On why gender employment equality in Britain has stalled since the early 1990s

Razzu, Giovanni and Singleton, Carl and Mitchell, Mark

University of Reading, University of Edinburgh, University of Edinburgh

1 May 2018

On why gender employment equality in Britain has stalled since the early 1990s

Giovanni Razzu[*] Carl Singleton[†] Mark Mitchell[‡]

May 2018
Link to latest version

Abstract

Using over four decades of British micro data, this paper asks why progress in closing the gender employment rate gap has stalled since the early 1990s. We find that how partner characteristics affected women's likelihood of employment explain most of the gap's shift in trend. Instead, changes to the structure of employment both between and within industry sectors impacted the gap at approximately constant rates throughout the period. There is evidence that continuing improvements in women's employment when they had children or higher qualifications worked towards narrowing the gap, even after progress overall had stalled.

Keywords: gender employment gaps; structural change; micro time series dataset
JEL codes: E24; J16; J21

[*]University of Reading, g.razzu@reading.ac.uk
[†]University of Edinburgh, carl.singleton@ed.ac.uk; with thanks to the Economic and Social Research Council (UK) for funding support under Grant No. ES/J500136/1.
[‡]University of Edinburgh, mark.mitchell@ed.ac.uk

1 Introduction

A long-term perspective over the Twentieth century shows the scale of the increase in British women's labour market participation and employment. When combined with the falling attachment to the labour force among men, these trends have resulted in a significant narrowing of the gender employment rate gap since the beginning of the last century.[1] However, in more recent decades a different but similarly stark pattern emerges: this progress towards equality between men and women appears to have stalled.

Figure 1A shows the absolute gap between male and female UK employment rates from 1971 to 2016. The overall reduction from thirty-nine to ten percentage points did not happen at the same rate throughout this period. The gap narrowed by around one percentage point per year between 1971 and 1990, but on average only a third of a point in the subsequent twenty-five years. This picture is not unique to the United Kingdom. For eighteen industrialised countries, including the United States, Canada and most of Western Europe, Figure 1B plots the relative changes in the female shares of total employment for two consecutive periods (before and after 1990) against each other. Most countries lie below the forty-five-degree line, implying a slowing down in the rate at which gender employment gaps have fallen. The relative stall in progress in Great Britain is slightly more pronounced than the average among these countries, but less than in the United States, Canada, Sweden, Denmark and Norway, though the Scandinavian countries had relatively smaller gaps left to close by 1990.

FIGURE 1: Gender employment gaps in the UK and elsewhere since the 1970s

(a) The UK working-age gender employment rate gap, 1971-2016

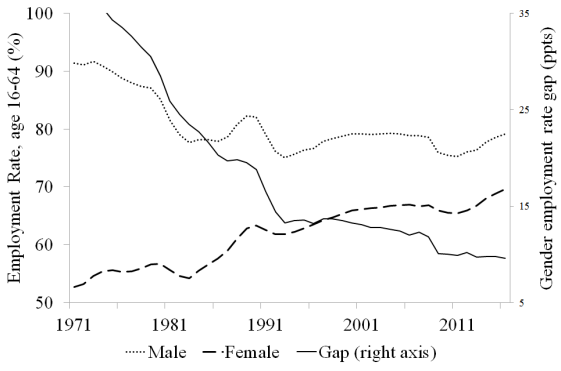

(b) The evolution of female employment shares across countries: 15 years before and after 1990

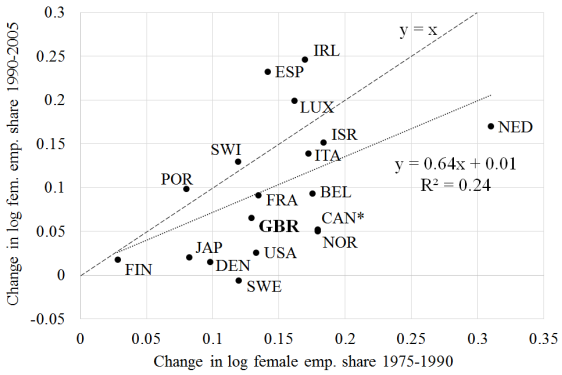

Notes.- (a) author calculations using ONS Labour Market Statistics (accessed 10/11/2017). Gap is expressed as male minus female. (b) author calculations using ILOSTAT labour market indicators (accessed 10/11/2017). *Canada is 1976-1990 due to lack of data availability for 1975.

[1] Throughout what follows and unless stated otherwise, gap refers to the percentage point difference of the female employment rate form the male counterpart.

Our aim in this paper is to shed light on why the gender employment gap in the UK since the early 1990s has not continued to narrow as quickly as it did in the preceding decades. No study we are aware of has addressed this issue head-on. Numerous studies, particularly using US data, have looked at the trend in female employment over a long period of time, but these, although extremely important, represent evidence on just one side of the gender employment gap. Here we also look at the patterns in male employment to develop a set of facts and conclusions about the evolution of the UK gender gap.

We identify three broad perspectives, or sets of evidence, from the existing literature on female employment in the long-run, which we will describe in more detail in Section 2. The first emphasises the periodical evolution of female work. For example, the distinct 'U-shaped' development in US female participation (e.g. Goldin, 1990), the discontinuous impacts produced by the Second World War (e.g. Goldin and Olivetti, 2013), and the important role played by specific birth cohorts (e.g. Joshi et al., 1985 for the United Kingdom and Goldin and Mitchell, 2017 for the United States). The second perspective highlights specific factors, either distinctly or in combination, concerning how men and women interact with the labour market. These factors include the roles of education and human capital, hours of work and labour supply, technological change and its impact on household work, fertility, social attitudes towards women's employment, and the role of legislation and policies, such as Equal Pay, taxation and parental leave. The third perspective emphasises the impacts of changes in the economic structure of countries, specifically the changes in the services shares of output and employment (e.g. Ngai and Petrongolo, 2017). In this paper, we provide new evidence on the roles played by individual worker characteristics and industry structural changes on employment equality progress in the UK since the early 1970s.

We do so using micro data from representative samples of British households. We study a measured fall in the gender employment rate gap of seventeen percentage points between 1973 and 1989, and a further fall by 2005 of only seven points. We first ask whether the changing distribution of some observable characteristics among men and women can account for the time series of employment rates and the gender gap. We find that the largest part of the stall in progress is accounted for by changes in employment rates related to what we call partner characteristics, which include a partner's employment status and their education levels. The likelihood of employment for a man living in a single person household or with a non-employed partner fell consistently since 1973. But, whereas the increased employment rate of women whose partners were also working accounted for over eighty percent of the rise in aggregate female employment before 1989, this increase slowed dramatically thereafter, and accounted for the majority of why the gender gap stalled. These results suggest that the

gender pattern of work in Great Britain is closely related to trends in the share of prime working-age men and women living in households with no work. Gender equality in employment was driven by large improvements in the within household distribution of work, at the expense of greater polarisation of work between households.

It is plausible that the historical rise in the services sectors has been to the comparative advantage of women and has reduced gender work equality in Britain. Confirming this, we find that between-industry shifts in employment accounted for an increase in the female share of overall employment at approximately the same rate before and after the early 1990s. Within-industry shifts in the types of workers employed have tended not to favour female employment but, similarly to the between-industry effects, have done so at about the same rate over the past four decades. Therefore, the pace of industrial structural change in Britain does not account for why employment equality progress has stalled. We also show that the overall slowdown in closing the employment gap has coincided with worsening representation of women within several industry sectors, such as in Manufacturing and Banking & finance, after we condition on the types of worker characteristics, excluding gender, that these sectors have tended to employ.

The remainder of the paper proceeds as follows: Section 2 reviews the existing evidence; Section 3 describes how changes to the composition of employment among men and women, and who is working, accounted for the decline and stall in the British employment rate gap; Section 4 analyses what role the pace of industry structural change has had; Section 5 summarises and further discusses our main findings. Further information is presented in the Online Appendix describing the main source of data, how we constructed a sample of households for analysis and some additional results.

2 Reviewing the evidence on gender employment gaps in the long run

The evidence is fairly limited on how gender gaps in the employment rates have evolved over time. However, wider insights can be drawn from a large literature which focuses on how female employment has evolved in the long run. In what follows, we structure a review of this evidence around the three broad perspectives described above, which are nonetheless to some extent interrelated. Before proceeding, we note that most of this literature concerns the US experience, with there being far less evidence for the United Kingdom.

Historical development and cohort-based explanations

In Great Britain, the female participation rate in both 1850 and 1960 was around thirty-five to forty percent (Joshi et al., 1985; Costa, 2000; Olivetti, 2014). In the more commonly studied United States, in 1890 only eighteen percent of women worked for pay, increasing to twenty-six and sixty percent by 1940 and 2000 respectively (Olivetti, 2014).

Goldin (1990, 2006) distinguishes three phases in the evolution of US female participation over the last two centuries. The first phase, from the late Nineteenth century to 1920, was characterised by the expansion of single women's employment, with married women's work being at this time still stigmatised. The second phase, from 1920 to 1950, witnessed growth in married women's paid employment, facilitated by increases in office work, new technologies, both in work and the home, and wider school and university enrolment. The third phase, from 1950 to 1970, saw the expansion of married women's work at an even greater rate. Within this long period, the 1940s are considered by some as a watershed period in the United States: a decade of huge changes in the labour force. However, others have argued that the female employment surge was only temporary and, as soon as the War was over, women returned to exclusive household work.[2]

Goldin (1990)'s long-term analysis of US women's work also represents a helpful introduction to cohort-based analyses. For US women born before the 1950s, she shows that labour force participation over the life cycle rose steeply from their twenties to late forties, and then decreased in their early fifties. But this changed for the women born since the mid-1950s, who experienced a higher participation rate overall, but with a dip in the middle (the so-called 'sagging middle'), and a phasing out that began at a later age than for previous cohorts. Goldin and Mitchell (2017) suggest that the sagging middle of the life cycle is due to the greater employment rates of women in their twenties, together with delayed childbearing and the impacts of motherhood. In another prominent US cohort-based analysis, Attanasio et al. (2008) study the life-cycle labour supply of three birth cohorts in the US: the Elizabeth Dole (1930s), Hillary Clinton (1940s) and Ophra Winfrey (1950s) generations. While the Clintons worked more than the Doles, both cohorts' participation was low during childbearing years, which was no longer the case in the Winfrey cohort. The authors' analysis suggests that the increased participation over the life cycle for the 1950s cohort was accounted for in a large part by decreased childcare costs relative to wages.

[2]Goldin (1991) gives a comprehensive summary of the watershed and the revisionist approaches to the role of the Second World War on female labour force participation. The two perspectives are not exclusive (Goldin and Olivetti, 2013): women's labour supply was altered by the War, particularly for those who entered white-collar occupations, but there is variance between groups of women in the persistence of effects from this large shock to the economy and society.

Evidence for the United Kingdom is more sparse. Joshi et al. (1985)'s noteworthy analysis of birth cohorts working between 1850 and 1980 found that, in post-war Britain, the increase in female participation, until the 1970s, was mainly among married women aged over thirty-five, and that cohort-specific factors explain most of the secular increase in participation. Gregg et al. (2007) analyse the role of relevant policy and legislative changes during the period 1974-2000, isolating those birth cohorts whose mothers experienced significant increases in employment. They find that an important role is played by maternity rights, which affected mothers' decisions to return to work.

Explanations based on specific factors

The majority of the literature has focussed on the role of a range factors in explaining the the evolution of female employment over time and the differential experiences of some cohorts. These factors include: increased education; reduced hours; increased real wages; technological change; decreased fertility; increased awareness of the role of women from changing social attitudes or the expansion of the feminist movement; the household division of labour; the income and labour market positions of partners. The impact of these factors has more often been studied in isolation rather than in combination. In an example of the latter, however, Eckstein and Lifshitz (2011) analyse the trends in US female employment between 1964 and 2007 with a dynamic model including education, earnings, fertility, marriage and divorce. They find that more education accounted for around a third of the increase in female employment, while the rise in real wages and narrowing of the gender wage gap accounted for a further twenty percent. Both a decline in fertility and increase in the prevalence of divorce accounted for small amounts, while approximately forty percent of the increase remains unexplained by observed characteristics. For the UK, Gomulka and Stern (1990) study the evolution of employment between 1970 and 1983. They conclude that changes in behaviour, as opposed to changes in individual characteristics such as age and education, played a major role. They suggest this was possibly affected by the introduction from the 1970s of more stringent employment and equalities legislation. Gutiérrez-Domènech and Bell (2004) extend this analysis to the period 1984-2002. They also find that in the 1980s changes in behaviour contributed to female participation growth. Overall, they suggest that periods of greater female participation changes are those in which changes in behaviour have a positive and significant impact.

The role of human capital in long-term labour market trends has been the subject of a vast literature (see for example Acemoglu and Autor, 2011). Most relevant to the gender-based discussion here, Olivetti (2006) shows that the substantial increase in US married women's hours of market work between 1970 and 1990 can mostly

be explained by greater returns to labour market experience. Her estimates of the gender-specific human capital production function show that women's marginal returns to experience increased by twenty-five percent across the decades, compared with a six percent increase for men. When applied within a model which features human capital accumulation, home production and childcare provision for married couples, these relative changes in returns to experience can explain ninety-six percent of the observed increase in married women's hours of work over their life-cycles.

