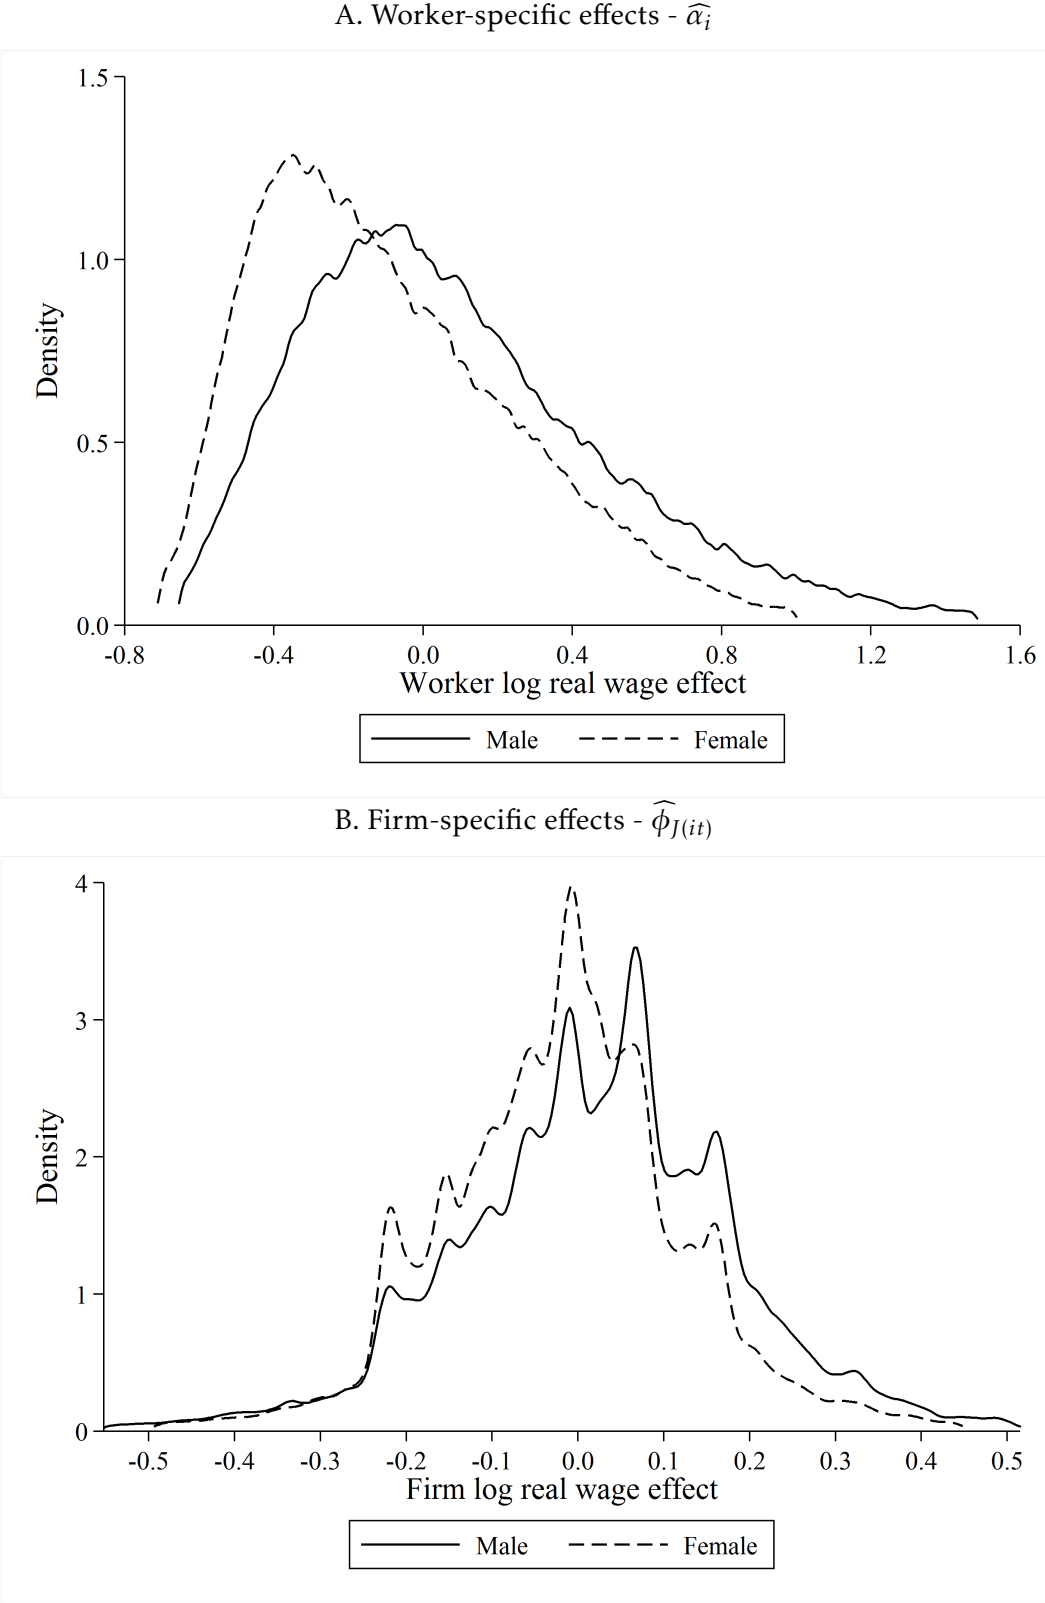
Notes.- wage effects estimated as per regression model (1) and column (3) of Table B1, with overall mean values then subtracted. £2002. Both male and female kernel densities were estimated with a bandwidth of one log point. The top and bottom one percent of the estimated effects (not gender-specific) are not displayed.

the female mean values for each set of fixed effects from the equivalent male values. These gender gaps (y-axis) are then plotted in Figure 4 for each overall employee wage percentile (x-axis). Initially focusing near the median wage, men and women, who are earning the most UK-typical rates of hourly pay, are on average working for firms with similar levels of firm-specific wage premiums. However, among those in relatively low-paid UK jobs, below the twentieth percentile, the gender gap in firm fixed effects is negative. This implies that to obtain the same pay as the men at the bottom of the wage distribution, UK women are working in relatively 'better' or higher-wage firms than those men. However, this gender gap is reversed among those earning close to the ninetieth percentile. Throughout the employee wage distribution, men earning the same as women have higher contributions to their wages from worker fixed factors, especially among the lowest and highest earners. These factors could include general human capital (transferable across firms, occupations and industry sectors), accumulated through education choices or work experience, but could also include the effects of wage discrimination.

Besides Cardoso et al. (2016), the other studies which have estimated a role of worker-firm sorting in gender pay gaps use a Blinder-Oaxaca type decomposition of the 'Raw' or observed gap. These estimates are also based on AKM-type wages model estimates (see for example Card et al., 2016, among the other studies discussed and summarised below in Table 6). We can present a measure which is more comparable with this wider literature, based on our Full model estimates:

$$\frac{E\left[\widehat{\phi}_{J(it)}|\ i \in M\right] - E\left[\widehat{\phi}_{J(it)}|\ i \in F\right]}{E\left[w_{it}|\ i \in M\right] - E\left[w_{it}|\ i \in F\right]} = \frac{0.036}{0.223} = 16\%\ , \tag{7}$$

where M and F denote the sets of men and women in the Analysis sample. This measure says that the allocation of men and women to the estimated firm fixed effects can account for 3.6 log points, or 16%, of the raw UK gender pay gap of 22.3 log points. However, this measure is somewhat less nuanced than that provided by the Gelbach decomposition. Comparing the numerator in (7) with equation (6), this measure does not first condition on the role played by the allocation of men and women to any observable fixed and time-varying worker and firm characteristics. Therefore, the estimate of $\widehat{\delta_f}$ from the Gelbach decomposition has, in our view, a more relevant interpretation: it gives an estimate of how much who works for whom matters, after conditioning on the fact that men, women, jobs and firms have different observable characteristics, which are also relevant for explaining a large part of the raw pay gap. In some sense, it measures the residual role of which workers work for which firms.

FIGURE 4: Gender gap in the contribution of estimated worker- and firm-specific wage effects throughout the overall employee hourly wage distribution
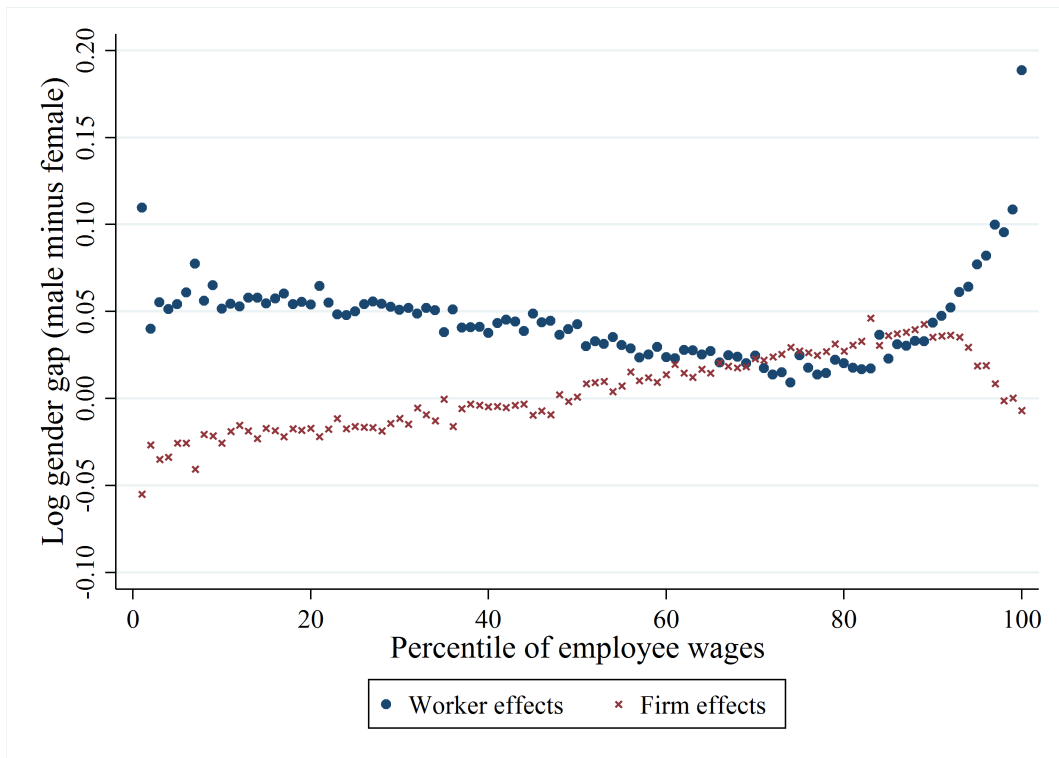
Notes.- wage effects estimated as per regression model (1) and column (3) of Table B1. £2002.
How figure constructed: all employees (non-gendered) are grouped in real hourly wage percentile bins; "Worker effects" here plots the difference between the average male and female estimated worker-specific wage effects within each percentile bin; "Firm effects" plots the equivalent gender difference by percentile bin of the overall employee wage distribution. Summing the two components within a percentile, and subtracting that value from zero, would give the gender gap contributed by the time-varying observable characteristics at each employee wage percentile.
Interpretation: a negative value displayed here for the "Firm effects" implies that women in that portion of the overall employee wage distribution tend to be employed by higher-wage firms than their male counterparts, on average.

