



Costa-Dias, M. S. R. D., Parodi, F., & Joyce, R. (2021). The gender pay gap in the UK: Children and experience in work. *Oxford Review of Economic Policy*. <https://doi.org/10.1093/oxrep/graa053>

Publisher's PDF, also known as Version of record

License (if available):
CC BY

Link to published version (if available):
[10.1093/oxrep/graa053](https://doi.org/10.1093/oxrep/graa053)

[Link to publication record in Explore Bristol Research](#)
PDF-document

This is the final published version of the article (version of record). It first appeared online via Oxford Academic at <https://doi.org/10.1093/oxrep/graa053>. Please refer to any applicable terms of use of the publisher.

University of Bristol - Explore Bristol Research

General rights

This document is made available in accordance with publisher policies. Please cite only the published version using the reference above. Full terms of use are available: <http://www.bristol.ac.uk/red/research-policy/pure/user-guides/ebr-terms/>

The gender pay gap in the UK: children and experience in work

Monica Costa Dias,^{*} Robert Joyce,^{**} and Francesca Parodi^{***}

Abstract: In this study, we document the evolution of the gender pay gap in the UK over the past three decades and its association with fertility, examining the role of various differences in career patterns between men and women and how they change with the arrival of the first child. We show that differences in accumulated years of labour market experience play an important role, while differences in industry, occupation, and job characteristics explain less, conditional on working experience. We develop an empirical wage model to estimate the causal effect of working experience on the wages of women. Estimates from this model are then used to simulate two counterfactual scenarios in which women who are employed always work full-time, or women's rates of both part-time and full-time work are the same as men's. We find that differences in working experience explain up to two-thirds of the gender pay gap of college graduates 20 years after the first childbirth, and that the gap is largely driven by differences in full-time experience. The role of working experience is more moderate for individuals with no college education, but it can still account for about one-third of the overall long-term gender wage gap.

Keywords: labour supply, wage dynamics, gender pay gap, human capital

JEL classification: D10, J16, J18, J24

I. Introduction

Gender differences in earnings are essentially universal across countries. Within the developed world those gaps have tended to fall greatly over the last century, although progress has stalled in recent decades and the gaps remain sizeable (Goldin, 2014; Blau and Kahn, 2017). The opening of the pay gap happens gradually over the course of life and is strongly related with the arrival of children and its strong negative effects on the

^{*}University of Bristol, Institute for Fiscal Studies, and University of Porto; e-mail: monica.costa-dias@bristol.ac.uk

^{**}Institute for Fiscal Studies; e-mail: robert_j@ifs.org.uk

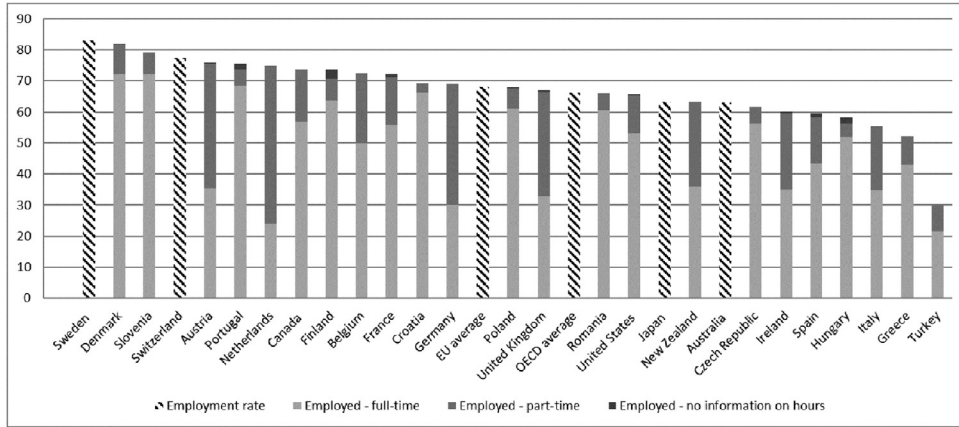
^{***}Collegio Carlo Alberto, University of Turin, Institute for Fiscal Studies; e-mail: francesca.parodi@carloalberto.org

We are thankful to Richard Blundell, Alex Bryson, Stefania Innocenti, Heather Joshi, Almudena Sevilla, and one anonymous referee for the many insightful comments and suggestions. All errors remain our own. We gratefully acknowledge financial support from the Joseph Rowntree Foundation, the Economic and Social Research Council (ESRC Centre for the Microeconomic Analysis of Public Policy at IFS—grant number ES/M010147/1—and grant numbers ES/K00624X/1 and ES/N015304/1), and the European Union Horizon 2020 programme on Dynamics of Inequality Across the Life-Course (Dial ES00142).

doi:10.1093/oxrep/graa053

© The Author(s) 2021. Published by Oxford University Press.

This is an Open Access article distributed under the terms of the Creative Commons Attribution License (<http://creativecommons.org/licenses/by/4.0/>), which permits unrestricted reuse, distribution, and reproduction in any medium, provided the original work is properly cited.

Figure 1: Maternal employment rates across OECD countries

Note: Data from OECD Family Database. Shows PT–FT (part time–full time) employment rates (%) of women aged 15–64 with at least one child aged 0–14 in 2014 or latest available year.

hourly wage rates, employment rates, and hours of paid work of women (Paull, 2008; Adda *et al.*, 2017; Kleven *et al.*, 2019a).

Countries differ significantly in the relative importance of these determinants of the gender pay gap. Figure 1 shows the wide cross-country variation in the employment and working hours of mothers. For instance, Scandinavian and northern European countries have the highest maternal employment rates, almost 30 percentage points higher than the employment rates in some of the southern European countries, such as Greece and Italy, and over 50 percentage points higher than employment rates in Turkey. The employment rates of mothers in English-speaking countries tend to fall in the middle of the table, close to the average for OECD and EU countries. In some countries, including Germany, the Netherlands, Austria, and the UK, a large fraction of working mothers do short working hours, while in other places, including the Scandinavian countries, most mothers in work do full-time hours.

Recent research (Kleven *et al.*, 2019a) also documented that in countries where mothers stay longer out of work after childbirth, such as the US and UK as compared to the Scandinavian countries, the drop in women’s earnings relative to men’s persists for longer after childbirth. This is partly the mechanical effect of mothers continuing to work less than fathers do many years after their first child is born, and partly a consequence of the price of labour as measured by the wage rate per hour being differentially affected by parenthood for mothers and fathers.

In this paper we investigate some of the drivers of the gender wage gap among parents. We document that the gap in hourly wages does not open immediately after childbirth, contrary to what happens to earnings, but instead widens gradually as the child grows up. We then consider two mechanisms that can partly account for that gradual but permanent opening of the gender wage gap. First, taking time away from paid work and working short hours—two prevalent features of maternal, but not paternal, labour supply—may inhibit the wage growth of women through skill depreciation, slow accumulation of skills, or lack of career progression, hence denting their future earnings capacity persistently. Second, mothers and fathers sort into different types of jobs—in

different industries, occupations, and offering a different set of amenities such as flexibility in working hours—perhaps because of how difficult or easy it is to combine such jobs with the responsibilities of parenthood. If mothers cluster into jobs that can easily be coordinated with their childcare responsibilities, even if at the cost of higher pay, while fathers assume the main bread-winner role and sort into high-paying jobs, the resulting diverging careers could be reflected in diverging wage profiles. These two mechanisms are likely to lead to a gradual accumulation of persistent losses in the wage rates of mothers as compared to fathers and contribute to explaining the widening of the gender wage gap after childbirth. They are also likely related. Indeed, mothers who take time away from paid work may see their prospects of progressing to higher paying jobs curtailed, at least partly as they fail to accumulate the working experience they need to progress.

Using longitudinal data for the UK spanning the 27 years that start in 1991, we assess the quantitative importance of these mechanisms. We find that ‘experience capital’ gained in work has the potential to account for a large share of the gender wage gap, including the way that it evolves over the life cycle. In contrast, conditionally on accumulated working experience, the different job characteristics (including industry and occupation) in which fathers and mothers concentrate do not further explain much of the opening of the gap after childbirth. This is not to say that sorting into jobs with different characteristics is not an important driver of the gender gap in wages. Instead, our result reflects that gender differences in job sorting happen together with the expansion in the gender experience gap, as mothers taking time away from paid work fail to progress in their careers and may even move down the job ladder.

Building on this evidence, we set up and estimate an empirical model of wage dynamics that allows us to assess the causal impact of accumulated work experience capital on the wage progression of women over their life-cycle. Our model formalizes, in a flexible way, the intertemporal connection between hourly wages and labour supply at the extensive and intensive margins. In particular, it accounts for the possibility that working part-time hours impacts hourly wages not only today, but also in future years by having implications for wage progression that are different from those of full-time work. It also allows for depreciation in experience capital during periods of no work. Given the importance of education in driving labour market outcomes, we fully interact our model with education attainment. We consider three education levels, corresponding to less than high-school qualifications, high-school qualifications, and university degree. To deal with endogenous participation and hours selection, we adopt a control function approach exploiting the many policy reforms that changed the incentives to work of mothers over the period covered by the data. For exclusion restrictions we use simulated family disposable income variables that are constructed for different working hours of the woman (Blundell *et al.*, 1998).

We find that working full-time hours is a key determinant of wage progression, particularly for women earlier in the career. Moreover, there is a clear positive education gradient on the impact of working full-time hours on wage growth. This means that educated women have more to gain from working full-time hours but also that they are the ones who have most to lose from not doing so. In contrast, we find no statistically significant impact of working part-time hours.

Our model can be used to simulate the role of experience losses in explaining the opening of the gender wage gap after the birth of the first child. We do so in two counterfactual simulations: the first simulation imposes that employed women always work full-time, and the second simulation further imposes that women are paid for their work at the same rate of men. For university graduates we find that losses in working experience after the birth of the first child can explain up to two-thirds of the opening of the gap in wages when the child reaches 20 years of age. This is not merely because mothers spend more time out of the labour force after childbirth than fathers; it is also because, when in paid work, they are likely to work shorter hours, and part-time work shuts down progression in hourly wages. The role of working experience is more moderate for individuals with no university education, but it can still account for about one-third of the overall gender wage gap 20 years after childbirth.