The development of labour-saving technology in the household sector, such as washing machines and microwave ovens, can explain more than half of the rise in US female labour force participation between 1900 and 1980 (Greenwood et al., 2005). Technological progress in the home has also been shown as a significant driver of the decline in marriage and rise in divorce in the United States, which have had additional effects on US employment gaps and income inequality, especially between households (Greenwood et al., 2016). Other advances in the household affecting the time costs of women's reproductive role and their comparative advantage in childcare, for example instant (baby) formula, were probably significant in married women's increasing participation in the mid-part of the Nineteenth century (Albanesi and Olivetti, 2016).

The extent to which the rise in female participation can be explained by increases in relative real wages has also been well-studied. Smith and Ward (1985) find that rising real wages accounted for sixty percent of the total expansion in the US female labour force during the 1950-1980 period. Layard et al. (1980) estimate labour supply elasticities from the 1974 UK General Household Survey, concluding that real wages explained about one third of the increase in female participation between 1973 and 1977. Related is the evidence on the so-called added-worker effect, whereby women might theoretically increase their labour supply due to a reduction in their partner's income (Layard et al., 1980; Attanasio et al., 2005; Benito and Saleheen, 2013).[3] This has been a popular explanation for the rise in US female participation in the 1970s and 1980s, since this period also coincided with a significant decline in aggregate male wage growth. However, Juhn and Murphy (1997) show that the individual-level evidence during this period tells a different story: the most rapid increases in female US labour supply during this period were among women married to men with relatively high wages, and high-wage men also experienced greater wage increases over this period; conversely women married to low-wage men slowed their rate of labour force entry in the 1970s and 1980s, just as their husband's wage growth also declined rapidly during this period.

[3]The empirical evidence for this effect mostly relates to short-term changes within households following transitory economic shocks and the business cycle (see for UK specific evidence of this effect: Razzu and Singleton, 2016, Bryan and Longhi, 2017 and Bredtmann et al., 2017), rather than explaining the long-term trends in overall employment that we focus on here.

The impact of legislation and other institutional developments has also been widely studied. These have affected, for example, fertility rates, medical advances and the tendency of women to marry later (Goldin and Katz, 2002), to plan childbearing and their career (Bailey, 2006), and to reconcile work and motherhood (Albanesi and Olivetti, 2016). Specifically, Olivetti and Petrongolo (2017) find little compelling evidence on the impact of parental leave on gender gaps in a set of OECD countries, while finding stronger evidence on the positive impact of early childhood spending and in-work benefits. For Britain, Gregg et al. (2007) conclude that it was the interaction between legislative measures and an increase in relative earnings, along with the reduction in taxation of part-time work, that determined the overall transformation of married women's employment between 1974 and 2000. These effects on the British labour market were not uniform but concerned mostly women with employed partners and relatively high levels of education.

A common theme of these explanations and the related empirical evidence is that the issue of gender labour supply cannot be separated from what happens at the household level. Indeed, some of the most prominent research on labour supply has emphasised the household context in which work and consumption decisions take place (e.g. Blundell and Macurdy, 1999; Doepke and Tertilt, 2016). Moreover, the development of women's employment during most of the Twentieth century was a story of increased *married* labour force participation driving the overall level. A more recent literature has suggested that the economic nature of households, and especially the differences between them, has changed markedly in recent decades. Gregg et al. (2010) document the conflicting trends between individual and household labour market outcomes for the United States, Britain, Germany, Spain and Australia. In these countries, individual worklessness rates have fallen over the past two decades, while household-based rates have not. For the UK, Gregg and Wadsworth (2001, 2008) find that the most important contribution over time has come from work polarisation across households. For a given employment rate there have been large increases in the prevalence of dual-earner and no-earner couples over time. Jobs appear to have transferred to better-educated women, while older and less-educated men have lost out. These groups of men and women do not typically share the same households.

Explanations based on changes in industry structure and the demand for skills or occupations

Finally, another perspective emphasises the role of changes in the economic structure of countries on gender equality. Echoes of this can be traced as far back as Adam Smith's argument that the development of commerce and industry would improve women's economic position by reducing the social position of the military (Dimand

et al., 2004). More recently it would be the decline of jobs in manufacturing and the growth of the services sectors that plays to the comparative advantage of women, resulting in a relative increase in the demand for female labour (Goldin, 2006; Ngai and Petrongolo, 2017). This comparative advantage can be explained by the services sectors being associated with the more intensive use of 'brain' and interpersonal skills or tasks, as opposed to those associated with 'brawn', and therefore by these sectors creating jobs that are considered more suitable for women who aspire to work. Olivetti and Petrongolo (2014) and Olivetti and Petrongolo (2016) show that the process of structural change, through differences and changes in the services share of the economy, can account for more than half of the cross-country variation in women's relative labour market outcomes.

Stalling gender equality progress

Few studies focus on explaining the stalling of progress since the 1990s. Goldin (2006) points out that whether or not the trend in US women's participation has levelled off depends on the age group examined. The evidence is that women in their thirties continue to be committed to the labour force, probably because "they were in the labor force for longer before they married, invested more in formal and informal training, and shaped their identities before having children". Blau and Kahn (2007)'s analysis of US married women's labour supply finds that most of the change since the 1990s is explained by the dramatic fifty percent reduction in women's own wage labour supply. Considering that a husband's own wage elasticity did not change substantially in this period, they conclude that married women's and married men's responses to wages were becoming more similar.

Aaronson et al. (2006) show how the lower US participation rate since 2000 is in stark contrast to the increasing trend of the previous three decades. They suggest the stall since the mid-1990s in the progress of higher female participation with each successive birth cohort was associated with cohorts of women born since the 1950s not having substantially greater attachment to work than those before them. This is confirmed by Lee (2014), who finds this was the case regardless of education, childbearing, caring responsibilities and marital status. Juhn and Potter (2006) also suggest that the female rate is unlikely to rise above the level seen in the 1990s, unless there are Government or private sector-led initiatives, such as the expansion of income tax credits, changes to social security that increase the retirement age and employer-led initiatives to improve flexible working. Fortin (2015) suggests that gender role attitudes explain a third of the US levelling off: post-1966 birth cohorts hold more traditional gender role attitudes than the baby-boomer cohort, particularly regarding housewifery. Macunovich (2009), also for the United States, finds that the reversal in the pattern of ever greater female employment rates has not only been

among highly educated women with young children, but has been more general and especially marked among single women, women with no children, and women with more than sixteen years of education (college graduates), who have high marriage and fertility rates.

In summary, substantial explanations are still missing on the stalled progress since the 1990s while most research has focused on the longer term increase in female employment. Regarding the latter, three explanations appear to be relatively prominent: the household dimension, the increase in real wages and the potential role played by structural changes to the economy and the associated increase in the employment share of the services sectors. In what follows we provide further evidence on two of these for the United Kingdom. Furthermore, unlike the employment rate, the decreasing gender gaps in hours worked per person and in real wages have not stalled in recent decades in Great Britain, as demonstrated by Figure 2. It is beyond the scope of this paper to explore more precisely these differences in the trends of other aggregate gender gaps in the British labour market. This is not through lack of interest, but because of the paucity and inconsistency of historical earnings data. Nonetheless, these differences on the surface allow us to tentatively conclude that the factors which accounted for the stalled progress in closing the employment rate gap have acted mostly on the extensive margin of labour supply.

FIGURE 2: Mean hours worked per person and mean log real hourly wages, with gender gaps

(a) Real hourly wages, 1975-2005

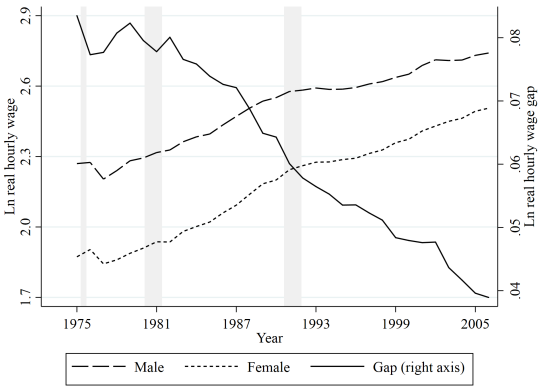

(b) Hours worked per person. 1973-1998

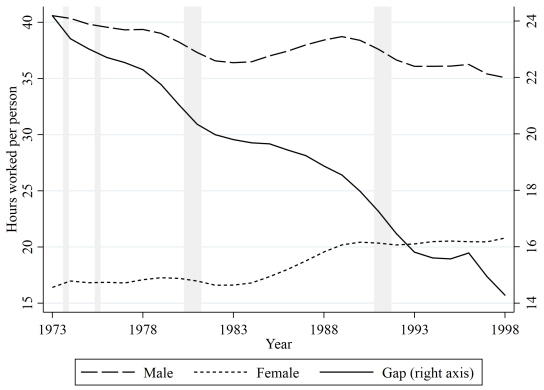

Notes.- the series for "Real hourly wages" are taken from Elsby et al. (2016), who calculated these mean values using the representative British New Earnings Survey Panel Dataset for all employees aged 25-59 in April, excluding the top and bottom 1% of male and female hourly wages in each year. Real wages were expressed in April 2012 pounds using the ONS RPIX prices index. "Hours worked per person" are author calculations using General Household Survey, 1972-1998, ages 25-54 only: the same sample of individuals and households as described throughout Sections 3 & 4. Working hours were not consistently recorded after 1998 when compared with before. See Online Appendix Table A2. Gaps are expressed as male minus female. Shaded areas denote official periods of UK recessions.

3 Did the British employment gap stall for some but not for others?

We use microdata from Great Britain's General Household Survey (GHS) covering the period 1972-2006. This was a repeated annual cross-section study, collected as a stratified random sample, and contains information for households and all their individual members.[4] We exploit a time series (or pseudo-cohort) version of the dataset for 1972-2004 created by the UK's Office for National Statistics (ONS). We further extend this time series to 2006 and add some information from the annual datasets. We only use variables that change minimally over the period in question, and so can be considered approximately as part of a consistent time series dataset.

The main variable of interest is employment status. The personal characteristics we focus on as potential explanatory variables are age groups [6 categories], highest qualification [4], number of dependent children [3], whether or not at least one of those children is an infant (aged less than five years: pre-compulsory UK school age) [2], and marital/partner status [2].[5] For those who are married, cohabiting, or in a same-sex couple, the partner's employment status [2] and education [4] are also considered. Once these characteristics are fully interacted, besides age groups which are only considered later, there are 180 distinct 'types', or groups, of workers. For example, one of these distinct types contains married men/women, who have a degree and two non-infant children, whose wives/husbands are employed and have further education qualifications. The final analysis dataset we use includes partnered and single individuals of prime working age (aged 25-54) with non-missing values for all explanatory variables, including partner characteristics. This contains approximately 7-12 thousand observations per year. To improve the accuracy of our estimates, we pool the annual samples in each rolling three-year period: i.e. estimates in what follows for 1973 are based on the pooled sample for 1972-1974. Online Appendix A provides further details on the GHS, its design, how our sample was constructed and how all variables were derived, including all of their categories. Online Appendix Figure A1 shows that, for the resultant sample, the trends we observe in the male and female employment rates and, especially, the level of the gap between them match those published by the ONS for the full UK working-age population over the period 1972-2006.

[4]We favour this source over the Labour Force Survey (LFS), from which official UK time series of aggregate employment are derived, because the information collected by the GHS is more detailed and more consistent over time, with fewer changes in the sampling methodology or the survey questions, especially concerning the key variables of interest here.

[5]The values in square brackets give the number of derived categories for each of these variables which we use in the following analysis.

What a gender-blind distribution of work could look like

To illustrate how the employment rate gap has changed over time, in the context of shifting demand for worker characteristics, we begin with a motivating and descriptive representation of what a 'gender-blind' distribution of work could look like in Great Britain. Let the observed employment rate of a worker with gender $g \in \{m, f\}$ and type of personal characteristics i be e_{ig}. The share of the male or female population who are type i is given by α_{ig}. Terms referring to non-gendered employment rates or shares of the population omit the g subscript. A starting point for a *counterfactual* gender-blind distribution of work is one where only differences in population type shares, given the non-gendered employment rate for each type, should determine the observed gender gap:

$$Gap_{Cf} = \sum_i (\alpha_{im} - \alpha_{if}) e_i \ . \tag{1}$$

We then describe an *adjusted* gap, subtracting this counterfactual from the *actual* value, thus addressing how much of he gap in any period is not accounted for by the different distribution of worker types among men and women:

$$Gap_{Adj} = Gap - Gap_{Cf} \ . \tag{2}$$

In this way, we describe what the gap would be if the type-specific employment rates of men and women were the same. Or, in other words, what the gap would be if it was determined only, for example, by higher employment rates for those with degrees, or for those who are married, and so on.

Figure 3 shows these measures of the UK actual, counterfactual and adjusted gaps over four decades since 1973 (panel b), along with the associated employment rate series of men and women (panel a). The actual employment rate gap declined consistently throughout the entire period, falling from thirty-eight to fifteen percentage points between 1973 and 1993. Thereafter, it levelled off, decreasing by only a further three points by 2005. During this later period, female employment continued to increase, but male employment reversed its previous trend. The counterfactual gap also decreased from around ten points in 1973 to five points in 1993: if men and women had gender-blind employment rates, we would have expected to measure a decrease in the gender gap between 1973 and 1993 which was less than a third of that actually observed. The overall non-gendered composition of British employment moving through this period tended to decreasingly favour the personal characteristics held by men relative to women. Not only do gender differences in workers' type-specific employment rates account for most of the gender gap throughout these four decades, but they also account for the majority of its decrease too. Since 1993, they also account for the entirety of why progress in closing

the gap has stalled: despite accounting for a minority of the actual gap's decrease before 1993, the counterfactual accounts for all the decrease since then to 2005. This suggests that the changing overall composition of work, and the non-gendered supply or demand of personal characteristics to the labour market, is unlikely to explain why progress in closing the gender gap has stalled more recently. Instead, this motivates us to look for evidence of changes in the gender-specific trends of labour supply behaviour.