3.1 The role of firms within industry sectors, part- or full-time work, and age groups

We can use our estimates of systematic firm wage premiums, from the Full regression model, to explore the distributions of men and women over relatively high- and low-paying firms, for different groups of jobs in the labour market. Table 2 shows percentiles of $\widehat{\phi}_{J(it)}$ by gender within these various employee groups. Appendix Figures C1-C3 represent the equivalent set of results as kernel density estimates. First, we look within the private and public sectors. In the latter, the distributions of men and women over firm wage premiums are more similar than in the whole labour market, with only a modest tendency of men to be working in firms with relatively very high wage premiums. The median firm fixed effect received by a man in the public sector is 2 log points higher than the median value received by a woman. The distribution of firm fixed effects in the private sector is more dispersed than in the public sector. The gender gap in firm fixed effects is low at the bottom of the male and female

16

private sector distributions. But the median male firm wage premium in the private sector is 9 log points higher than the female median value, though this gap narrows moving towards the top of the fixed effects distributions.

The 'Other industry sectors' group is dominated by the public sector, so the gender distributions of firm fixed effects are similar to those in the public sector. The male and female distributions over the firm wage premiums are also relatively similar within the 'Manufacturing' and 'Financial services' industry groups. However, the gender differences in the 'Non-financial (sales) services' group, which is dominated by retail and hospitality services, are starker. The seventy-fifth and ninetieth percentiles of the firm wage premiums received by men in this sector are as much as 10-11 log points higher than the equivalent female figures.

We also look at differences in where men and women work within part- and full-time employment (Appendix Figure C2). In part-time employment, men are more likely to be employed by a firm with a very low wage premium than women are. Conversely in full-time employment, men are more likely to be employed by a firm with a high wage premium than women, but no more so than in the overall gender distributions over the estimated firm fixed effects.

Finally, we look within age groups. For those aged 25-34, there are only small differences between men and women in the likelihoods of working for firms with relatively high or low wage premiums. In this age group, the median firm fixed effect received by a man is 2 log points higher than for women. These results show some similarities with Manning and Swaffield (2008), who found that in the first years after entering the labour market the occupation allocation of men and women did not account for the widening pattern of the pay gap thereafter. However, our findings suggest it is also the gender allocation to firms which does not generate large wage gaps among younger workers. But gaps in representation do open up for the 35-44 age group, and the differences between men and women in what type of firm they work for then persist up to age 64. The gender gap in the median firm fixed effect is around 5 log points between ages 35 and 54. The gender representation gaps also widen with age, with the ninetieth percentile of the firm fixed effects received by men aged 35-64 being around 7 log points higher than the equivalent figure among female employees aged 35-64. Therefore, there is some evidence that the extent to which men disproportionately work for high wage firms could increase with age. This would be consistent with recent findings from the United States (Goldin et al., 2017). Though we must caveat this result here, since the shortness of the sample period studied means we cannot robustly disentangle the birth cohort effects from the life-cycle.

TABLE 2: Gender gaps at percentiles of estimated firm-specific wage effects ($\widehat{\phi}_{J(it)}$) by sector, part- vs. full- time and age groups

	Percentiles				
	10th	25th	50th	75th	90th
Overall (all employee-years)	0.01	0.03	0.04	0.06	0.06
Private	0.01	0.06	0.09	0.05	0.04
Public	0.00	0.01	0.02	0.03	0.04
Manufacturing	0.04	0.03	0.02	0.02	0.02
Non-financ. (sales) services	-0.01	0.02	0.03	0.10	0.11
Financial services	0.00	0.00	0.00	0.00	0.03
Other ind. sectors	0.01	0.02	0.03	0.03	0.05
Part-time	-0.02	-0.03	-0.02	0.00	0.01
Full-time	0.00	0.01	0.03	0.04	0.05
Age25-34	-0.01	0.00	0.02	0.04	0.04
Age 35-44	0.02	0.03	0.05	0.06	0.07
Age 45-54	0.02	0.04	0.05	0.07	0.07
Age 55-64	0.02	0.04	0.03	0.05	0.07

 Notes.- effects estimated as per regression model (1) and column (3) of Table B1. The derivation of the industry groupings used here is based on the SIC2003 and described in Appendix A.

4 Robustness and Further Discussion

This section considers how robust our main results are within sub-samples of employee-year observations. We also look at how much of the adjusted pay gap could be accounted for not only by gender-firm sorting but also by gender-occupation sorting. Finally, we compare our results to what other studies of the pay gap in other countries have found, regarding the importance of whom men and women work for.

4.1 Before and after the 2008/9 recession

As Figure 1 shows, the UK gender pay gap decreased by approximately 10-12 log points between 2002 and 2016. It is possible that some of this change could be accounted for by a decline in the extent of gender-firm sorting. This period also included the

2007-08 financial crisis and a deep recession. Following 2008 there was a fall in British employee wage inequality, though there was also an increase in the variance of firm average wages (Schaefer and Singleton, 2017). For firms and employees who remained economically active during and since the financial crisis, the downturn could have systematically affected the wages they paid and received. Furthermore, the recession substantially changed the composition of UK employment and production compared with before (Blundell et al., 2014). Therefore, to check whether the main results generalise throughout our sample period, we re-estimate the Full wages model for two sub-periods, allowing all of the parameters to change: first 2002-07, before the financial crisis; and second 2008-16. We then apply the decomposition of the adjusted gender pay gap within these sub-periods. Table 3 presents the results, where we repeat for comparison in column (3) our main findings for the whole sample period. The estimate of the adjusted pay gap is 15.1 log points in the first sub-period, compared with 13.6 log points in the second. But, even though the pay gap decreased markedly since 2002, and despite the UK's Great Recession and the following recovery having had impacts on relative gender outcomes (Razzu and Singleton, 2016), the estimated contribution from gender-firm sorting was around 1 log point in both sub-periods. In this regard, our main findings generalise throughout 2002-16, when we allow for changes through the period in the composition of employment and wages over firms and workers, and allow for changes in the estimated wage premiums for observable and unobservable factors.

TABLE 3: Decomposition results for the log gender pay gap: Analysis sample sub-periods

	2002-07 (1)	2008-16 (2)	2002-16 (3)
Worker effects - $\widehat{\delta_a}$	0.143 (0.942)	0.126 (0.928)	0.137 (0.943)
Gender-firm sorting - $\widehat{\delta_f}$	0.009 (0.058)	0.010 (0.072)	0.008 (0.057)
Adjusted gap - $\widehat{\delta_y}$ or $\widehat{\lambda}$	0.151	0.136	0.145
Observed gap (mean)	0.267	0.201	0.223

Notes.- see Table 1. Column (3) here presents the same results as Column (1) of Table 1, repeated for comparison. Column (1) presents equivalent results to Column (3), where the sample period is reduced to only employee-year observations in 2002-07 (N=541,346). Similarly, Column (2) presents results where the sample period is reduced to only employee-year observations in 2008-16 (N=1,002,805).

4.2 Occupations, firms and the pay gap

Firms can be described by the collections of tasks, jobs and occupations carried out by their employees. For any given firm this description can change over time. Employees

often change occupations when they switch jobs, either within or between firms. For these reasons, it is important that we account for occupational wage premiums in the Basic and Full regression models. All the models estimated in Section 3 included controls for occupations at the 2-digit level of the ONS Standard Occupational Classification (SOC). This amounts to 26 occupation groups. Information for employees on their 3-digit occupational classification is also available in the ASHE. In our baseline results, we prefer using the 2-digit groups because there is greater scope for measurement error (spurious occupation switching) when using a more detailed classification. Nonetheless, in column (1) of Table 4, we confirm that our main findings are robust to this modelling choice, when we instead measure the adjusted pay gap controlling for 92 3-digit occupations groups. The contribution from the gender-firm sorting effect decreases to 4% of the adjusted pay gap.

In column (2) of Table (4), we present decomposition results comparable to our main results, except here we do exclude occupations from the Basic model and $\widetilde{\mathbf{X}}$, and account for a 'gender-occupation sorting' effect in a similar manner as in Cardoso et al.[16] Column (3) shows the results from applying the same decomposition as in column (2), but with 3- instead of 2-digit occupation effects estimated in a variant of the Full wages model. The measured contribution to the adjusted pay gap from gender-firm sorting in these decompositions increases to 1.2-1.3 log points, compared with our main findings of 0.8 log points. Although UK occupations receive very different rates of pay, gender-occupation sorting contributes only one percentage point to the pay gap. Together, who men and women work for and what occupations they are in accounts for 13-14% of an overall adjusted UK gender pay gap of 17 log points.

4.3 Comparing our main results with other studies and countries

Several recent analyses have looked at the extent to which the sorting of men and women over firms contributes to the pay gaps of other countries. As emphasised throughout our previous discussion, the most comparable set of results to our own are from Cardoso et al. (2016) for Portugal. The observable pay gap over their sample period was a similar level to that studied here in the United Kingdom. They found that approximately a fifth of an adjusted Portuguese hourly wage gap was accounted for by gender-firm sorting, with a further fifth accounted for by gender-job-title sorting.[17] These estimates are an order of magnitude greater than what we find for the United Kingdom. One explanation is that firms could potentially matter less for overall UK wage variance than in Portugal.

---

[16]Cardoso et al. (2016) excluded covariates for job titles (a very specific classification of occupations) from their Basic wages model, and then used the Gelbach (2016) decomposition to simultaneously account for contributions to an adjusted pay gap from estimated sets of worker, firm and job-title fixed effects. Because firms are a collection of occupations, and because we have a smaller sample of the employees within firms, we prefer to adjust the pay gap for occupational wage premiums before measuring the contribution from gender-firm sorting.

[17]These job-titles in the Portuguese data refer to somewhere in the order of 30,000 different groups.

TABLE 4: Decomposition results for the log gender pay gap: robustness to occupational classification and accounting for the contribution from occupation-specific wage effects

	Main results	Firm and occupation sorting	
	w. 3-digit occs	2-digit	3-digit
	(1)	(2)	(3)
Worker effects - $\widehat{\delta_a}$	0.126	0.148	0.146
	(0.960)	(0.871)	(0.856)
Gender-firm sorting - $\widehat{\delta_f}$	0.005	0.012	0.013
	(0.040)	(0.070)	(0.074)
Gender-occ. sorting - $\widehat{\delta_o}$		0.010	0.012
		(0.060)	(0.070)
Adjusted gap - $\widehat{\delta_y}$ or $\widehat{\lambda}$	0.131	0.170	0.170
Observed gap (mean)	0.223	0.223	0.223

Notes.- see Table 1. Column (1) provides equivalent results as Column (1) in Table 1, except here the occupation controls included in the Basic and Full model were at the 3- rather than 2-digit SOC level. Column (3) then gives results where we exclude the 3-digit controls from the Basic model, but use the Gelbach decomposition to account for how much of the Adjusted pay gap is contributed by the gendered distribution over these occupation effects, as estimated in the Full model, alongside the worker- and firm-specific effects. Column (2) here is the equivalent of Column (3), except for 2-digit occupation effects.

