

Who works for whom and the UK gender pay gap

Sarah Louise Jewell

Giovanni Razzu

Carl Singleton*

Published at: *British Journal of Industrial Relations*. (2019; In Press)

<https://doi.org/10.1111/bjir.12497>

Abstract

This study reports novel facts about the UK gender pay gap. We use a representative, longitudinal and linked employer-employee dataset for 2002-16. Men's average log hourly wage was 22 points higher than women's in this period. We find 16% of this raw pay gap is accounted for by estimated firm-specific wage effects. This is almost three times the amount explained by gender occupation differences. When we decompose a preadjusted measure of the pay gap, we find less than 1 percentage point or a 6% share is accounted for by the gender allocation across high and low wage firms. In other words, only a small share of what is traditionally referred to as the 'unexplained' part of the pay gap is explained by the differences between men and women in whom they work for.

Keywords: gender wage gap; firm-specific wages; occupation premiums

*s.l.jewell@reading.ac.uk, g.razzu@reading.ac.uk & c.a.singleton@reading.ac.uk: Department of Economics, University of Reading, UK.

We are grateful for helpful comments and advice from two anonymous referees, Salvatore Lattanzio, Daniel Schaefer and Ian Walker, as well as from participants at the 2018 WPEG (University of Sheffield) and WOLFE (University of York) conferences, and at the 2019 annual conferences of the Royal Economic Society, Scottish Economic Society and European Society for Population Economics.

This work is mostly based on the Annual Survey of Hours and Earnings dataset (Crown copyright 2017), having been funded, collected and deposited by the Office for National Statistics (ONS) under secure access conditions with the UK Data Service (SN:6689). The UK Data Service agrees that the results are non-disclosive, and cannot be used to identify a person or organisation. The use of these data does not imply the endorsement of the data owner or the UK Data Service at the UK Data Archive in relation to the interpretation or analysis of the data. This work uses research datasets which may not exactly reproduce National Statistics aggregates.

Singleton thanks the Economic and Social Research Council (UK) for funding support under Grant No. ES/J500136/1.

1 Introduction

The gap between the average hourly pay of all male and female employees in the United Kingdom stood at 17.1% in April 2018.¹ This study contributes to the literature on the pay differences between men and women by assessing what role firm-specific wage premiums had in the UK's pay gap between 2002 and 2016.

To do so, we estimate a wages model which allows for worker and firm-specific fixed wage effects, following the contributions of [Abowd et al. \(1999\)](#) and [Abowd et al. \(2002\)](#). In our analysis sample, the observed mean log hourly pay gap (male minus female) among employees aged 25-64 was 22.3 points.² We find that 16% of this gap is accounted for by estimated firm-specific wage effects, implying that men have a greater tendency than women to work for firms which on average pay higher wage premiums to their employees. Although this may not appear to be a large share, to put it into perspective, the equivalent contribution from the fact that men and women work in different occupations, which also have different wage premiums, is just 6%. However, the vast majority of the pay gap is explained by the fixed characteristics of workers, which affect their wages irrespective of what jobs they are in, such as their education, preferences and work experience.

We also look at the importance of how workers are allocated to firms in terms of the adjusted gender pay gap, which is typically reported alongside the raw gap by commentators and policy makers, since it can account for differences in the other observable characteristics of workers or jobs relevant for wages, such as age, tenure, occupation, industry sector and full-time status. We do this by applying the [Gelbach \(2016\)](#) decomposition. Out of an adjusted mean log hourly pay gap of 14.5 points, we find that an estimated share of 6% is contributed by the conditional allocation of employees over firm-specific wage premiums, with the remainder accounted for by the fixed characteristics of workers, which are transferred across firms.

These findings contribute to a vast literature which 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 categories: productivity, preferences and discrimination, which are all interrelated ([Altonji and Blank, 1999](#)). With the diminishing gender gaps in education and labour market participation in the majority of developed countries, the importance and focus on explanations from the first category, especially human capital-based ones, has lessened. Pay gaps nonetheless persist and are pervasive. More recent work has looked to worker preferences and psychological attributes for explanations, due to their impacts on productivity, choices and beliefs (see for reviews [Croson and Gneezy, 2009](#); [Bertrand, 2011](#); [Azmat and Petrongolo, 2014](#)). The role of firms cuts across these sets of explanations. Early work found that US women were more likely to work for low wage firms than men, and vice versa regarding high wage firms (e.g. [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 high wage firms ([Goldin et al., 2017](#); [Barth et al., 2017](#)).

The only studies 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 by [Mumford and Smith \(2007, 2009\)](#) and [Drolet and Mumford \(2012\)](#). These authors used cross-sectional linked employer-employee data to disentangle the influence of observable employer and employee characteristics. They showed that the proportions of women relative to men in occupations and firms of different types did account for part of the UK gender pay gap.

However, these studies, as well as those which have looked at other countries using a similar method, were hampered by being unable to address simultaneously the unobservable fixed heterogeneity over workers and firms in the determination of wages, i.e. they lacked longitudinal linked employer-employee datasets. Therefore, their results cannot be directly compared with what we find here.

Similarly, our main findings are not directly comparable with the majority of the recent UK-focused gender pay gap literature, which has mostly used sources of longitudinal household survey data. This literature could not control robustly for the potential influence of how male and female employees were allocated across firm types, which could systematically differ by gender and be correlated with the other determinants of wages. It is plausible that omitting this factor, which generally explains a significant fraction of overall UK wage variation, could have confounded previous results. For example, [Costa Dias et al. \(2018\)](#) demonstrate the importance of accumulated years of work experience and working hours 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. But it is well-known that the measured returns to tenure and work experience are likely to be upwards biased unless the unobserved worker-firm-match quality is controlled for ([Topel, 1991](#)). Our main findings show that the way in which employees are assigned to firms can explain only a small fraction of the preadjusted UK hourly pay gap. Therefore, we add support to some of the previous conclusions about the determinants of UK gender pay differences, which were based on household survey data from a similar period, but which could not have addressed the potential influence of how workers were matched to firms.

More direct comparisons of our main findings are possible with a few recent studies of the gender pay gaps in other countries (e.g. [Card et al., 2016](#); [Sorkin, 2017](#); see Section 5.4 for a full comparison with several relevant studies). In particular, [Cardoso et al. \(2016\)](#) (henceforth CGP) looked at how much of the Portuguese hourly wage 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 wage gap over three decades. Our methodological approach is close to that of CGP, 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 first adjusting for the influence of both time-varying and fixed observable wage determinants, and not just the former as in the case of CPG.

There is an active policy context related to this study. UK law recently changed such that all British employers with at least 250 employees must annually publish their own gender pay gaps.³ Despite the potential significant economic costs to firms of complying with this legislation, there is no robust evidence, which is representative of the whole UK labour market, that shows the gender pay gap is not an issue of differences in pay between firms rather than within them.⁴ It will be impossible to address this evidence gap using the pay gap data collected from firms under the new legislation. Quite simply, there is no robust way to address how much of the pay gaps reported by firms, and the differences therein between firms, are explained by workforce composition. Therefore, with this study we address a 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 gender pay gap.

The remainder of the paper is organised as follows: Section 2 describes the data and our longitudinal employer-employee sample construction; Section 3 outlines the methods used to decompose measures of the gender pay gap. Section 4 presents the main results; Section 5 discusses the robustness of our findings and compares these UK results with those from other countries; and Section 6 concludes.

2 Data

The main data source is the Annual Survey of Hours and Earnings (ASHE), 2002-16, which is based on a 1% random sample of UK employees, drawn from HM Revenue and Customs Pay As You Earn (PAYE) records, collected and administered by the [Office for National Statistics \(2019\)](#). Questionnaires are sent to employers, who are legally required to complete these with reference to payrolls for a period in April. Because the randomness of the sample in every year is based on all the individuals having the same last two-digits for their personal lifetime National Insurance number, this dataset can be viewed as a panel of employees without attrition and with replacement.⁵

2.1 Analysis sample construction

Particularly valuable for our analysis are the longitudinal identifiers for individuals and enterprises contained within the ASHE. 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 for every year between 2002 and 2016, which we describe in more detail in [Online Appendix A](#).

We include only the main job observation of an individual in any year, which must not be at a trainee or an apprentice level, and must not have incurred a loss of pay in the reference period for whatever reason. If main job markers are missing from the ASHE, we impute these based on the job with the highest working hours. 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. 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. This measure includes any incentive pay or premiums for working nights, weekends or during public holidays. 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.⁶

We only 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, so that we only study the clear majority who would have completed full-time education by this age in the UK. Therefore, any worker fixed wage effects estimated in our regression models should account for most of the wage heterogeneity over employees associated with human capital accumulation before they have begun to engage fully with the labour market. Our analysis uses the following 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 [Online Appendix A](#). We drop the small number of employee-year observations which have missing values for any of these variables. This results in what we call the ‘Whole ASHE’ sample.

All our main results, however, use what we call the ‘Analysis’ sample. This is an 87% sub-sample of the Whole ASHE. It consists of the largest connected set, or mobility group, of workers and firms, since the method we apply only allows for the comparison of any estimated firm-specific wage effects within

a connected set of firms and workers.⁷ The largest connected set contains 1.71 million employee-year observations, with 131,903 men and 124,501 women represented. The median and mean number of years that men and women appear as employees in the sample are 10 and 9, respectively. The sample covers 86,779 different firms, where 18,404 are present in 2002, and thereafter, approximately 4-6 thousand new firms enter the sample 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.⁸ Online Appendix Table A1 shows the distribution of the number of employers worked for by employees over the whole 2002-16 period. Just less than half (48%) of the men and women represented in the sample had only one employer over this period, the majority had at least two employers and around a fifth had more than two. Table 1 compares some descriptive statistics over employee-year observations between the Analysis and Whole ASHE samples. The 87% sub-sample is similar in most regards, though in the largest connected set the average firm size (actual number of employees) and its variance is higher. This implies that larger firms are marginally over-represented in our analysis compared with the UK employee population.

2.2 Descriptive evidence on gender gaps

For all employees and each year in 2002-16, mean and median values of the raw (observed/unadjusted) gender wage gap are shown in Figure 1. In the Whole ASHE sample, the mean value declined from around 28 log points in 2002 to 16 points in 2016. The median gap similarly declined throughout this period.⁹ In the Analysis sample, both the mean and median wage gaps are consistently higher across all years than in the Whole ASHE sample, by approximately one log point. Online Appendix Figure A1 shows that both male and female average log wages are higher in the Analysis sample, though more so for male wages. Figure 1B also shows wage gap series derived using the British Household Panel Survey (BHPS, 2002-08) and Understanding Society Survey (USS, 2009-16), using a method and sample selection as comparable as possible to our Whole ASHE measures (see Online Appendix A). These other datasets, used widely in the literature (e.g. Costa Dias et al., 2018), 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. Compared with the ASHE, the BHPS/USS data have a smaller sample size and feature attrition from the panel of households and families in the survey.

We also estimate the kernel densities of log real hourly wages for the male and female observations in the Analysis sample, shown by Figure 2. The female wage distribution is more concentrated toward values of the UK's real National Minimum Wage, whereas observations with very high hourly rates of pay are dominated by men. Online Appendix Figure 2 demonstrates that the gender-specific wage distributions are practically unchanged when we move between the Whole ASHE sample and the largest connected set of employees and firms therein.

There is also evidence in these data of substantial gender segregation of employees across firms. For men in the sample, at their firms the female employee share is on average just 31%. However, for women, the share of their coworkers who are female is 71%. Mumford and Smith (2009) found similar statistics

for the UK focusing on 2004 only. The extent of gender segregation is greater in the Whole ASHE sample, suggesting that smaller firms are substantially more gender segregated than larger firms (see Table 1).

3 Method

First we estimate a so-called AKM-type wages model, which features both worker and firm fixed effects (Abowd et al., 1999). We use this to look at how much of the raw wage differences between male and female UK employees in 2002-16 are accounted for by firm-specific wage effects. We then go a small step further and apply a decomposition method suggested by Gelbach (2016). This method can address the role in the pay gap of whom employees work for, after conditioning on, or adjusting for, the influence of some observable worker, firm and job characteristics. In other words, it provides a way to decompose what is traditionally known as the ‘adjusted’ or residual wage gap, typically estimated using cross-sectional datasets, into contributions which can be addressed using longitudinal data. CGP have identified and described thoroughly the usefulness of this method for studying the role of firm fixed factors in the adjusted gender pay gap. Therefore, our treatment here is concise. However we suggest one expansion on CGP. 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 productivity.