Some of the channels we investigate have been studied in the large empirical literature on gender gaps. For instance, Bertrand *et al.* (2010) identify career interruptions and working hours as key drivers of the pay gap among MBA graduates; Goldin, (2014) discusses the high penalty from career breaks and flexibility in some high-wage occupations as drivers of the persisting gender wage gap; Blau and Kahn (2017) find that the three most important factors explaining the current gender wage gap in the US are occupation, industry, and experience; Kleven *et al.* (2019b) point to occupation, sector, and firm choices as mechanisms driving child penalties in Denmark. Our study adds to this literature by considering the causal impact of current working choices on wage growth and how these effects cumulate over time to explain the gradual opening of the gender wage gap after childbirth. We also document that other mechanisms, such as occupation, industry, or specific job characteristics, are comparatively less important once cumulative differences in working experience are accounted for.

One other set of papers uses structural models of the life-cycle to study female labour supply and wages. Among others, Olivetti (2006) builds a life-cycle model with human capital accumulation and home production to investigate the change in women's supply of hours of market work; Attanasio *et al.* (2008) construct a life-cycle model of female participation decisions and savings with a process for human capital accumulation that accounts for periods out of work; Blundell *et al.* (2016) estimate a dynamic model of employment, human capital accumulation, and savings for women and use it to analyse the effects of welfare policy in the UK; Adda *et al.* (2017) model fertility and labour supply decisions of women over the life-cycle to quantify the long-run impact of fertility incentivizing policies. We depart from these papers by adopting a more flexible empirical model that is nevertheless well suited to quantify the role of present and past experiences in work in explaining female wage dynamics. We then use our framework to study the extent to which our mechanisms explain gender differences in pay among parents.¹

The rest of the paper is organized as follows. Section II describes the data and defines the main variables. Section III shows descriptive evidence on the evolution of the gender pay gap and its relation to childbirth and to the different working lives of mothers and fathers. It then shows that opening differences in experience capital are

¹ There is a large set of papers looking into other potential drivers of gender gaps that are further away from the focus of this study, including discrimination and pre-market factors such as education (Altonji and Blank, 1999) or sorting and bargaining (Card *et al.*, 2015).

strongly related with the growing gender pay gap after childbirth. Section IV outlines the econometric model and estimation technique we use to identify the causal impact of working experience on the wages of women. Section V reports the empirical estimates and section VI discusses the model predictions by simulating counterfactual wage gaps if women were to work at the same rate as men do. Section VII suggests policy implications. Finally, section VIII draws some concluding remarks.

II. Data

We use two data sources for our empirical analysis. All analysis requiring longitudinal data is based on the UK Household Longitudinal Study (UKHLS). This is a combination of two panel studies of families with the same structure and overlapping samples, the British Household Panel Survey (BHPS) and the Understanding Society (USoc).² The UKHLS has been following the lives of families and their offshoots since 1991. The survey started with a representative sample of 5,050 households living in Great Britain; it was later replenished in 1997 and 2001 with 1,000 households from the former European Community Household Panel, in 1999 with two samples of 1,500 households each from the Welsh and Scottish extensions, and in 2001 with the first set of households from Northern Ireland, totalling 1,900. It has undergone a large expansion and restructuring in 2009 that was marked with the adoption of its new label, USoc. Some 40,000 new families were added from across the UK. Except for attrition—which was particularly large in the transition between the two studies in 2009, affecting just over 30 per cent of the existing sample—all household members in the original samples remain in the sample until the end of the period, which for this paper is 2017. Other individuals have also been added to the sample, as they formed families with original members of the panel or were born into them. In our analysis, we exclude families living in Northern Ireland as the region is only sampled from 2001 onwards.

All members of the household aged 16 and above are interviewed yearly, and a large set of demographic, education, and labour-market information is recorded. The UKHLS also collects historical information documenting the history of full-time and part-time employment and the socio-economic environment of the parental household of each interviewee.

We use the sub-sample of individuals observed during their main working years, between the ages of 20 and 55, after they left full-time education. In each annual wave, the UKHLS collects information on employment status, usual weekly working hours, paid and unpaid overtime, and gross pay including for overtime. We discretize the distribution of working hours into three bins that we label as no-work (4 hours per week or fewer), part-time work (5–24 weekly hours), and full-time work (25 and more weekly hours). Those who report being disabled or long-term sick are dropped from the sample.

Table 1 shows the sample sizes in UKHLS for our population of interest. Overall, our sample includes almost 150,000 observations of over 27,000 individuals. Among these, over 7,000 were first interviewed in the BHPS period, prior to 2009, of which just

² University of Essex. Institute for Social and Economic Research, NatCen Social Research, Kantar Public. (2017). Understanding Society: Waves 1–7, 2009–16 and Harmonized BHPS: Waves 1–18, 1991–2009. [data collection]. 9th Edition. UK Data Service. SN: 6614.

Table 1: UKHLS—sample sizes and distribution of education

	Men	Women	All
Sample size: number of individuals			
All	11,899	15,144	27,043
In BHPS	3,112	3,985	7,097
In BHPS and USoc	630	778	1,408
Sample size: number of observations			
All	63,159	84,062	147,221
In BHPS	23,841	31,732	55,573
In BHPS and USoc	8,825	11,069	19,894
Median duration of observation spells (years)			
All	5	6	6
In BHPS	7	8	8
In BHPS and USoc	18	19	18
Distribution of education			
GCSEs	0.460	0.474	0.468
A-levels	0.404	0.417	0.412
University degree	0.136	0.109	0.121

Notes: The row labels 'all', 'in BHPS', and 'in BHPS and USoc' stand for, respectively, the entire UKHLS sample used in this study, the subset of individuals who enter the sample prior to 2009, and the subset of individuals who enter the sample prior to 2009 and remain until after 2009 or later.

over 1,400 are present in both periods. The median length of the observation period is 6 years, though it is longer for those who enter the sample earlier during the BHPS period, particularly if followed into the USoc period.

We consider three education groups: GCSEs, representing those who leave education at 16 without completing high-school education; A-levels, representing those with a high-school diploma or equivalent; and degree, representing those who graduate from university (3-year degree). The distribution of education in the population is shown at the bottom of [Table 1](#). The largest group is that with lower education attainment, and this is true for both men and women. Men in our age group and time window are more likely to have a degree than women, but, as we show, this does not hold in the more recent periods. Only about 12 per cent of our sample has a degree.

Wages are measured in hourly rates built as the ratio of the total gross weekly pay by total hours, both measures including paid and unpaid overtime. We remove aggregate wage growth from the wage rate and trim it at percentiles 2 and 99 from below and above, respectively, to limit the impact of measurement error in wages and hours. All results in monetary quantities are in 2016 prices.

The historical labour supply information is collected in waves 1992, 2001, 2002, 2009, and 2013. It exists for over 70 per cent of our sample. We use the history of employment and working hours to construct, for these years, two experience variables that measure accumulated experience time in part-time and full-time jobs since the beginning of the working life. We then complete the experience variables over the entire observation window using year-on-year information on employment spells and hours.

We complement the UKHLS with data from the British Labour Force Survey (LFS). It is used by the Office for National Statistics to produce official quarterly labour market statistics and has the advantage of being a much larger sample than the UKHLS. We use it for two purposes. First, to show some descriptive evidence of labour market

trends. The LFS has a larger sample than the UKHLS, which enables us to characterize trends more accurately when the longitudinal feature of the data is not required. We find similar but more volatile patterns when reproducing the figures on UKHLS data. And second, to construct two measures of ‘women/family friendliness’ of jobs at the industry or occupation levels, that we then merge with the UKHLS data. Specifically, we merge in information on: the percentage of female employees by 3-digit occupation from the SOC classification; the percentage of employees doing part-time hours by occupation at the same level of disaggregation; and the percentage of female managers by 2-digit industry from the SIC classification.

III. Descriptive analysis

(i) Differences in wages between men and women over time

We start by looking at the long-term trends in the gender wage gap in the UK. As for many other developed economies, gender wage disparities in the UK remain high despite a steady convergence over time since the 1980s. [Figure 2](#) plots the average hourly wages of male and female employees over time according to the LFS.³ It also plots, in black and on the right-hand axis, the proportional difference between the two. The gap has decreased over the last 20 years from almost 30 per cent in 1993. Currently, the average female employee earns around a fifth less per hour than the average male employee.

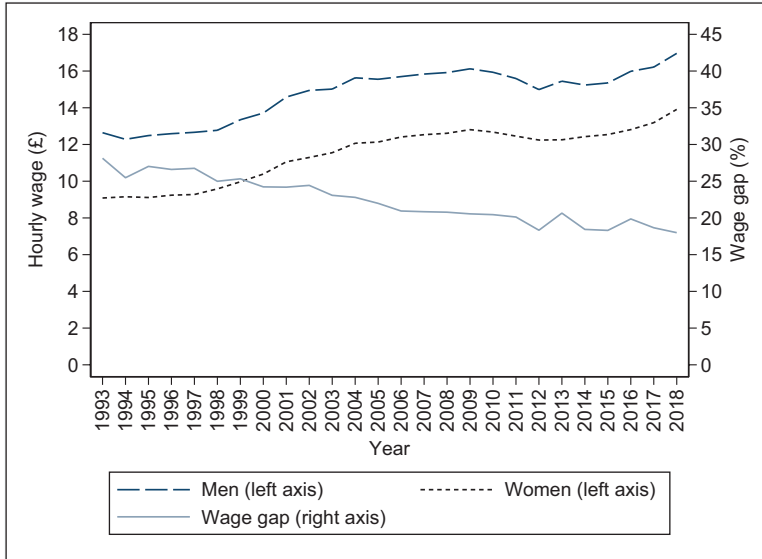
This wage gap is what it says on the tin: the difference between average female wages and average male wages. It is not a ‘like-for-like’ comparison between otherwise-identical workers or jobs. One reason why wage differentials between men and women might have changed over this period is that their relative levels of education have also changed. This is actually an important aspect to take into account in interpreting the declining wage gap over time.