There is some evidence for this, with estimates from an AKM-type model using these Portuguese data showing that as much as 20% of hourly wage variance there is accounted for by establishment effects (Card et al., 2016, 2018), which compares with about 13% from the estimates of UK firm effects here. Another explanation for the difference in magnitude between the UK and Portuguese estimates could relate to the sets of covariates included in the Basic wages models. In Cardoso et al., the Basic model included only covariates for age and age squared, tenure and tenure squared, and year effects. The 'raw' observed mean pay gap in their sample is 24 log points, and the adjusted gap estimated from their Basic model is 23 log points. To provide a closer comparison, we also consider a Gelbach decomposition with a similarly limited set of covariates in our Basic wages model.[18] We apply this decomposition to the same estimates obtained from the Full wages model used in our main results. Our measure of the adjusted pay gap in this more Basic wages model is now 21 log points, which is similarly close to the observed UK sample average value, 22 log points, as it was for Portugal in Cardoso et al. We decompose this measure into four components (column (1) of Table 5): the role of worker effects, gender-firm sorting, gender-occupation sorting, and the contribution from the allocation of men and women over the additional time-varying covariates we included

[18]Though we continue to exclude a linear age term from the model, which we always omitted because of the controls for birth cohorts in other specifications.

in our original Basic model (a cubic firm size polynomial, full-time work dummy, cubic terms in age and tenure). The estimate we find for the share of the adjusted pay gap from gender-firm sorting is then 16%. The share contributed by gender-occupation sorting is 6%, and the contribution from the time-varying covariates no longer included in the Basic model is -8%, with the remainder coming from the role of the worker fixed effects. Our estimate of the UK gender-firm sorting effect would be biased upwards by omitting important time-varying covariates from the Basic wages model.

TABLE 5: Decomposition results for the log gender pay gap: a more Basic wags model and the gross weekly earnings gap

	A more Basic model (1)	Gross weekly earnings (2)	Main results (hourly wages) (3)
Worker effects - $\widehat{\delta_a}$	0.183 (0.864)	0.161 (0.933)	0.146 (0.943)
Gender-firm sorting - $\widehat{\delta_f}$	0.034 (0.159)	0.012 (0.067)	0.008 (0.057)
Gender-occ. sorting - $\widehat{\delta_o}$	0.012 (0.056)		
Other observable chars. - $\widehat{\delta_x}$	-0.017 (-0.079)		
Adjusted gap - $\widehat{\delta_y}$ or $\widehat{\lambda}$	0.211	0.173	0.145
Observed gap (mean)	0.223	0.498	0.223

Notes.- see Table 1. Column (3) presents the same results as Column (1) of Table 1, here repeated for comparison. Column (1) applies the Gelbach decomposition to the Full model assuming a more Basic wages model. The contribution presented from 'Other observable chars.' gives the contribution, to the associated estimate of the adjusted pay gap, from the time-varying covariates included in the Full model but not in this more Basic wages model, as described in the text. Column (2) use the exact same Analysis sample of employee-years and firms, and the same regression models, as used to estimate our main results, except the dependent variable of these models is now the natural logarithm of real gross weekly wages, excluding overtime.

Table 6 summarises the main results from five other studies which are closely comparable to our own. All of these use a common methodological approach. Specifically, the approach taken by Card et al. (2016) (henceforth CCK), with data from Portugal for the period 2002-2009, has since been followed closely by Bruns (2016) for Germany, Gallen et al. (2017) for Denmark, Sorkin (2017) for the United States, and Coudin et al. (2018) for France. CCK and these other studies estimate AKM-type models, where all coefficients, including the firm fixed effects, can vary by gender. CCK then propose a novel way to further decompose the contribution to the raw pay gap from gender differences in the estimated firm fixed effects, into what they call sorting and

bargaining components. The sorting component relates to differences in the firms that men and women work for. The bargaining component is derived from any differences in the relative firm-specific premiums men and women would receive upon switching between the exact same firms. This decomposition depends on how the separately estimated sets of effects by gender are normalised. CCK find overall that 21% of the Portuguese average pay gap is accounted for by estimates of the gender-firm fixed effects. Decomposing this further, they show that the sorting and bargaining components can account for around 15% and 5% of the overall raw pay gap, respectively. Using French data for the two decades between 1995 and 2014, Coudin et al. (2018) found a positive role for the sorting component of 11% in the raw French pay gap, but no role for the bargaining component. Bruns (2016) and Gallen et al. (2017) also found that the sorting component dominated any role for gender wage bargaining effects in recent decades in Germany and Denmark. For the reasons we explained before, our results for the contribution to the adjusted, or conditional, UK hourly pay gap from gender-firm sorting are not directly comparable to the results from these studies. However, our estimate that 16% of the raw pay gap is accounted for by differences in the allocation of men and women to firm fixed effects, shown by equation (7), is roughly comparable to the total 21% figure from CCK, with the caveat that our estimate was obtained without allowing any gender differences in model parameter estimates. In terms of an estimate of how much the unconditional sorting over firms contributes to the raw gender pay gap, these UK data do not generate an outlier compared with other studies and countries.

Besides our preference to account for whether gender-firm sorting contributes to a measure of the adjusted pay gap, there are two other reasons why we favoured the Gelbach decomposition to CCK's approach here. First, the principal motivation of CCK is more specific than our own. They motivate their decomposition by a need to reconcile what they identify as two competing strands of literature on how firms might generate pay gaps: through a sorting or bargaining channel. In comparison, for the United Kingdom, we are only motivated by measuring if gender-firm sorting matters. Second, our sample size is much smaller than CCK's, though it is still comparable to the sample size in most of the studies mentioned above which have followed them. That said, identification of the AKM-model depends on connected sets and mobility groups. In these UK data, if we were to estimate gender-specific AKM-type models, as the CCK approach requires, we would begin to rapidly forego the representativeness of the sample and the precision with which any firm fixed effects can be estimated.

It is important to distinguish what type of pay is being studied when comparing estimates on the importance of firms in gender pay gaps. Of all the studies summarised in Table 6, Sorkin (2017) for the United States finds the greatest role for sorting effects. Following a similar approach to CCK, he finds that a quarter of the raw US overall pay gap can be accounted for by the sorting channel. However, this US study looks at earnings rather than wages, since hours data were not available, unlike for Portugal

in CCK's study. Sorkin suggests that there is a possibility of high- and low-hours firms, and that US men and women may be sorted on this basis, which could explain why his estimate for the sorting channel is greater than found by CCK. There is also recent evidence from the United States that there are increasing wage returns from working longer hours, which men are disproportionately more likely to do (Goldin, 2014; Cortes and Pan, 2016). If longer working hours are associated with greater rates of pay in the UK, and men are sorted towards those firms which offer longer hours, then we would expect estimates of the gender-firm sorting effect to be larger when we decompose the weekly earnings gap. To explore the importance of the UK pay measure, we replicate our main results, using the same exact Analysis sample of employee-year and firm observations, and with the same Full and Basic wages models, but instead of log hourly employee wages being the dependent variable, we replace this with log gross weekly earnings. The raw gender gap in weekly earnings in the Analysis sample is 49.8 log points. The estimate of the covariate-adjusted earnings gap is 17.3 log points. Applying the Gelbach decomposition to this value, 1.2 log points, or a share of 6.7%, is accounted for by the gender-firm sorting effect. Therefore, the estimated influence of where men and women work on the adjusted weekly earnings gap is of a similar magnitude as it is for the hourly wage gap, despite the far greater quantity of hours worked by UK men than women (see Appendix Table A1). We also find that of the raw earnings gap, 6.7 log points, or a share of 13.5%, is accounted for by the gender difference in the estimated firm fixed effects from the weekly earnings Full wages model. Therefore, at least so far as the UK is concerned, there is no evidence that the extent to which men and women are sorted towards especially high- or low-hours firms accounts for a substantial part of why the weekly earnings gap is far greater than the hourly wage gap.

5 Conclusion

Although the UK gender pay gap has decreased steeply in recent decades, it remains large and significant. A growing literature now highlights the importance of where people work, and especially for whom they work, in shaping wage inequality. In this study we asked how much of an adjusted UK hourly gender pay gap, between 2002 and 2016, is accounted for by the differences between men and women in whom they work for. We found that the contribution from this gender-firm sorting effect is small: in the order of one percentage point, or just 6% of the adjusted pay gap. In other words, if overnight all gender inequality was wiped out in terms of the firms that people work for, conditional on the existing observable differences in worker, job and firm characteristics, which are also relevant for wages, the gender pay gap would only decrease by this small amount. The clear majority of what explains the pay gap shows up within firms. Future research should focus on identifying the worker and firm behaviours which can explain this.

TABLE 6: Summary of results from elsewhere on the share of the gender pay gap accounted for by where men and women work

	Country, Period	Raw Pay Gap (log points)	Share of raw/adj. gap acc. for by gender-firm sorting	Additional notes
Cardoso et al. (2016)	Portugal, 1986-2008	24	19% (adj.)	Gelbach decomposition of worker, firm and job-title fixed effects; 18% of adj. gap from job-title-sorting; private sector representative sample of 28 million obs.; age 18-64; real hourly earnings included benefits and overtime pay.
Card et al. (2016)	Portugal, 2002-09	23	15% (raw)	Gender-specific AKM-type models used to estimate gender-firm specific pay premium, then Blinder-Oaxaca-type decomposition of estimated firm effects into separate sorting and 'bargaining' components (5% of raw gap); sample of 14m observations; age 19-65.
Bruns (2016)	Germany, 1995-2008	25 (2001-08)	25-31% (raw)	Method follows Card et al. (2016); -5-0% of raw gap from bargaining component; sample of 2.5 million workers in 90,000 establishments; age 20-60; no hours data; top-coded data.
Gallen et al. (2017)	Denmark, 1980-2010	19 (2000-09)	14% (raw)	Method follows Card et al. (2016); estimated sorting effect varies over time, being only 3% in 1980-89; sorting effects at establishment level; annual earnings data and no hours for full-time workers.
Sorkin (2017)	United States, 2000-08	33	25% (raw)	AKM-type decomposition, similar to Card et al. (2016); sample of 500 million people-years; age 18-61; US-LEHD data does not contain hours worked so uses annual earnings.
Coudin et al. (2018)	France, 1995-2014	17	11% (raw)	Method follows Card et al. (2016); sorting (10.6%) and bargaining (-2.4%) component; connected sample of 100,000 workers in 90,000 firms; private sector employees only; hourly wage data.