FIGURE 3: Gender employment rates and gap, age 25-54.

(a) Male & female

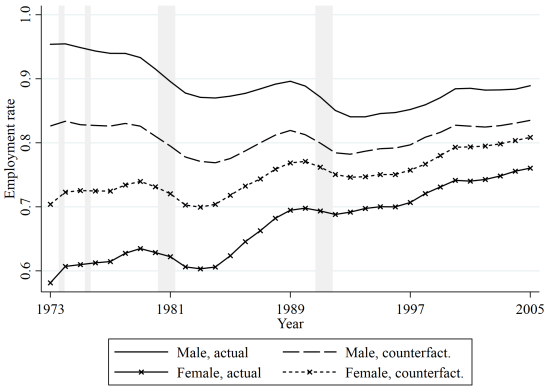

(b) Gender gap

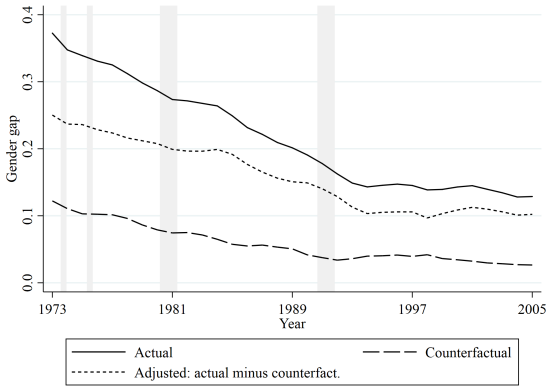

Notes.- author calculations using General Household Survey, 1972-2006. Gap is expressed as male minus female. The *counterfactual* should be interpreted as the level or gap which would occur only due to the different relative prevalence of characteristics among men and women; i.e. not because similar men and women have different employment rates. Shaded areas denote official periods of UK recessions.

These general measures and our discussion thereof offer little in the way of any specific explanation for why progress appears to have stalled since the early 1990s. We attempt to address this in what follows. Also, this gender-blind analysis is limited by the extent to which the values e_i, being functions of both e_{if} and e_{im}, provide a reasonable way to construct a counterfactual. By construction, the actual employment rate for each worker type is endogenous to the disproportionate demand and supply of male and female work throughout the economy. The distribution of work across personal characteristics could in fact be very different in a truly gender-blind economy.

Worker characteristics and household employment

We improve on the analysis above to assess how the employment rate gap has changed over time and to disentangle the roles of the gender type-specific employment rates and the composition of types in the population. We do this in two steps. First, we carry out a shift-share analysis of male and female rates, and consequently the gap, which sheds some light on what factors can account for the latter's trends. We then apply a regression-based decomposition of the changes in male and female employment rates and their gap through time. This latter approach considers each

set of worker and household types jointly, compared with the shift-share analysis, which can only address in turn the unconditional contribution from the changes in the employment rates associated with each mutually exclusive set of characteristics. To focus on accounting for why progress has stalled, we split our analysis period into equal length sub-periods, thus decomposing changes relative to initial values in years $x \in \{(19)73, 89\}$.

Shift-share analysis

Starting with the shift-share analysis, the change in the employment rate gap can be written as

$$\Delta_{t-x} Gap_t = \sum_i \Bigg[\underbrace{(\Delta_{t-x}\alpha_{imt})e_{imx} - (\Delta_{t-x}\alpha_{ift})e_{ifx}}_{\text{Composition shifts}} + \underbrace{\alpha_{imx}(\Delta_{t-x}e_{imt}) - \alpha_{ifx}(\Delta_{t-x}e_{ift})}_{\text{Employment rate changes}} \Bigg] + \text{ Interactions }, \tag{3}$$

where t denotes some year subsequent to the initial value and Δ_{t-x} denotes the difference operator relative to that initial period.[6] The 'Composition shifts' represent how much of the change in the gap between years x and t is due to changes in the share of the male and female population who are of type i. The 'Employment rate changes' represent how much of the change in the gap is accounted for by differences in the employment rates change of males and females.

Table 1 reports the results of this decomposition. The last two columns account for how the gap declined by seventeen percentage points between 1973 and 1989 but only by seven points between 1989 and 2005. Composition shifts account for a third of the change in the gap in the second period, compared with around an eighth in the first period. Employment rate changes account for the clear majority of why the gap declined in the first period, and the lack of similar changes in the second period accounts for why its decline slowed during this time. This slowdown appears to be driven by a decrease of eleven and an increase of seven percentage points in the male and female rates respectively in the first period, compared with analogous decreases of only four and zero percentage points in the second period. This evidence suggests that it was these latter changes, as opposed to those in the composition of the

[6] We admit the interaction terms rather than re-scaling initial values to their averages over the periods. We favour this approach because the interpretation and graphical representation of what each component accounts for is more straightforward.

prime-working-age population, that accounted for the declining pace of the gender gap decrease from the 1990s onwards.

TABLE 1: Shift-share decomposition of changes in overall employment rates and the gender gap, 1973-89 and 1989-2005 (ppts)

	Male		Female		Gap	
	73-89	89-05	73-89	89-05	73-89	89-05
Composition shifts	0.8	2.3	3.1	4.8	-2.3	-2.5
Employment rates	-10.9	-3.8	7.4	0.0	-18.3	-3.8
Interactions	4.3	0.8	0.8	1.7	3.5	-0.9
Total (actual)	-5.8	-0.7	11.3	6.6	-17.1	-7.3
Employment rates:*						
By highest qualification:						
None	-8.7	-2.1	3.3	-2.1	-12.0	0.1
Any below A-level	-1.6	-1.1	3.0	1.4	-4.6	-2.5
A-level or prof.	-0.4	-0.5	1.0	0.6	-1.4	-1.1
Degree	-0.1	-0.2	0.1	0.1	-0.2	-0.3
By dependent children:						
No children	-4.0	-2.7	1.1	-1.2	-5.0	-1.5
>0, no infant	-3.6	-0.6	2.9	-0.5	-6.5	0.0
1-2, incl. infant	-1.6	-0.5	3.0	1.9	-4.6	-2.3
3+, incl. infant	-1.7	-0.1	0.5	-0.1	-2.2	0.1
By partner status:						
No partner	-7.0	-1.9	-0.3	-0.6	-6.7	-1.3
Partner employed	-1.5	-0.7	9.6	1.4	-11.1	-2.1
Partner non-employed	-2.3	-1.2	-1.9	-0.8	-0.4	-0.4

* Contributions from changes only in the employment rates associated with each mutually exclusive set of characteristics; i.e. each set of contributions sums to the overall contribution for the time period from 'Employment rates' stated above.
Notes.- author calculations using General Household Survey, 1972-2006, ages 25-54 only. Gap expressed as male minus female (changes).

We can also address how changes in employment rates by highest qualification account for the decline and stall in the gender gap. Table 1 shows how between 1973 and 1989 the fall in the gap was mostly accounted for by changes in the employment rates of men and women with no qualifications, increasing the overall female rate by three percentage points, and decreasing the male rate by nine points. However, between 1989 and 2005 both male and female employment rates for those with no qualifications decreased by two percentage points. Therefore, the majority of the slowdown in closing the gender gap is accounted for by the employment rates of men and women with no qualifications equally worsening in recent decades, whereas previously women's increasing relative employment in this group had rapidly contributed to closing the overall gap.

Looking at the contribution of employment rate changes solely by the individual's dependent children type, much of the reduction in the pace at which the gap declined between the two periods can be accounted for by the changes for those with no children or no infant, and especially for women in these categories. The direction of the change in employment rates they experienced over the period 1973-1989 was reversed in 1989-2005: the increasing employment rate of women with no, or non-infant children initially contributed to decrease the gender gap, but from 1989 to 2005 the reduction in the employment rate of these types reduced the overall female rate.

In terms of 'partner status', it appears to have been the large increase in the employment of women with employed partners in the first period, relative to the second, that determined the stall in the closing of the gap. Over half the contribution of employment rate changes in 1973-1989 was driven by changes in the rate of employment of these women, whereas in 1989-2005 this proportion fell to just over one third. The reduction in the employment rate of men with no partner in 1973-1989 accounted for roughly two thirds of the reduction in the overall male employment rate in 1973-1989, compared with just less than a half in 1989-2005.[7]

Regression-based decomposition

In Equation 3, and the results of Table 1, the impact of employment rate changes was decomposed by highest qualification, number of dependent children and partner status in turn. As a result, this exercise only allows for relative comparisons of the impact of employment rate changes within each mutually exclusive group. Here we quantify more formally how much of the change in male and female employment rates and, consequently, the gap between them, can be jointly explained by changes to the composition of the workforce, or by changes to values which associate worker characteristics with employment. To do so we use the Blinder-Oaxaca decomposition as extended for use in binary-response models by, for example, Fairlie (1999, 2006) and Bauer and Sinning (2008). In contrast to the shift-share analysis above, this decomposition allows us to simultaneously quantify the impact of differences in the responses of employment from multiple type characteristics, for 1973-1989 and 1989-2005, on the change in the gender employment rate gap over the same periods.

With our dataset being a pseudo panel, as opposed to a true panel, it is not entirely possible to disentangle the changing preferences and behaviours of cohorts from the

[7] Online Appendix Figure B1 depicts the male and female employment rates and the gap between them implied, relative to the base years of 1973 and 1989, only by changes in the distribution of single or partner households among men and women. Panel c therein shows that these implied changes in the gap are small relative to each base year, suggesting that it was not simply changes in the prevalence of single or couple households that drove the gap's pattern over time. Figures B2-B4 also provide analogous illustrations of changes for each of the mutually exclusive sets of characteristics described in Table 1 respectively.

true relationships between observables and employment. For instance, even if our samples in 1973 and 1989 were identical based on observables, it may be that the choices of individuals across time periods differ due to shifts in tastes and/or values. This means we are unable to say, for example, that any estimated coefficient effects are entirely due to exogenous changes in labour market supply behaviour. Having available a true longitudinal panel would not fully resolve this limitation, as it would result in an ageing sample. In an ideal scenario, but one which is impossible, we would be able to observe individuals in each period as each worker type. As we are concerned with the relationship between aggregate employment and individual characteristics, having a nationally representative sample in each time period is beneficial. It is also a possibility that large demographic shifts, for example in the prevalence of early retirement or school leaving ages, could result in our samples differing systematically over time. Focusing on a sample of prime working-age individuals in each period, should help to alleviate this problem by excluding those at the extremes of the labour force age distribution.

We first estimate logit models of male and female employment in selected years. The dependent variables considered in these models, contained in X_{gt}, are dummies for individual's five-year banded age [6 categories], partner characteristics (the interaction of whether or not single and their partner's employment status and highest qualification level) [9], dependent children (0, 1-2, 3+, interacted with whether any of the children is an infant) [5], and highest qualification [4], with all variables constructed as described in Online Appendix A Table A2. From each of these models we store the vector of coefficient estimates, $\hat{\beta}_{gt}$, which represent the change in probability of employment associated with a change in each variable in X_{gt}. We then decompose the change in the gender employment rate gap between the initial year x and time t as

$$\begin{aligned}\Delta_{t-x}Gap_t = &\underbrace{\left\{E_{\beta_{mx}}[e_{mt}|X_{mt}] - E_{\beta_{mx}}[e_{mx}|X_{mx}]\right\} - \left\{E_{\beta_{fx}}[e_{ft}|X_{ft}] - E_{\beta_{fx}}[e_{fx}|X_{fx}]\right\}}_{\text{Composition}} \\ &+ \underbrace{\left\{E_{\beta_{mt}}[e_{mx}|X_{mx}] - E_{\beta_{mx}}[e_{mx}|X_{mx}]\right\} - \left\{E_{\beta_{ft}}[e_{fx}|X_{fx}] - E_{\beta_{fx}}[e_{fx}|X_{fx}]\right\}}_{\text{Coefficients}} \\ &+ \text{Interactions}\,,\end{aligned} \tag{4}$$

where $E_{\beta_{mx}}[e_{mt}|X_{mt}]$ is the expected male employment rate conditional on the composition observed in period t, but evaluated with the parameters that associate observables with employment in period x. We use parameter estimates $\hat{\beta}_{gt}$ from each gender-year employment model as estimates of β_{gt}. The 'Composition' effect gives the part of the change in the gender employment rate gap explained purely by changes in the distribution of observables in X_{gt}. The 'Coefficients' effect is the amount

explained by changes in the response of employment rates, $\hat{\beta}_{gt}$, acting via the initial distribution of observable characteristics. With the latter effect derived directly from the coefficients of individual year-gender employment models, any differences in their impact across time periods may capture either changing labour supply or demand behaviour.