[Figure 3](#) shows a rapid increase in the level of education of the working population over the past 20 years. The take-up of education happened mostly at the top, with university graduation rates more than doubling over the period, while the proportion of individuals leaving school with minimal qualifications dropped strongly to compensate. The increase in education was faster for women, who overtook men by the late 2000s and are now more likely than men to have a university degree. Because graduates tend to earn more than non-graduates, these differential trends in educational attainment have contributed to reducing the gender wage gap.

[Figure 4](#) shows the evolution of the gender wage gap as a proportion of male earnings over time, by education. For those with GCSE-level qualifications, this plot confirms that indeed the gender pay gap has fallen over the past two decades. However, it reveals no clear downward trend for the higher education groups. As a result, there

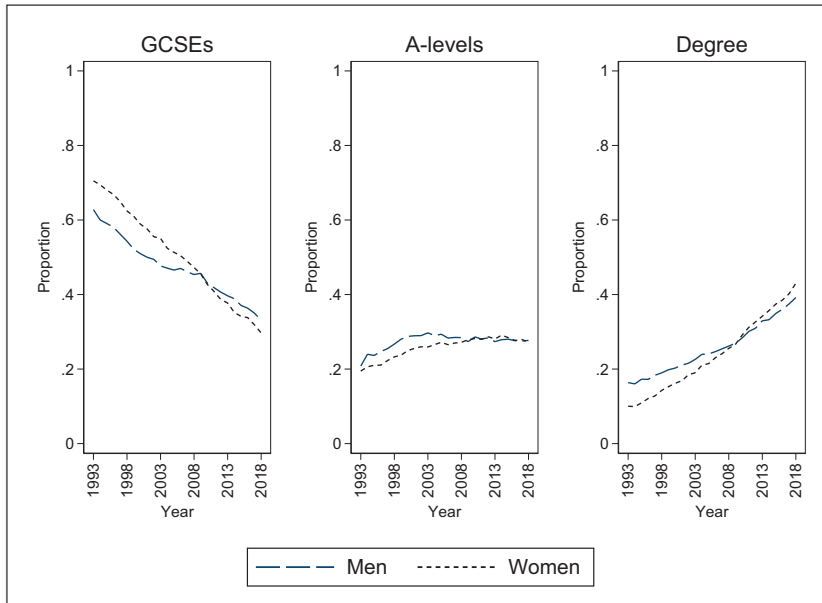
³ LFS data is used to produce the time trends presented in this section because its larger sample size captures the trends more accurately; results obtained using UKHLS data produce similar but more irregular patterns.

Figure 2: Trends in real hourly wages and the gender wage gap

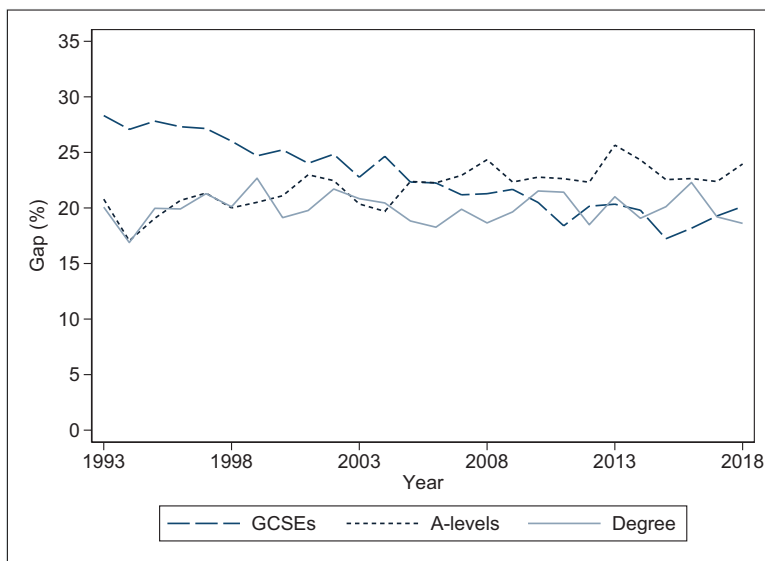


Notes: LFS 1993–2018. Real wage rates per hour in 2016 prices; observations in the top one and bottom two percentiles of the wage distribution by gender and year are excluded. Wage gap measured in proportion to male wages.

Figure 3: Trends in education attainment, by gender



Source: LFS 1993–2018.

Figure 4: Trends in the gender wage gap by education

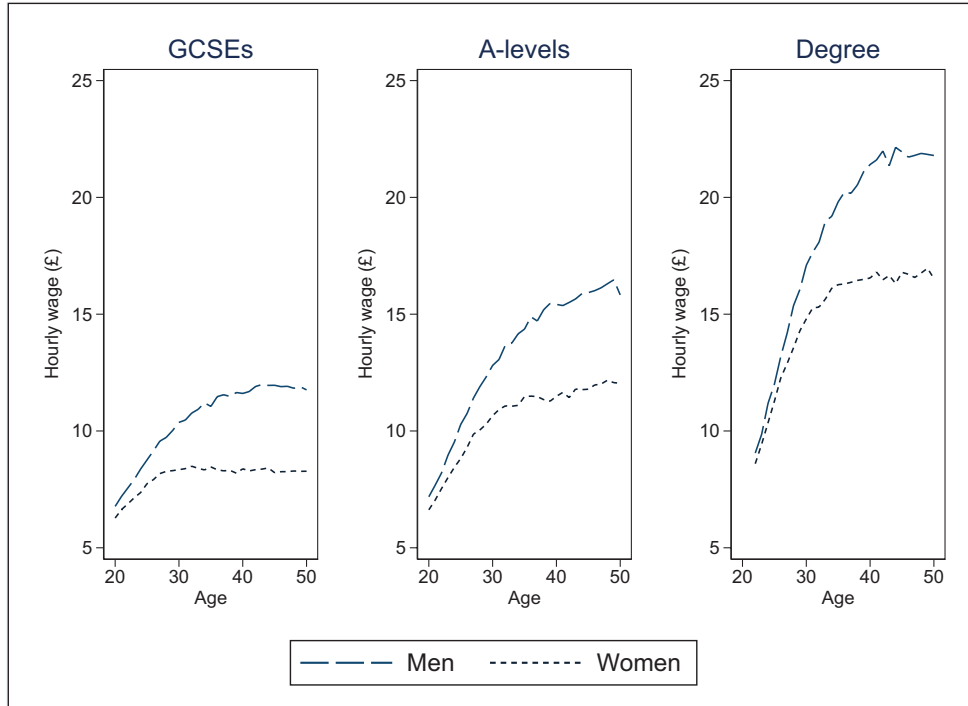
Notes: LFS 1993–2018. Wage gap measured in proportion to male wages.

has been a notable change in the nature of the gender wage gap. The gap used to be bigger (in proportional terms) for those with less formal qualifications than for university and high-school graduates, whereas the reverse is now the case. In summary, the fall in the overall gender wage gap over the past 20 years has been driven mostly by the lowest-educated individuals, and by an increase in the number of women who are highly educated.

(ii) Children, career patterns, and the gender wage gap over the life-cycle

A crucial starting point for disentangling the drivers of wage differences between men and women, which simple aggregate figures miss, is that those differences evolve over the life-cycle. This in turn is highly related to the arrival of children and changes in labour market behaviour associated with that. Figure 5 shows how average wages for male and female employees relate to their age (pooling those observed at the relevant ages between the start of 1993 and the middle of 2017). Notice that the sets of individuals who are employed at each age are different, so it is possible, for example, that women with low levels of experience return to employment in their 40s, thereby dragging down average female wages at that age. Wages are shown in 2016 constant-wage terms, so the increasing profile with age means that wages increase over the course of life by more than would be expected simply due to economy-wide growth.

The figure shows that wages typically increase with age throughout their 20s, for both men and women, which is consistent with the returns to additional experience being especially high for those with little experience. These returns look higher for graduates: their wage profile is especially steep throughout their 20s and, for men, well beyond that.

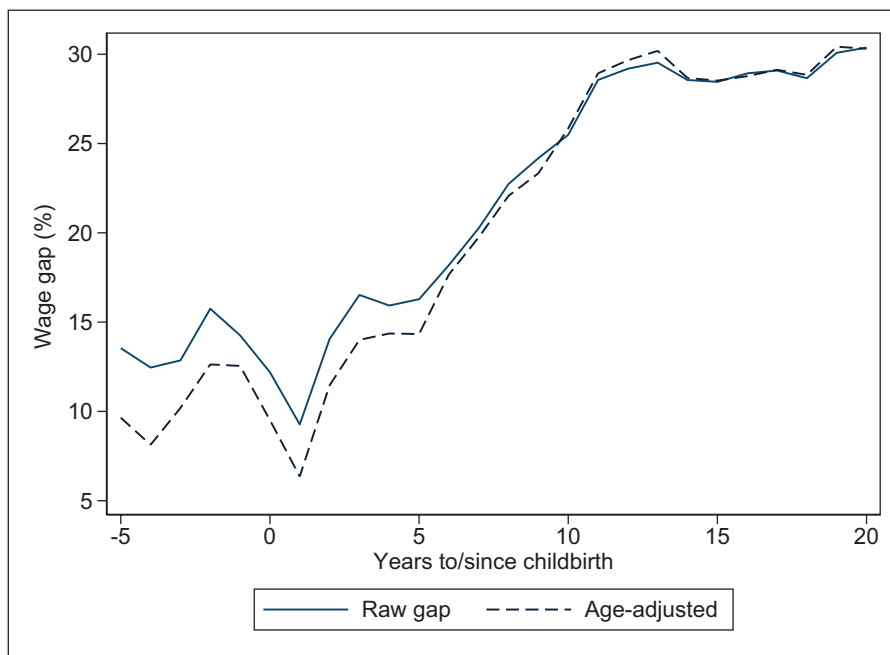
Figure 5: Mean hourly wages across the life-cycle by gender and education

Notes: LFS 1993 to 2018. Wage rates per hour in constant 2016 wage levels; observations in the top one and bottom two percentiles of the wage distribution by gender and year are excluded.