3.1 The AKM-type Full model of wages

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 model is given by:

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

or in stacked matrix notation as:

$$\mathbf{y} = \mathbf{X}\boldsymbol{\beta} + \mathbf{A}\boldsymbol{\alpha} + \mathbf{F}\boldsymbol{\phi} + \boldsymbol{\varepsilon} . \quad (1)$$

In this equation, \mathbf{y} gives an $(N \times 1)$ 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 $(N \times k)$ matrix \mathbf{X} and the vector \mathbf{x}_{it} contain k time-varying covariates, with associated $(k \times 1)$ coefficient vector $\boldsymbol{\beta}$. The $(N \times P)$ and $(N \times L)$ 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 $(P \times 1)$ vector $\boldsymbol{\alpha}$. Similarly, the firm fixed effects, $\phi_{J(it)}$, are contained in the $L \times 1$ 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 $E[\boldsymbol{\varepsilon} | \mathbf{X}, \mathbf{A}, \mathbf{F}] = 0$.

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 or she works, and in

whatever job. The firm fixed effects, $\phi_{J(it)}$, measure relative wage premiums, which employees receive upon switching firms. The Full 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. As per an [Oaxaca \(1973\)](#) decomposition, their contribution to the raw pay gap can be measured simply as:

$$E_{it} \left[\hat{\phi}_{J(it)} \mid i \in M \right] - E_{it} \left[\hat{\phi}_{J(it)} \mid i \in F \right], \quad (2)$$

where M and F denote the sets of men and women in N , respectively.

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 connected sets of workers and firms ([Abowd et al., 2002](#)). This required connectedness does not depend on gender, since we do not allow the firm fixed effects to be gender specific. As already mentioned, in our analysis we restrict attention only to the largest identified connected set in the panel dataset. Nonetheless, this means estimating a large number of coefficients, $(k + P + L - 2)$, for the Full model.

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](#)). The additive separability of worker and firm fixed effects is a strong assumption, though one we can check later by considering whether a model which replaces these with worker-firm match-specific fixed effects improves the fit to the data. 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 ϵ 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 more detail in the context of our data and application in Online Appendix B, finding that in general the data do support it. However, this is not to say that the mobility implied by AKM-type models needs to be completely random. Worker-firm matching based on the worker or firm fixed components of wages, any relevant time-varying observable characteristics controlled for, such as whether a job is full-time, and non-wage factors, like a firm's location, can all be correlated with mobility without bias in the model's coefficient estimates.

It is a well-known feature of AKM-type models that estimates of the worker and firm fixed effects are estimated with considerable imprecision ([Card et al., 2018](#)). This is sure to be the case here when we only have data on average for around 1% percent of the employees in a firm in any given year, and only for a relatively small number of years in total. This can cause issues in some applications, especially those that seek to decompose the variance of wages. However, as noted by CGP, both the [Oaxaca](#) and [Gelbach](#) decompositions should satisfy large sample properties, since they use the full sets of estimated fixed effects. The only concern would be if the sampling errors in the estimated worker and firm fixed wage effects systematically differed by gender, though we have no reason to suspect this.

3.2 The Gelbach decomposition

To demonstrate our application of the [Gelbach](#) decomposition, it is useful if we re-write the Full model:

$$\mathbf{y} = \mathbf{X}\boldsymbol{\beta} + \underbrace{\mathbf{g}\tilde{\lambda} + \mathbf{W}\boldsymbol{\beta}_w + \mathbf{A}\tilde{\boldsymbol{\alpha}}}_{=\mathbf{A}\boldsymbol{\alpha}} + \underbrace{\mathbf{Z}\boldsymbol{\beta}_z + \mathbf{F}\tilde{\boldsymbol{\phi}}}_{=\mathbf{F}\boldsymbol{\phi}} + \boldsymbol{\varepsilon} , \quad (3)$$

where \mathbf{g} contains a dummy variable for whether or not an individual is male, and $\tilde{\lambda}$ measures the wage gap conditional on the other factors in the model. The $(N \times p)$ and $(N \times l)$ matrices \mathbf{W} and \mathbf{Z} contain p and l time-invariant observable worker and firm characteristics, respectively. Therefore, $\tilde{\boldsymbol{\alpha}}$ and $\tilde{\boldsymbol{\phi}}$ measure the effects on employee wages of all unobserved (or residual) fixed worker and firm factors. But the coefficients contained in $\{\tilde{\lambda}, \boldsymbol{\beta}_w, \tilde{\boldsymbol{\alpha}}, \boldsymbol{\beta}_z, \tilde{\boldsymbol{\phi}}\}$ cannot be all 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 model is given by:

$$\mathbf{y} = \mathbf{g}\lambda + \tilde{\mathbf{X}}\tilde{\boldsymbol{\beta}} + \mathbf{e} . \quad (4)$$

The adjusted gender pay gap in (4) is given by the coefficient λ . The covariates are contained in the $(N \times [k + p + l])$ matrix $\tilde{\mathbf{X}} = [\mathbf{X}, \mathbf{W}, \mathbf{Z}]$. 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 model. In a standard exploration of the 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 job is in the private sector. This is the main way we differ from the application of the [Gelbach](#) decomposition by CGP, who only included a more limited set of k time-varying covariates in their Basic model. The vector \mathbf{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 model, Equation (1), with the $(1 \times N)$ row vector:

$$\begin{aligned} \boldsymbol{\gamma}' &= \left(\mathbf{g}' \left[\mathbf{I} - \tilde{\mathbf{X}}(\tilde{\mathbf{X}}'\tilde{\mathbf{X}})^{-1}\tilde{\mathbf{X}}' \right]' \left[\mathbf{I} - \tilde{\mathbf{X}}(\tilde{\mathbf{X}}'\tilde{\mathbf{X}})^{-1}\tilde{\mathbf{X}}' \right] \mathbf{g} \right)^{-1} \mathbf{g}' \left[\mathbf{I} - \tilde{\mathbf{X}}(\tilde{\mathbf{X}}'\tilde{\mathbf{X}})^{-1}\tilde{\mathbf{X}}' \right]' \\ &= (\boldsymbol{\Gamma}'\boldsymbol{\Gamma})^{-1} \boldsymbol{\Gamma}' , \end{aligned} \quad (5)$$

where \mathbf{I} is the $(N \times N)$ identity matrix. Pre-multiplying some variable, contained within an $(N \times 1)$ 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 $\tilde{\mathbf{X}}$: i.e. it ‘adjusts’ the observed gender gap. If we replaced $\tilde{\mathbf{X}}$ in (5) with a vector of 1’s, then pre-multiplying \mathbf{y} by $\boldsymbol{\gamma}'$ would give the actual mean gender pay gap observed in the data. Using (5), we can write the following decomposition equation for the components of the adjusted pay gap:

$$\underbrace{\boldsymbol{\gamma}'\mathbf{y}}_{\hat{\delta}_y = \hat{\lambda} - \text{adjusted gender pay gap}} = \underbrace{\boldsymbol{\gamma}'\mathbf{A}\tilde{\boldsymbol{\alpha}}}_{\hat{\delta}_a - \text{dist. worker effs}} + \underbrace{\boldsymbol{\gamma}'\mathbf{F}\tilde{\boldsymbol{\phi}}}_{\hat{\delta}_f - \text{gender-firm sorting}} + \underbrace{\boldsymbol{\gamma}'\boldsymbol{\varepsilon}}_{=0} , \quad (6)$$

noting that $\widehat{\mathbf{A}\alpha} = \mathbf{A}\widehat{\alpha} - \mathbf{W}\beta_w$ etc. The term on the left-hand-side of (6) is an estimate of the adjusted gender pay gap, and is equivalent to the least squares estimate $\widehat{\lambda}$ of the Basic model (4). The first term on the right-hand-side of (6), $\widehat{\delta}_a$, 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’.¹⁰ We can also write this effect as follows:

$$\widehat{\delta}_f = E_{it} \left[\widetilde{\phi}_{J(it)} \mid i \in M, \widetilde{\mathbf{X}}\widetilde{\beta} \right] - E_{it} \left[\widetilde{\phi}_{J(it)} \mid i \in F, \widetilde{\mathbf{X}}\widetilde{\beta} \right]. \quad (7)$$

The final term in (6) is equal to zero by the standard assumptions of the AKM-type Full model: $E[\widehat{\epsilon} \mid \widetilde{\mathbf{X}}\widetilde{\beta}] = 0$. If in the expression for γ' we replace all occurrences of $\widetilde{\mathbf{X}}$ with \mathbf{X} (i.e. if we assume that the Basic model does not include any time-invariant factors: \mathbf{W} and \mathbf{Z} are equal to $\mathbf{0}$), then Equation (6) accounts for the role of all estimated worker or firm time-invariant heterogeneity, as per the application by CGP: the gender-firm sorting effect becomes

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

4 Main Results

4.1 Estimating the Full model & decomposition of the raw pay gap

Table 2 summarises the estimation of the AKM-type wages model, described by Equation (1). The time-varying observable characteristics in \mathbf{X} are: year effects, squared and cubed employee age, cubic polynomials for employee tenure and firm size, and dummy variables for full-time and occupations (2-digit classification: 26 categories). Differences in employee education, general human capital and work experience prior to entering the sample period are captured by the worker fixed effects. Columns (1) and (2) show statistics over male and female worker-year observations, respectively, whereas Column (3) combines both genders. The correlation of the estimated worker and firm fixed effects is negative for men and women, and only slightly positive overall. However, this is almost sure to be substantially biased downwards given the nature of our sample and limited mobility bias (Andrews et al., 2008, 2012). One key assumption of the AKM modelling framework is the additive separability of the worker and firm fixed effects. As a rough test of this assumption, studies compare the fit of the AKM-type wages model with an equivalent model which admits instead worker-firm match-specific fixed effects (e.g. Card et al., 2013). When we apply this test here, we only find a small improvement in model fit, with the adjusted R^2 increasing by 2 percentage points to 0.92, suggesting that the additive separability of the worker and firm fixed effects is a workable assumption with these data.

For completeness, we also show the results from decomposing the variance in worker wages in the Analysis sample into component shares accounting for the role of worker and firm fixed wage effects, their covariance, and other factors in the model.¹¹ We don’t dwell on these results due to the well-known sources of finite sample bias on these estimated shares, other than to highlight that the firm effects share overall is 12%. 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 $covar(w_{it}, \hat{\phi}_{J(it)})/var(w_{it}) = 13.2\%$.¹² Systematic differences in firm wage premiums do account for a sizeable fraction of overall UK wage dispersion. Figures 3A and 3B emphasise this further, by plotting kernel density estimates of $\hat{\alpha}_i$ and $\hat{\phi}_{J(it)}$, respectively, by gender. The distribution of the firm-specific wage premiums received by male employees in the UK labour market is visibly more positively skewed than the female distribution, suggesting that men are disproportionately more likely to be employed in high-wage firms when compared with women.

The raw mean gender pay gap among all employees in the Analysis sample is 22.3 log points. Besides CGP, the other studies which have estimated a role of worker-firm sorting in gender pay gaps focus on an *Oaxaca* decomposition of the ‘Raw’ or observed gap. These estimates are also based on AKM-type model estimates (see for example Card et al., 2016, among the other studies discussed and summarised below in Table 6). Table 2 summarises the results of an *Oaxaca* decomposition for the raw UK pay gap. The vast majority, 84%, is accounted for by the differences in estimated worker fixed effects. However, the different allocation of men and women over the estimated firm fixed effects in the UK labour market can account for 3.6 log points or 16% of the raw gender pay gap. The role of the different occupations which men and women work in is smaller than the firm-specific component, accounting for 1.3 log points or 6% of the raw gap. The residuals and remaining observable characteristics, including tenure, age and full-time status, all together contribute -1.4 log points or -6% to the pay gap, i.e. they on average favour women.