The results of the double (time and gender) Blinder-Oaxaca decompositions are summarised in Table 2 and Figure 4. Focusing on the last two columns of Table 2, which show the change in the gender gap, the overall Coefficients effect, as opposed to Composition, account for the vast majority of the decline in the employment gap and the stall in progress since the early 1990s. Between 1973 and 1989, these effects made up over ninety percent of the decline in the gap, in contrast to seventy-five percent between 1989 and 2005. More specifically, what appears to be driving both the decline and stalled progress is the Coefficients effect for partner characteristics, accounting for almost twelve percentage points, or around seventy percent, of the total decline of the gap between 1973-1989. In the second period, however, this effect accounts for none of the decrease in the employment rate gap. For men, changes in the regression model coefficients relating their partners' characteristics to their own employment likelihood reduced the overall male employment rate by nine and eight percentage points in the first and second periods, respectively. Analogous changes for women, however, resulted in an increase of three percentage points in the female employment rate in the first period, but a decrease approximately equal to that of the male rate in the second period. This reversal between periods would suggest a significant change in how the characteristics of women's partners affected the likelihood they were employed.

Table 2 also shows the results from decomposing further the joint Coefficients effect from partner characteristics. Again, focusing on the last two columns, the difference between the two periods in the specific Coefficient effect of having a non-employed or employed partner is what drove the overall change in the gap. The Coefficient effect for women whose partner is employed are in stark contrast over time. In the period 1973-1989, this effect, representing changes in how having an employed partner predicted an individual's own employment likelihood, accounted for four, over a third, of the eleven percentage point increase in the female employment rate. In the following years to 2005, however, this effect accounted for a seven percentage point reduction in the female employment rate. The analogous Coefficient effects reduced the male employment rate over both periods. But across the two periods, there is a large reduction in the male coefficient effect from having a non-employed partner. Between 1973 and 1989, this particular effect accounted for a fall in the male employment rate of over five percentage points. Thereafter to 2005, the same Coefficient effect only explains a further two percentage point decrease in the male employment rate.

TABLE 2: Blinder-Oaxaca type decomposition of changes in overall employment rates and the gender gap, 1973-89 and 1989-2005 (ppts)

	Male		Female		Gap	
	73-89	89-05	73-89	89-05	73-89	89-05
Composition:						
Age/Cohort	0.1	-0.1	0.6	-0.2	-0.5	0.1
Partner charcs	1.0	0.9	-1.9	-1.4	2.9	2.4
Dependent chld.	0.2	0.0	3.0	0.3	-2.8	-0.3
Highest qual. (self)	1.0	2.0	3.0	5.7	-1.9	-3.7
Coefficients:						
Age/Cohort	-5.1	1.2	0.3	3.0	-5.5	-1.9
Partner charcs	-8.8	-8.3	3.0	-8.4	-11.8	0.1
Single	*-1.3*	*-1.1*	*-2.2*	*-2.5*	*0.9*	*1.4*
Non-employed partner	*-5.6*	*-1.8*	*-0.4*	*-1.3*	*-5.2*	*-0.5*
Employed partner	*-2.1*	*-4.2*	*4.3*	*-7.0*	*-6.4*	*2.8*
Partner quals.	*0.2*	*-1.3*	*1.3*	*2.3*	*-1.1*	*-3.5*
Dependent chld.	1.0	1.5	-1.2	1.8	2.2	-0.4
Highest qual. (self)	2.6	1.6	3.5	5.0	-0.9	-3.3
Interactions	2.1	0.5	1.0	0.8	1.1	-0.4
Total (actual)	-5.8	-0.7	11.3	6.6	-17.1	-7.3

Notes.- author calculations using General Household Survey, 1972-2006, ages 25-54 only. The contribution from individual covariates are grouped as follows – 'Age/cohort': dummy variables for age groups; 'Partner chars': marital status interacted with dummy variables for a partner's employment status and their highest qualification; 'Dependent chld.': interaction of dummy variables for number of children and whether or not there is an infant in the household; 'Highest qual. (self)': individual's highest qualification level.

Both highlighted factors, which appear to have shifted the male and female employment trends, raise concerns of potentially important omitted variables. Some predictors of an individual's employment status are potentially missing from our analysis, such as disability status. These missing worker characteristics could be more or less correlated over time with those we are able to include in the regression models. For example, it may be the case that a large proportion of men have partners who are non-employed because of high levels of household wealth. If there were large increases over time in the number of men with wealthy, non-employed partners, then the estimated employment effects of having a non-employed partner will be overestimated, and the associated Coefficients effects we account for in the decomposition would then be exaggerated between periods.

In addition, Table 2 shows that a large proportion of the overall change in the employment rate gap is attributable to the Coefficients effect of what we label 'Age/Cohort': around a third and a quarter in 1973-1989 and 1989-2005 respectively. However, in this decomposition of employment rate changes over time, it is not

possible to disentangle the role of the life-cycle from birth cohort effects. However, we can see that our estimates of these effects did not have a major role in explaining why the gap's progress has stalled, compared with the much greater role of the partner characteristic effects.

There is also evidence in Table 2 that improved prospects for women who were highly qualified and who had dependent children continued to close the gender employment rate gap beyond 1990. The last two columns of the table show that changes in the coefficients relating these characteristics to employment reduced the gap by a larger amount in the second period relative to the first. Specifically, in 1973-1989 Coefficients effects accounted for a reduction in the likelihood a woman with dependent children being in employment, however in 1989-2005 they led to a two percentage points increase in the female employment rate. Also, while changes in the coefficients relating women's highest qualifications to employment resulted in an increased likelihood of employment in both 1973-1989 and 1989-2005, in the latter period the overall effect was one and a half percentage points larger. The Coefficients effect in the case of qualifications, however, explains roughly the same proportion of the total change in the female employment rate as the related Composition effect. This suggests that both changes in the distribution of qualifications among women, and how they affect the likelihood of female employment, have been important in continuing to close the overall employment rate gap.

Figure 4 also depicts results from the decomposition of the change in the employment rate gender gap using, for illustrative purposes, four equally spaced time sub-periods: 1973-1981, 1981-1989, 1989-1998, and 1997-2005. It shows the reduction in the total gender employment rate gap in each period. It also highlights the relative significance of Coefficients relative to Composition effects, with the former being smaller in magnitude in all but the last period when the majority of the slowdown occurred. For completeness, Online Appendix Figure B5 provides equivalent illustrations for changes in the male and female employment rates individually.

4 Does the pace of industry structural change account for stalling progress?

As summarised before, several studies have attributed the historical pace of structural change - typically describing the rise and fall of industry shares of aggregate output, total hours worked or employment - to trends in gender labour market outcomes in the UK and elsewhere. In a general equilibrium framework, changes to industry structure, the demand for specific worker types by those industries, and

FIGURE 4: Blinder-Oaxaca type decomposition of the change in the gender employment rate gap, 1973-81, 1981-89, 1989-97 and 1997-2005

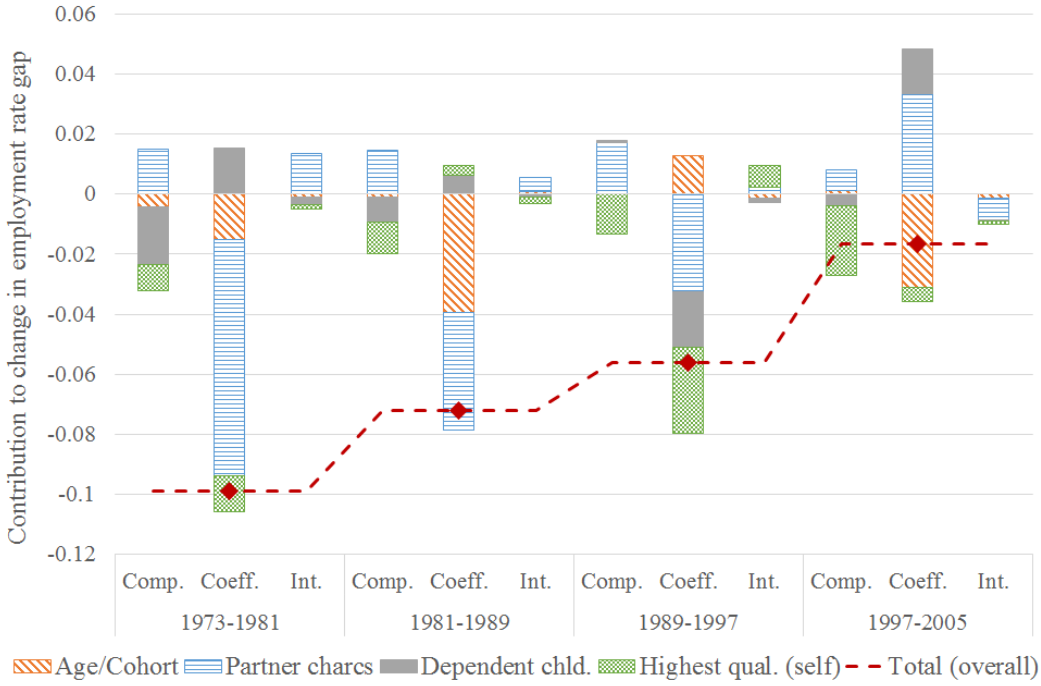

Notes.- author calculations using General Household Survey, 1972-2006, ages 25-54 only. Estimates are obtained from regressions of male and female employment likelihood on individual characteristics, and first using these to decompose changes between years for each gender, and then presenting how each of these together imply contributions to the change in the employment rate gap, stated as male minus female. See Online Appendix Figure B5 for a similar representation of the results for male and female changes. See notes to Table 2 for a description of the grouped sets of covariates defined by each legend item.

industry- or worker-type-specific wages, could all be co-determined with the supply of employment by different worker types (see for examples Olivetti and Petrongolo, 2014 and Ngai and Petrongolo, 2017). However, the analysis of the previous section broadly takes this co-determination as a given, and instead only accounted for how the changing composition of worker types and their associated employment rates might explain the stalled progress in gender equality in Britain since the early 1990s. In this section we look at the other side of the coin, accounting for the changing pace of industrial transformation in the British labour market, approximately taking as given how the prevalence of worker types in employment has also changed over time. Before proceeding, we should note the possibility that changes in male and female labour supply have in fact driven the changing industry composition of British employment. However, as argued by Lee and Wolpin (2006) and Olivetti and Petrongolo (2014), the existing evidence to date suggests this is not a major issue, and that it is a reasonable assumption that the growth in the services sector in developed economies was driven by demand factors associated with technical change, and that labour supply factors were approximately neutral in determining the industry mix of employment.

We classify workers across eight major industry sectors, which we define to give approximately consistent categories over time. However, unlike for the measures of worker heterogeneity discussed above, there are clear breaks in the time series of how British industry has been classified within the GHS. For the period between 1972 and 1980, the survey used the 1968 Standard Industrial Classification (SIC-UK). For the years between 1981 and 1994 it used the 1980 SIC-UK, and for the remaining years in our sample it used the 1992 SIC-UK, which, for comparability, we nonetheless convert to the 1980 classification using ONS correlations. To be confident that the seams between classification schemes are not accounting for any of our findings, we study the role of industrial change in British employment within these three separate periods, and then compare the results across the periods. [8]

Like Olivetti and Petrongolo (2016) and Ngai and Petrongolo (2017), who characterise trends in the female share of hours worked in several countries, we quantify how much of the changes in the female share of total employment are accounted for by changes between and within British industry sectors, using a shift-share analysis. Let F_{kt} be the share of all employment in sector k which is female in some period t. The share of sector k in total employment is represented by λ_{kt}. A shift-share decomposition of the change in the whole economy's female employment share, since initial period $x \in \{(19)73, 82, 94\}$ is then given by

$$\Delta_{t-x}F_t = \sum_k \left[\underbrace{(\Delta_{t-x}\lambda_{kt})F_{kx}}_{\text{Ind. composition shifts}} + \underbrace{\lambda_{kx}(\Delta_{t-x}F_{kt})}_{\text{Within-ind. gender shares}} + \underbrace{(\Delta_{t-x}\lambda_{kt})(\Delta_{t-x}F_{kt})}_{\text{Interactions}} \right]. \quad (5)$$

The results of this decomposition are presented in Table 3. The first component of the above, which we label 'Ind. composition shifts', gives the change in the female employment share since each initial period arising from changes only in the industry composition of employment: i.e. holding constant the composition of employment *within* each sector, but allowing only changes in the composition of employment *between* sectors. For the periods 1973-9 and 1982-93, approximately half of the rise in the female employment share was accounted for by changes in the industry composition of British employment. We also see in Table 3 that for the later period 1994-2005, these between-industry shifts in employment accounted for a rising female employment share at approximately the same rate as before. However, this was after progress in closing the gender gap had stalled, and so these changes then accounted for more than the actual rise in the female share of employment over that period. Based on this evidence, it appears a slowdown in the pace of industrial structural change does

[8] Online Appendix Table A2 contains further details on how we classified British industry for the purposes of the following discussion.

not account for the stalled progress in gender employment equality since the early 1990s. Therefore, it is unlikely that factors such as technical change could explain the stalled progress. Instead, as Figure 5 also shows for changes in the female share of employment throughout each of the three sub-periods, the slowdown in what we label as changes in 'Within industry gender shares' must account for the stalled progress in gender employment equality.

TABLE 3: Shift-share decomposition of changes in the female share of total employment, 1973-79, 1982-93 and 1994-2005 (ppts)

	73-79	82-93	94-05
Industry composition shifts	1.2	2.6	2.0
Within-industry gender shares:			
Employment-education gender shares	4.0	5.7	1.3
Industry education demand shifts	-1.6	-2.0	-1.9
Subtotal (incl. interactions)	0.9	2.6	-0.7
Interactions	0.1	0.1	0.4
Total (actual)	2.2	5.2	1.7

Notes.- author calculations using General Household Survey, 1972-2006, ages 25-54 only. See Online Appendix Table B3 for a further disaggregation of the 'Within-industry gender shares' component by industry sector.