References

Abowd, J., R. H. Creecy, and F. Kramarz. 2002. "Computing Person and Firm Effects Using Linked Longitudinal Employer-Employee Data." longitudinal employer-household dynamics technical papers, Center for Economic Studies, U.S. Census Bureau.

Abowd, J., F. Kramarz, and D. Margolis. 1999. "High Wage Workers and High Wage Firms." *Econometrica*, 67(2): 251–334.

Altonji, J., and R. Blank. 1999. "Race and gender in the labor market." In *Handbook of Labor Economics*. Eds. by O. Ashenfelter, and D. Card, 3, Part C Elsevier, , 1st edition, Chap. 48 3143–3259.

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

Barth, E., A. Bryson, J. Davis, and R. Freeman. 2016. "It's Where You Work: Increases in the Dispersion of Earnings across Establishments and Individuals in the United States." *Journal of Labor Economics*, 34(S2): S67 – S97.

Barth, E., S. Pekkala Kerr, and C. Olivetti. 2017. "The Dynamics of Gender Earnings Differentials: Evidence from Establishment Data." NBER Working Papers 23381, National Bureau of Economic Research, Inc.

Bayard, K., J. Hellerstein, D. Neumark, and K. Troske. 2003. "New Evidence on Sex Segregation and Sex Differences in Wages from Matched Employee-Employer Data." *Journal of Labor Economics*, 21(4): 887–922.

Bertrand, M. 2011. "New Perspectives on Gender." In *Handbook of Labor Economics*. Eds. by O. Ashenfelter, and D. Card, 4 Elsevier, , 1st edition, Chap. 17 1543 – 1590.

Blau, F. D. 1977. *Equal pay in the office*. Lexington Books.

Blau, F., and L. Kahn. 2017. "The Gender Wage Gap: Extent, Trends, and Explanations." *Journal of Economic Literature*, 55(3): 789–865.

Blundell, R., C. Crawford, and W. Jin. 2014. "What Can Wages and Employment Tell Us about the UK's Productivity Puzzle?." *The Economic Journal*, 124(576): 377–407.

Bruns, B. 2016. "Assessing the role of workplace heterogeneity in recent trends of the gender wage gap." Working Papers 2016006, Berlin Doctoral Program in Economics and Management Science (BDPEMS).

Card, D., A. R. Cardoso, J. Heining, and P. Kline. 2018. "Firms and Labor Market Inequality: Evidence and Some Theory." *Journal of Labor Economics*, 36(S1): S13 – S70.

Card, D., A. R. Cardoso, and P. Kline. 2016. "Bargaining, Sorting, and the Gender Wage Gap: Quantifying the Impact of Firms on the Relative Pay of Women." *The Quarterly Journal of Economics*, 131(2): 633–686.

Card, D., J. Heining, and P. Kline. 2013. "Workplace Heterogeneity and the Rise of West German Wage Inequality." *The Quarterly Journal of Economics*, 128(3): 967–1015.

Cardoso, A. R., P. Guimaraes, and P. Portugal. 2016. "What drives the gender wage gap? A look at the role of firm and job-title heterogeneity." *Oxford Economic Papers*, 68(2): 506–524.

Cortes, P., and J. Pan. 2016. "When Time Binds: Returns to Working Long Hours and the Gender Wage Gap among the Highly Skilled." IZA Discussion Papers 9846, Institute for the Study of Labor (IZA).

Costa Dias, M., R. Joyce, and F. Parodi. 2018. "The gender pay gap in the UK: children and experience in work." IFS Working Paper 18/02, The Institute for Fiscal Studies.

Coudin, E., S. Maillard, and M. To. 2018. "Family, firms and the gender wage gap in France." IFS Working Paper 18/01, The Institute for Fiscal Studies.

Croson, R., and U. Gneezy. 2009. "Gender Differences in Preferences." *Journal of Economic Literature*, 47(2): 448–74.

Eeckhout, J., and P. Kircher. 2011. "Identifying Sorting – In Theory." *Review of Economic Studies*, 78(3): 872–906.

Faggio, G., K. G. Salvanes, and J. van Reenen. 2010. "The evolution of inequality in productivity and wages: panel data evidence." *Industrial and Corporate Change*, 19(6): 1919–1951.

Fortin, N. M., B. Bell, and M. Böhm. 2017. "Top earnings inequality and the gender pay gap: Canada, Sweden, and the United Kingdom." *Labour Economics*, 47(C): 107–123.

Gallen, Y., R. V. Lesner, and R. M. Vejlin. 2017. "The labor market gender gap in denmark: Sorting out the past 30 years." IZA Discussion Papers 10789, Institute for the Study of Labor (IZA).

Gelbach, J. B. 2016. "When Do Covariates Matter? And Which Ones, and How Much?." *Journal of Labor Economics*, 34(2): 509 – 543.

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

Goldin, C., S. Pekkala Kerr, C. Olivetti, and E. Barth. 2017. "The Expanding Gender Earnings Gap: Evidence from the LEHD-2000 Census." *American Economic Review*, 107(5): 110–14.

Groshen, E. 1991. "The Structure of the Female/Male Wage Differential: Is It Who You Are, What You Do, or Where You Work?." *Journal of Human Resources*, 26(3): 457–472.

Guimaraes, P., and P. Portugal. 2010. "A simple feasible procedure to fit models with high-dimensional fixed effects." *Stata Journal*, 10(4): 628–649.

Lopes de Melo, R. 2018. "Firm Wage Differentials and Labor Market Sorting: Reconciling Theory and Evidence." *Journal of Political Economy*, 126(1): 313 – 346.

Manning, A., and J. Swaffield. 2008. "The gender gap in early-career wage growth." *Economic Journal*, 118(530): 983–1024.

Mueller, H. M., P. P. Ouimet, and E. Simintzi. 2017. "Wage Inequality and Firm Growth." *American Economic Review*, 107(5): 379–83.

Mumford, K., and P. N. Smith. 2007. "THE GENDER EARNINGS GAP IN BRITAIN: INCLUDING THE WORKPLACE." *The Manchester School*, 75(6): 653–672.

Mumford, K., and P. N. Smith. 2009. "What determines the part-time and gender earnings gaps in Britain: evidence from the workplace." *Oxford Economic Papers*, 61(suppl1): i56–i75.

Office for National Statistics. 2017. "Annual Survey of Hours and Earnings, 1997-2016: Secure Access. [data collection]." 9th Edition SN: 6689, UK Data Service.

Razzu, G., and C. Singleton. 2016. "Gender and the business cycle: An analysis of labour markets in the US and UK." *Journal of Macroeconomics*, 47(PB): 131–146.

Schaefer, D., and C. Singleton. 2017. "Recent Changes in British Wage Inequality: Evidence from Firms and Occupations." 2017 Meeting Papers 459, Society for Economic Dynamics.

Snell, A., P. Martins, H. Stüber, and J. Thomas. 2018. "Bias in Returns to Tenure When Firm Wages and Employment Comove: A Quantitative Assessment and Solution." *Journal of Labor Economics*, 36(1): 47 – 74.

Song, J., D. J. Price, F. Guvenen, N. Bloom, and T. von Wachter. 2016. "Firming Up Inequality." unpublished manuscript.

Sorkin, I. 2017. "The Role of Firms in Gender Earnings Inequality: Evidence from the United States." *American Economic Review*, 107(5): 384–87.

Topel, R. 1991. "Specific Capital, Mobility, and Wages: Wages Rise with Job Seniority." *Journal of Political Economy*, 99(1): 145–76.

Weichselbaumer, D., and R. Winter-Ebmer. 2005. "A Meta-Analysis of the International Gender Wage Gap." *Journal of Economic Surveys*, 19(3): 479–511.

Appendix A. Further description of the data and sample construction

This section provides additional details regarding the datasets used and how we have constructed sub-samples and derived variables thereof. The relevant documentation and variable descriptions attached to these datasets are publicly available from the UK Data Service. The ONS has also published various documents concerning the data quality and consistency of the ASHE.

We focus on methodological details in the period 2002-16. Throughout these years, the ASHE is intended to be a true random sample of all employees in employment, irrespective of employment status, occupation, size of employer etc. Given the legal obligation of employers to respond using payrolls, the ASHE has always had a consistently high response rate. There is no cumulative attrition from the panel. Any individual in the sampling frame not included in the ASHE, in any year, for whatever reason, remains in the sampling frame for the following year. Conditional on a hundred percent response to the survey, the ASHE would be a truly random one percent sample of employees: all with a National Insurance number which has a numerical part ending in 14. However, there are two major sources of under-sampling, both occurring if individuals do not have a current tax record. This can happen for some individuals who have recently moved job, or for those who earn very little (mostly working part-time), and who are therefore not paying income tax or National Insurance when their employers are looked up by the ONS. For either of these reasons, a worker would not have (yet) been assigned to an administrative employer reference, known in the UK as a PAYE number (Pay As You Earn). Therefore, the statistical authority would not be able to find an employer address to send the survey questionnaire to. From 2004 the ASHE aimed to increasingly sample some of these employees previously under-represented. It added supplementary data for individuals without a PAYE reference, and attempted to represent the employees whose jobs would have changed between the determination of the sampling frame each January and the reference period in April. Nevertheless, the ASHE datasets remain representative of the employee payments and hours worked in the UK. We view the dataset as providing on average an approximate one percent sample of the employees within all UK firms, as a repeated snapshot every April.

From 2005 a new survey questionnaire was introduced for the ASHE, which was intended to reduce the latitude for employers' own interpretations of what was being asked of them. From 2007 there were further notable methodology changes. Before, occupations were classified as follows: if the respondent stated an employee's job had not changed in the past year, then the previous year's occupational classification was applied - otherwise, it was manually coded. Afterwards an automatic coding, text recognition, tool was used. "The effect of using ACTR was to code more jobs into higher paying occupations. The jobs that tended to be recoded into these higher paying occupations generally had lower levels of pay than the jobs already coded to those occupations. Conversely, they tended to have higher levels of pay than the other jobs in the occupations that they were recoded out of. The impact of this was to lower the average pay of both the occupation group that they had moved from and that they had moved to." From 2007 the sample size of the ASHE was reduced by 20 percent, with reductions targeted at those industries exhibiting the least variation over time in earnings patterns.

We use the ASHE annual cross-sections for each year from 2002 to 2016 and construct a panel as follows: in case of multiple jobs per individual, we exclude non-main jobs. In case of missing main job markers, we impute these based on the job with the highest working hours. In a next step, we link employees across consecutive years based on their unique personal identifiers. We can also impute missing enterprise reference numbers (entrefs) backwards, since the ASHE contains a variable which indicates whether an employee is holding the same job as in the last reference period. After linking two consecutive years in this way, we use local unit identifiers to impute missing entrefs across individuals within the same year (the ONS states that the local unit identifiers are not consistent across years, but rather they are created to identify establishments within years). We continue to update missing entrefs in this way back to and including 2002. While for the years 2003-16 we are only imputing values for a couple of missing entrefs per year,

in 2002 we add a large proportion. We could also impute entrefs in this way for 1998-2001, but the sample then becomes increasingly unrepresentative of the UK employee population.