4.1.1 Firms and the gender pay gap throughout the employee wage distribution

We construct a further graphical representation, which depicts the role of the estimated firm wage premiums on the hourly wage gap throughout the overall employee wage distribution, looking beyond the mean. In Figure 4, we collect all employees, both men and women, 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 male and 3,102 female observations. We then look up the estimates of the 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 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 approximately the fortieth 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 in firm-specific wage premiums is reversed among employees working in relatively high-paid jobs, between approximately the median and at least the ninety-fifth percentile of the overall wage distribution: when comparing men and women earning the same but relatively high rates of pay, men are more likely to work for firms which provide high wage premiums. 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.

4.1.2 Firm wage premiums and gender within industry sectors, part- or full-time work, and age groups/cohorts

We 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 within different groups of jobs. In other words, we ask whether men are more or less likely to be employed by relatively low or high-wage firms than women, depending on the type of job or worker. Table 3 summarises the gender differences in $\hat{\phi}_{J(it)}$ within these groups, showing the gaps at various percentiles of the gender-specific distributions over these estimated values, i.e. within a group of jobs we compare the median estimate among men with the median among women. Online Appendix Figures C1-C3 represent the equivalent results as kernel density estimates. Table 3 also shows the contribution of the differences in estimated firm fixed effects to the raw mean gender wage gap within these different groups of jobs, with the comparable result for all jobs of a 16% share repeated in the first row.

First, we look within the private and public sectors. In the latter, the distributions of men and women over the estimated firm wage premiums are more similar than in the whole labour market, with only a modestly greater tendency of men to be working in firms with relatively very high wage premiums (Figure C1B). 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 (Figure C1A). Men and women are approximately equally likely to receive very low firm wage premiums if they are working in the private sector. 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 allocation of workers to firm-specific wage effects contributes a greater amount and share of the raw mean wage gap in the private sector, 18%, than in the public sector, 11%.

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 (Figure C1F). The male and female distributions over the firm wage premiums are also relatively similar within the ‘Manufacturing’ and ‘Financial services’ industry groups (Figures C1C&E), despite in the latter group of jobs there being evidence of large gaps at the mean and at the top of the wage distribution (Healy and Ahamed, 2019). However, the gender differences in the ‘Non-financial (sales) services’ group, which is dominated by retail and hospitality services, are starker (Figure C1D). 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. Within this group, the estimated firm fixed effects contribute 15% of the raw mean wage gap.

We also look at differences in where men and women work within part-time and full-time employment (Figure C2). In part-time employment, men are more likely to be employed by a firm with a very low wage premium than women. Conversely in full-time employment, men are more likely to be employed by a firm with a high wage premium than women. However, the raw mean wage gaps within part-time and full-time work are similar, as are the contributions to these gaps from the 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 (Figure C3A). In this age group, the median firm fixed effect received by a man is 2 log points higher than for women.¹³ 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 (Figures C3B-D). 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, though as an estimated share of the raw mean wage gap this contributes a consistent share over the life-cycle. 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 birth cohort effects from the life-cycle. Related, we cannot explore with these data how any gender-firm sorting interacts with the child penalty for female earnings, which has been shown to account for a large part of the gap in earnings (e.g. Costa Dias et al., 2018; Kleven et al., 2019).

4.2 Applying the Gelbach decomposition to the adjusted pay gap

Table 4 presents our main results from applying the Gelbach decomposition method in different ways to the same Full model regression estimates summarised by Table 2. In our preferred results, given by Column (3), the adjusted pay gap is 14.5 log points, obtained as the least squares estimate $\hat{\lambda}$ from the Basic model, described by Equation (4), with the following controls in $\tilde{\mathbf{X}}$: 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). This estimate simply demonstrates the importance of these few covariates in accounting for gender pay differences. Of this 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, $\hat{\delta}_f$. In other words, in a counterfactual where men and women were identically distributed across firms and their associated wage premiums, conditional on some of 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 by the estimates of unobserved fixed worker factors, which by definition are transferable across firms, and include differences in education, work experience and preferences.

This measure is more nuanced than that provided by an Oaxaca decomposition. It first conditions 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 $\hat{\delta}_f$ from the Gelbach decomposition has a very relevant interpretation: it gives an estimate of how much who works for whom matters, after first 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.

Column (2) of Table 4 takes a step backwards from the preferred results in Column (3) to a ‘More Basic’ model, showing the results of the adjusted pay gap decomposition when we exclude the time-invariant worker and firm observable factors from our preferred 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 7.4% share. 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.

Column (1) takes a step even further backwards from the preferred results to a ‘Most Basic’ model. In this case the adjusted pay gap is 21.1 log points, only a small reduction compared with the raw pay

gap. This estimate is obtained by controlling only for year effects, squared employee age and a quadratic polynomial for employee tenure. We decompose this measure into four components: 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 in Column (2) (a cubic firm size polynomial, full-time work dummy, cubic terms in age and tenure). The estimate 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. Therefore, our estimate of the UK gender-firm sorting effect could be substantially biased upwards by omitting important time-varying covariates from the Basic model.

Finally, in Column (4) we present decomposition results comparable to our preferred results, except here we do exclude occupations from the Basic model and $\tilde{\mathbf{X}}$, and account for a gender-occupation sorting effect in a similar manner as in CGP. The measured contribution to the adjusted pay gap from gender-firm sorting in these decompositions increases to 1.2 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% of an overall adjusted UK gender pay gap of 17 log points.

5 Robustness and Further Discussion

This section considers some robustness checks on our main results, summarised in Table 5. We also compare our results to what studies of the pay gaps in other countries have found, regarding the importance of whom men and women work for.

5.1 Before and after the 2008/9 recession

As Figure 1 shows, the raw 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 in 2008-09. 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, 2019). 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 model for two sub-periods, allowing all of the parameters to change: first 2002-07, before the financial crisis and recession; and second 2010-16, afterwards. We then apply the decomposition of the adjusted gender pay gap within these sub-periods with our preferred version of the Basic model.

The estimate of the adjusted pay gap is 15.1 log points in 2002-07, compared with 13.4 log points in 2010-16 (Columns (1) & (2), Table 5). 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.

5.2 Occupational classification

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 4 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 (3) of Table 5, we confirm that our preferred results 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 (4) of Table 5, we also show comparable estimates to Column (4) of Table 4. With more detailed occupation controls, the contributions to the adjusted pay gap from gender-firm and gender-occupation sorting increase marginally to 1.3 and 1.2 log points, respectively.

The ASHE does not contain information on employees' education, as discussed before. However, the influence of education levels or qualifications obtained before the age of 25 on wages should here be captured by the estimated worker fixed effects, and so omitting these variables from the model ought to not significantly affect the raw wage gap decomposition results. However, moving between less and more detailed occupational controls offers suggestive evidence of how robust our decomposition of the adjusted wage gap is to this omission from the wages model. The more specific estimates of occupational wage premiums plausibly correlate more strongly with employee education levels and qualification types. Therefore, comparing the results using the 3-digit and 2-digit classifications suggests that the relative importance of the gender-firm sorting effect may not be greater if we could pre-adjust the wage gap for workers' education differences.

5.3 Labour market experience

The results described so far do not address explicitly the importance of accumulated life-time work experience in determining gender wage differences. For the raw pay gap decomposition, we do not consider this omission much of a concern, since the influence of pre-2002 work experience on wages will be captured by the worker fixed effects and employee tenure is included in the wages model. However, we can construct a proxy time-varying measure of labour market experience using the New Earnings Survey Panel Dataset 1972-2016 (Office for National Statistics, 2017), which contains the ASHE job observations but does not contain firm identifiers. This measure is simply taken as the number of years a worker appeared in the NESPD/ASHE up to and including each year. This is likely to be a poor indicator of cumulative labour market experience, especially from a gender perspective, given it does not address hours of work and will be underestimated for workers with frequent moves between firms due to the survey methodology. When we add a cubic polynomial in this proxy measure of experience within the Full wages model, our preferred results are practically unchanged. The estimated contribution of firm sorting to the raw gender wage gap

remains 3.6 log points, or a 16.0% share, the adjusted wage gap is reduced marginally to 14.3 log points, and the gender-firm sorting effect contributes 0.8 log points to this, or a 5.4% share (Column 5, Table 5).

5.4 Comparing with other studies and countries and other robustness

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 CGP for Portugal. The observable pay gap over their sample period was a similar level to the one 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, where these job titles refer to somewhere in the order of 30,000 different groups. These contributions are an order of magnitude greater than what we find for the United Kingdom in our preferred results. One explanation is that firms could potentially matter less for overall UK wage variance than in Portugal. 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 models. In CGP, 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. Therefore, the estimates from our Most Basic model provide a closer comparison (Column (1), Table 4). Our measure of the adjusted pay gap in this case is 21 log points, which is similarly close to the observed UK sample average value, 22 log points, as it was in Portugal for CGP. As already discussed, the estimate we find for the share of the adjusted pay gap from gender-firm sorting in this specification is 16%, with a further 6% accounted for by gender-occupation sorting. In the UK case, these greater estimates for the importance of firms in the adjusted pay gap are accounted for by not conditioning on whether employees work full-time.

Table 6 summarises the main results from five other studies which are closely comparable to this one. All use a common methodological approach. Specifically, the approach taken by Card et al. (2016) (henceforth CCK), using data from Portugal for the period 2002-2009, has since been followed closely by Bruns (2019) for Germany, Gallen et al. (2019) for Denmark, Sorkin (2017) for the United States, Coudin et al. (2018) for France and Casarico and Lattanzio (2019) for Italy. 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 decompose further 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 that 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 (2019) and Gallen

[et al. \(2019\)](#) also found the sorting component dominated any role for gender wage bargaining effects in recent decades in Germany and Denmark. [Casarico and Lattanzio \(2019\)](#) find the bargaining component in Italy can account for a third of an estimated overall 30% influence of firm-specific wage premiums between 1995 and 2015, which was also increasing in importance over this period. 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 is roughly comparable to the total 21% figure from CCK, with the caveat that this estimate was obtained without allowing any gender differences in the model parameter estimates.

As a further robustness check, we re-estimate the Full model, allowing the parameters on all time-varying observable characteristics to vary with gender, but keeping the firm fixed effects non-gender-specific. Estimates of this model are summarised in Online Appendix Table D1. Compared with our preferred specification, the model fit is approximately unchanged, as is the share of the raw pay gap accounted for by the firm fixed wage effects at 16%. The worker fixed effects, however, account for a smaller share when the parameters on observable characteristics are allowed to vary by gender. Overall, 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 did not apply fully CCK's approach here. First, the principal motivation of CCK is more specific than our own. They motivate their decomposition by an apparent need to reconcile what they identify as two competing strands of literature on how firms might generate pay gaps: through a sorting or a bargaining channel. In comparison, for the United Kingdom, we are only motivated by measuring whether gender-firm sorting matters. Second, our sample is a much smaller share of the population than CCK's, though it is comparable to the population share in most of the studies mentioned above which have since followed CCK. That said, identification of the AKM-model depends on connected sets and mobility groups. In these UK data, if we were to estimate fully gender-specific AKM-type models, as the CCK approach requires, we would begin to forego rapidly the representativeness of the sample and further diminish 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 unavailable, 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 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](#); [Cortés and Pan, 2019](#)). 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 chosen 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 models, but instead of log hourly employee wages we use

log gross weekly earnings as the dependent variable. The raw gender gap in weekly earnings in the Analysis sample is 49.8 log points. Of this, 6.7 log points or a 14% share is accounted for by the estimated firm fixed effects (see Online Appendix Table D2). 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 6.7% share is accounted for by the gender-firm sorting effect (Column (6), Table 5). 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 (Table 1). 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. There is, however, an important caveat to this result; the records in the ASHE given by employers are only for paid hours of work. With these data we cannot rule out the possibility that there is gender-firm sorting with regard to unpaid hours of work, which, if related to hourly wages or weekly earnings, could still account for part of the gender pay gap.

6 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 the gender pay gap, between 2002 and 2016, was accounted for by the differences between men and women in whom they worked for. For the raw mean wage gap of 22.3 log points, we found that 16% was explained by the differences in the estimated firm-specific wage premiums earned by men and women. This was almost three times the amount accounted for by the differences in the occupations worked in by men and women. We also looked at the role of firms in accounting for an adjusted measure of the gender pay gap of 14.5 log points. We found that the contribution from the gender-firm sorting effect was small after adjusting for the influence of other observable characteristics of workers, jobs and firms: around one percentage point or just 6% of the adjusted pay gap is accounted for by gender-firm sorting. In other words, if overnight all gender inequality was wiped out in terms of the firms that people worked for, conditional on the existing observable differences in relevant worker, job and firm characteristics remaining, such as tenure, full-time status and occupations, then the gender pay gap would only have decreased by this small amount.