FIGURE 5: Female share of total employment: levels implied by changes to the overall industry distribution of employment, or the gender shares within industries, since 1973/82/94

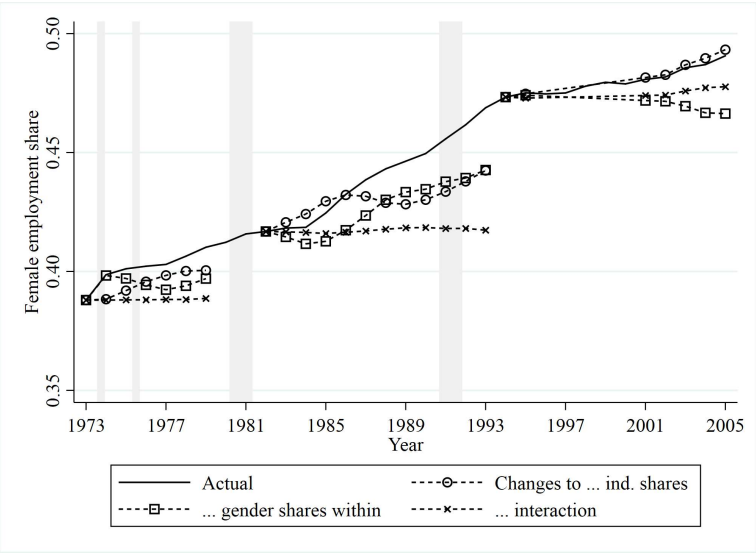

Notes.- author calculations using General Household Survey, 1972-2006, ages 25-54 only. Shaded areas denote official periods of UK recessions.

However, the role of factors such as technical change cannot be completely ruled out from the analysis above. For instance, if firms have tended to employ relatively

more highly qualified workers over time, possibly due to technical change, then we would expect this to have negatively affected the female share of employment, since men make up a larger proportion of those with higher qualifications. In order to assess this, we can rewrite the contribution in (5) from the changes in within-industry gender shares for any industry sector k as

$$\lambda_{kx}(\Delta_{t-x}F_{kt}) = \sum_i \Bigg[\underbrace{\lambda_{kx}\mu_{ikx}(\Delta_{t-x}F_{it})}_{\text{Emp. type gender shares}} + \underbrace{\lambda_{kx}(\Delta_{t-x}\mu_{ikt})F_{ix}}_{\text{Within-ind. emp. type shifts}} + \underbrace{\lambda_{kx}(\Delta_{t-x}\mu_{ikt})(\Delta_{t-x}F_{it})}_{\text{Interactions}} \Bigg], \tag{6}$$

where μ_{ikt} now refers to the share of worker type i in the total employment of sector k, and F_{it} refers to the female share in the total employment of worker type i. What we label 'Within-ind. emp. type shifts' now accounts for the role of changes to the types of workers typically employed by the firms in each industry sector. Specifically, we focus on the role of education levels in the workforce over time, i.e. suppose that i denotes the four levels of an individual's highest qualification. Table 3 shows how these 'Industry education demand shifts' do in fact account for a fall in the female share of overall employment in all three sub-periods. Similar to the contribution from between-industry employment shifts, this within-industry component would have accounted for a change in the female employment share at approximately the same rate throughout our sample period. This further emphasises that a change of pace in whatever the labour demand factors might be is highly unlikely to explain the stalled progress in the gap, either through such a slowdown affecting the extent to which employment has moved between sectors, or through how its composition has changed within. Instead, changes to 'Employment-education gender shares' accounted for large rises in the female share of employment before the early 1990s, but only a relatively small amount since. Overall then, this decomposition shows that other factors, as discussed in the preceding section, account for why recent progress in closing the gap has stalled in Britain.

Have employment gaps continued to close in some sectors but not others?

The previous results are largely silent on the heterogeneity in both the level and changes of the gender gaps within industry sectors. To shed light on this we return to our approach of describing what gender-blind distributions of work could have looked like. In a similar way as before, we derive the female employment shares within each industry sector that we would predict based on those sectors' gender-blind 'demand' for worker characteristics, using the same set of 180 different worker types as before in Section 3. For example, we consider what would be the

female share of the Construction sector if it employed men and women with no qualifications, who are married, with no children etc. in the same proportion as there are such men and women in the prime working-age population, without any regard for gender. We then subtract these predicted, or conditional, gender-blind female shares from their actual levels in each year, such that a negative value for this difference indicates that women were disproportionately under-represented in that sector. We interpret this as a measure of relative female under/over-representation, conditioning on the composition of worker types employed by each sector. Table 4 shows these differences for the eight major industry sectors and the years 1982, 1993 and 2005, as well as the changes between these years. As to be expected, we see that the difference between the actual female share and its predicted gender-blind level within some sectors is large, such as Construction, where it stood at minus thirty-five percentage points in 1982, falling further over the period. In four sectors there were improvements in conditional female representation for the period 1982-93, which were partly or completely reversed thereafter in the period 1993-2005: Primary divisions, Engineering, Distribution and Banking & finance. In Construction and Other Manufacturing female representation worsened in both sub periods. In Transport & communications female representation worsened in the first period but improved in the second, and in Other services there were small increases in female over-representation in both periods. In general, despite differences in the patterns between sectors, female representation within the most male dominated sectors, as well as in Banking & finance, appears to have worsened since the early 1990s, whereas there were improvements in the decade before. This tells us that the stalled progress in closing the employment gap is also symptomatic of women's conditional representation in certain industry sectors worsening, despite the closing of gaps with men in terms of some observable characteristics overall, such as education levels.

TABLE 4: Difference between actual and predicted 'gender-blind' female employment shares by industry (ppts)

				Change	
	1982	1993	2005	82-93	93-05
Primary divisions, 0-2	-27.9	-23.9	-31.4	4.0	-7.5
Engineering, 3	-24.5	-24.2	-27.2	0.4	-3.0
Other manufacturing, 4	-7.9	-10.0	-16.0	-2.1	-6.0
Construction, 5	-34.6	-36.4	-37.0	-1.7	-0.6
Distribution, 6	2.5	3.6	1.3	1.1	-2.3
Transport & comms, 7	-27.2	-27.8	-24.9	-0.6	2.9
Bank & finance, 8	-1.9	-0.3	-4.5	1.6	-4.2
Other services, 9	13.7	14.5	14.8	0.8	0.3

Notes.- author calculations using General Household Survey, 1981-83, 1992-93, 2004-2006, ages 25-54 only. Industry divisions listed are defined by SIC-UK 1980.

5 Summary and further discussion

In several developed countries progress towards gender equality in the labour market has slowed in recent decades. Progress in this regard stalled for British employment after the early 1990s. We have attempted to answer why this happened, using four decades of household micro data. Factors such as technical change, which we observe through their effects on the mixing of employment and worker types between and within industry sectors, do not appear to account for the relative changes in the trends of male and female prime-working-age employment. Similarly, the composition of individual characteristics between and among men and women does not address the stalled progress. Instead we find that changes in the employment rates associated with some individual characteristics, particularly those relating to partners or the lack thereof, account for why the employment rate gap has stalled. Before the shift in trend, the relatively rapid rise in gender equality had tended to take place within households, while employment inequality across households grew: in the employment pool more highly educated women, whose partners tended to also be in paid work, replaced less educated men, whose wives or partners did not see their own likelihoods of working increasing much at all.

A natural question to ask is whether this stalled progress is the new norm: can we expect the gender employment gap to close any time soon? We have shown that between-industry shifts in employment have consistently tended to account for a reduction in the British gender employment gap. We also found that changes in the observable composition of worker types filling jobs within industry sectors, with regards the levels of education, have over the same period consistently worked against closing the gap. However, the contribution from between-industry shifts was of a magnitude greater. But it is only practical to envisage that these shifts in employment cannot be limitless. The last UK Census in 2011 showed that eighty-one percent of the workforce was engaged in the services sectors, an increase of six percentage points compared with 2001, though the preceding decades had seen even larger relative increases in the services share of employment.[9] If the pace of structural change in Britain permanently slows, or that change becomes less friendly to women in terms of the types of jobs created and destroyed, then it is hard to see a path to the employment gap closing significantly further.

Weinberg (2000) showed that most of the rise in US demand for female workers since the 1970s could be accounted for by increases in computer use within and across industries and occupations. In this way, the computerisation of some tasks could explain the increased substitutability between high-skilled women and low-skilled men in the production process, which Olivetti and Petrongolo (2014) have argued

[9]Source: "Five facts about... the UK service sector", ONS Digital (2016): visual.ons.gov.uk/...

could be important in explaining gender wage and hours gap trends. However, in developed economies such as Britain's, computer adoption has been widespread throughout, if not completed. Instead, in the manufacturing sector there is now an emerging shift beyond just automation and computerisation to the so-called Industry 4.0, which defines a more pronounced use of cyber-systems and data analytics.[10] If this evolution in manufacturing technology takes hold, then it will arrive alongside even greater demand for very-high-skilled workers, particularly those with higher qualifications in Science, Technology, Engineering and Mathematics (STEM). It is well-known that women are under-represented in these subject areas among graduates (Card and Payne, 2017). So, while the initial computerisation of production led to the substitution of high-skilled women for low-skilled men, the most likely future trends will favour greater demand for the types of skills predominantly held by highly-educated men. This, alongside our related finding that female representation within the most dominated male sectors appears to have worsened since the early 1990s, suggests that if UK policy makers want to do more about gender employment equality, then they could look to redress the gender imbalance in education subject choices (Morgan and Carrier, 2014).

As summarised before, many previous studies have made clear that the long-run trends and recent patterns in female work have very much been a married women's affair. In looking ahead, however, it seems important to relate this more directly to the findings on household work polarisation of Gregg and Wadsworth (2001) and Gregg and Wadsworth (2008), since the household dimension of work sheds light on issues that a singular focus on individual employment rates does not. In our own analysis, we observe that the trends of household polarisation place a set of bounds within which gender employment rates have simultaneously evolved. The substantial growth in female employment before the early 1990s subsequently slowed down as it approached the limit imposed by these bounds; i.e. once there had already been very large increases in the employment of women whose partners were also very likely to work. The implication then is that little more progress could be achieved in the aggregate gap from further increases in the employment of highly-educated married women. For the gap to fall further, the bounds placed by the extent of the household polarisation of work imply that the labour market needs to see more women moving into work whose partners are not working or, paradoxically, men whose partners do work to stop working themselves.

[10]Industry 4.0 was a term used originally by the German Government as part of its high-tech strategy in the late 2000s. The "4.0" implies a fourth industrial revolution, after the preceding three: mechanisation, mass production/assembly line and computerisation/automation.

References

Aaronson, S., B. Fallick, A. Figura, J. Pingle, and W. Wascher. 2006. "The Recent Decline in the Labor Force Participation Rate and Its Implications for Potential Labor Supply." *Brookings Papers on Economic Activity*, 37(1): 69–154.

Acemoglu, D., and D. Autor. 2011. *Skills, Tasks and Technologies: Implications for Employment and Earnings*. 4 of Handbook of Labor Economics, Chap. 12 1043–1171 Elsevier.

Albanesi, S., and C. Olivetti. 2016. "Gender Roles and Medical Progress." *Journal of Political Economy*, 124(3): 650–695.

Attanasio, O., H. Low, and V. Sãnchez Marcos. 2008. "Explaining Changes in Female Labor Supply in a Life-Cycle Model." *American Economic Review*, 98(4): 1517–1552.

Attanasio, O., H. Low, and V. Sanchez-Marcos. 2005. "Female Labor Supply As Insurance Against Idiosyncratic Risk." *Journal of the European Economic Association*, 3(2-3): 755–764.

Bailey, M. J. 2006. "More Power to the Pill: The Impact of Contraceptive Freedom on Women's Life Cycle Labor Supply." *The Quarterly Journal of Economics*, 121(1): 289–320.

Bauer, T., and M. Sinning. 2008. "An extension of the Blinder-Oaxaca decomposition to nonlinear models." *AStA Advances in Statistical Analysis*, 92(2): 197–206.

Benito, A., and J. Saleheen. 2013. "Labour Supply as a Buffer: Evidence from UK Households." *Economica*, 80(320): 698–720.

Blau, F. D., and L. M. Kahn. 2007. "Changes in the Labor Supply Behavior of Married Women: 1980-2000." *Journal of Labor Economics*, 25(3): 393–438.

Blundell, R., and T. Macurdy. 1999. "Chapter 27 - Labor Supply: A Review of Alternative Approaches." 3 of Handbook of Labor Economics Elsevier, 1559–1695.

Bredtmann, J., S. Otten, and C. Rulff. 2017. "Husband's Unemployment and Wife's Labor Supply: The Added Worker Effect across Europe." IZA Discussion Papers 11087, Institute for the Study of Labor (IZA).

Bryan, M., and S. Longhi. 2017. "Couples' Labour Supply Responses to Job Loss: Growth and Recession Compared." *The Manchester School* n/a–n/a.

Card, D., and A. A. Payne. 2017. "High School Choices and the Gender Gap in STEM." Working Paper 23769, National Bureau of Economic Research.

Costa, D. L. 2000. "From Mill Town to Board Room: The Rise of Women's Paid Labor." *Journal of Economic Perspectives*, 14(4): 101–122.