The gender wage gap is small or non-existent at around the time of labour market entry and it barely widens up to the mid-20s, particularly for college graduates. The gap then opens up more from around the late 20s and gets gradually wider over the next 20 years of the life-cycle. This is because male wages continue to increase, especially for the highly educated, while female wages completely flatline on average.

The opening of the gender wage gap when people reach their late 20s is likely related to the arrival of children. Figure 6 shows this explicitly by plotting the wage gap not by age, but by time to or since the birth of the first child in a family (where zero is the year in which that child is born). There is, on average, a wage gap of over 10 per cent even before the arrival of the first child. A small part of this gap is simply due to age differences—men tend to be slightly older than women when the first child arrives—though the age-adjusted line on Figure 6 still yields a wage gap of 7–12 per cent in the 5 years preceding the first child. A key feature of the patterns shown in the figure is that the gap appears fairly stable until the child arrives, and is small relative to what follows. After the child arrives, there is a gradual but continual rise in the wage gap over the following 12 years, until it reaches around 30 per cent of male wages.

While we measure wages on an hourly basis, and hence differences in working hours cannot directly explain the gender pay gap described above, different working patterns may lead to different hourly pay for more subtle reasons related to productivity, career progression, or other job characteristics. Figure 6 suggests that changes in women's

Figure 6: Gender wage gap by time to/since birth of first child

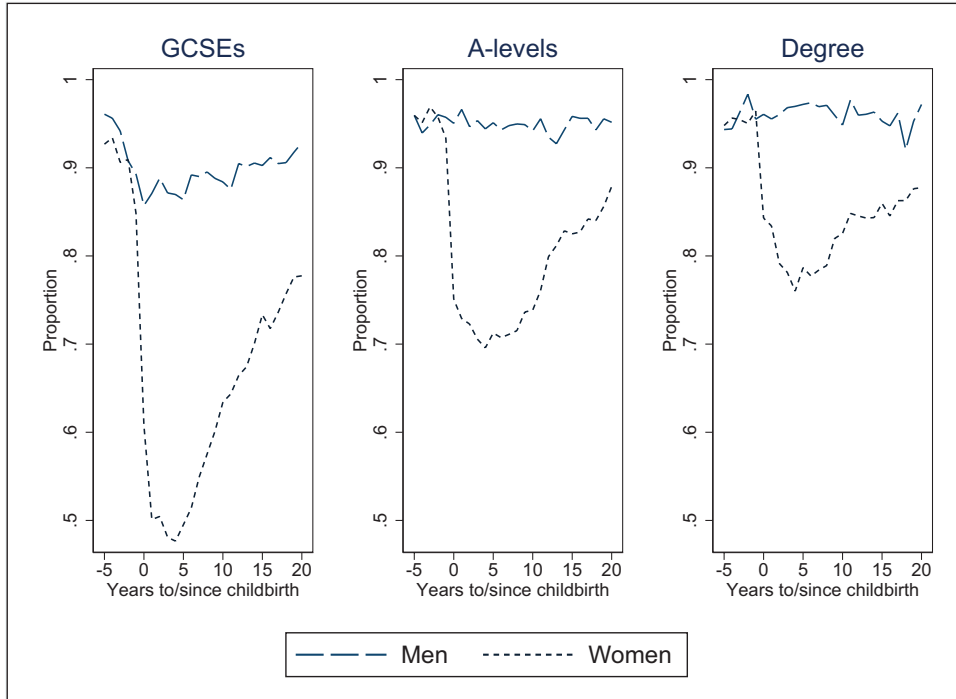
Notes: UKHLS 1991 to 2017. Wage gap measured in proportion of male wages. Observations in the top one and bottom two percentiles of the wage distribution are excluded. The age-adjusted series shows the gap obtained from wage rates net of education-specific age effects.

working patterns after the arrival of children may well be important in explaining this wage gap. The crucial observation is that the wage gap opens up gradually—not in any sudden jump—after the first child arrives and continues to widen for many years after that point. This pattern would be consistent with a gender gap in the level of labour market experience following the same basic shape as the gender gap in wages: relatively stable in the years before childbirth, growing incrementally for many years after that point, before eventually largely stabilizing once more. The next three figures show this.

Figure 7 shows the employment rates of men and women by the time to or since the birth of the first child. Before the arrival of the first child, it is difficult to discern any differences between the employment rates of men and women. However, when the first child arrives, a large employment gap opens up immediately: many women leave paid employment at this point, while any employment responses by men look tiny in comparison, and non-existent for high-school and college graduates.⁴

The employment response among the lowest-educated women is more than double the response among other women. Between one year before and one year after the birth of the child, women's employment rates drop by 30 percentage points (ppts) for those with GCSEs, 13ppts for those with A levels, and 9ppts for graduates. The other important feature of Figure 7 is that, once the employment gap opens up after the arrival

⁴ For the purpose of measuring employment, maternity or paternity leave is treated as being in paid work.

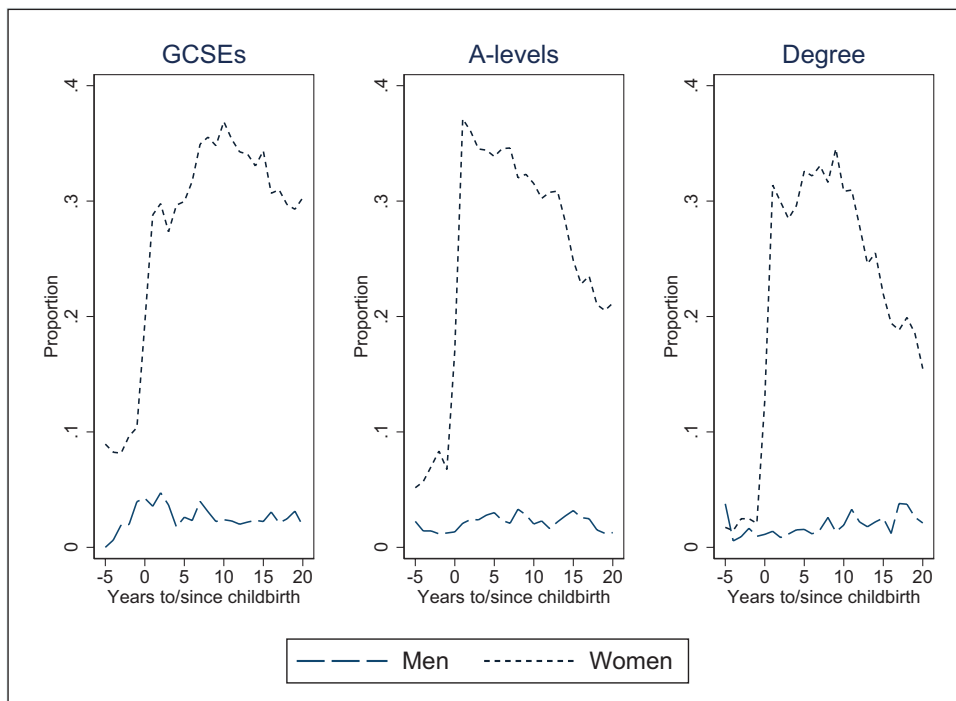
Figure 7: Employment rates of men and women

Notes: UKHLS 1991 to 2017.

of the first child, it persists. Women's employment rates do start to rise again once the first child is around school age, but they remain below male employment rates for the full 20 years shown. Hence, the gap in time spent in paid work keeps growing year on year for a long time after the first child arrives.

Figure 8 shows that not only do many women move out of paid employment altogether after having their first child, but many others move to part-time work (recall that part-time is defined as working 5–24 hours per week). Again, the male rate of part-time employment looks essentially unaffected by the arrival of the first child, and the gap that opens up is persistent: women are still significantly more likely to be in this kind of work than men when their first child reaches adulthood.

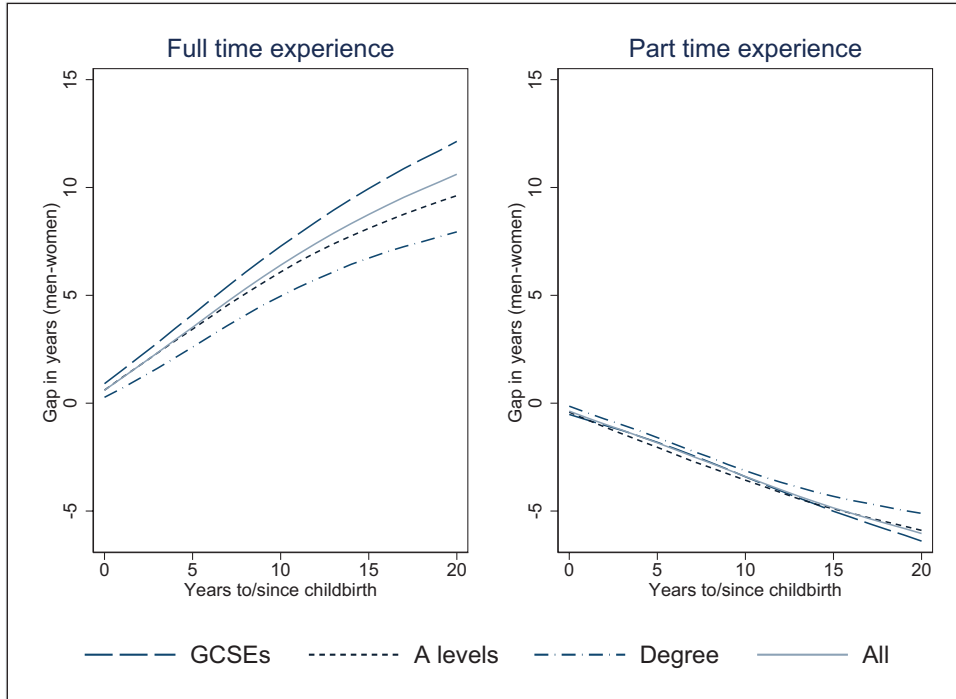
Figure 9 shows the direct implications of these patterns: a steadily increasing gap in accumulated labour market experience after the arrival of the first child. By the time their first child is aged 20, women have on average been in paid work for 3 years less than men, comprising 10 years less full-time paid work and 7 years more part-time paid work. The gap is larger still for the low-educated. Previous research (Blundell *et al.*, 2016) tells us that the 3 years less spent in any form of paid work understates the gender differences in accumulated human capital and that it is the 10-year gap in full-time experience that is more relevant. This is because it is only full-time paid work which seems to have substantial benefits in terms of the accumulation of experience that allows workers to command higher wages in future. We confirm this in the new analysis summarized in the next section, and examine the implications of the lack of wage progression in part-time work for the gender wage gap.