The clear majority of what explains the pay gap shows up within firms. Future research should focus on identifying the particular worker and firm behaviours which can explain this. Other recent work has used longitudinal employer-employee linked data to assess the gender differences in how far between firm moves matter for life-cycle earnings growth, compared with within firm progression (e.g. [Barth et al., 2017](#)). Combining such an analysis with specific life-cycle events, notably childbirth, could provide a deeper understanding of the firm's role in gender earnings differences. [Coudin et al. \(2018\)](#) have made progress on this using French data, finding that the employment and mobility choices of mothers are affected after childbirth, in turn affecting gender earnings differences persistently thereafter. Unfortunately, we are not aware of any UK datasets presently available which would allow a similar study.

It would be a significant challenge for the validity of our analysis and conclusions if worker mobility between firms was related to firm-specific components of wage growth, rather than to the level of wage premiums on offer. The robustness checks in Online Appendix B and similar checks in the related literature summarised by Table 6 have suggested that this may not be an issue. However, this deserves more careful

attention using larger samples of firms' employees and longer periods of study, which future access to UK sources of administrative earnings data would allow. Similarly, a further limitation of the present study is that we have not allowed for the possibility that firm-specific wage premiums might differ between men and women within firms. If in fact they do, then this would suggest the presence of differences in bargaining power or systematic discrimination between men and women within firms. We are hopeful that these issues can be addressed in the years ahead, when the UK authorities follow the examples of other countries and open up more administrative data for research.

Notes

¹Excluding 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 8.6% in 2018. The median gap for all employees, including part-time, was 17.9% in 2018. Source: ONS, ASHE Total Table 1.6a, Hourly Pay excluding overtime; [\[External link\]](#)

²This is a standard way to measure wage gaps in the economics literature, which in effect takes less account of the long upper tails of the wage distribution, particularly for men. The arithmetic mean log wage is equivalent to the log geometric mean wage. Exponentiating the figure in this case, the sample's male geometric mean wage is 25% greater than the female value.

³Legislation titled “The Equality Act 2010 (Gender Pay Gap Information) Regulations 2017”; [\[External link\]](#). Public sector employers throughout the UK already had duties to report pay gaps, following from the Equality Act 2010. There was significant interest in and scrutiny of the reported figures among the British public and media. See for example “Gender pay gap: More than 500 firms reveal their figures”, 6 January, 2018, [\[External link\]](#); “Gender pay gap: Men still earn more than women at most firms”, 21 February, 2018, [\[External link\]](#).

⁴There have been several regulatory impact assessments on the policy proposal, since before the enactment of the Equality Act 2010. For example, the Regulatory Impact Assessment of April 2016, relating to this new legislation, suggested seemingly conservative estimates of the net economic costs to private sector employers of £3.8million per annum, [\[External link\]](#).

⁵A National Insurance number gives individuals the right to work in the UK and identifies them for income tax purposes. The two main reasons why an individual might not be observed in some year of the ASHE 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. The panel goes back as far as 1968, though firm identifiers are generally unavailable before 2002.

⁶Obtained from UK National Statistics, accessed 24/4/2017.

⁷[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.”

⁸Throughout 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.

⁹These levels and trends are similar to what one can obtain from UK official National Statistics aggregates, contained historically within the ONS “Patterns of Pay” series, which is based on the ASHE, [\[External link\]](#).

¹⁰We 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. [Beckhout 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 wage effects.

¹¹For the male and female decompositions, the covariance of the residuals with all the estimated factors is not strictly zero, since the model is estimated gender-blind.

¹²This will include any upwards bias from our small sample setting and estimation errors for the fixed effects ([Krueger and Summers, 1988](#)). Although comparisons are 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](#)).

¹³These results show 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.

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.
- Andrews, M. J., L. Gill, T. Schank, and R. Upward.** 2008. "High wage workers and low wage firms: negative assortative matching or limited mobility bias?" *Journal of the Royal Statistical Society Series A*, 171(3): 673–697.
- Andrews, M., L. Gill, T. Schank, and R. Upward.** 2012. "High wage workers match with high wage firms: Clear evidence of the effects of limited mobility bias." *Economics Letters*, 117(3): 824–827.
- 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., 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.** 2019. "Changes in Workplace Heterogeneity and How They Widen the Gender Wage Gap." *American Economic Journal: Applied Economics*, 11(2): 74–113.
- 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.
- Casarico, A., and S. Lattanzio.** 2019. "What Firms Do: Gender Inequality in Linked Employer-Employee Data." Cambridge Working Papers in Economics 1966, Faculty of Economics, University of Cambridge.

- Cortés, P., and J. Pan.** 2019. “When Time Binds: Substitutes for Household Production, Returns to Working Long Hours, and the Skilled Gender Wage Gap.” *Journal of Labor Economics*, 37(2): 351–398.
- 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.
- Drolet, M., and K. Mumford.** 2012. “The Gender Pay Gap for Private-Sector Employees in Canada and Britain.” *British Journal of Industrial Relations*, 50(3): 529–553.
- Eeckhout, J., and P. Kircher.** 2011. “Identifying Sorting – In Theory.” *Review of Economic Studies*, 78(3): 872–906.
- Gallen, Y., R. V. Lesner, and R. Vejlin.** 2019. “The labor market gender gap in Denmark: Sorting out the past 30 years.” *Labour Economics*, 56(C): 58–67.
- 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.
- Healy, G., and M. M. Ahamed.** 2019. “Gender Pay Gap, Voluntary Interventions and Recession: The Case of the British Financial Services Sector.” *British Journal of Industrial Relations*, 0(0): .
- Kleven, H., C. Landais, J. Posch, A. Steinhauer, and J. Zweimüller.** 2019. “Child Penalties Across Countries: Evidence and Explanations.” NBER Working Papers, National Bureau of Economic Research, Inc.
- Krueger, A., and L. Summers.** 1988. “Efficiency Wages and the Inter-industry Wage Structure.” *Econometrica*, 56(2): 259–93.
- 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.
- 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.
- Oaxaca, R.** 1973. “Male-Female Wage Differentials in Urban Labor Markets.” *International Economic Review*, 14(3): 693–709.
- Office for National Statistics.** 2017. “New Earnings Survey Panel Dataset, 1975-2016: Secure Access.” 7th Edition SN: 6706, UK Data Service.

- Office for National Statistics.** 2019. “Annual Survey of Hours and Earnings, 1997-2016: Secure Access. [data collection].” 14th 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.** 2019. “Recent Changes in British Wage Inequality: Evidence from Large Firms and Occupations.” *Scottish Journal of Political Economy*, 0(0): .
- 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.

TABLE 1: 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
Share female at firm	0.31	0.71	0.27	0.74
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
<i>N</i>	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.

TABLE 2: Summary of estimated Full model with two-way fixed effects & decomposition of raw gender pay gap: 2002-16

	Male (1)	Female (2)	Total (3)
St. dev. of log wages - $std_{it}(w_{it})$	0.55	0.49	0.53
<i>N</i> : worker-years	824,806	888,326	1,708,132
<i>P</i> : workers	131,903	124,501	256,404
<i>F</i> : firms			86,779
St. dev. worker effects - $std_{it}(\hat{\alpha}_i)$	0.45	0.38	0.43
St. dev. firm effects - $std_{it}(\hat{\phi}_{J(it)})$	0.2	0.17	0.18
St. dev. observables - $std_{it}(\mathbf{x}'_{it}\hat{\beta})$	0.51	0.43	0.48
Correlation - $corr_{it}(\hat{\alpha}_i, \hat{\phi}_{J(it)})$	-0.022	-0.013	0.004
Adjusted R^2			0.903
RMSE			0.165
<i>Variance shares</i> ($X/var_{it}(w_{it})$):			
Worker effects - $var_{it}(\hat{\alpha}_i)$	0.68	0.60	0.65
Firm effects - $var_{it}(\hat{\phi}_{J(it)})$	0.13	0.12	0.12
Covariance - $2covar_{it}(\hat{\alpha}_i, \hat{\phi}_{J(it)})$	-0.01	-0.01	0.00
Residuals - $var_{it}(\hat{\epsilon}_{it})$	0.07	0.09	0.02
Other	0.13	0.19	0.21
<i>Raw gender wage gap decomp. (shares):</i>			
Raw gap - $E_{it}[w_{it} i \in M] - E_{it}[w_{it} i \in F]$			0.223
Worker - $E_{it}[\hat{\alpha}_i i \in M] - E_{it}[\hat{\alpha}_i i \in F]$			0.187 (0.84)
Firm - $E_{it}[\hat{\phi}_{J(it)} i \in M] - E_{it}[\hat{\phi}_{J(it)} i \in F]$			0.036 (0.16)
Occupations			0.013 (0.06)
Other			-0.014 (-0.06)

Notes.- author calculations using the ASHE 2002-16, all employees age 25-64. £2002. Pay excludes overtime. Gap is male minus female. Estimated Full model includes covariates in \mathbf{x}_{it} for year fixed effects, squared and cubed terms for employee age, a cubic polynomial for employee tenure, a cubic polynomial for firm size (n. of employees) and a dummy variable for whether a worker was employed full-time.

TABLE 3: Gender gaps at percentiles of estimated firm-specific wage effects ($\hat{\phi}_{J(it)}$) & contribution of firms to the raw mean wage gap by sector, part- vs. full-time and age groups/cohorts

	Percentiles of firm effects					Decomposition of raw mean gap (log points)		
	10th	25th	50th	75th	90th	Wage gap	Firm contrib.	Firm share of gap
Overall (all employee-years)	0.01	0.03	0.04	0.06	0.06	22.29	3.61	0.16
Private	0.01	0.06	0.09	0.05	0.04	29.43	5.27	0.18
Public	0.00	0.01	0.02	0.03	0.04	20.79	2.35	0.11
Manufacturing	0.04	0.03	0.02	0.02	0.02	24.59	2.63	0.11
Non-financ. (sales) services	-0.01	0.02	0.03	0.10	0.11	27.25	4.19	0.15
Financial services	0.00	0.00	0.00	0.00	0.03	28.23	-0.24	-0.01
Other ind. sectors	0.01	0.02	0.03	0.03	0.05	18.12	2.98	0.16
Part-time	-0.02	-0.03	-0.02	0.00	0.01	14.04	2.27	0.16
Full-time	0.00	0.01	0.03	0.04	0.05	14.01	2.13	0.15
Age 25-34	-0.01	0.00	0.02	0.04	0.04	7.93	1.52	0.19
Age 35-44	0.02	0.03	0.05	0.06	0.07	25.97	4.24	0.16
Age 45-54	0.02	0.04	0.05	0.07	0.07	30.98	4.83	0.16
Age 55-64	0.02	0.04	0.03	0.05	0.07	24.88	3.73	0.15

Notes.- effects estimated as per regression model (1) and Table 2. The derivation of the industry groupings used here is based on the SIC2003 and described in Online Appendix A. For the ‘Overall’ results, see also Table 2 and Figure 3B. For complete representations of the distributions of $\hat{\phi}_{J(it)}$ for the different groups of jobs, see Online Appendix Figures C1-C3.

TABLE 4: Main decomposition results for the adjusted log gender pay gap

	Most Basic model (1)	More Basic model (2)	Preferred results (3)	Firm & occ. sorting (4)
Worker effects - $\hat{\delta}_a$	0.183 (0.864)	0.138 (0.927)	0.146 (0.943)	0.148 (0.871)
Gender-firm sorting - $\hat{\delta}_f$	0.034 (0.159)	0.011 (0.074)	0.008 (0.057)	0.012 (0.070)
Gender-occ. sorting - $\hat{\delta}_o$	0.012 (0.056)			0.010 (0.060)
Other observable chars. - $\hat{\delta}_x$	-0.017 (-0.079)			
Adjusted gap - $\hat{\delta}_y$ or $\hat{\lambda}$	0.211	0.149	0.145	0.170
Observed gap (mean)	0.223	0.223	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.

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

Column (1) applies the [Gelbach](#) decomposition to the Full model assuming a most Basic 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 most Basic model, as described in the text.