Dimand, R. W., E. L. Forget, and C. Nyland. 2004. "Retrospectives: Gender in Classical Economics." *Journal of Economic Perspectives*, 18(1): 229–240.

Doepke, M., and M. Tertilt. 2016. "Chapter 23 - Families in Macroeconomics." 2 of Handbook of Macroeconomics Elsevier, 1789–1891.

Eckstein, Z., and O. Lifshitz. 2011. "Dynamic Female Labor Supply." *Econometrica*, 79(6): 1675–1726.

Elsby, M., D. Shin, and G. Solon. 2016. "Wage Adjustment in the Great Recession and Other Downturns: Evidence from the United States and Great Britain." *Journal of Labor Economics*, 34(S1): S249 – S291.

Fairlie, R. W. 1999. "The Absence of the African-American Owned Business: An Analysis of the Dynamics of Self-Employment." *Journal of Labor Economics*, 17(1): 80–108.

Fairlie, R. W. 2006. "An Extension of the Blinder-Oaxaca Decomposition Technique to Logit and Probit Models." IZA Discussion Papers 1917, Institute for the Study of Labor (IZA).

Fortin, N. M. 2015. "Gender Role Attitudes and Women's Labor Market Participation: Opting-Out, AIDS, and the Persistent Appeal of Housewifery." *Annals of Economics and Statistics*(117/118): 379–401.

Goldin, C. 1990. *Understanding the Gender Gap: An Economic History of American Women*. In , NBER Books(gold90-1): National Bureau of Economic Research, Inc.

Goldin, C. 1991. "The Role of World War II in the Rise of Women's Employment." *American Economic Review*, 81(4): 741–756.

Goldin, C. 2006. "The Quiet Revolution That Transformed Women's Employment, Education, and Family." *American Economic Review*, 96(2): 1–21.

Goldin, C., and L. F. Katz. 2002. "The Power of the Pill: Oral Contraceptives and Women's Career and Marriage Decisions." *Journal of Political Economy*, 110(4): 730–770.

Goldin, C., and J. Mitchell. 2017. "The New Life Cycle of Women's Employment: Disappearing Humps, Sagging Middles, Expanding Tops." *Journal of Economic Perspectives*, 31(1): 161–82.

Goldin, C., and C. Olivetti. 2013. "Shocking Labor Supply: A Reassessment of the Role of World War II on Women's Labor Supply." *American Economic Review*, 103(3): 257–262.

Gomulka, J., and N. Stern. 1990. "The Employment of Married Women in the United Kingdom 1970-83." *Economica*, 57(226): 171–199.

Greenwood, J., N. Guner, G. Kocharkov, and C. Santos. 2016. "Technology and the Changing Family: A Unified Model of Marriage, Divorce, Educational Attainment, and Married Female Labor-Force Participation." *American Economic Journal: Macroeconomics*, 8(1): 1–41.

Greenwood, J., A. Seshadri, and M. Yorukoglu. 2005. "Engines of Liberation." *The Review of Economic Studies*, 72(1): 109–133.

Gregg, P., M. Gutiérrez-Domènech, and J. Waldfogel. 2007. "The Employment of Married Mothers in Great Britain, 1974-2000." *Economica*, 74(296): 842–864.

Gregg, P., R. Scutella, and J. Wadsworth. 2010. "Reconciling workless measures at the individual and household level. Theory and evidence from the United States, Britain, Germany, Spain and Australia." *Journal of Population Economics*, 23(1): 139–167.

Gregg, P., and J. Wadsworth. 2001. "Everything You Ever Wanted to Know about Measuring Worklessness and Polarization at the Household Level but Were Afraid to Ask." *Oxford Bulletin of Economics and Statistics*, 63(0): 777–806.

Gregg, P., and J. Wadsworth. 2008. "Two Sides to Every Story: Measuring Polarization and Inequality in the Distribution of Work." *Journal of the Royal Statistical Society. Series A (Statistics in Society)*, 171(4): 857–875.

Gutiérrez-Domènech, M., and B. Bell. 2004. "Female labour force participation in the United Kingdom: evolving characteristics or changing behaviour?" Bank of England working papers 221, Bank of England.

Joshi, H. E., R. Layard, and S. J. Owen. 1985. "Why Are More Women Working in Britain?" *Journal of Labor Economics*, 3(1): S147–S176.

Juhn, C., and K. M. Murphy. 1997. "Wage Inequality and Family Labor Supply." *Journal of Labor Economics*, 15(1, Part 1): 72–97.

Juhn, C., and S. Potter. 2006. "Changes in Labor Force Participation in the United States." *Journal of Economic Perspectives*, 20(3): 27–46.

Layard, R., M. Barton, and A. Zabalza. 1980. "Married Women's Participation and Hours." *Economica*, 47(185): 51–72.

Lee, D., and K. I. Wolpin. 2006. "Intersectoral Labor Mobility and the Growth of the Service Sector." *Econometrica*, 74(1): 1–46.

Lee, J. Y. 2014. "The Plateau in U.S. Women's Labor Force Participation: A Cohort Analysis." *Industrial Relations: A Journal of Economy and Society*, 53(1): 46–71.

Macunovich, D. J. 2009. "Reversals in the Patterns of Women's Labor Supply in the U.S., 1976-2009." IZA Discussion Papers 4512, Institute for the Study of Labor (IZA).

Morgan, S., and H. Carrier. 2014. *The Transition from Education to Work: A Focus on Subject Choices*. Chap. 4, Gender Inequality in the Labour Market in the UK Oxford University Press.

Ngai, L. R., and B. Petrongolo. 2017. "Gender Gaps and the Rise of the Service Economy." *American Economic Journal: Macroeconomics*, 9(4): 1–44.

Office for National Statistics. Social and Vital Statistics Division. 2007a. "General Household Survey, 2005. [data collection]. 2rd Edition. UK Data Service. SN: 5640."

Office for National Statistics. Social and Vital Statistics Division. 2007b. "General Household Survey: Time Series Dataset, 1972-2004. [data collection]. UK Data Service. SN: 5664."

Office for National Statistics. Social and Vital Statistics Division. 2009. "General Household Survey, 2006. [data collection]. 3rd Edition. UK Data Service. SN: 5804."

Olivetti, C. 2006. "Changes in Women's Hours of Market Work: The Role of Returns to Experience." *Review of Economic Dynamics*, 9(4): 557–587.

Olivetti, C. 2014. "The Female Labor Force and Long-Run Development: The American Experience in Comparative Perspective." In *Human Capital in History: The American Record*. National Bureau of Economic Research, Inc, 161–197.

Olivetti, C., and B. Petrongolo. 2014. "Gender gaps across countries and skills: Demand, supply and the industry structure." *Review of Economic Dynamics*, 17(4): 842–859.

Olivetti, C., and B. Petrongolo. 2016. "The Evolution of Gender Gaps in Industrialized Countries." *Annual Review of Economics*, 8(1): 405–434.

Olivetti, C., and B. Petrongolo. 2017. "The Economic Consequences of Family Policies: Lessons from a Century of Legislation in High-Income Countries." *Journal of Economic Perspectives*, 31(1): 205–30.

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.

Smith, J. P., and M. P. Ward. 1985. "Time-Series Growth in the Female Labor Force." *Journal of Labor Economics*, 3(1): S59–S90.

Weinberg, B. A. 2000. "Computer Use and the Demand for Female Workers." *Industrial & Labor Relations Review*, 53(2): 290–308.

On why gender employment equality in Britain has stalled since the early 1990s

Online Appendix

Giovanni Razzu[†] Carl Singleton[‡] Mark Mitchell[§]

May 2018

Appendix A. Further description of the data and sample construction

The General Household Survey (GHS) was a nationally representative survey of individuals living in private households in Great Britain, carried out annually by the Office for National Statistics (ONS) between 1972 and 2006. Before 1988 and after 2004 the GHS was conducted over calendar years, whereas from 1988 to 2004 the survey reported on a financial-year basis, running annually from April to March. The survey was designed to collect consistent and nationally representative information on social, demographic and economic characteristics of households over time. As such, the survey questionnaire and variables contained in the GHS were subject to minimal year-to-year changes. To obtain a representative sample of Great Britain, the survey used a probability stratified two-stage sample design, with households stratified first by geographical location and then by economic indicators, such as the proportion of households with no car. The survey run was only broken in 1997/98 and 1999/2000 for review. The former break led to changes thereafter in the survey design, which did not significantly impact any of the key individual nor household variables we use here. The latter break was to review the sampling procedure to continue to maintain representativeness. From 2006 the GHS content was redesigned, as per EU requirements, to contain comparable data to the Survey of Income and Living Conditions (EU-SILC) and became the General Lifestyle Survey (GLS).

We use the GHS Time Series data set, which contains a set of key variables from each annual wave of data between 1972 and 2004. It was originally referred to as the GHS Pseudo-Cohort Dataset. The time series dataset was created by the ONS to analyse changes in social and economic inequalities over those three decades or so. To correct for small discrepancies between years, all variables were adjusted to match their most recent version. For example, all 'Marital Status' variables prior to 1992 were adjusted to contain a category for 'Cohabiting', due to its inclusion as a possible response from that year onwards. An indicative reference for the GHS time series dataset is as follows: Office for National Statistics. Social and Vital Statistics Division. (2007). *General Household Survey: Time Series Dataset, 1972-2004*. [data collection]. UK Data Service. SN: 5664. We favour the GHS Time Series dataset over other annual nationally representative surveys, such as the UK Annual Population and Labour Force Surveys, due to its detail at the household level and its consistent design over a long period of time. We extend the time series dataset by adding some variables from the annual cross-sectional datasets for 1972-2004, such as the UK Standard Industrial

[†]University of Reading, g.razzu@reading.ac.uk
[‡]University of Edinburgh, carl.singleton@ed.ac.uk
[§]University of Edinburgh, mark.mitchell@ed.ac.uk

Classification (SIC) Job Industry Division of individuals' main job, and by also adding observations and years from the 2005 and 2006 datasets.

To create a consistent and practical analysis dataset, we drop all observations with missing values for any of the variables later used to define a worker's type, employment status or industry of employer. Partners are matched by identifying compatible pairs of individuals who share the same household and family identifiers, and report to be similarly married, cohabiting or in a same-sex couple. In the tiny number of cases where partners could not be definitively traced within the datasets, these households and individuals were dropped. 'Partner characteristics' are then defined simply as the characteristics of the person with whom each of these individuals was matched. The GHS classifies a family as *"A married or opposite sex cohabiting couple on their own; or a married or opposite sex cohabiting couple, or a lone parent, and their never-married children (who may be adult), provided these children have not children of their own"*. The analysis sample therefore represents an approximate random and representative sample of people aged 25-54 who were living in Britain each year between 1972 and 2006. Table A1 documents a decline in the analysis sample size over the period both for men and women. However, this is not due to any restrictions we place on the sample or increasing survey non-response, but rather a combination of financial constraints faced by the ONS and the use of different sampling procedures over time. In all the analysis described in the main text, we pool the sample for each three-year period to improve accuracy, such that estimates of the employment rate for 1973 are based on the 1972-1974 pooled samples of the GHS. As shown in Figure A1, the trends in both British male and female employment rates aged 25-54, and the gap between them, in the resulting sample of the GHS, match those used by the ONS as UK National Statistics for the working-age population over the same period, derived from the Labour Force Survey. Not only the trends, but the cyclical patterns in the employment rates estimated from the two separate sources are also comparable.

Every man and woman in our analysis sample is assigned to one of 180 mutually exclusive worker types or groups, based on the complete interaction of a set of observable characteristics. For example, married men/women, who have a degree, two non-infant children, a wife/husband who is employed and who also has A-levels as their highest qualification, are assigned to a specific worker type for that exact combination of characteristics. Table A2 below contains notes on the variables from both the GHS Time Series (TS) and annual cross-sectional datasets which were used to define these individual worker types.

TABLE A1: Sample number of observations

	Male	Female	Total
1972	5,835	6,017	11,852
1973	5,659	5,933	11,592
1974	4,900	5,101	10,001
1975	5,488	5,694	11,182
1976	5,369	5,626	10,995
1977	5,534	5,664	11,198
1978	5,387	5,521	10,908
1979	5,173	5,317	10,490
1980	5,239	5,311	10,550
1981	5,456	5,579	11,035
1982	4,517	4,696	9,213
1983	4,274	4,467	8,741
1984	4,195	4,317	8,512
1985	4,359	4,472	8,831
1986	4,515	4,705	9,220
1987	4,580	4,748	9,328
1988	4,391	4,601	8,992
1989	4,902	5,080	9,982
1990	4,588	4,754	9,342
1991	4,785	5,013	9,798
1992	4,769	5,113	9,882
1993	4,665	4,947	9,612
1994	4,026	4,376	8,402
1995	3,966	4,384	8,350
1996	3,691	4,011	7,702
1998	3,384	3,720	7,104
2000	3,373	3,726	7,099
2001	3,691	4,046	7,737
2002	3,432	3,850	7,282
2003	4,023	4,413	8,436
2004	3,355	3,776	7,131
2005	4,576	5,080	9,656
2006	3,275	3,767	7,042

Notes.- author calculations using General Household Survey, 1972-2006.