We only keep observations for individuals aged 25-64 in each period, and which have not been marked as having incurred a loss of pay in the reference period through absence, employment starting in the period, or short-time working, and which are marked as being at an adult rate of pay (i.e. dropping trainees and apprenticeships). This is practically the same filter applied by the ONS in their annually published results on UK "Patterns of Pay" using the ASHE. We drop observations with missing records for basic hours, gross weekly earnings, or hourly wage rates. Basic hours are intended by the survey to be a record for an employee in a normal week, excluding overtime and meal breaks. Gross weekly pay is the main recorded value in the survey, and from this overtime records are subtracted. Hourly wage rates are then derived by the ONS through dividing by basic hours worked. We drop observations with over a hundred or less than one basic hour worked, as these could reflect measurement error and the mistaken inclusion of overtime in the usual hours of work. Full-time is defined as working over thirty basic hours in a week. But there are a tiny number of discrepancies in some years, we believe relating to teaching contracts, where the definition applied by the ONS differs. We however recode these such that for all observations the thirty hours threshold applies. To further address some potential for measurement error, we drop observations whose derived hourly rate of pay, excluding overtime, is less than eighty percent of the applicable National Minimum Wage (NMW) each April, allowing for the different age-dependent rates of the NMW over time (which, in this application, is always the highest adult rate, given that we restrict our attention to individuals at least 25 years old). We set this threshold lower than the NMW to avoid dropping observations where employers have rounded pay figures about the NMW, where the degree of rounding could vary with the actual value of the NMW, a behaviour on the part of employers which has been hypothesised by the ONS.

To create a tenure variable, we use the recorded employment start date of individuals. The ASHE contains information on when an employee started working for an enterprise from 2002 onwards. We drop a tiny number of unrealistic entry dates, where the start date lies in the future, or where it implies an employee started working aged fifteen or younger. There are some inconsistencies across years in these records. First, an employee can be employed by the same enterprise for three consecutive years, holding the same job, but the start dates recorded in the first and third years, though identical, can vary from the second. In this case we update the "one-off" deviation with the value of the previous year. Second, if we observe an employee in a chain of consecutive years in the same firm, holding the same job, but the start dates differ for some years, then we impute the earliest date available. Finally, we use the employment start date to impute a tiny number of missing entrefs for employees backwards to 2002 again. This allows us to not have to observe employees in a chain of consecutive years to make imputations. Again, we then use within-year local unit identifiers to update longitudinal entrefs within a year, for a handful of employees with missing entrefs.

The ASHE contains the number of employees in an enterprise as listed in the administrative Inter-Departmental Business Register (IDBR). A small minority of employees in the same enterprise and year have missing or varying values for this variable. We impute the same value for all employees within year and enterprise as the modal value for the firm. For 2002-10 occupations are classified using the four-digit SOC2000, and for 2011-16 using the SOC2010. We use both classifications in our analysis, rather than cross-walking. When we use 2-digit occupations as control variables, the base or excluded category is 41 – Administrative Occupations. When we use 3-digit occupations in our robustness checks, the excluded category is 211 – Science Professionals. To derive a firm's time -invariant industry classification, we first convert ONS Standard Industrial Classification (SIC) 2007 to 2003, using files made available by the UK Data Service. This conversion uses the 2008 Annual Respondents Dataset, where both classifications were applied, and where any 2007 code mapping to multiple 2003 codes was decided using whichever of the two bore a greater share of economic output. We then take the modal SIC2003

section (one-digit) classification of the firm in the sample. We then group industry sectors as follows: Manufacturing/Construction/Engineering, or just "Manufacturing," is given by SIC2003 sections C-F; Retail/Wholesale/Services or just "Non-financial (sales) services", is given by G-H; 'Financial services' is given by J-K; and Primary/public/other services, or just "Other", is given by A-B, I and L-Q. The Manufacturing sector is the excluded category in the regression models. We assign each firm to the public or private sector using their modal value in the sample of the ASHE variable "idbrsta", which records the legal status of the enterprise according to the IDBR. We assign "Private" to be private companies, sole proprietors and partnerships, with everything else being Public, including central government and local authorities. We derive an individual's birth cohort by taking their modal value of the dataset year minus their age within the sample, and in the regression models the excluded category is the earliest cohort. The excluded year effect is for 2002.

The household survey-based pay statistics in Figures 1 and A1 are derived using the longitudinal British Household Panel Survey (2002-08) and its successor the Understanding Society Survey (2009-16). They use the Great Britain and Northern Ireland samples and waves, but without any other boost samples. Years refer to tax years (April-March). Some individuals (less than 1%) were interviewed twice in these periods, in which case we use their first earnings response. The hourly wage is estimated from responses by employees aged 25-64 for monthly earnings in their main job. It is derived by taking usual monthly pay, converting this to a weekly figure (multiplying by 3/13), and then dividing by the sum of usual normal and usual overtime weekly hours. Only observations with usual weekly hours between 1 and 100 hours were used. Hourly wages below 0.8 of the applicable National Minimum Wage rate were dropped. Any individuals with missing values for pay, hours, age, sex or interview date were excluded from the statistics.

TABLE A1: Descriptive statistics for employees: comparison of the Analysis sub-sample and the Whole ASHE, 2002-16

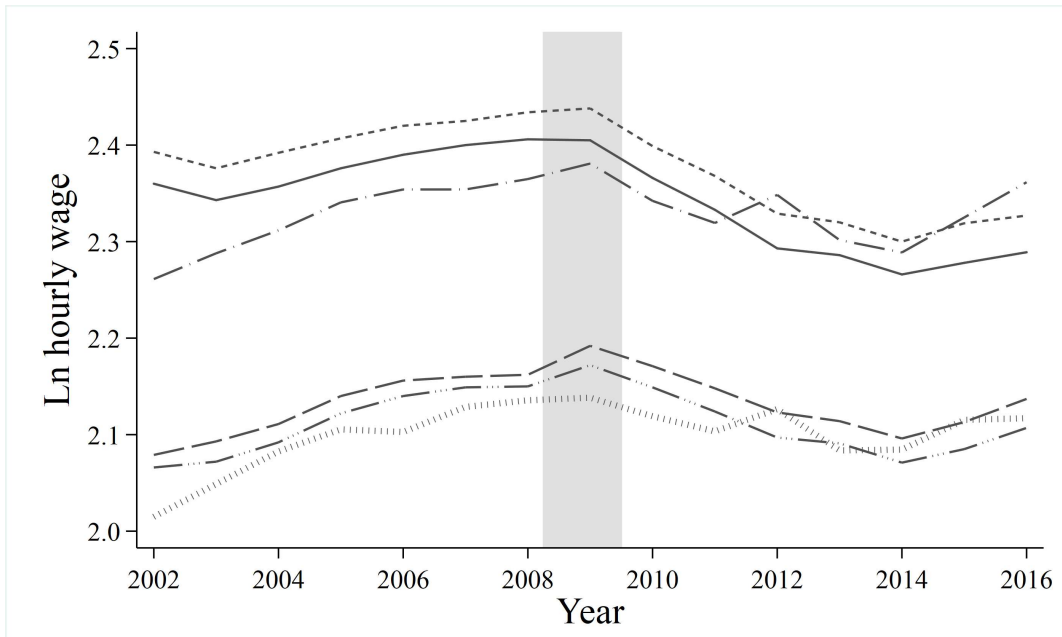|	Analysis		Whole ASHE	
	Male (1)	Female (2)	Male (3)	Female (4)
Mean ln real hourly wage (2002 prices)	2.44	2.21	2.41	2.20
Median ln real hourly wage	2.38	2.13	2.34	2.11
St. dev. of ln real hourly wage	0.55	0.49	0.55	0.48
Mean ln real weekly earnings	6.02	5.52	5.99	5.49
Share employed full-time	0.92	0.59	0.91	0.58
Mean usual weekly hours if full-time	38.5	36.8	38.7	36.9
Mean usual weekly hours if part-time	19.5	19.5	19.4	19.4
Share in private sector	0.71	0.48	0.74	0.53
Mean age (years)	42.6	42.9	42.8	42.9
St. dev. of age	10.4	10.2	10.6	10.3
Mean firm size (n. of employees)	18,343	20,245	15,574	18,049
St. dev. of firm size	42,410	42,110	39,617	40,251
N	824,806	883,326	971,830	990,997

Notes.- author calculations using the ASHE 2002-16, all employees age 25-64. £2002. Pay and hours worked excludes overtime. See Figure 1 and text for details of sample and variable construction.

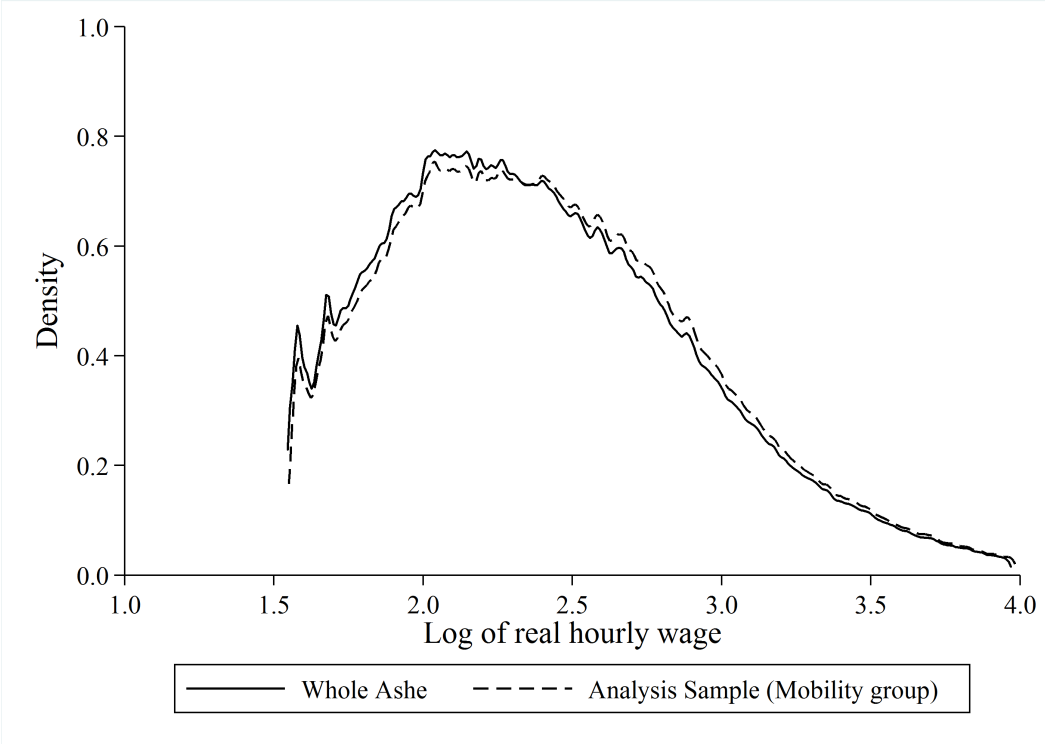FIGURE A1: Mean and median log real hourly pay of UK men and women, 2002-16