Column (2) gives results where the assumed Basic regression model (4) *does not* include time-invariant employee or firm characteristics; i.e. $\tilde{\mathbf{X}}$ from (4) is identical to \mathbf{X} in the Full model.

Column (3) gives results where the assumed Basic regression model (4) *does* include time-invariant employee or firm characteristics; i.e. $\tilde{\mathbf{X}}$ from (4) is *not* identical to \mathbf{X} in the Full model.

Column (4) gives results where we exclude the 2-digit occupation controls from the Basic model estimated as per column (3), but use the [Gelbach](#) decomposition to account for how much of the Adjusted pay gap is contributed by the gendered distribution over these effects, as estimated in the Full model, alongside the worker- and firm-specific effects.

TABLE 5: Decomposition results for the log gender pay gap: robustness checks

	Sub-periods		3-digit occ. class.		Experience ctrls	Weekly earnings
	2002-07 (1)	2010-16 (2)	(3)	(4)	(5)	(6)
Worker effects - $\hat{\delta}_a$	0.143 (0.942)	0.125 (0.932)	0.126 (0.960)	0.146 (0.856)	0.135 (0.946)	0.161 (0.933)
Gender-firm sorting - $\hat{\delta}_f$	0.009 (0.058)	0.009 (0.068)	0.005 (0.040)	0.013 (0.074)	0.008 (0.054)	0.012 (0.067)
Gender-occ. sorting - $\hat{\delta}_o$				0.012 (0.070)		
Adjusted gap - $\hat{\delta}_y$ or $\hat{\lambda}$	0.151	0.134	0.131	0.170	0.143	0.173
Observed gap (mean)	0.267	0.191	0.223	0.221	0.223	0.498

Notes.- see Table 4.

Column (1) presents equivalent results to Column (3) in Table 4, where the sample period is reduced to only employee-year observations in 2002-07 ($N=541,346$).

Column (2) similarly presents results where the sample period is reduced to only employee-year observations in 2010-16 ($N=771,742$).

Column (3) provides equivalent results as Column (3) in Table 4, except here the occupation controls included in the Basic and Full model were at the 3- rather than 2-digit SOC level.

Column (4) provides equivalent results as Column (4) in Table 4, but here using the 3-digit SOC level.

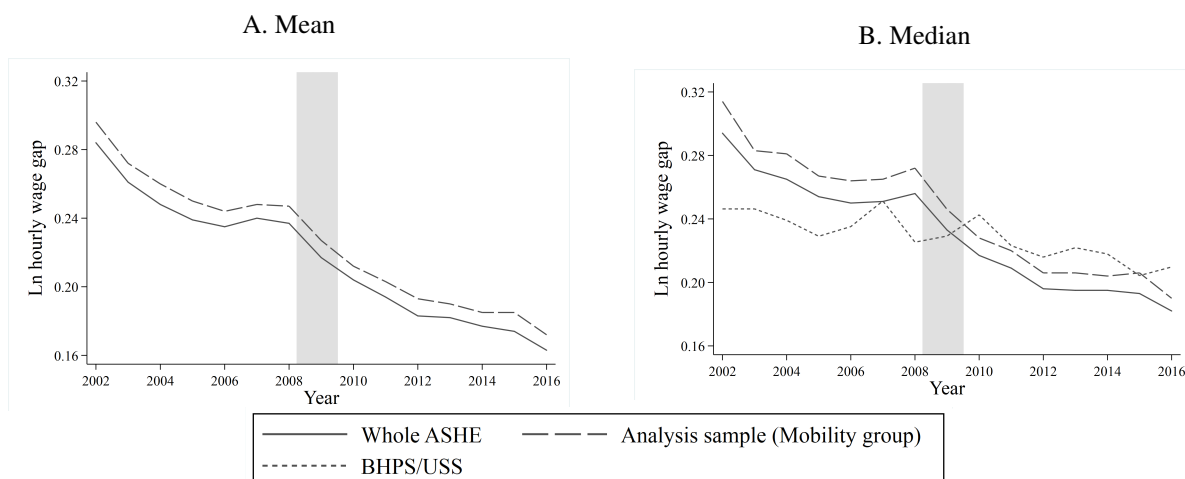
Column (5) provides equivalent results as Column (3) in Table 4, but adds a cubic polynomial for a proxy measure of cumulative labour market experience to the wages models, described in the text.

Column (6) use the exact same Analysis sample of employee-years and firms, and the same regression models, as used to estimate the preferred results, except the dependent variable is the natural logarithm of real gross weekly wages(see also Table D2).

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

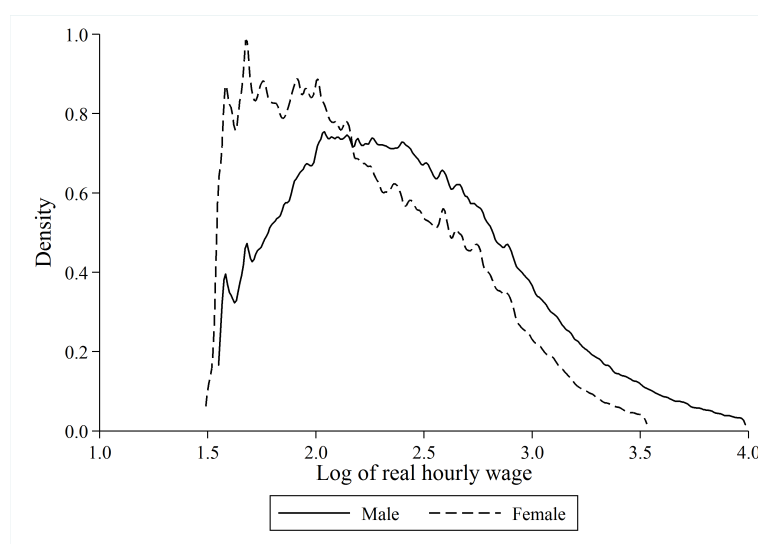
	Country, Period	Raw Pay Gap (log points)	Share of raw/adj. gap acc. for by gender-firm sorting	Additional notes
Cardoso et al. (2016)	Portugal, 1986-2008	24	19% (adj.)	Gelbach decomposition of worker, firm and job-title fixed effects; 18% of adj. gap from job-title-sorting; private sector representative 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 (2019)	Germany, 1995-2008	25 (2001-08)	25-31% (raw)	Method follows Card et al. (2016) ; -5% to 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. (2019)	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.
Casarico and Lattanzio (2019)	Italy, 1995-2015	17	30% (raw)	Method follows Card et al. (2016) ; sorting 2/3 and bargaining 1/3 of overall role of firm wage premiums; connected sample of 20.5 million workers in 1.2 million firms; private sector employees only; weekly wage data.

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



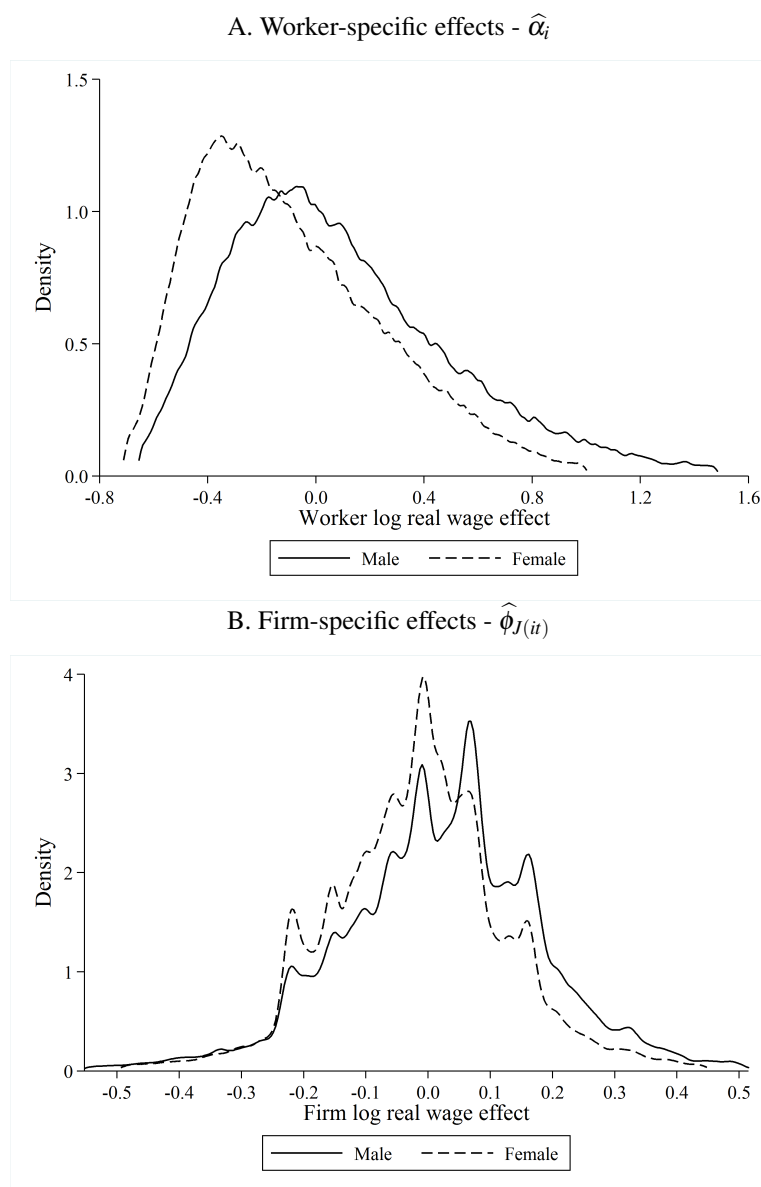
Notes.- author calculations using the ASHE 2002-16 and BHPS/USS 2002-2016, all employees aged 25-64. Pay excludes overtime. Gap is male minus female. “Analysis sample” is an 87% sub-sample of the “Whole ASHE” and represents statistics using only jobs in the largest connected set of workers and firms,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

FIGURE 2: Male and female distributions of employee real hourly wages



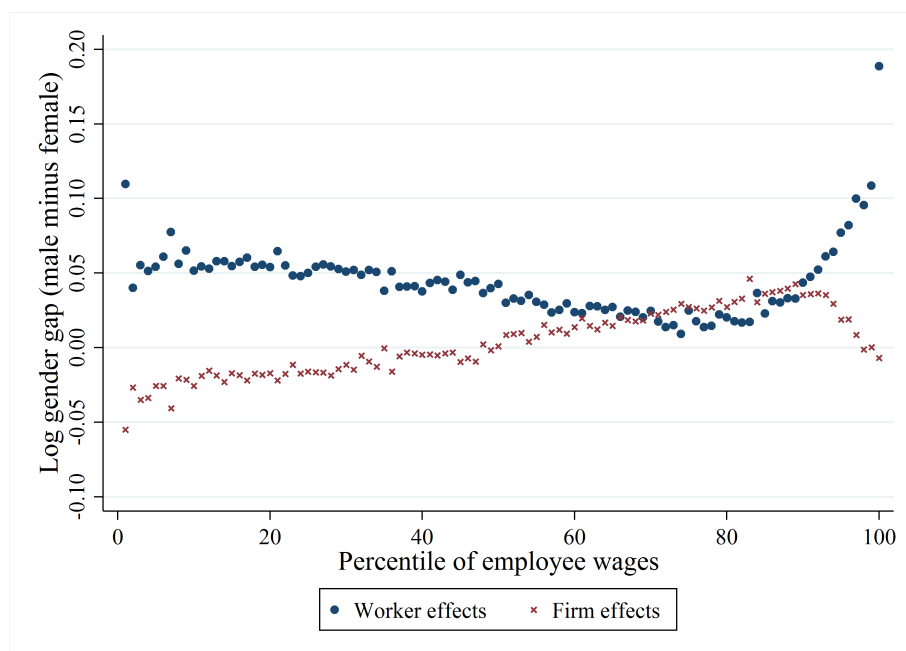
Notes.- see Figure 1. £2002. Uses the “Analysis sample”. Both male and female kernel densities were estimated with a bandwidth of one log point. The top and bottom one percent of male and female hourly earners are not displayed.

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



Notes.- wage effects estimated as per regression model (1) and Table 2, 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.

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 Table 2. £2002.

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.