TABLE A2: Notes on time series variables used and derived from the General Household Survey, 1972-2006

	GHS variables	*Notes*	*Derived Categories*
Age	TS: *page*[1] 2005/6: *age*	Individuals were assigned to five-year age groups.	25-29, 30-34, 35-39, 40-44, 45-49, 50-54.
Highest Qual.	TS: *pedfull* 2005/6: *edlev00*	The 19 and 13 possible levels of education were respectively reduced to 4 major UK levels, with these partly determined by what is most comparable over time in the UK.	None, Any less than A-level, A-level or professional (High school/ secondary), Degree (College/ tertiary).
Dependent Chld.	TS: *pn0to4*, *pn5to15* 2005/6: *n0to4*, *n5to15*	Children defined as being aged less than 16 years. Dependency determined by household relationship variables. Variables give the number of children an individual has aged 0-4, and 5-15 respectively.	None, 1-2, 3+
Infant	TS: *pn0to4* 2005/6: *n0to4*	Infants are defined as being aged 0-4 years. Indicator created of whether or not a household contained an infant.	Infant in household, or not.
Employment	TS: *pdvilo3a* 2005/6 *dvilo3a*	International Labour Organization (ILO) definition of employment status. 2 categories generated from the original 3: In employment, unemployed, economically inactive.	In employment, or not.
Marital Status	TS: *pdvmardf* 2005/6: *dvmardf*	2 categories were derived from the original 7. Before 1992 there were only five categories: those for 'Cohabiting' and 'Same Sex Couple' were added in 1992 and 1996 respectively.	Married/ Cohabiting/ Same Sex couple; Single/ Widow/ Divorced/ Separated
Job Industry	1972-80: *indust*	Standard Industrial Classification (SIC) UK 1968 orders. The 24 orders were matched to their respective SIC-UK 1980 industry divisions using ONS correlations.[2] These 10 divisions were then reduced to 8 by grouping the 3 smallest.	Primary Orders (1, 2, 4, 18)[3], Engineering (5-9), Other Manufacturing (3,10-16), Construction (17), Distribution (20), Transport & Communications (19), Banking & Finance (21), Other Services (22-24).
	1981: *sicclass* 1982: *sic* 1983-93: *sicr* 1994: *indmain*	SIC-UK 80 industry divisions. These 10 divisions were reduced to 8 by grouping the 3 smallest.	Primary Divisions (0,1,2), Engineering (3), Other Manufacturing (4), Construction (5), Distribution (6), Transport & Communications (7), Banking & Finance (8), Other Services (9).
	1995-97: *indmain* 1998-2006: *indstry1*	SIC-UK 92 industry section sub-classes In 1995/66. These were matched to the suitable SIC-UK 80 industry division using ONS correlations. From 1998-2006 the GHS carried out this conversion. In both cases the 10 industry divisions were reduced to 8 by grouping the smallest 3.	As above.
Hours of Work	1972-1998: *workhrs*	Number of hours usually worked per week excluding overtime. Post 1998 it was not possible to separately identify usual and overtime hours worked.	

[1] Variables preceded by TS are from GHS Time Series dataset, and those by a year from the corresponding cross section.
[2] For the appropriate documentation, see UK Standard Industrial Classification of Economic Activities Archive
[3] Text is the derived category and numbers in parentheses are the SIC-UK orders/divisions that were grouped.

FIGURE A1: Comparison of GHS sample (GB) employment rates and gender gap with ONS National Statistics (UK), 1971-2016

(a) Male & female

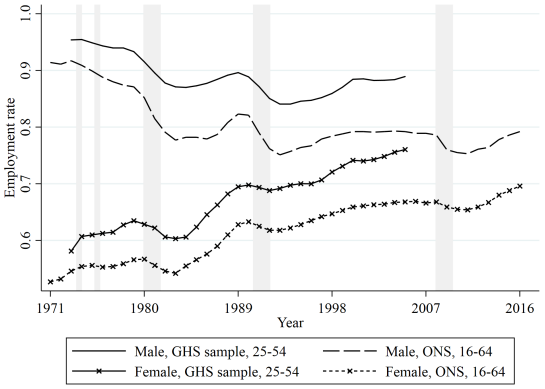

(b) Gender gap

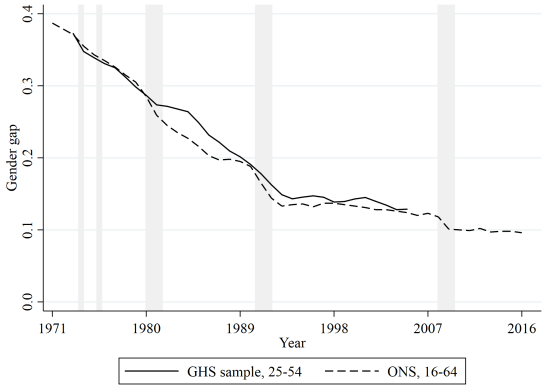

Notes.- author calculations using General Household Survey, 1972-2006 and ONS Labour Market Statistics (accessed 31/05/2017). Gap is expressed as male minus female. Shaded areas denote official periods of UK recessions.

FIGURE A2: The distribution of work across households, 1973-2005

(a) Single-adult hholds

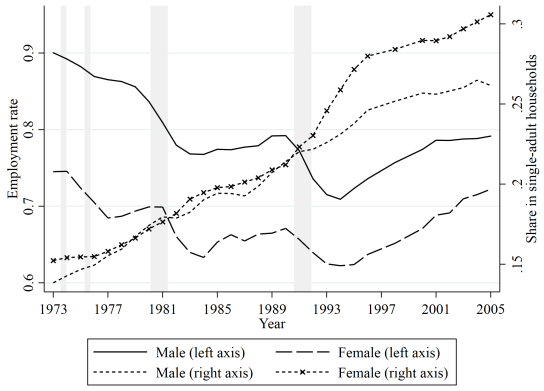

(b) Partner hholds - Male

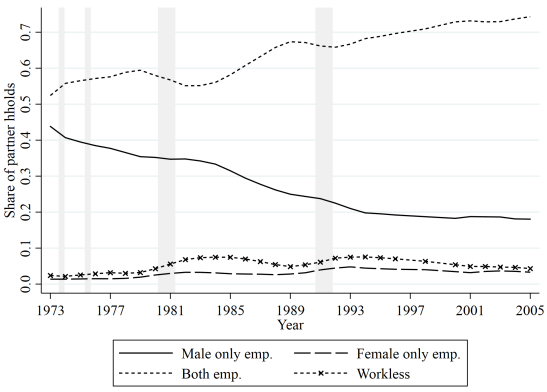

(c) Partner hholds - Female

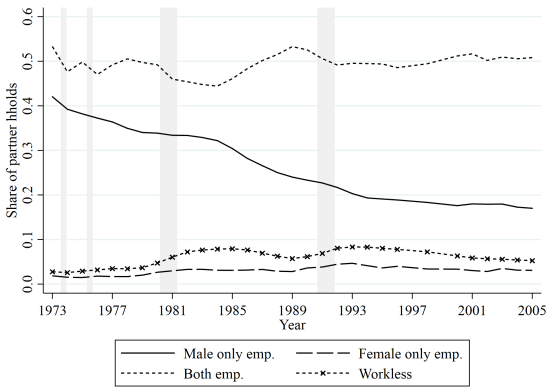

Notes.- author calculations using General Household Survey, 1972-2006. Shaded areas denote official periods of UK recessions.

Appendix B. Additional figures and tables

FIGURE B1: Total employment rates and the gender gap implied by the changes (*only*) in the prevalence of single adult or couple households since 1973/89.

(a) Male

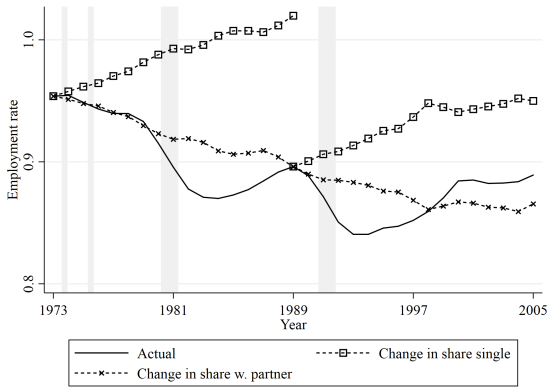

(b) Female

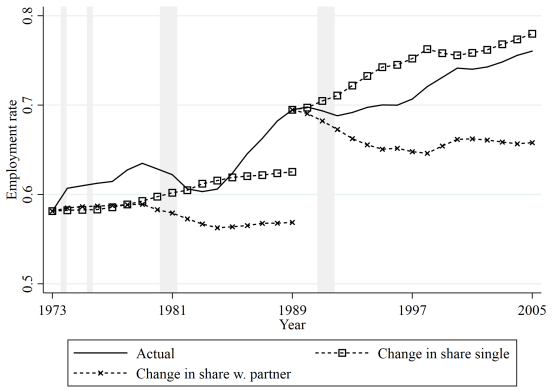

(c) Gap

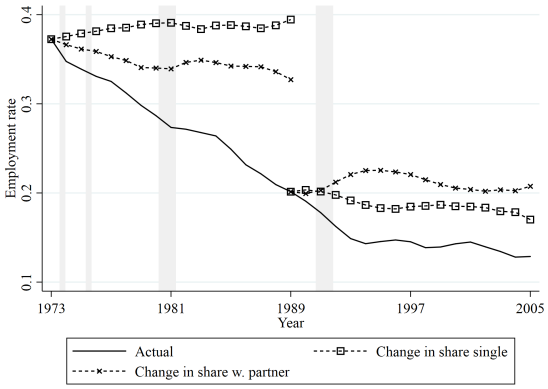

Notes.- author calculations using General Household Survey, 1972-2006. Gap is expressed as male minus female. Each series gives the total employment rate implied by only the cumulative change, relative to 1973/89, in the share of single adult or couple households. Shaded areas denote official periods of UK recessions.

FIGURE B2: Total employment rates and the gender gap implied by the changes in either population composition or the employment rates within groups, by HIGHEST QUALIFICATION, since 1973/89.

(a) Male

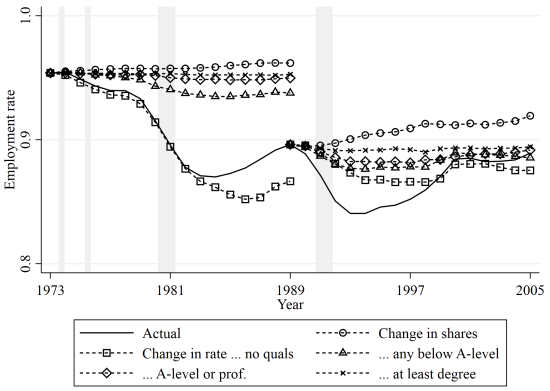

(b) Female

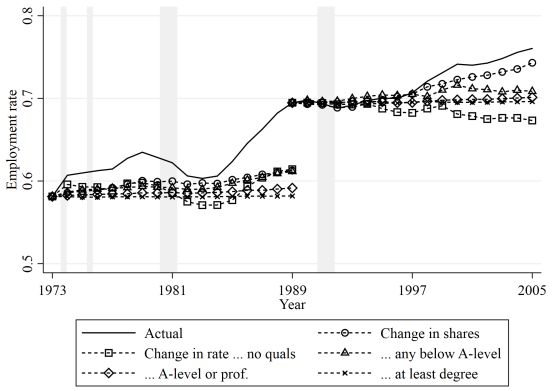

(c) Gap

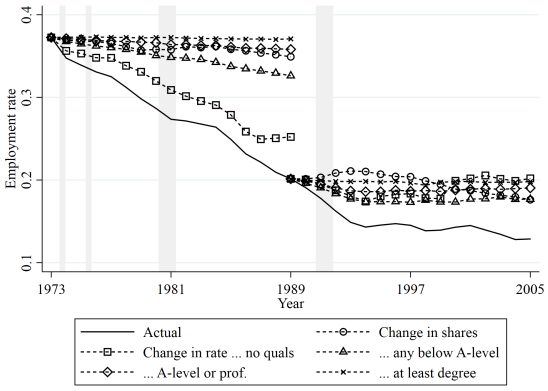

Notes.- author calculations using General Household Survey, 1972-2006. Gap is expressed as male minus female. Each series gives the total employment rate implied by only the cumulative change, relative to 1973/89, in the employment rate (or gap) from only observed changes in type shares or the employment rates for some types.

FIGURE B3: Total employment rates and the gender gap implied by the changes in either population composition or the employment rates within groups, by NUMBER & AGE of DEPENDENT CHILDREN, since 1973/89.