Figure 8: Proportion of all men and women in jobs of no more than 25 hours per week

Notes: UKHLS 1991 to 2017.

There are certainly other factors, besides levels of experience in paid work, that may be affected by childbirth and that may contribute to differences in wages between men and women. One possibility is that women undertake different kinds of work upon becoming mothers, potentially in different sectors of the economy. Such changes in job characteristics could be related to their wages for a number of reasons. For example, priorities or constraints could change around the time that children arrive such that women move towards occupations in which the benefits are less skewed towards wages and more towards other factors, such as flexibility. It could also be that a concentration of women in certain occupations or industries allows employers to exercise market power in order to hold wages down if, for example, they know that many of those women have limited ability or desire to search for alternative employment because they are time-constrained or want to work close to home. These different kinds of mechanisms linking occupation, industry, or other job characteristics to the gender wage gap would have very different implications for policy, and it is beyond the scope of this work to disentangle them (and there are many other possibilities besides the examples given). But what we can do is to provide a sense of their likely importance in accounting for the evolution of the gender wage gap.

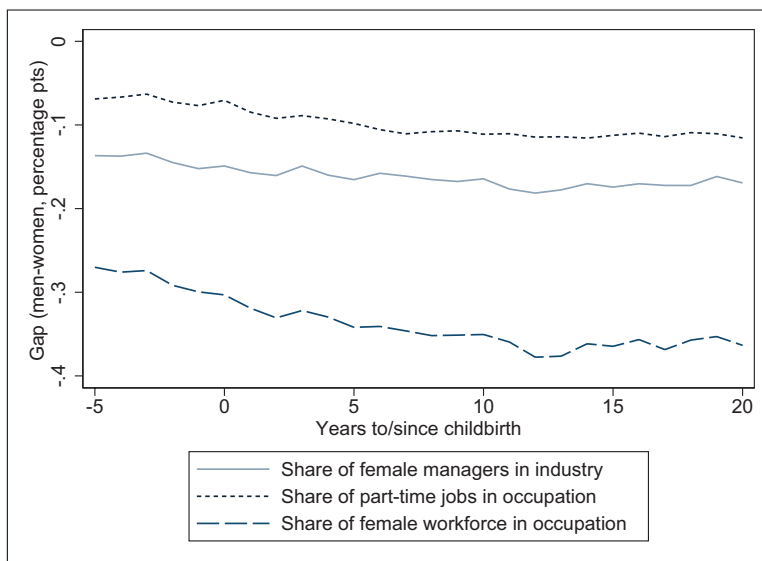
Figure 10 summarizes three example differences between the occupations and industries that women and men work in, and how these differences evolve around the time of childbirth. We take the occupation or industry that each worker is in, and map this to the composition of the workforce in that occupation or industry (computed from

Figure 9: Gender gap in years spent working full-time and part-time

Notes: UKHLS 1991 to 2017. The figure cumulates the gender gaps in years of work shown in Figures 6 and 7, and therefore does not include any differences in experience that already exist more than 5 years before the birth of the first child; these earlier differences are negligible.

the LFS). As time goes on, women who have children tend (relative to men who have children) to concentrate increasingly in female-dominated occupations, occupations in which part-time work is relatively common, and sectors in which female managers are relatively common. To that extent, there are similarities with the evolution of the gender wage gap—which also grows over the life-cycle, as we have seen. However, a closer look reveals a caveat to that: whereas the gender wage gap is fairly stable in the years before childbirth and then begins to gradually increase from the time of the first child, occupation and industry differences between men and women seem to be on a more uniformly increasing trajectory that starts a few years before the birth of the first child. This may in part be due to job changes in anticipation of having children. But it casts some doubt on the ability of these occupation or industry differences to explain powerfully the shape of the gender wage gap around the birth of the first child.

So far we have highlighted numerous factors that can play a role in driving the gender wage gap that persists in the UK labour market: education, labour supply choices along the intensive and extensive margins and the resulting work experience accumulation patterns, and characteristics of specific occupations/sectors, jobs, and working arrangements. We now perform a more comprehensive decomposition exercise that allows us to quantify the association between the gender wage gap and these factors. To do so, we estimate a set of wage regressions. We start from a baseline specification that only accounts for demographic

Figure 10: Gender gaps (men minus women) in characteristics of occupation and industry

Notes: UKHLS 1991 to 2017.

controls (age and region). The gender pay gap conditional on these variables is captured by the coefficient of a female dummy. We then proceed to richer specifications by progressively controlling for a number of differences between male and female workers: education and education interacted with age, part-time and full-time experience polynomials interacted with education, industry (2-digit SIC2007 codes) and occupation (3-digit SOC2010 codes), as well as other characteristics of the job including current working hours, indicator for public sector, and a self-reported measure of flexible working arrangements.

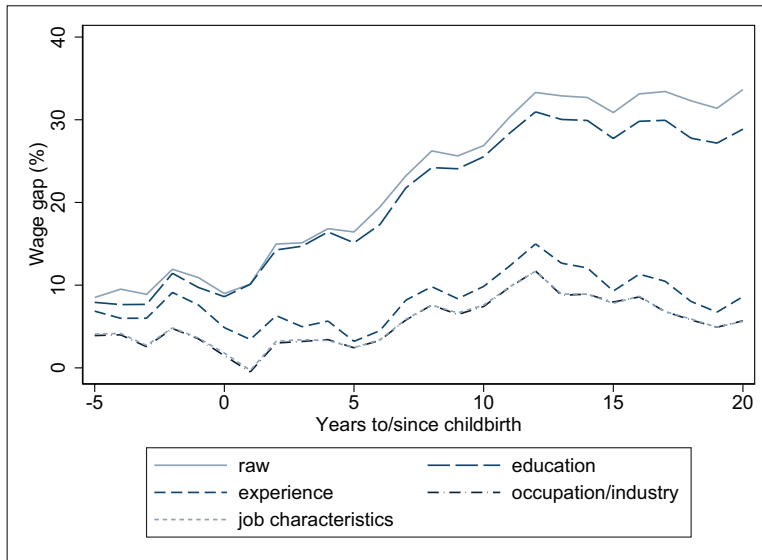
Table 2 and Figure 11 displays the estimate coefficients for the female dummy and how it varies as more detailed information about individual characteristics, working history, and the characteristics of the jobs is added to the regression model. The first row in the table shows that the raw gap of 22 log points is only mildly reduced by accounting for gender differences in educational attainment. Gender differences in experience show by far the strongest impact in reducing the gap (compare the estimate in column 3 to those in columns 1 and 2). After controlling for experience in full-time and part-time jobs, the gender gap in wages drops to just below 9 log points. After that, differences in industry and occupation can further reduce the gap by another 2 log points (column 4), and other job characteristics have no further impact (column 5).

Figure 11 shows these results in more detail, by years to/since first childbirth. It highlights that experience has the potential to account for a large amount of the gender wage gap, including the way that it evolves over the life-cycle—albeit still leaving a substantial part of the gender wage gap unexplained (and our causal analysis in the next section confirms this). Industry and occupation differences, by contrast, seem to explain far less.

These results show that opening differences in working experience are likely to be a major driver of the opening differences in the gender pay gap over the course of life. They do not dismiss other drivers as unimportant. Instead, they suggest that the

Table 2: Log hourly wage regressions, all education levels

	(1)	(2)	(3)	(4)	(5)
Female	-0.223*** (0.003)	-0.194*** (0.003)	-0.087*** (0.003)	-0.067*** (0.003)	-0.069*** (0.003)
Age, region	yes	yes	yes	yes	yes
Education	no	yes	yes	yes	yes
Experience	no	no	yes	yes	yes
Occupation, industry	no	no	no	yes	yes
Job characteristics	no	no	no	no	yes
N			75,482		

Figure 11: Gender wage gap by time to/since birth of first child, controlling for association between wages and individual and job characteristics

Note: Estimates use UKHLS 1991 to 2017.

divergence in the careers of men and women happen together with the relative losses in working experience that women go through after childbirth. For instance, one would expect that working experience facilitates moving up the occupational ladder as workers gain the skills necessary for promotions. The consequence of this parallel movement is that, conditionally on working experience, the role of other drivers of wages is less visible. We therefore focus on working experience in what follows, to examine its causal role on gender differences in wages.

IV. Model and estimation

Our estimates so far do not allow us to quantify how the gender wage gap opens with the arrival of children and the role of working experience in driving it because the

different processes through which men and women select into work confound these effects. If high-ability women on a steeper wage curve are more likely to remain in work after childbirth than those on lower earnings, while male selection into work is not affected by fatherhood, then our measures of how the gender pay gap opens after childbirth and of the role of working experience in driving that gap will be biased. Indeed, we can see some of that selection in action when contrasting the labour supply behaviour of mothers of young children across education groups, with more-educated women being less likely to interrupt their working careers. In this section we discuss our empirical strategy to identify the causal impact of experience on wages by explicitly accounting for endogenous selection into work. Our model is for women, as variation in experience conditional on age is more meaningful for them. In our simulation exercise later in the paper, we control for pre-birth wages to control for permanent differences in wages between mothers that do and do not interrupt work after childbirth.

(i) Regression model of wages

To measure the causal role of working experience in driving the wages of women, we specify and estimate a simple but flexible model of experience capital accumulation and wages. In all that follows, we fully interact all aspects of the model with education and, hence, consider separately the choices and outcomes of the three education groups that we have studied so far. To simplify the notation, we omit the education index but it is implicit in all parameters below.