Who works for whom and the UK gender pay gap

Online Appendix

Sarah Louise Jewell Giovanni Razzu Carl Singleton

Corresponding author: c.a.singleton@reading.ac.uk: Department of Economics, University of Reading, UK.

Published at: *British Journal of Industrial Relations*. (2019; In Press)

<https://doi.org/10.1111/bjir.12497>

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 of around 55%. The response rate for employee jobs is around 60%, implying that for those firms who do respond they tend to provide complete responses for all their employees in the sampling frame. There is no cumulative attrition from the panel. Any individual in the sample 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 three major sources of under-sampling, one due to firms not responding, and the other two 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." In 2007 and 2008 the target 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. However, the 1% of population sampling frame was reinstated from 2009 onward.

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. However, all other payments received in the period are included within this gross pay, including incentive-related pay and any premiums for weekend or night work. 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.

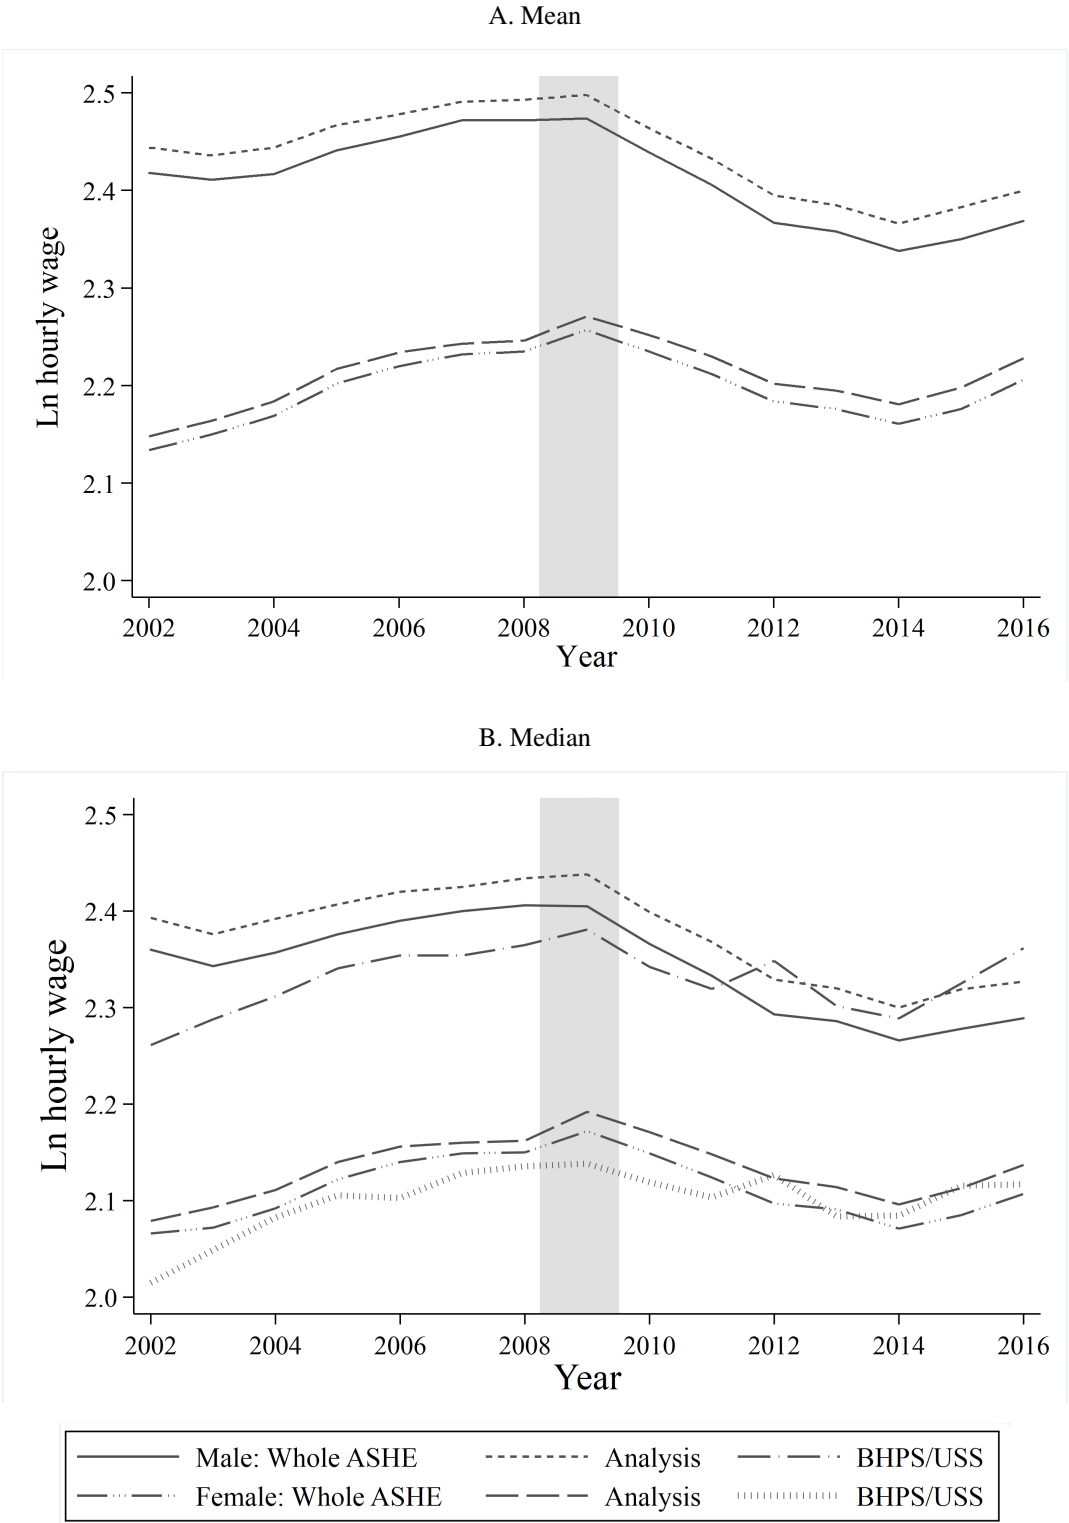
The household survey-based pay statistics in Figure 1 and A1 are derived using the longitudinal British Household Panel Survey (2002-08) and its successor the Understanding Society Survey (2009-16). They use the Great Britain and Northern Ireland samples and waves, but without any other boost samples. 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: Distribution of the number of different jobs held by workers in the Analysis sample (%), 2002-2016

Number of jobs	Worker-year weighted			Worker weighted		
	Male (1)	Female (2)	Total (3)	Male (4)	Female (5)	Total (6)
1	40.93	40.95	40.94	48.17	48.20	48.18
2	32.61	32.60	32.60	30.52	30.56	30.54
3	16.90	17.01	16.96	14.13	14.16	14.14
4	6.75	6.72	6.73	5.20	5.14	5.17
5	2.12	2.06	2.09	1.52	1.49	1.51
6	0.54	0.54	0.54	0.37	0.38	0.37
7+	0.15	0.12	0.13	0.10	0.08	0.09
<i>N / P</i>	824,806	883,326	1,708,132	124,501	131,808	256,304

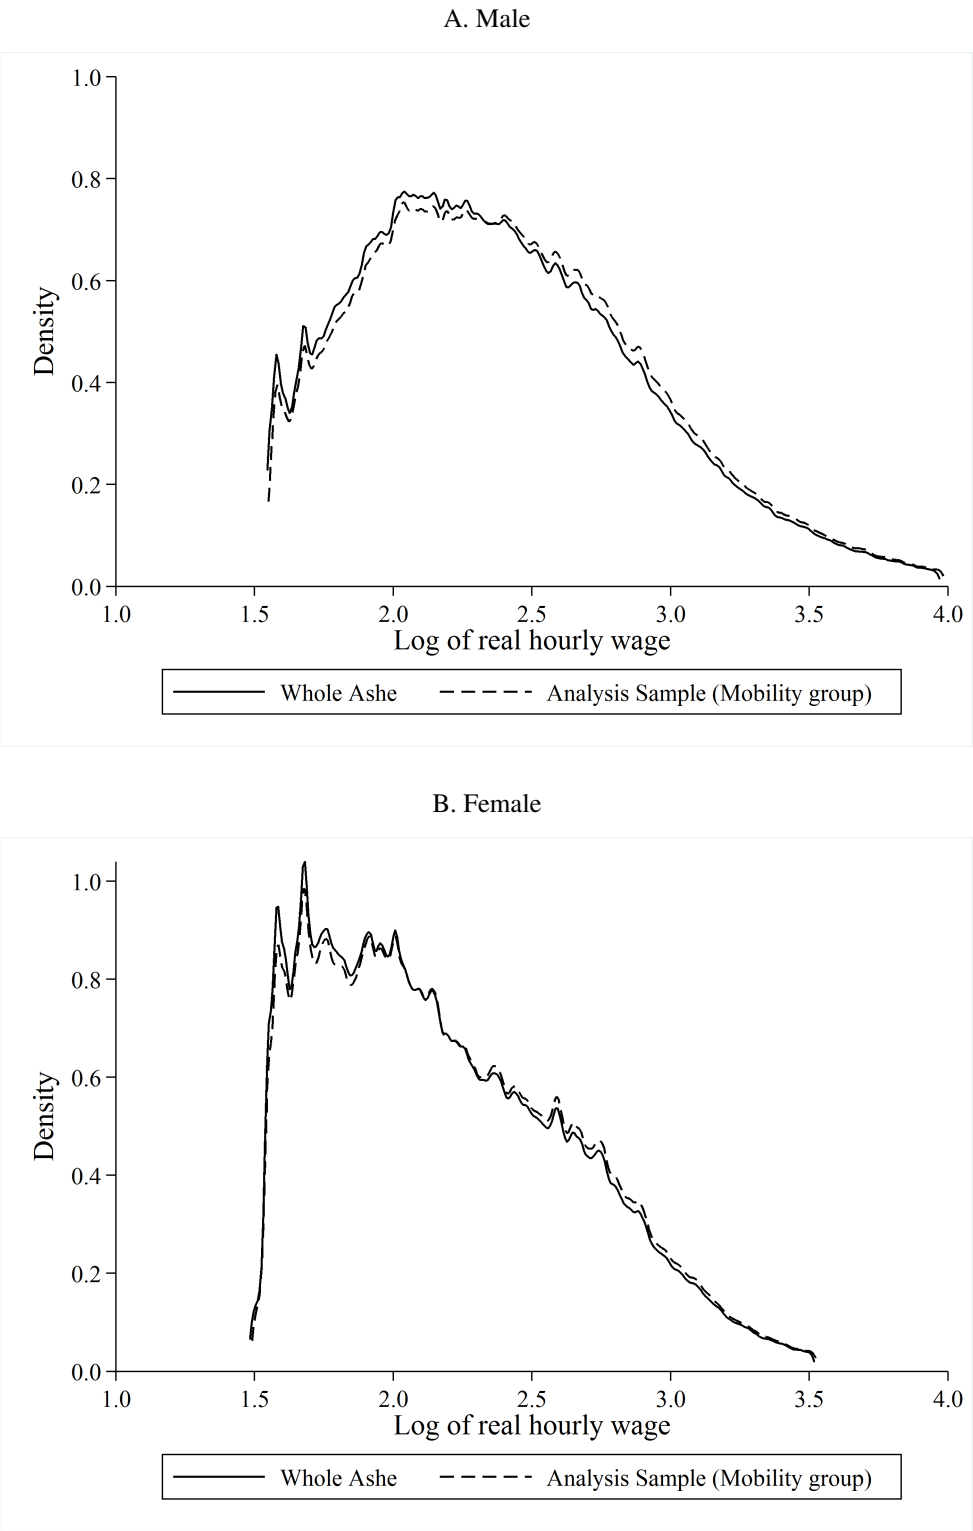
Notes.- author calculations using the ASHE 2002-16, all employees age 25-64.