(a) Male

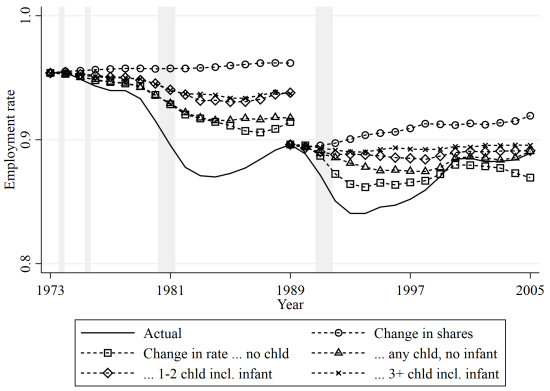

(b) Female

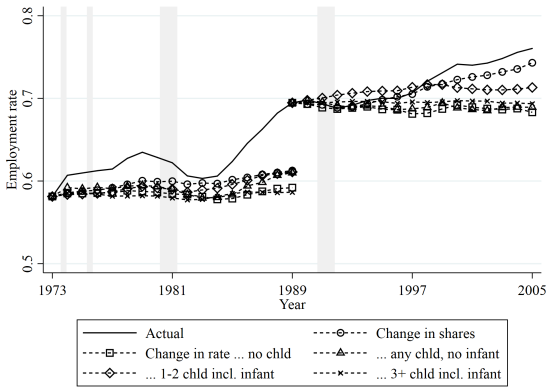

(c) Gap

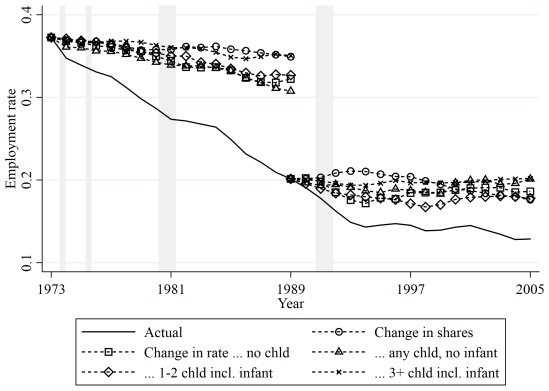

Notes.- see Figure B2, except analysis is a shift-share decomposition over fertility types

FIGURE B4: Total employment rates and the gender gap implied by the changes in either population composition or the employment rates within groups, by PARTNER'S EMPLOYMENT STATUS, since 1973/89.

(a) Male

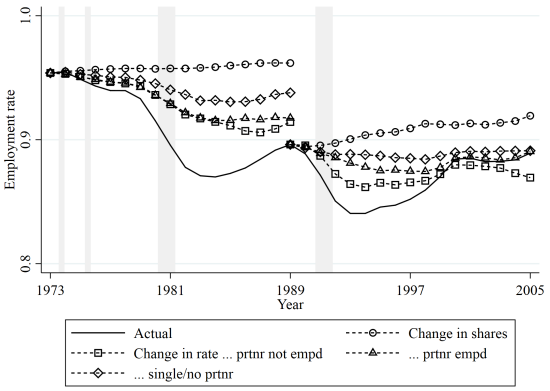

(b) Female

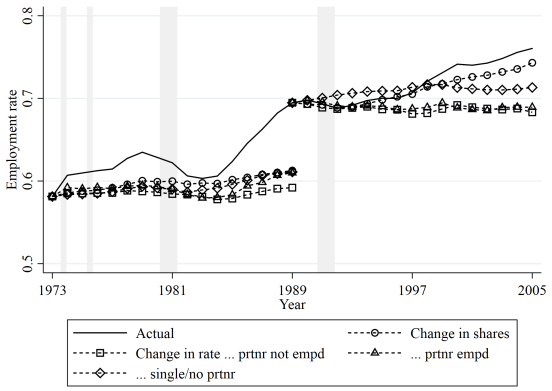

(c) Gap

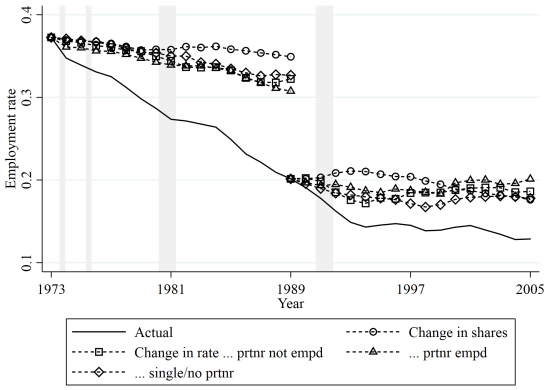

Notes.- see Figure B2, except analysis is a shift-share decomposition over partner employment status.

FIGURE B5: Blinder-Oaxaca type decompositions of the change in the male and female employment rates, 1973-81, 1981-89, 1989-97 and 1997-2005

(a) Male

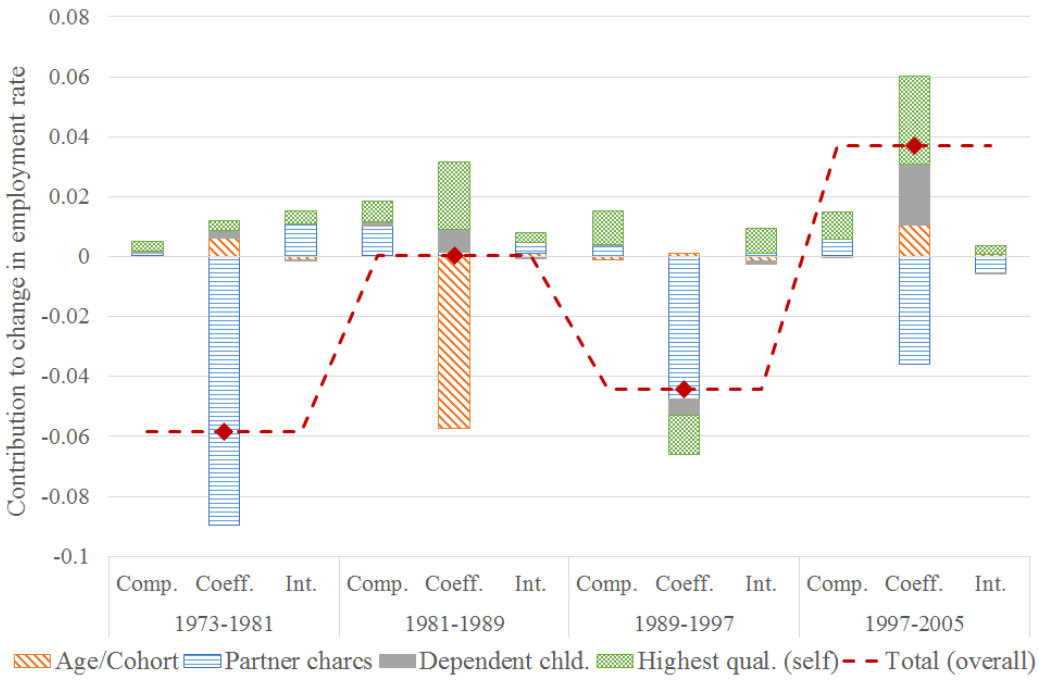

(b) Female

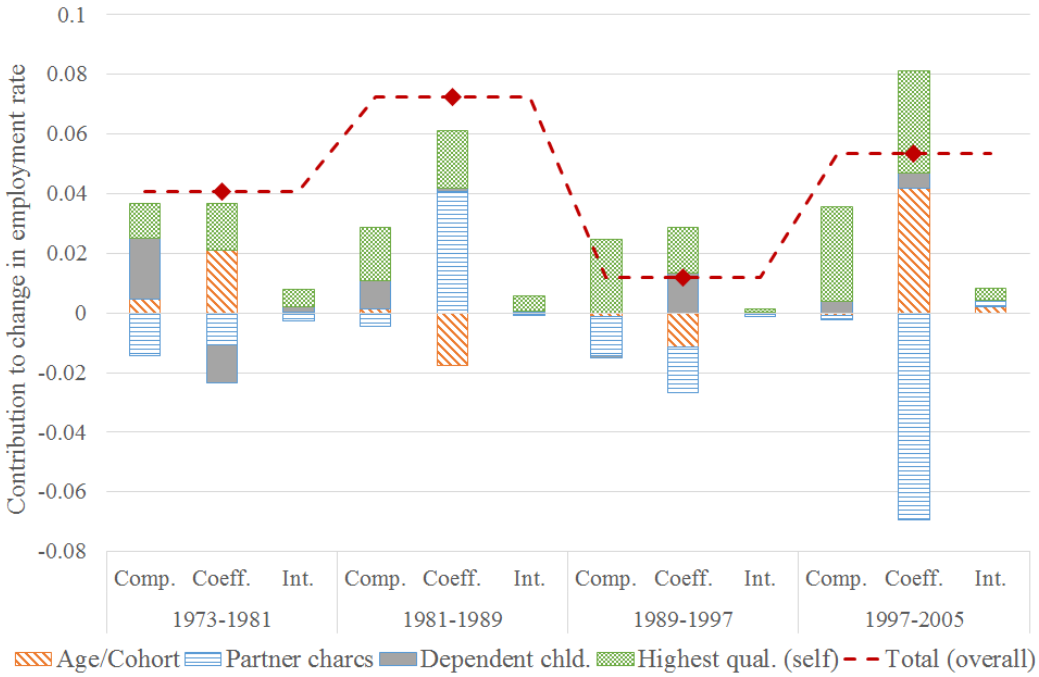

Notes.- see Figure 4.

TABLE B1: Blinder-Oaxaca type decompositions of changes in male employment rates, 1973-81, 1981-89, 1989-97 and 1997-2005 (ppts)

	Time period			
	1973-81	1981-89	1989-97	1997-2005
Composition:				
Age/Cohort	0.000	0.000	-0.001	0.000
	(0.000)	*(0.001)*	*(0.001)*	*(0.001)*
Partner charcs	0.000	0.010	0.004	0.006
	(0.001)	*(0.002)*	*(0.002)*	*(0.002)*
Dependent chld.	0.001	0.001	0.000	0.000
	(0.001)	*(0.001)*	*(0.000)*	*(0.001)*
Highest qual. (self)	0.003	0.007	0.011	0.009
	(0.001)	*(0.001)*	*(0.003)*	*(0.003)*
Sub-total	0.005	0.019	0.014	0.015
	(0.001)	*(0.003)*	*(0.003)*	*(0.003)*
Coefficients:				
Age/Cohort	0.006	-0.057	0.001	0.011
	(0.007)	*(0.021)*	*(0.008)*	*(0.008)*
Partner charcs	-0.089	0.002	-0.048	-0.036
	(0.012)	*(0.021)*	*(0.011)*	*(0.014)*
Dependent chld.	0.003	0.007	-0.005	0.020
	(0.006)	*(0.012)*	*(0.006)*	*(0.007)*
Highest qual. (self)	0.003	0.022	-0.013	0.029
	(0.003)	*(0.013)*	*(0.005)*	*(0.007)*
Sub-total	-0.077	-0.026	-0.065	0.024
	(0.006)	*(0.006)*	*(0.006)*	*(0.005)*
Interactions:				
Sub-total	0.014	0.007	0.007	-0.002
	(0.003)	(0.002)	(0.003)	(0.002)
Total change:	-0.058	0.001	-0.044	0.037
	(0.006)	(0.005)	(0.006)	(0.005)
N. obs. initial year	16,349			
N. obs. end year	15,212	13,881	7,075	11,206

Notes.- author calculations using General Household Survey, 1972-2006, ages 25-54 only. The contribution from individual covariates are grouped as follows – 'Age/cohort': dummy variables for age groups; 'Partner chars': marital status interacted with dummy variables for a partner's employment status and their highest qualification; 'Dependent chld.': interaction of dummy variables for number of children and whether or not there is an infant in the household; 'Highest qual. (self)': individual's highest qualification level. Standard errors in parentheses are estimated using clustering on the 180 individual worker group types defined in the main text.

TABLE B2: Blinder-Oaxaca type decomposition of changes in female employment rates, 1973-81, 1981-89, 1989-97 and 1997-2005 (ppts)

	Time period			
	1973-81	1981-89	1989-97	1997-2005
Composition:				
Age/Cohort	0.005	0.001	-0.001	-0.001
	(0.003)	*(0.002)*	*(0.002)*	*(0.003)*
Partner charcs	-0.014	-0.005	-0.013	-0.001
	(0.002)	*(0.001)*	*(0.003)*	*(0.002)*
Dependent chld.	0.020	0.010	0.000	0.004
	(0.012)	*(0.008)*	*(0.006)*	*(0.007)*
Highest qual. (self)	0.012	0.018	0.025	0.032
	(0.003)	*(0.004)*	*(0.006)*	*(0.005)*
Sub-total	0.022	0.024	0.010	0.034
	(0.011)	*(0.010)*	*(0.007)*	*(0.008)*
Coefficients:				
Age/Cohort	0.021	-0.018	-0.012	0.042
	(0.017)	*(0.011)*	*(0.022)*	*(0.017)*
Partner charcs	-0.011	0.041	-0.015	-0.069
	(0.024)	*(0.015)*	*(0.038)*	*(0.024)*
Dependent chld.	-0.013	0.001	0.013	0.005
	(0.008)	*(0.006)*	*(0.026)*	*(0.007)*
Highest qual. (self)	0.016	0.019	0.015	0.034
	(0.005)	*(0.006)*	*(0.020)*	*(0.011)*
Sub-total	0.013	0.044	0.002	0.012
	(0.007)	*(0.009)*	*(0.010)*	*(0.008)*
Interactions:				
Sub-total	0.005	0.005	0.000	0.008
	(0.003)	(0.003)	(0.004)	(0.003)
Total change:	0.041	0.073	0.012	0.054
	(0.013)	(0.013)	(0.011)	(0.009)
N. obs. initial year	17,051			
N. obs. end year	15,586	14,435	7,731	12,623

Notes.- see Table B1

TABLE B3: Contributions from within-sector gender share changes to the overall change in the female share of total employment, 1982-93 and 1994-2005 (ppts)

	82-93	94-05
Primary divisions, 0-2	0.69	-0.36
Engineering, 3	0.40	-0.01
Other manufacturing, 4	-0.14	-0.64
Construction, 5	0.31	0.14
Distribution, 6	0.15	-0.47
Transport & comms, 7	0.18	0.18
Bank & finance, 8	0.31	-0.28
Other services, 9	0.71	0.71
Total	2.6	-0.7

Notes.- author calculations using General Household Survey, 1981-2006, ages 25-54 only. Industry divisions listed are defined by SIC-UK 1980.