A. Mean

B. Median

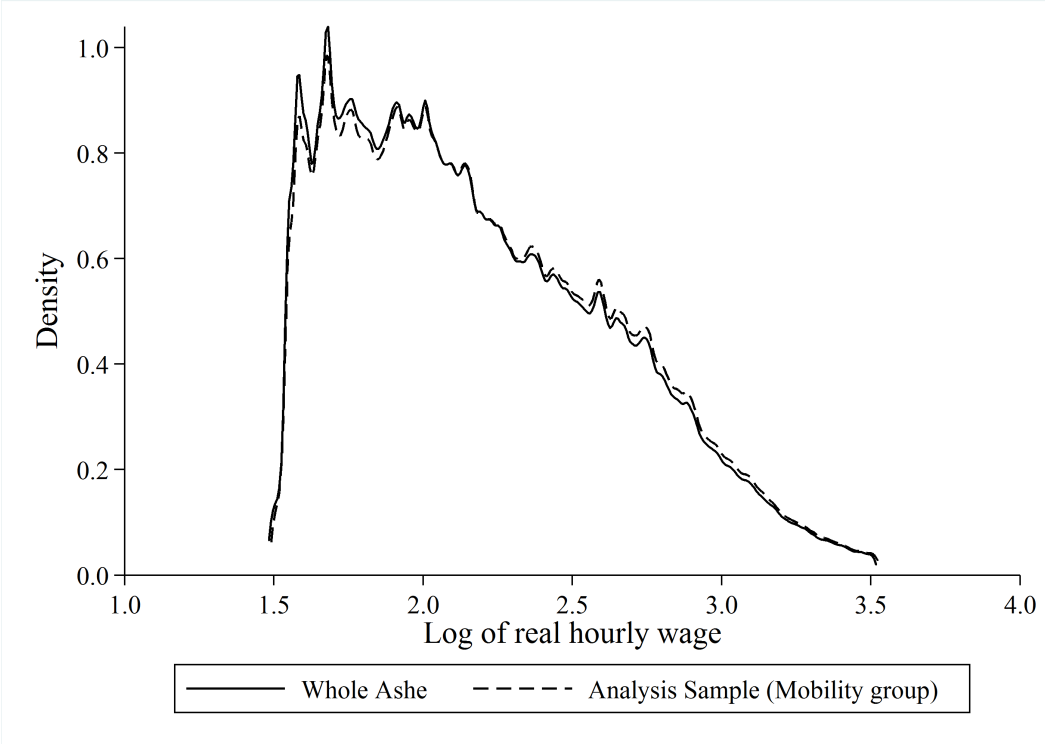|			
Male: Whole ASHE	Analysis	BHPS/USS	
Female: Whole ASHE	Analysis	BHPS/USS	
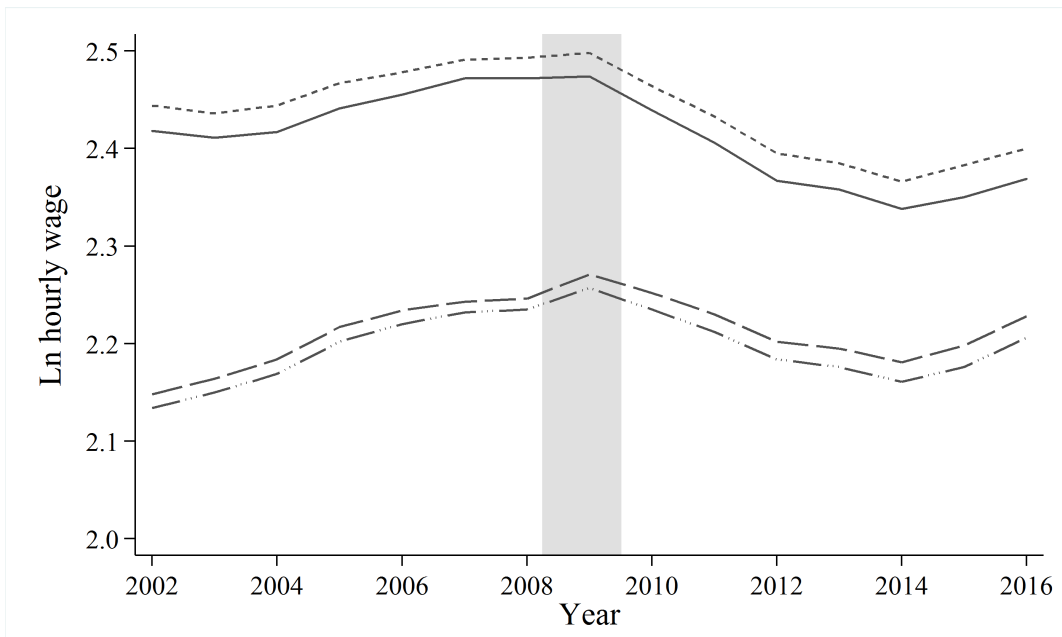
Notes.- see Figure 1.

FIGURE A2: Distributions of male and female real hourly pay: comparison of Whole UK-representative ASHE sample with the Analysis sub-sample used for the main results

A. Male

B. Female

Notes.- see Figures 1 & 2

Appendix B. Robustness of the regression model

A key assumption of the AKM-type models is that the mobility of workers is exogenous to the unobserved time-varying heterogeneity of the worker-firm match: for example, match-specific shocks to wages and productivity, which could induce workers to switch between firms. Card et al. (2013, 2016) and Card et al. (2018) have demonstrated how to carry out a simple and transparent test of this assumption, which we apply to our Analysis sample from the ASHE. This test takes the form of an event study of how wages change when employees switch between firms. If the assumptions of the AKM-type Full wages model are correct, as described in the main text, that firms pay proportional wage premiums to all of their employees, then we would expect to observe that employees who switch from firms where their co-workers are relatively low-paid (in the economy) would then experience wage increases after moving to firms where their new co-workers were relatively high-paid, and vice versa. The model also predicts that the wage gains and losses for employees moving in the different directions between any two firms should be symmetric. Furthermore, wages before and after switching should be relatively stable, i.e. firm switching is not driven by a deterioration or an expected future growth in match-specific quality.

We construct this event study using the Analysis sample of employees as follows. First, we assign to each employee, in every year, the mean wage among all of his co-workers in the same firm and year (excluding himself). Then, year-by-year, we assign each employee to the quartile of how relatively well or poorly paid on average those co-workers were, compared with all the other employees and their respective co-workers in the economy. Then we select a sample of employee-firm switching events from within the Analysis sample. For a switch to be included in the event study, an employee must be observed working in the same firms in the two years before and after their switch took place. We refer to such employees as 'switchers'. Throughout the whole sample period, this gives us a sample of 21,455 switching events. We then define 16 sub-types of event, defined by the co-worker wage quartiles of the switchers before and after they switched firms. For each of these event types, and the employee-employer relationships represented by each, we then compute the mean log real hourly wages of the switchers, conditional on the number of years before or since switching. These statistics are displayed in Figure B1. For example, in the first panel below, we plot how, on average, real hourly wages evolved for switchers who were originally in the lowest quartile for average co-worker real hourly wages, before they then switched to a different firm. As predicted by the model, there is a step-change increase in individual employee wages for those who moved to a firm where their new co-workers were then relatively high-paid. Whereas, those who moved between firms where their co-workers were similarly low-paid experienced no substantial wage increases. Similarly, in the last panel below, we can see that among those employees who switched away from firms with co-workers who were relatively well-paid, individuals then on average experienced larger real hourly wage decreases if they switched to firms with relatively low-paid co-workers. In other words, Figure B1 shows that employees do experience step-change wage increases (decreases) when they switch to firms where their new co-workers are more (less) high-paid than their old co-workers. Furthermore, the magnitude of these observed average employee wage changes upon firm switching display some symmetry, as is also predicted by the Full wages model. Also, note that none of the series in any panel of Figure B1 cross, nor are there noticeably different trends before and after in the average employee wages across the different types of switching (high to low, high to high, low to low etc.).