Following Blundell *et al.* (2016), wages are modelled as a simple function of accumulated human capital k . The hourly wage rate w of woman i at age t is simply:

$$\ln w_{it} = \alpha_i + \ln(k_{it} + 1) + u_{it} \quad (1)$$

where α_i is an individual-specific wage level per unit of experience capital and u_{it} is a time-varying idiosyncratic wage shock. The latter may include a persistent and a transitory component.

We assume that human capital is accumulated in work and consider two positive labour supply points, corresponding to part-time and full-time hours. In line with our previous analysis, these correspond to weekly hours between 5 and 25, and 26 or more, respectively; anyone working 4 or fewer hours per week is considered as not working. Formally, the human capital process is:

$$\ln(k_{it} + 1) = \ln(k_{it-1} + 1) - \delta + \pi(e_{it-1})P_{it-1} + \varphi(e_{it-1})F_{it-1} \quad (2)$$

and we set the initial value k_{i1} to zero. We allow for skills to depreciate in each period at a rate δ and to accumulate while in work at rates π or φ that depend on whether the woman was working part-time or full-time in period $t - 1$. The rate of human capital accumulation also depends on how many years of working experience the woman has accumulated by the end of period $t - 1$, denoted by e_{t-1} . We differentiate between experience in part-time and full-time work, denoted by e^P and e^F respectively. So $e = (e^P, e^F)$ and the experience accumulation process is simply

$$e_{it}^P = e_{it-1}^P + P_{it-1} \quad \text{and} \quad e_{it}^F = e_{it-1}^F + F_{it-1} \quad \text{with starting values } e_{i1}^P = e_{i1}^F = 0.$$

The parameters (π, φ) are specified as follows:

$$\pi(e_{it}) = \pi_1 + \pi_2 e_{it}^P + \pi_3 e_{it}^F$$

$$\varphi(e_{it}) = \varphi_1 + \varphi_2 e_{it}^P + \varphi_3 e_{it}^F$$

where (π_j, φ_j) for $j = 1, 2, 3$ are unknown parameters that we estimate. This setting allows for decreasing returns to additional periods of work if π and φ are decreasing functions of (e^P, e^F) .

Our goal here is to consistently estimate the dynamic returns to work and to working hours in order to assess their combined role in driving the gender pay differentials. To do so, we re-write the wage equation (1) in first differences and replace the growth in experience capital using equation (2) and the expressions for (π_j, φ_j) . This eliminates human capital, which is not directly observed in the data, and yields our regression equation:

$$\Delta \ln = \omega_{it} = -\delta + (\pi_1 + \pi_2 e_{it}^P + \pi_3 e_{it}^F) P_{it-1} + (\psi_1 + \psi_2 e_{it}^P + \psi_3 e_{it}^F) F_{it-1} + \Delta u_{it} \quad (3)$$

Equation (3) shows that the returns to a year of work for workers with no past work experience (π_1, φ_1) cannot be distinguished from the depreciation rate. We, therefore, estimate

$$\tilde{\pi}_1 = \pi_1 - \delta \text{ and } \tilde{\psi}_1 = \psi_1 - \delta$$

(ii) Estimation

The direct estimation of equation (3) using the sample of women in continuous work will result in biased estimates of the unknown parameters for three main reasons. The first is the classical employment selection problem, which in our case applies to two consecutive periods. Second, lagged hours are likely correlated with the unobserved term Δu if, as usually accepted, hours of work respond to contemporaneous wage shocks. And third, accumulated years of work may also be correlated with Δu for two main reasons. The first is more obvious: Δu may have a long memory and hence contain information driving past labour supply—something that would happen if, as is frequently assumed, the permanent wage shock is auto-regressive. The second potential source of dependence arises from the conditioning on current and past employment. If, for instance, more experienced workers are more likely to keep working through periods of low unobserved wage growth than less experienced workers, Δu_t and e_{t-1} will be correlated in the sample of workers even if they are not in the full sample.

The discussion above highlights the fact that we have various sources of bias that need to be accounted for in estimating the wage regression (3), including the extensive and intensive margins of labour supply and the endogeneity of years of experience in part-time and full-time work. To deal with these, we extend the control function method of Heckman (1979). The details of our method can be found in the on-line [Appendix](#). Here we describe the sources of exogenous variation that we use to construct the control functions that are then included in the regression equation (3).

The first set of excluded variables includes three accounts of the disposable income of each family in each year, by whether the woman did not work, worked part-time hours, or worked full-time hours. These instruments are simulated on a micro-simulation tool that contains a fine description of the direct taxes and benefits in the UK, and how they changed over the entire observation period.⁵

They are meant to capture how the sequence of tax and welfare reforms changed the incentives to work of women in different family arrangements. To ensure that the exogenous policy variation is separated from other potentially endogenous sources of variation, we use predicted female wages to calculate her gross earnings by working hours. The predictions are based on a full set of time and age dummies, fully interacted with education. We also set the earned income of the spouse to zero. Finally, we net out aggregate time effects and fixed family demographic effects. The remaining variation in the simulated instruments captures only how policy reforms differentially affect the disposable income of different types of families over time, should the woman not work, work part-time hours, or work full-time hours (see Blundell *et al.*, 1998). The simulated disposable incomes are complemented with a set of instruments more closely related to accumulated years of working experience, a cubic polynomial in age and the ages of the youngest and oldest children.

All regressions include an additional set of exogenous covariates, including regional dummies and a second order polynomial in two background indices that summarize family background.⁶ The first stage regressions also include indicators for the presence of children, presence of a partner, and whether the partner is working.

V. Results

We estimate our model using UKHLS data for the 1991–2016 period. Estimates of the first stage regressions can be found in the [Appendix to this paper, Tables 4–7](#). They show that the instruments we are using are strong drivers of the endogenous variables in most cases. As expected, residual simulated disposable income is a stronger determinant of employment and hours of work for women with basic education only, and has less power as a driver of the same outcome among university graduates. In addition, age and age of oldest child are strong predictors of accumulated experience. We test the strength of the set of instruments meant to capture exogenous variation for each of the endogenous variables, and find evidence in support of the instruments in

⁵ Fortax is a detailed tax and benefit micro-simulation tool that can be used to accurately predict the budget constraints families face by earned income. It accounts in detail for the tax and welfare system in place at each point in time and how they changed over the period of our data. More information can be found in Shephard (2009) and Shaw (2011).

⁶ The indices summarize some of the characteristics of the individual's parental home and are meant to capture permanent individual traits that drive productivity in the labour market and labour supply choices. They are the first two factors from a principal component analysis on a set of variables describing the socio-economic background of the woman. These include her parents' education and whether they were working when she was 16, whether she lived with both parents at that same age, books at home as a child, ethnicity, number of siblings, and sibling order.

Table 3: Wage regressions, women by education

	GCSEs		A levels		University	
	LRM (1)	CF (2)	LRM (3)	CF (4)	LRM (5)	CF (6)
Lagged part-time hours	-0.016 (0.010)	0.001 (0.018)	-0.017 (0.012)	0.024 (0.023)	0.004 (0.027)	-0.009 (0.064)
× lagged full-time experience	-0.000 (0.001)	0.000 (0.001)	-0.001 (0.001)	-0.001 (0.001)	-0.003 (0.002)	-0.003 (0.002)
× lagged part-time experience	-0.000 (0.001)	-0.002 (0.001)	-0.000 (0.001)	-0.000 (0.001)	-0.002 (0.002)	-0.003 (0.003)
Lagged full-time hours	0.058** (0.006)	0.036** (0.009)	0.066** (0.005)	0.061** (0.006)	0.078** (0.009)	0.079** (0.011)
× lagged full-time experience	-0.002** (0.000)	-0.001 (0.001)	-0.002** (0.000)	-0.002** (0.000)	-0.003** (0.000)	-0.003** (0.001)
× lagged part-time experience	0.000 (0.000)	-0.001 (0.001)	-0.001** (0.000)	-0.001 (0.001)	-0.001 (0.001)	-0.004* (0.002)
Observations	15,027	13,341	17,908	15,118	6,762	5,946
F-test for control functions		2.886		3.156		1.946
<i>p</i> -val		0.013		0.008		0.084

Notes: Estimates on UKHLS data for the 1991–2016 period. The column titles ‘LRM’ and ‘CF’ stand for linear regression model and control function, respectively. The *F* statistic is for the statistical significance of the control functions for employment, lagged employment, lagged hours, and lagged experience in PT and FT hours. All regressions also include a second order polynomial in the two background indices and regional dummies. Bootstrapped standard errors clustered at the individual level in parentheses under the estimated coefficients. * $p < 0.05$, ** $p < 0.01$.

all cases except for hours choices among college graduates (see *p*-val at the bottom of Table 5).

Using these estimates, we then construct the various control functions and include them in our regression model for wage growth. Estimates of the parameters of interest are shown in Table 3 for each of the three education levels. Columns 1, 3, and 5 display estimates by education from the linear regression model, without controlling for employment selection, endogenous hours, or accumulated experience; columns 2, 4, and 6 display the control function estimates. Clearly, controlling for endogenous selection and experience does not much affect the estimate of the experience effects. Figures at the bottom of the table show the *p*-value for the test of statistical significance of the set of control functions used to tackle endogeneity. They show marginally significant effects for the two lower education groups, but not for the top education group.

The results in row 4 suggest that full-time work has a strong positive impact on wages, and that this effect increases with education. Moreover, the numbers in rows 5 and 6 further suggest that the returns to one additional year of full-time work drop with full-time experience but not with part-time experience. These estimates are consistent with the view that the returns to work drop over the life-cycle as workers accumulate experience, leading to a concave wage profile over the course of working life. It is also evident from the figures in Table 3 that part-time work has little or no impact on wages. Indeed, estimates in rows 1 to 3 show small and statistically insignificant estimates of the impact of part-time work. These figures show that, on average, the wages of women working part-time hours stagnate.