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



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



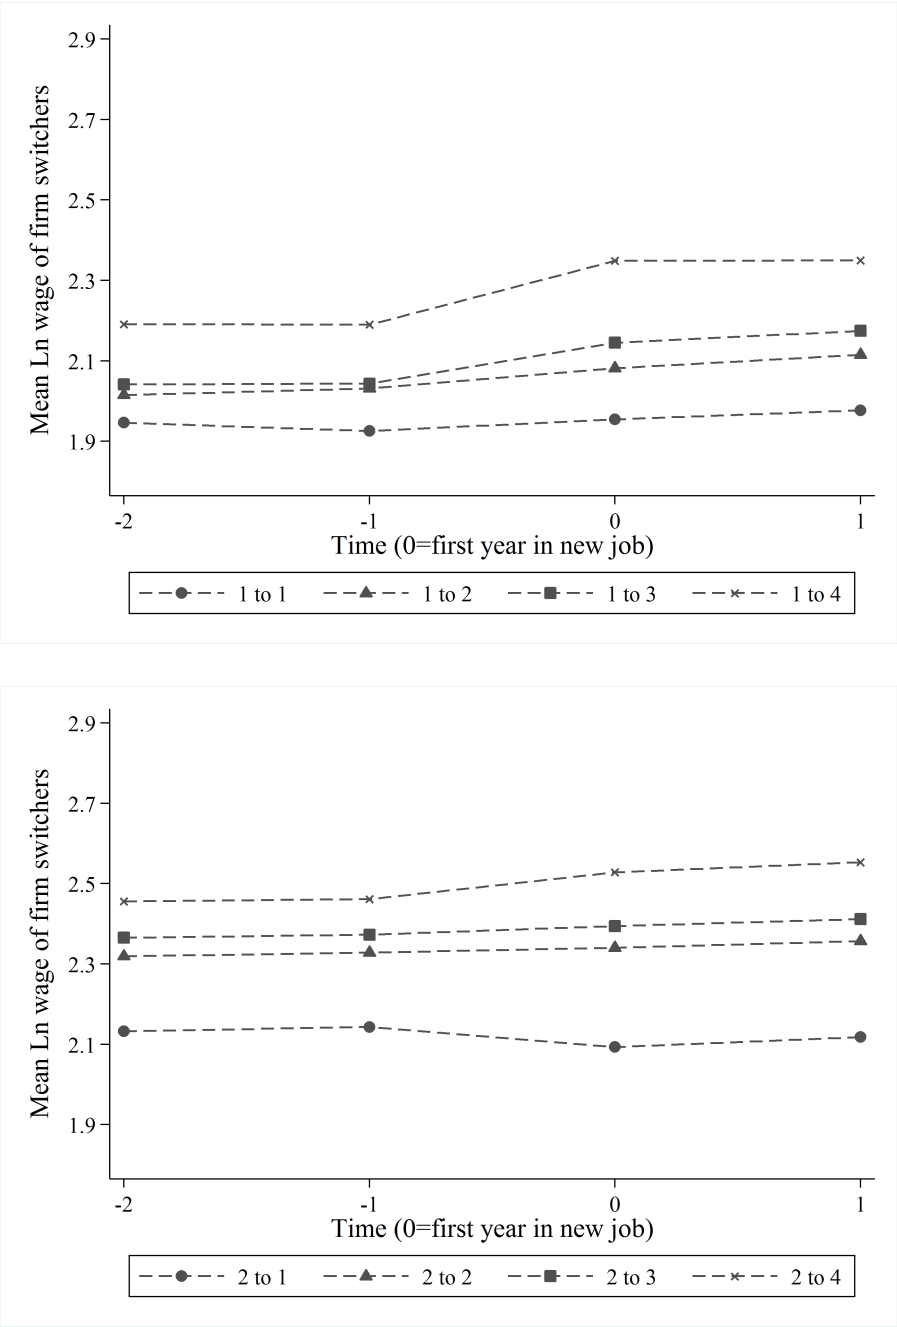
Notes.- see Figures 1 & 2

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 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.

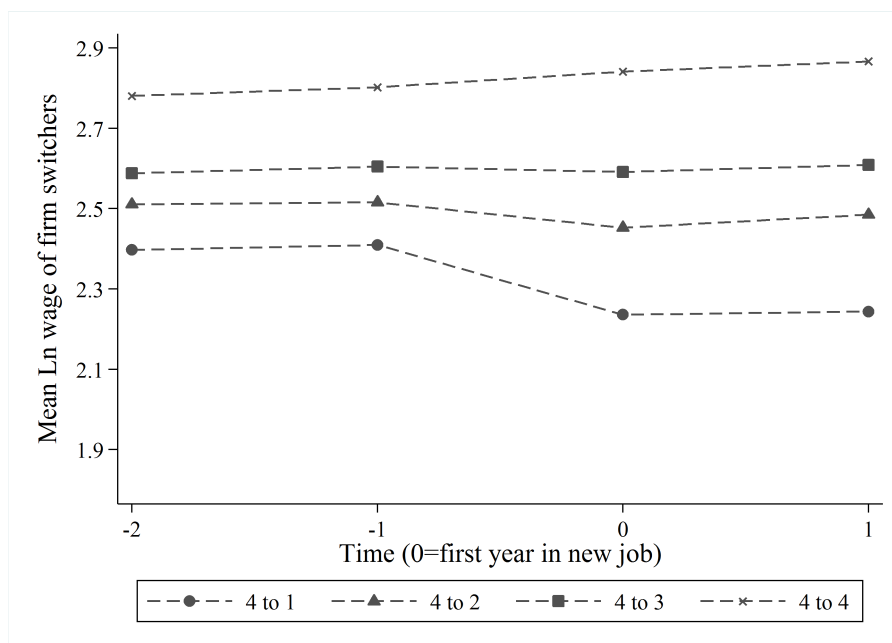
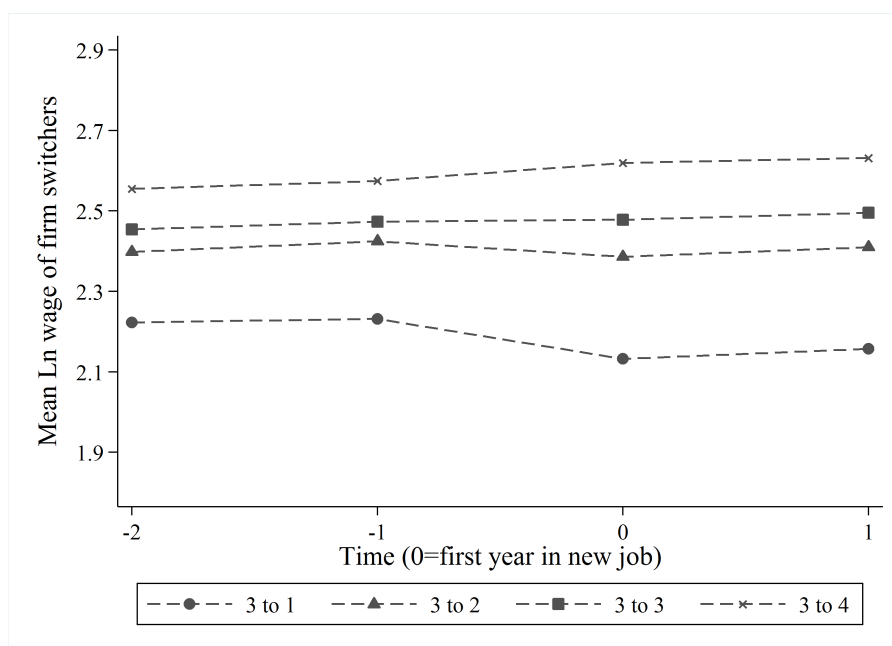
We construct this event study using the Analysis sample of employees as follows. First, we assign to each employee, in every year, the mean wage among all of his co-workers in the same firm and year (excluding himself). Then, year-by-year, we assign each employee to the quartile of how relatively well 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](#) and [Table B1](#). For example, in the first panel of the figure 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 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



(continued on the next page)

(continued from the previous page)



Notes.- “X to Y” indicates the quartile of co-worker wages for employees in their old firm (X), from which they switch to their new firm (Y). Each event series uses firm switches throughout the ASHE 2002-2016 Analysis sample period. See Table B1 for a summary of this data.

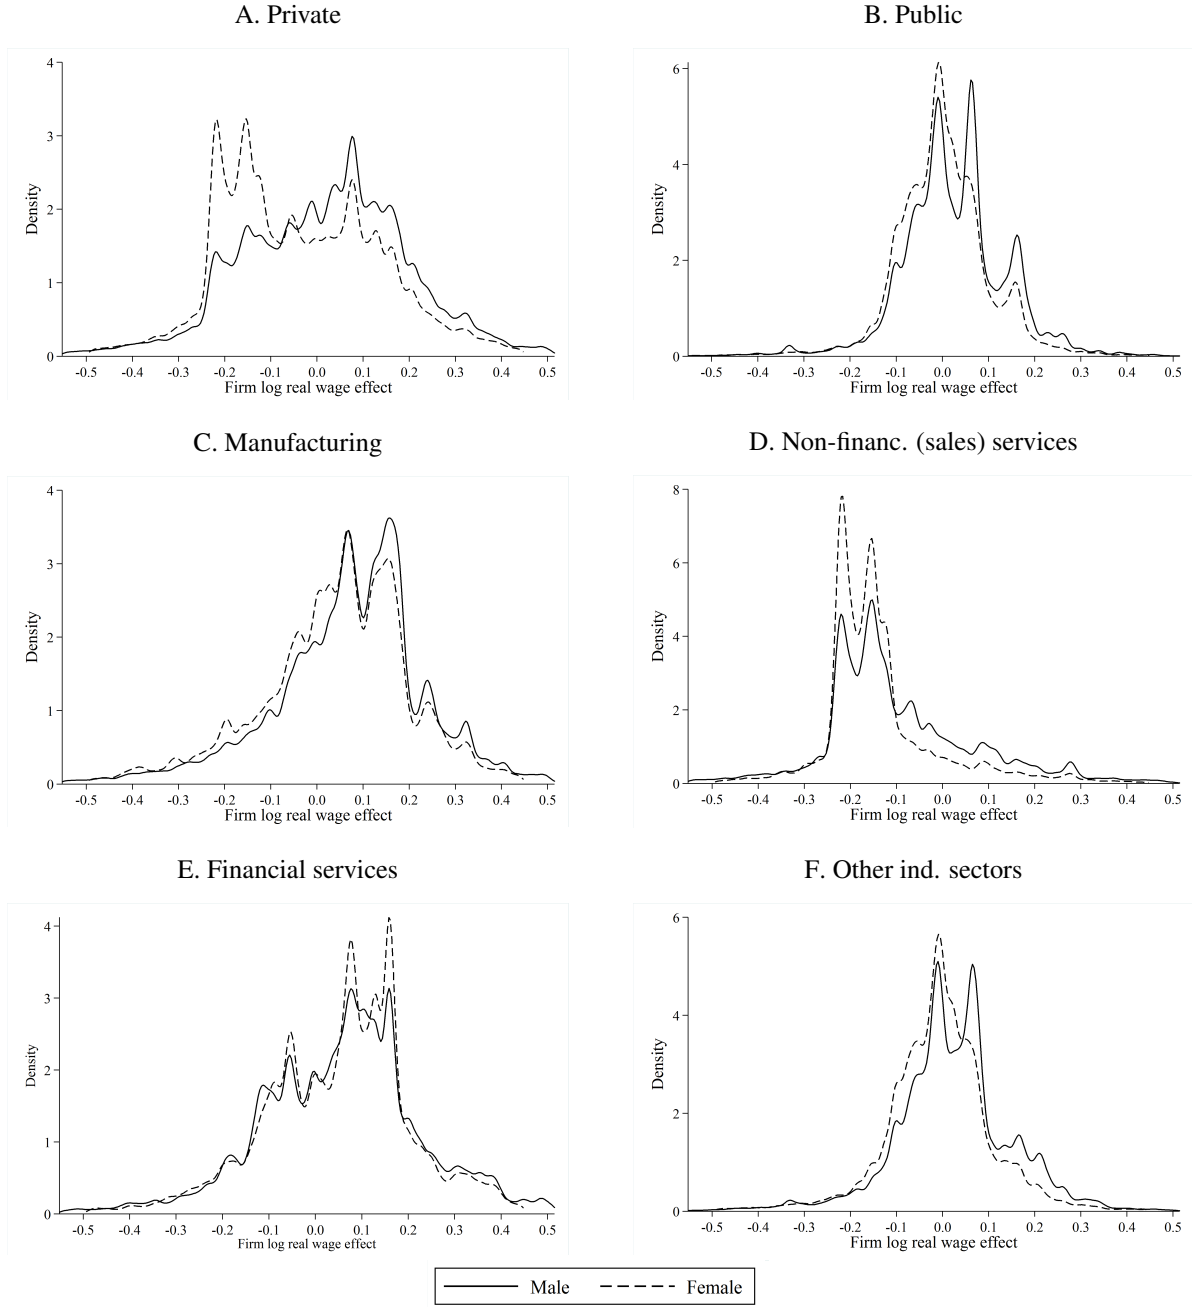
TABLE B1: Summary of event-study: Mean log wages of employees two years before switching firms and two years after, 2002-2016