FIGURE B1: An event-study of average employee log real hourly wages before and after switching firms, depending on the quartile of co-worker average wages in the old and new firms
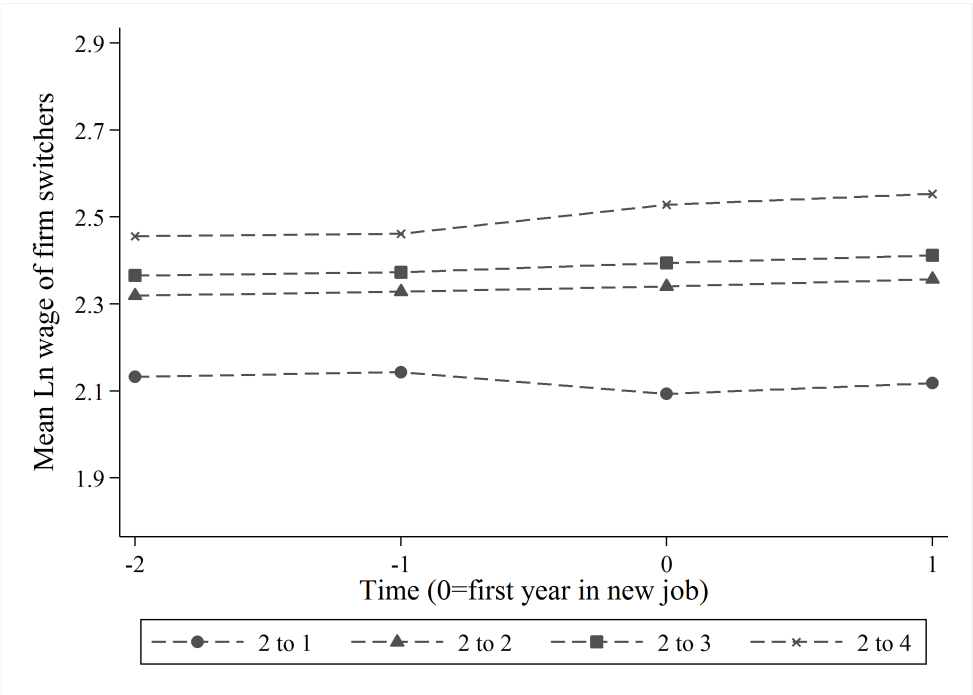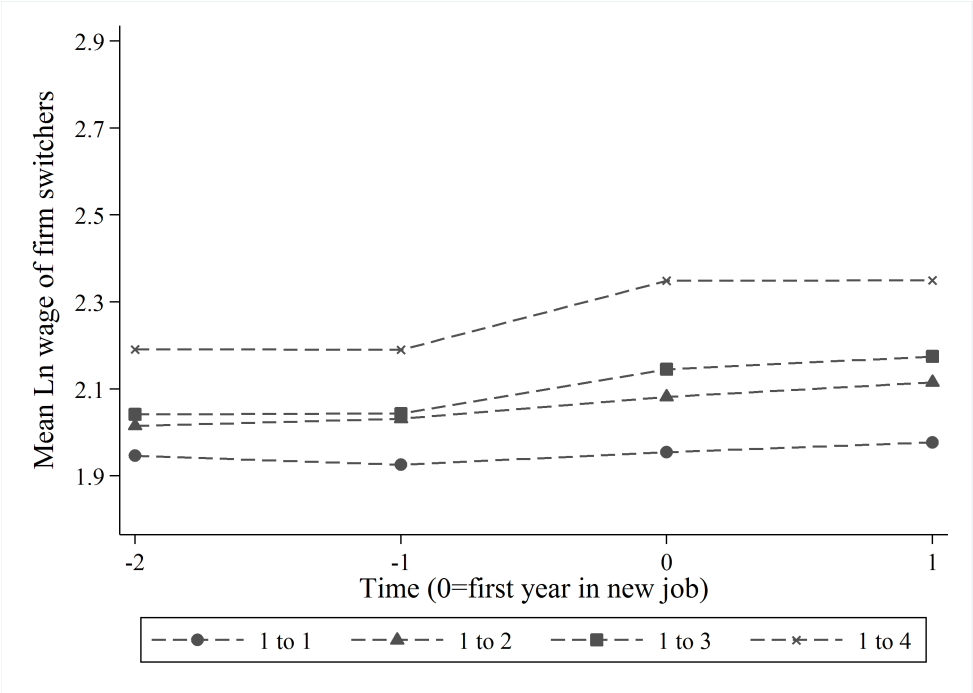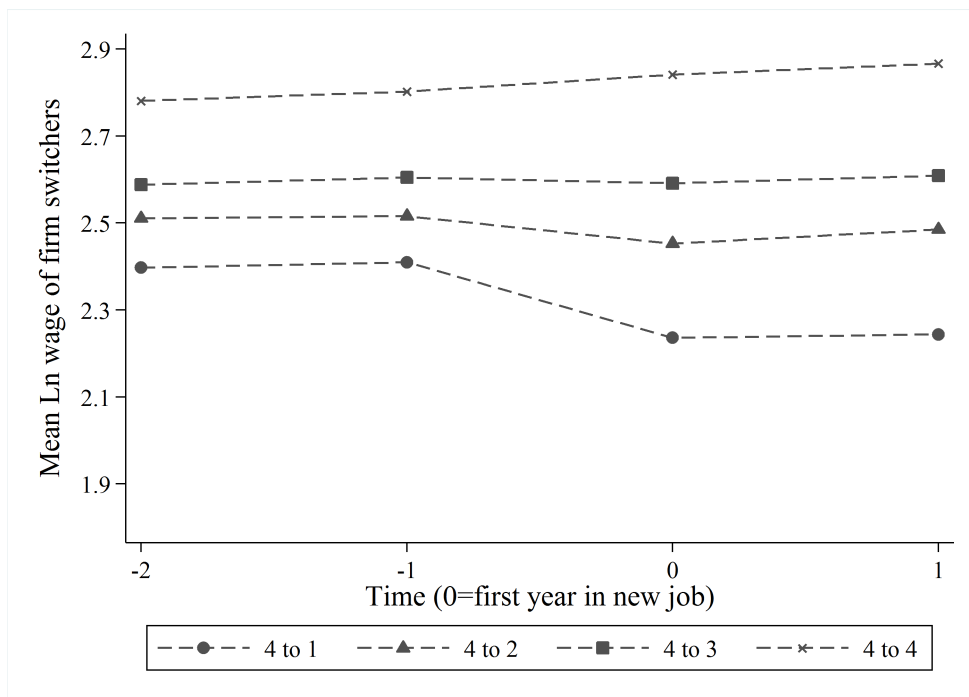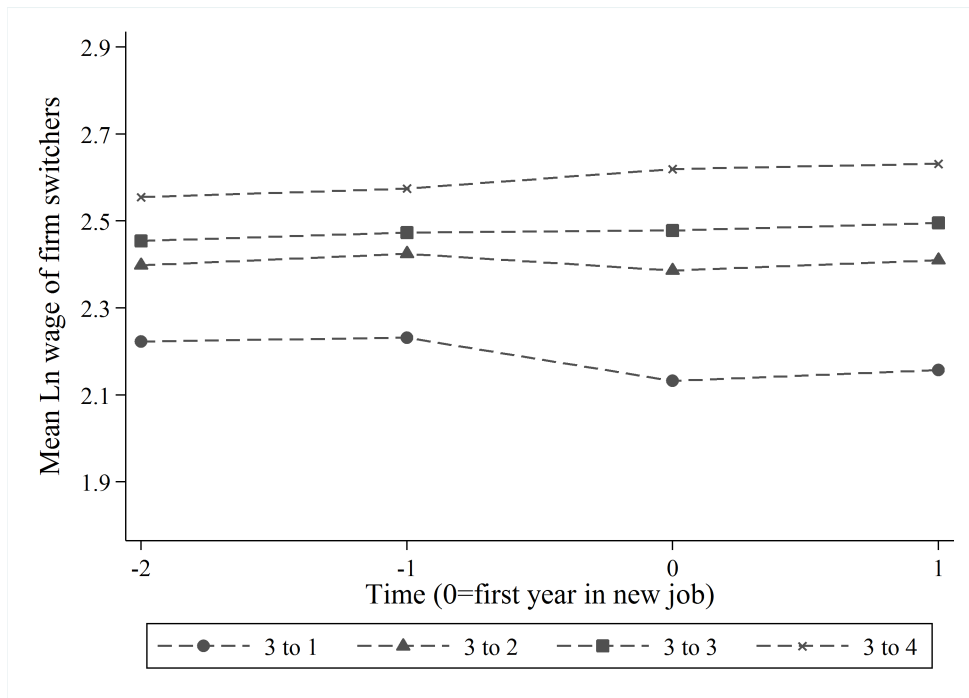
Notes.- "X to Y" indicates the quartile of co-worker wages for employees in their old firm (X), from which they move to their new firm (Y). Each event series uses firm switchers throughout the ASHE 2002-2016 Analysis sample period. The numbers of observed switching events studied are as follows: "1 to 1", 3,498; "1 to 2", 1,073; "1 to 3", 705; "1 to 4", 458; "2 to 1", 931; "2 to 2", 1,582; "2 to 3", 1,660; "2 to 4", 768; "3 to 1", 583; "3 to 2", 1,018; "3 to 3", 2,573; "3 to 4", 1,194; "4 to 1", 385; "4 to 2", 571; "4 to 3", 1,249; "4 to 4", 3,207.

TABLE B1: Regression model estimation results

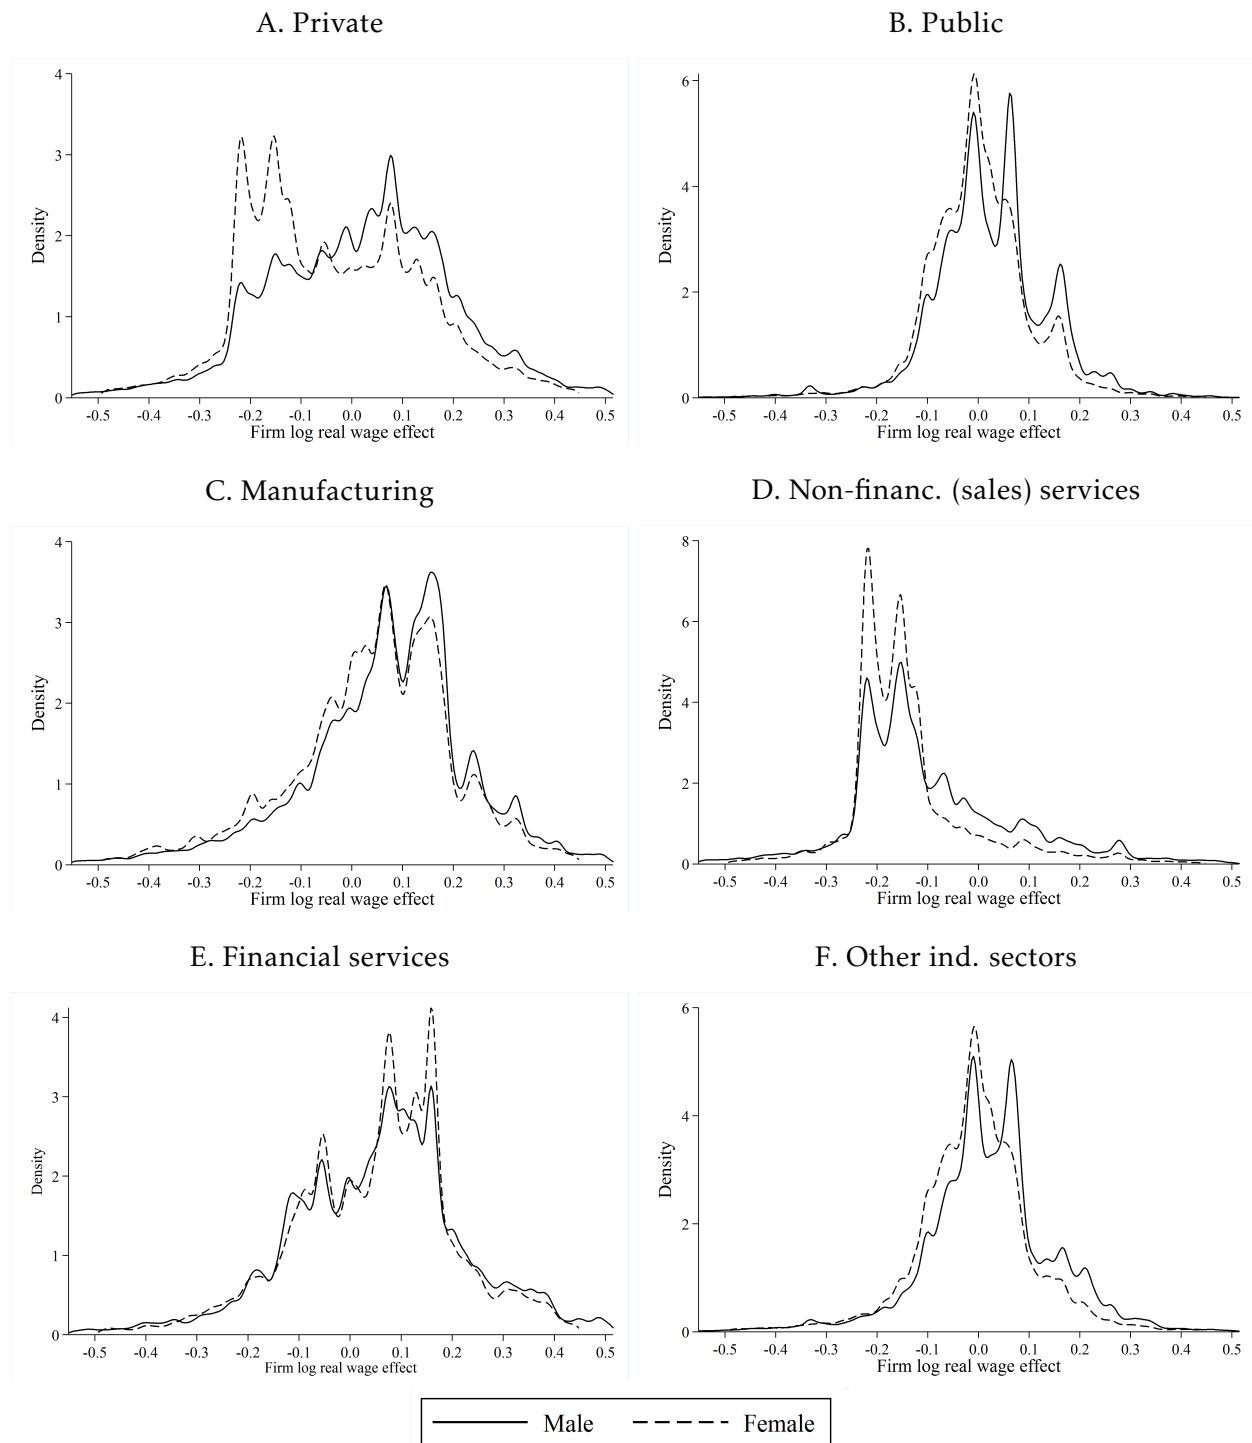|	Basic		Full
	(1)	(2)	(3)
Female	0.149	0.145	
	(0.003)	(0.003)	
Age (years) squared	$3.5x10^{-4}$	$7.9x10^{-5}$	$1.3x10^{-4}$
	$(1x10^{-5})$	$(2x10^{-5})$	$(2x10^{-6})$
Age (years) cubed	$-5x10^{-6}$	$-2x10^{-6}$	$-4x10^{-6}$
	$(1x10^{-7})$	$(2x10^{-7})$	$(3x10^{-8})$
Tenure (years)	0.021	0.022	0.019
	(0.001)	(0.001)	(0.000)
Tenure (years) squared	$-7x10^{-4}$	$-7x10^{-4}$	$-7x10^{-4}$
	$(4x10^{-5})$	$(4x10^{-5})$	$(1x10^{-5})$
Tenure (years) cubed	$-9x10^{-6}$	$-9x10^{-6}$	$-1x10^{-5}$
	$(7x10^{-7})$	$(8x10^{-7})$	$(2x10^{-7})$
Full-time	0.079	0.076	-0.042
	(0.004)	(0.004)	(0.0009)
Private sector		-0.009	
		(0.015)	
Non-financ. (sales) services		-0.172	
		(0.011)	
Financial services		0.075	
		(0.010)	
Other ind. sectors[*]		-0.031	
		(0.013)	
Constant	1.692	2.058	
	(0.009)	(0.040)	
Occupations	Yes	Yes	Yes
Firm size (n. emplys, cubic)	Yes	Yes	Yes
Year effects	Yes	Yes	Yes
Birth cohorts	No	Yes	No
Worker effects	No	No	Yes
Firm effects	No	No	Yes
R^2	0.549	0.566	0.923
N		1,708,132	