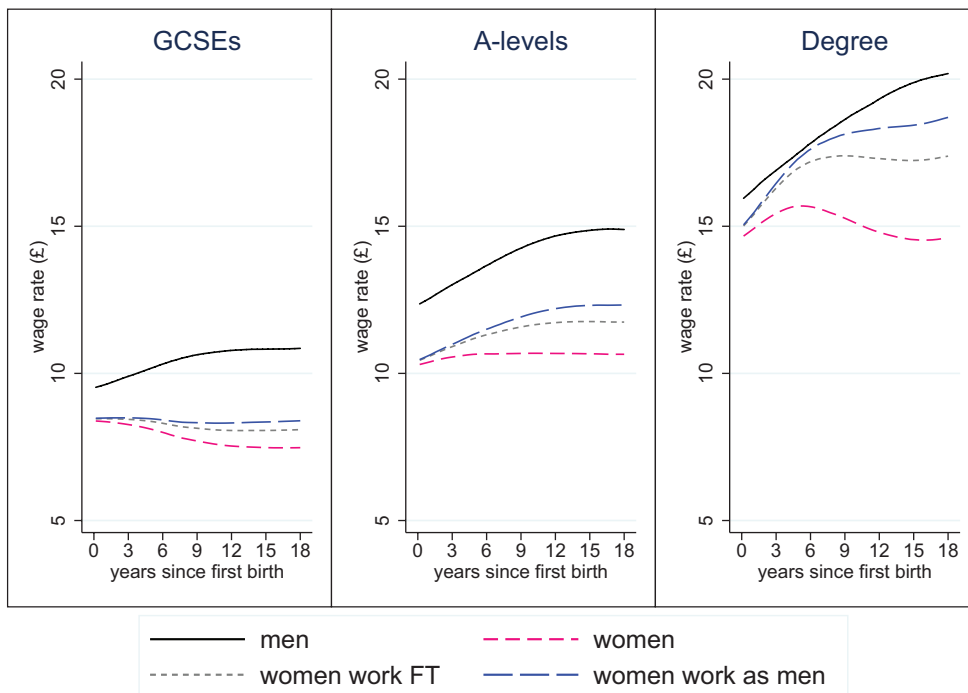
VI. Counterfactual simulations of the cumulative effect of employment and hours of work

We use the control function estimates of our wage equation to run two counterfactual experiments. First, we set all work to be full time. We do this by setting the value of part-time work for future wages to be equal to the estimated value of full-time work. And second, we assume that women work at the same rate as men with the same education and demographic characteristics. We then simulate the experience profiles of women under these two alternative scenarios by assuming that there is no depreciation of market skills during the periods when the woman is not doing paid work. Our counterfactual wages are constructed from observed wages by netting out the experience effect at observed levels of experience and adding the experience effects at the counterfactual levels of experience. Hence, the unobserved component of wages remains unaltered in our simulations.

The possibility that different women select into work as the child grows up is also not changed. All of these will still be reflected in the simulated profiles.

Figure 12 shows the results from these simulations for parents, by time since the birth of the first child and education. In each of the three panels, the top solid line and the bottom dashed line are the observed wage profiles of fathers and mothers, respectively. The plots show how these two lines move apart as the child grows, and they also highlight that the wages of women at best stagnate after childbirth. This holds true for all education groups. In relative terms, the pay differential when the child reaches 18 is

Figure 12: Counterfactual simulations—hourly wage rate by time since first birth



Notes: Based on BHPS data for the 1991–2008 period. 2016 wage levels.

remarkably similar across education groups. In all cases, women earn about 30 per cent less than men do at that point in the child's life.

The two intermediate lines in the graph show how employment and working hours after childbirth inhibit wage progression for women. The right-hand-side graph for college graduates suggests that working experience is a determinant factor in explaining the opening gap after childbirth. Given our estimates of the effect of working experience on wages, if college graduate women were to work full-time hours if they worked at all, the wage gap 18 years after childbirth could close up to 50 per cent. If, in addition, they were to work at the same rate as men do, the wage gap at the same age of the child could be further reduced to one-third of the observed level.

However, experience plays a less important role in determining the gender pay gap for workers who do not have a college degree. For these groups, gender differences in pay are comparatively large at childbirth and the gap in accumulated experience after that can only account for about one-third of the gap in pay when the child reaches 18 years of age. Other differences in job characteristics, firm characteristics, occupation, or in how wages are negotiated are likely to play a determinant role in driving gender pay differentials among workers with GCSE and A-level qualifications.

VII. Policy implications

We have shown in the previous section that UK women across all education levels face a so-called child penalty in wages of around 30 per cent with respect to men by the time their oldest child reaches age 18. This result is in line with Kleven *et al.* (2019a) who find long-run (10 years after the first birth) child penalties on earnings between 31 and 34 per cent in the UK and US as opposed to 21–26 per cent in Scandinavian countries and up to 51–61 per cent in Austria and Germany.

These earning gaps are by definition the combined effect of the gap in employment rates, the gap in hours of work, and the gap in hourly wage rates between women and men. Countries differ in the relative importance of these components in determining the overall gap in earnings accumulated by mothers before and after the birth of the first child, depending on the characteristics of their labour markets and on the policies in place. In the UK, as we have shown in our counterfactual simulations, the intensive margin of the labour supply gap (i.e. the prevalence of part-time work among women) drives a significant amount of the observed earnings gap over the life-cycle. The US overall earning gap is similar to the UK one, but the small fraction of women working part-time suggests that the participation and wage components are the relevant drivers there. One reasonable inference from this is that just encouraging or facilitating more mothers to remain in employment after childbirth is not enough to make large inroads into the gender wage gap unless the causes of poor pay progression in part-time work can be addressed. What we find in this study is that it is only full-time employment that appears to promote wage growth.

However, in Scandinavian countries part-time work is not very prevalent and female participation rates are higher than the OECD average, yet earning gaps for mothers still exist. In these countries the earnings gap is smaller than elsewhere and the labour supply gaps still explain a large proportion of it (see Kleven *et al.* (2019b) for estimates for Denmark). Interestingly, the wage gap also opens far less with motherhood in the

Scandinavian countries than it does in the UK. This evidence is consistent with our results that, in part, the opening of the gender wage gap among parents builds up as a consequence of the lost experience in work, which may embed lack of occupation progression and failure to move to the best-paying industries, firms, and jobs for mothers. Some southern European countries are found at the opposite extreme, with female employment rates significantly lower than the OECD average and especially large gender earnings gaps, probably driving down the apparent gender earning gaps due to more marked selection of higher-earning women into employment than in other countries. Austria and Germany, instead, are peculiar cases as they show extremely high and persistent earning penalties for women that are combined with high female participation but also especially high rates of part-time work. While all these results are for earnings rather than wages, they seem consistent with our findings that low levels of labour supply at both the intensive and extensive margins have long-term consequences for the earnings ability of mothers.

Some public interventions aimed to promote female employment and fertility have indeed been shown to be important in determining women's labour market outcomes. We can identify three major policy tools: (i) parental leave; (ii) child care provision; (iii) income support and tax incentives (Olivetti and Petrongolo, 2017). The design and mix of these policies vary widely across countries. Almost all developed countries have adopted paid maternity leave—the exception being the US. The average paid maternity leave across OECD countries is around 18 weeks, with a wide variation across countries in terms of both length and entitlements. Mothers are usually given from a minimum of 13–14 weeks leave around childbirth (e.g. in Norway and Germany) at a replacement rate of 100 per cent of previous earnings, up to a maximum of 9 months at an average of around a 30 per cent replacement rate (as in the UK).⁷ Maternity leave policies aim at preserving the labour market attachment of mothers despite temporary interruptions, but they might have undesired effects on women's wages and career progression as they create incentives for mothers to stay away from their jobs for longer and to create a longer career interruption than would otherwise have been the case. The net effect of parental leave on female employment has been the focus of several empirical studies. Among others, Lalive and Zweimüller (2009) found a significant short-run (first 3 years after childbirth) negative effect on female employment and earnings of an extension of the parental leave from 1 to 2 years that was implemented in Austria in 1990. These negative effects were larger for high-wage women. In 2007 Germany implemented a reform that linked maternity leave benefits to pre-birth income, thus incentivizing high-earning women to take up maternity leave. The reform resulted in delayed return to work and increased part-time employment in the medium run (Kluge and Schmitz, 2014).

The second class of policies, public provision of childcare, can be done through cash transfers or in-kind provision. Spending on childcare varies widely across countries. OECD figures show that spending on early childhood education and care range between almost 2 per cent of GDP in Iceland down to 0.1 per cent in Greece and Turkey; the UK spending at 0.7 per cent of GDP matches the OECD average (see Farquharson (2019)). Subsidized childcare programmes often have several aims, including promoting child development and functioning as transfers to parents, but in relation to labour

⁷ Figures from OECD Family Database.

markets their goal is to facilitate parental (and particularly maternal) labour supply. In countries where there is little private provision of childcare, publicly provided childcare provides an important substitute for parental (and particularly maternal) care. But even where there is a robust private childcare market in place, public childcare subsidies reduce the out-of-pocket costs parents face and, hence, the penalty they face on their effective hourly wage relative to childless workers.

Overall, empirical evidence on the impact of childcare subsidies on maternal labour supply has been mixed. Studies tend to find larger effects in countries where crowd-out is less of a concern (Bauernschuster and Schlotter, 2015); where pre-reform female labour supply is lower (Lundin *et al.*, 2008); and where initial childcare costs are higher and other barriers to female employment, such as low labour demand or more traditional gender attitudes, are lower (Lefebvre and Merrigan, 2008; Havnes and Mogstad, 2011; Cascio and Schanzenbach, 2013; Givord and Marbot, 2015; Nollenberger and Rodríguez-Planas, 2015). The size of the incentive also matters. For instance, Brewer *et al.* (2020) find that the provision of free childcare places only affects maternal labour supply significantly if the number of subsidized hours is large enough. Part-time places are unlikely to change mothers' choices as they do not offer enough flexibility for them to take up paid work. This non-linearity in the effects of interventions has rarely been studied and is key for optimal design.

Lastly, in many countries the tax and transfer system now provides some form of income top-up and tax incentive to low-earning parents. The UK's Working Tax Credit, following a major expansion of the tax credit system in the late 1990s, has been a prominent example (though it is in the process of being subsumed within universal credit, which is replacing and integrating multiple transfer payments), alongside the US's Earned Income Tax Credit. The aim of these policies is to top up the incomes of low-earning working families while also encouraging higher rates of employment among those families (especially single parents). There is good evidence for the UK that it achieved these goals (Brewer *et al.*, 2006; Blundell *et al.*, 2016).