Quartile to Quartile	Switches (1)	% of switches (2)	Years since switching				3-year wage change (7)
			-2 (3)	-1 (4)	0 (5)	1 (6)	
1 to 1	3,498	16.30	1.95	1.93	1.95	1.98	0.03
1 to 2	1,073	5.00	2.01	2.03	2.08	2.12	0.10
1 to 3	705	3.29	2.04	2.04	2.15	2.17	0.13
1 to 4	458	2.13	2.19	2.19	2.35	2.35	0.16
2 to 1	931	4.34	2.13	2.14	2.09	2.12	-0.01
2 to 2	1,582	7.37	2.32	2.33	2.34	2.36	0.04
2 to 3	1,660	7.74	2.37	2.37	2.39	2.41	0.05
2 to 4	768	3.5	2.46	2.46	2.53	2.55	0.10
3 to 1	583	2.72	2.22	2.23	2.13	2.16	-0.07
3 to 2	1,018	4.74	2.40	2.42	2.39	2.41	0.01
3 to 3	2,573	11.99	2.45	2.47	2.48	2.50	0.04
3 to 4	1,194	5.57	2.55	2.57	2.62	2.63	0.08
4 to 1	385	1.79	2.40	2.41	2.24	2.24	-0.15
4 to 2	571	2.66	2.51	2.52	2.45	2.48	-0.03
4 to 3	1,249	5.82	2.59	2.60	2.59	2.61	0.02
4 to 4	3,207	14.95	2.78	2.80	2.84	2.87	0.09

Notes.- author calculations using the ASHE 2002-16, all employees age 25-64. £2002. See Figure B1.

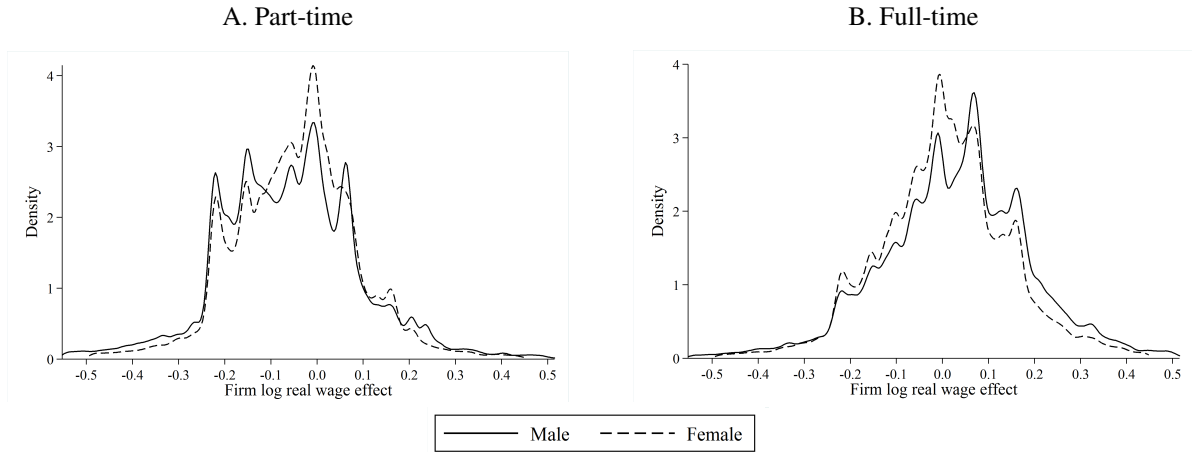
Appendix C. Additional figures

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



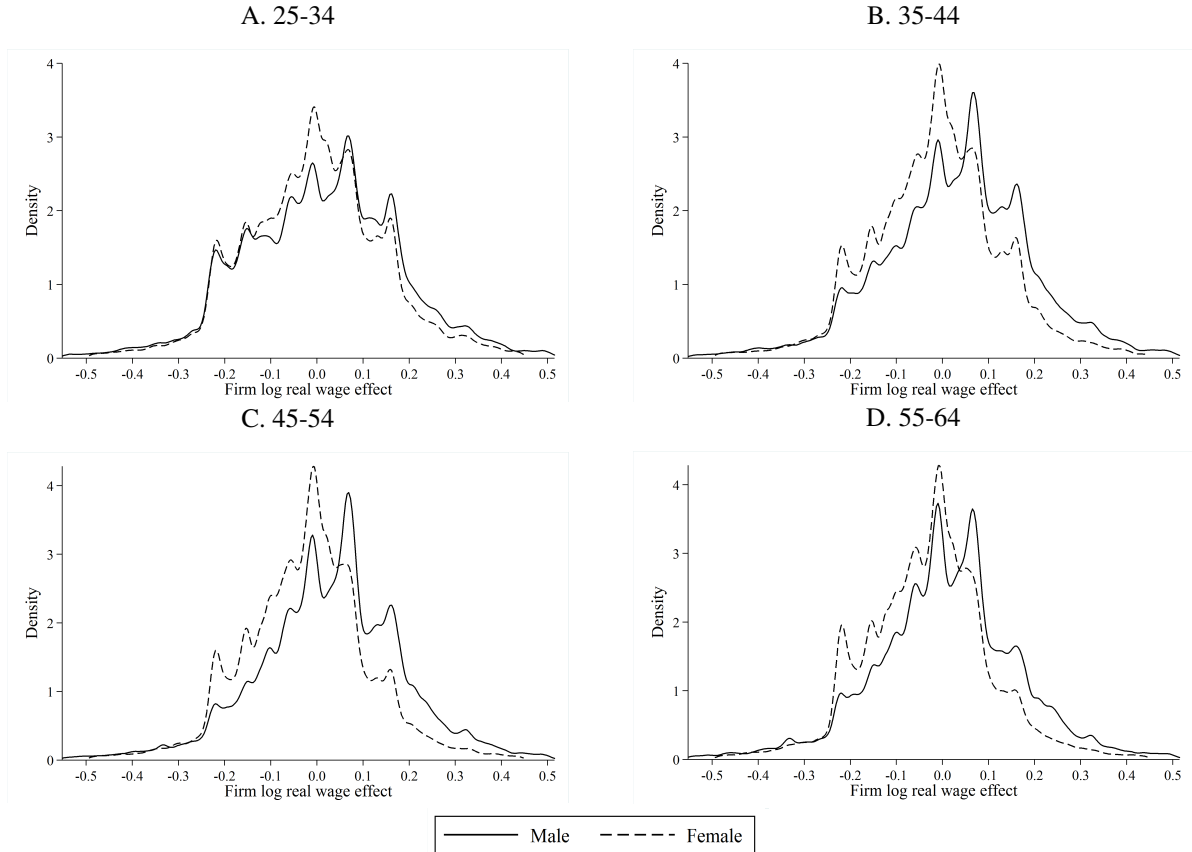
Notes.- see Figure 3. Both male and female kernel densities were estimated with a bandwidth of one log point. The top and bottom one percent of the overall set of estimated firm-specific effects are not displayed in any of the sub-figures. See Table 3 and Appendix A for descriptions of the industry groupings. 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 ($\hat{\phi}_{J(it)}$): employees in part- and full-time jobs



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 ($\hat{\phi}_{J(it)}$): age groups



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.

Appendix D. Additional tables

TABLE D1: Summary of estimated Full model with two-way fixed effects & decomposition of raw gender pay gap: gender-specific covariate effects

	Male (1)	Female (2)	Total (3)
St. dev. of log wages - $std_{it}(w_{it})$	0.55	0.49	0.53
<i>N</i> : worker-years	824,806	888,326	1,708,132
<i>P</i> : workers	131,903	124,501	256,404
<i>F</i> : firms			86,779
St. dev. worker effects - $std_{it}(\hat{\alpha}_i)$	0.46	0.37	0.42
St. dev. firm effects - $std_{it}(\hat{\phi}_{J(it)})$	0.20	0.17	0.18
St. dev. observables - $std_{it}(\mathbf{x}'_{it}\hat{\beta})$	0.53	0.46	0.51
Correlation - $corr_{it}(\hat{\alpha}_i, \hat{\phi}_{J(it)})$	-0.022	-0.010	0.004
Adjusted R^2			0.904
RMSE			0.164
<i>Variance shares</i> ($X/var_{it}(w_{it})$):			
Worker effects - $var_{it}(\hat{\alpha}_i)$	0.69	0.57	0.64
Firm effects - $var_{it}(\hat{\phi}_{J(it)})$	0.13	0.12	0.12
Covariance - $2covar_{it}(\hat{\alpha}_i, \hat{\phi}_{J(it)})$	-0.01	-0.01	0.00
Residuals - $var_{it}(\hat{\epsilon}_{it})$	0.07	0.09	0.02
Other	0.12	0.22	0.22
<i>Raw gender wage gap decomp. (shares):</i>			
Raw gap - $E_{it}[w_{it} i \in M] - E_{it}[w_{it} i \in F]$			0.223
Worker - $E_{it}[\hat{\alpha}_i i \in M] - E_{it}[\hat{\alpha}_i i \in F]$			0.174 (0.78)
Firm - $E_{it}[\hat{\phi}_{J(it)} i \in M] - E_{it}[\hat{\phi}_{J(it)} i \in F]$			0.036 (0.16)
Occupations			0.013 (0.06)
Observable / other			0.00 (0.00)

Notes.- author calculations using the ASHE 2002-16, all employees age 25-64. £2002. Pay excludes overtime. Gap is male minus female. Estimated Full model includes gender-specific covariates in \mathbf{x}_{it} for year fixed effects, squared and cubed terms for employee age, a cubic polynomial for employee tenure, a cubic polynomial for firm size (n. of employees) and a dummy variable for whether a worker was employed full-time.

TABLE D2: Summary of estimated Full model with two-way fixed effects & decomposition of raw gender pay gap: weekly earnings

	Male (1)	Female (2)	Total (3)
St. dev. of log earnings - $std_{it}(w_{it})$	0.63	0.76	0.74
<i>N</i> : worker-years	824,806	888,326	1,708,132
<i>P</i> : workers	131,903	124,501	256,404
<i>F</i> : firms			86,779
St. dev. worker effects - $std_{it}(\hat{\alpha}_i)$	0.47	0.48	0.49
St. dev. firm effects - $std_{it}(\hat{\phi}_{J(it)})$	0.26	0.27	0.27
St. dev. observables - $std_{it}(\mathbf{x}'_{it}\hat{\beta})$	0.60	0.72	0.71
Correlation - $corr_{it}(\hat{\alpha}_i, \hat{\phi}_{J(it)})$	-0.196	-0.187	-0.152
Adjusted R^2			0.902
RMSE			0.232
<i>Variance shares</i> ($X/var_{it}(w_{it})$):			
Worker effects - $var_{it}(\hat{\alpha}_i)$	0.56	0.40	0.44
Firm effects - $var_{it}(\hat{\phi}_{J(it)})$	0.17	0.13	0.13
Covariance - $2covar_{it}(\hat{\alpha}_i, \hat{\phi}_{J(it)})$	-0.12	-0.08	-0.07
Residuals - $var_{it}(\hat{\epsilon}_{it})$	0.08	0.09	0.04
Other	0.31	0.46	0.46
<i>Raw gender wage gap decomp. (shares):</i>			
Raw gap - $E_{it}[w_{it} i \in M] - E_{it}[w_{it} i \in F]$			0.498
Worker - $E_{it}[\hat{\alpha}_i i \in M] - E_{it}[\hat{\alpha}_i i \in F]$			0.245 (0.49)
Firm - $E_{it}[\hat{\phi}_{J(it)} i \in M] - E_{it}[\hat{\phi}_{J(it)} i \in F]$			0.067 (0.14)
Occupations			0.010 (0.02)
Other			0.176 (0.35)

Notes.- see Table 2.