Notes.- author calculations using the ASHE, all employees aged 25-64. Pay excludes overtime. See text for details of sample and variable construction. Column (1) gives results for the Basic regression model excluding covariates which are either time-invariant over employees or firms. Column (2) adds some of these time-invariant factors to the model. Column (3) is the estimated Full regression model given by equation (1).

Standard errors were estimated using firm-level clusters, and are shown in parentheses.

[*] The excluded grouped industry category is "Manufacturing."

Appendix C. Additional figures

FIGURE C1: Distribution of estimated firm-specific wage effects ($\widehat{\phi}_{J(it)}$): employees in the private and public sectors, and working within groups of industry sectors

A. Private

B. Public

C. Manufacturing

D. Non-financ. (sales) services

E. Financial services
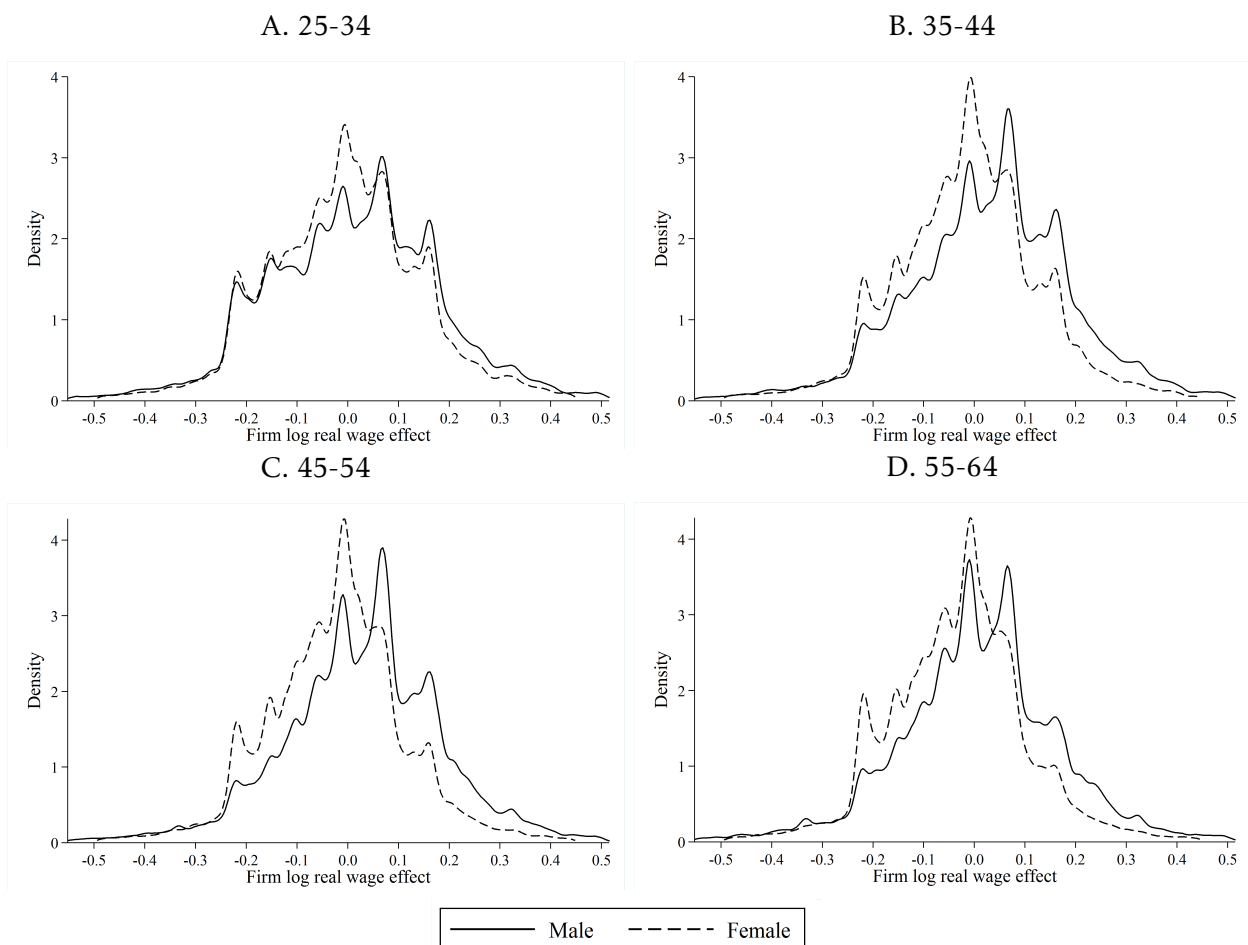
F. Other ind. sectors

Notes.- see Figure 3. Both male and female kernel densities were estimated with a bandwidth of one log point. The top and bottom one percent of the overall set of estimated firm-specific effects are not displayed in any of the sub-figures. See Table 2 and Appendix A for descriptions of the industry groupings. In (a), the sample size of employee-years is 580,000 male and 430,000 female; in (b) 240,000 and 460,000; in (c) 190,000 and 60,000; in (d) 130,000 and 150,000; in (e) 170,000 and 150,000; in (f) 330,000 and 530,000;

FIGURE C2: Distribution of estimated firm-specific wage effects ($\widehat{\phi}_{J(it)}$): employees in part- and full-time jobs
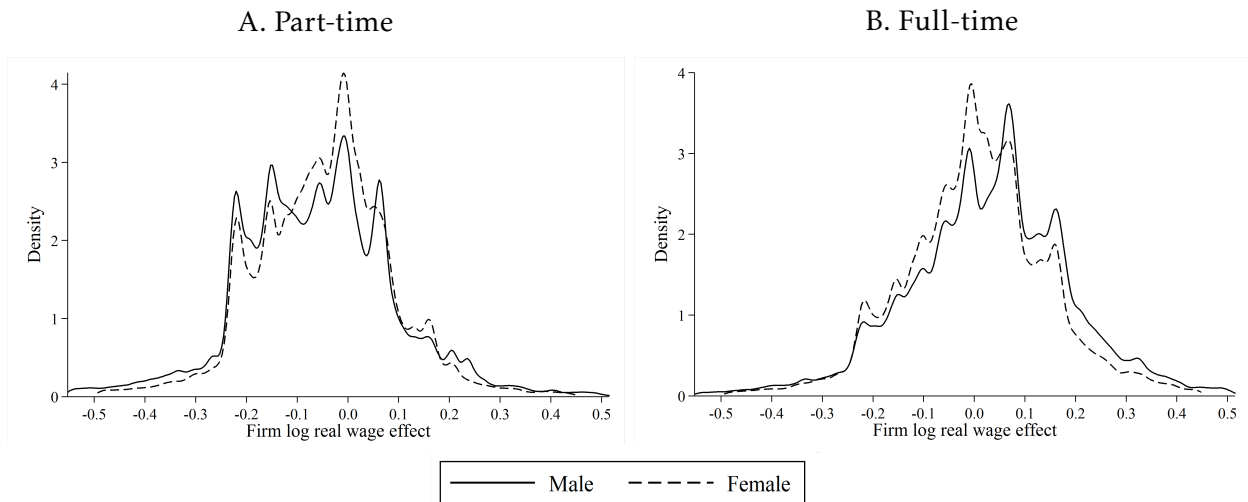
A. Part-time

B. Full-time

Notes.- see Figure 3 and C1. Both male and female kernel densities were estimated with a bandwidth of one log point. The top and bottom one percent of the overall set of estimated firm-specific effects are not displayed in any of the sub-figures. In (a), the sample size of employee-years is 70,000 male and 360,000 female; in (b) 760,000 and 520,000.

FIGURE C3: Distribution of estimated firm-specific wage effects ($\widehat{\phi}_{J(it)}$): age groups

A. 25-34

B. 35-44

C. 45-54

D. 55-64

Notes.- see Figure 3 and C1. Both male and female kernel densities were estimated with a bandwidth of one log point. The top and bottom one percent of the overall set of estimated firm-specific effects are not displayed in any of the sub-figures. In (a), the sample size of employee-years is 220,000 male and 230,000 female; in (b) 250,000 and 260,000; in (c) 230,000 and 260,000; in (d) 130,000 and 130,000.