A third, more profound long-term impact that one might expect to follow from this would be a steeper career trajectory as parents remain more attached to the labour market and build up more experience in work than they would otherwise have done: encouraging them into work would be merely a stepping stone to wage progression. Unfortunately this appears not to have been an outcome in the UK, due to precisely the kinds of wage dynamics uncovered in this paper. The wage returns to experience are generally quite low for those with lower education levels, and they are especially low for those doing part-time work. The means-tested nature of tax credits mean that they disproportionately impact those with lower education levels and they can often incentivize part-time work over full-time work, since they get withdrawn from families if earnings rise too high. Hence the long-term impacts of tax credits on wages have been very limited (Blundell *et al.*, 2016).

Since the nature of these wage dynamics has only been studied relatively recently, they have not yet been factored in to the design of in-work cash transfers. It may well be that their design can be refined given that we now know a lot more about the longer-term difference in payoffs between part-time and full-time work. That said, it should be noted that its impact on the gender wage gap would still be limited by the fact that means-tested transfers are targeted at those with low pay, who are also disproportionately concentrated in the lowest education group. For them, the impact of even

full-time work on future wages is relatively modest. It is in fact highly educated mothers for whom these experience-wage dynamics are largest and who fall furthest behind their male counterparts in the years after childbirth.

The empirical reduced form evaluations cited above shed light on the kinds of reforms most successful at addressing gender wage gaps. However, much of the evidence still focuses mainly on the short- and medium-run impact of policies, whereas the effects of these policies crucially depend on the incentives that women, and households in general, face over the entire life-cycle. Therefore, in order to assess the long-run impact of these policies, one must take into account working experience accumulation. The main contribution of this paper is precisely to provide a simple but flexible empirical framework that features this dynamic mechanism and that can be used to inform policy debate. Our results highlight, for instance, that policies incentivizing mothers of young children to remain actively in full-time work are likely to help their career progression and pay handsomely in the long term, particularly for those mothers with medium to high qualifications.

Taking a broader perspective, it is important to bear in mind that a general cross-country policy recommendation does not exist. Every country has its own mix of policies and each policy should not be analysed in isolation from the specific institutional background in which it is set as its effect depends on what other incentives and protections are there. Moreover, the interaction of public policies with social norms and attitudes towards mothers' participation in the labour market, that are country-specific and are evolving at different paces across countries, should not be overlooked. Finally, there could be relevant demand effects of public interventions aimed at reducing gender pay and employment gaps. For instance, policies that help bring up the labour supply of mothers may help contain statistical and other types of discrimination differentially affecting wage setting and career progression of men and women (Xiao, 2020).

VIII. Concluding remarks

Gender differences in rates of full-time and part-time paid work after childbirth are an important driver of differences in hourly wages between men and women. This is because they affect the amount and type of labour market experience that men and women build up, and this experience affects the hourly wage levels they can command. In this paper we show that differences in working experience are determinant in explaining the gender pay gap of college graduates, for whom they can explain up to two-thirds of the wage differences 20 years after childbirth. The role of experience in driving the gender wage differences of those with GCSE-level and A-level qualifications is more modest, accounting for about one-third of the gap 20 years after the first childbirth.

It is not only taking time out of paid work that matters; crucially working part-time after childbirth seems to hold back women's wages. This is because extra experience in full-time work leads to higher hourly wages, whereas extra experience in part-time work does not.

A key challenge for future research, then, is to understand why part-time work shuts down wage progression so much. There are a number of possibilities, including less training provision, missing out on informal interactions and networking opportunities,

and genuine constraints placed upon the build-up of skill by working fewer hours. Understanding this properly looks to be of great potential importance for policy-makers who want to address the gender wage gap. Of course, our results also suggest that an alternative (or complementary) focus would be on understanding the causes of gender differences in rates of full-time work in the first place, such as the division of childcare responsibilities.

Our results also show that closing gender gaps in rates of full-time and part-time paid work, or narrowing the difference between the impacts of full-time and part-time paid work on wage progression, cannot be expected to close the gender wage gap fully. This is especially relevant when thinking about the relationship between the gender wage gap and poverty: among lower-educated people, there is already a relatively substantial gender wage gap before the first child is born, and gender differences in full-time and part-time paid work in the subsequent 20 years explain only a minority of the gender wage gap that has built up by that point. Previous research suggests that other contributing factors could include women being less likely to work in more productive firms, less likely to successfully bargain for higher wages within a given firm, and more likely to enter family-friendly occupations over high-paying ones.⁸ Better understanding of mechanisms such as these, and their underlying causes, is another key priority for further research.

References

- Adda, J., Dustmann, C., and Stevens, K. (2017), 'The Career Costs of Children', *Journal of Political Economy*, **125**(2), 293–337.
- Altonji, J. G., and Blank, R. M. (1999), 'Race and Gender in the Labor Market', *Handbook of Labor Economics*, **3**, 3143–259.
- Attanasio, O., Low, H., and Sánchez-Marcos, V. (2008), 'Explaining Changes in Female Labor Supply in a Life-cycle Model', *American Economic Review*, **98**(4), 1517–52.
- Bauernschuster, S., and Schlotter, M. (2015), 'Public Child Care and Mothers' Labor Supply—Evidence from Two Quasi-experiments', *Journal of Public Economics*, **123**, 1–16.
- Bertrand, M., Goldin, C., and Katz, L. F. (2010), 'Dynamics of the Gender Gap for Young Professionals in the Financial and Corporate Sectors', *American Economic Journal: Applied Economics*, **2**(3), 228–55.
- Blau, F. D., and Kahn, L. M. (2017), 'The Gender Wage Gap: Extent, Trends, and Explanations', *Journal of Economic Literature*, **55**(3), 789–865.
- Blundell, R., Duncan, A., and Meghir, C. (1998), 'Estimating Labor Supply Responses Using Tax Reforms', *Econometrica*, 827–61.
- Costa Dias, M., Meghir, C., and Shaw, J. (2016), 'Female Labor Supply, Human Capital, and Welfare Reform', *Econometrica*, **84**(5), 1705–53.
- Brewer, M., Cattan, S., Crawford, S., and Rabe, B. (2020), 'Does More Free Childcare Help Parents Work More?', IFS Working Paper W20/9.
- Duncan, A., Shephard, A., and Maria Jose, S. (2006), 'Did Working Families' Tax Credit Work? The Impact of In-work Support on Labour Supply in Great Britain', *Labour Economics*, **13**(6), 699–720.

⁸ Others have looked at some of these issues. See, for example, Adda *et al.* (2017), Card *et al.* (2015).

- Card, D., Cardoso, A. R., and Kline, P. (2015), 'Bargaining, Sorting, and the Gender Wage Gap: Quantifying the Impact of Firms on the Relative Pay of Women', *The Quarterly Journal of Economics*, **131**(2), 633–86.
- Cascio, E. U., and Schanzenbach, D. W. (2013), 'The Impacts of Expanding Access to High-quality Preschool Education', Technical Report, National Bureau of Economic Research.
- Farquharson, C. (2019), 'Early Education and Childcare Spending', IFS Briefing Note 258.
- Givord, P., and Marbot, C. (2015), 'Does the Cost of Child Care Affect Female Labor Market Participation? An Evaluation of a French Reform of Childcare Subsidies', *Labour Economics*, **36**, 99–111.
- Goldin, C. (2014), 'A Grand Gender Convergence: Its Last Chapter', *American Economic Review*, **104**(4), 1091–119.
- Havnes, T., and Mogstad, M. (2011), 'No Child Left Behind: Subsidized Child Care and Children's Long-run Outcomes', *American Economic Journal: Economic Policy*, **3**(2), 97–129.
- Heckman, J. J. (1979), 'Sample Selection Bias as a Specification Error', *Econometrica*, **47**(1), 153–61.
- Kleven, H., Landais, C., and Søgaard, J. E. (2019b), 'Children and Gender Inequality: Evidence from Denmark', *American Economic Journal: Applied Economics*, **11**(4), 181–209.
- —Posch, J., Steinhauer, A., and Zweimüller, J. (2019a), 'Child Penalties Across Countries: Evidence and Explanations', in *AEA Papers and Proceedings*, vol. **109**, 122–6.
- Kluge, J., and Schmitz, S. (2014), 'Social Norms and Mothers' Labor Market Attachment: The Medium-run Effects of Parental Benefits', available at SSRN 2471521.
- Lalive, R., and Zweimüller, J. (2009), 'How Does Parental Leave Affect Fertility and Return to Work? Evidence from Two Natural Experiments', *The Quarterly Journal of Economics*, **124**(3), 1363–402.
- Lefebvre, P., and Merrigan, P. (2008), 'Child-care Policy and the Labor Supply of Mothers with Young Children: A Natural Experiment from Canada', *Journal of Labor Economics*, **26**(3), 519–48.
- Lundin, D., Mörk, E., and Öckert, B. (2008), 'How Far Can Reduced Childcare Prices Push Female Labour Supply?', *Labour Economics*, **15**(4), 647–59.
- Nollenberger, N., and Rodríguez-Planas, N. (2015), 'Full-time Universal Childcare in a Context of Low Maternal Employment: Quasi-experimental Evidence from Spain', *Labour Economics*, **36**, 124–36.
- Olivetti, C. (2006), 'Changes in Women's Hours of Market Work: The Role of Returns to Experience', *Review of Economic Dynamics*, **9**(4), 557–87.
- Petrongolo, B. (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.
- Paull, G. (2008), 'Children and Women's Hours of Work', *The Economic Journal*, **118**, F8–F27.
- Shaw, J. (2011), 'Fortax: UK Tax and Benefit System Documentation', Technical Report, IFS Working Papers.
- Shephard, A. (2009), *Fortax: Reference Manual*.
- Xiao, P. (2020), 'Wage and Employment Discrimination by Gender in Labor Market Equilibrium', mimeo.