

INTERNATIONAL MONETARY FUND

Are Low-Skill Women Being Left Behind? Labor Market Evidence from the UK

Era Dabla-Norris, Carlo Pizzinelli, and Jay Rappaport

WP/22/42

IMF Working Papers describe research in progress by the author(s) and are published to elicit comments and to encourage debate.

The views expressed in IMF Working Papers are those of the author(s) and do not necessarily represent the views of the IMF, its Executive Board, or IMF management.

**2022
FEB**



WORKING PAPER

IMF Working Paper
Asia and Pacific Department

Are Low-Skill Women Being Left Behind? Labor Market Evidence from the UK

Prepared by Era Dabla-Norris, Carlo Pizzinelli, and Jay Rappaport*

Authorized for distribution by Era Dabla-Norris

February 2022

IMF Working Papers describe research in progress by the author(s) and are published to elicit comments and to encourage debate. The views expressed in IMF Working Papers are those of the author(s) and do not necessarily represent the views of the IMF, its Executive Board, or IMF management.

ABSTRACT: Labor markets in the UK have been characterized by markedly widening wage inequality for low-skill (non-college) women, a trend that predates the pandemic. We examine the contribution of job polarization to this trend by estimating age, period, and cohort effects for the likelihood of employment in different occupations and the wages earned therein over 2001-2019. For recent generations of women, cohort effects indicate a higher likelihood of employment in low-paying manual jobs relative to high-paying abstract jobs. However, cohort effects also underpin falling wages for post-1980 cohorts across all occupations. We find that falling returns to labor rather than job polarization has been a key driver of rising inter-age wage inequality among low-skill females. Wage-level cohort effects underpin a nearly 10 percent fall in expected lifetime earnings for low-skill women born in 1990 relative to those born in 1970.

JEL Classification Numbers:	J21, J24, J31
Keywords:	Job polarization, Occupational choice, Life-cycle, Cohorts, Inter-generational inequality
Author's E-Mail Address:	edablanorris@imf.org ; cpizzinelli@imf.org ; jaywrappaport@gmail.com

* We thank Cristian Alonso, Vitor Gaspar, Klaus Hellwig, Rachel Ngai, Myrto Oikonomou, and participants in various IMF seminars for their comments.

Contents

1	Introduction	2
2	Relation to the literature	6
3	The data	7
4	Occupational shifts by cohorts	8
5	Empirical strategy	9
6	Results for employment propensities and wages	14
7	Decomposition exercises	18
8	Possible drivers of cohort effects	21
9	Conclusion	24
A	Additional Tables and Figures	29
A.1	Sensitivity analysis of baseline regression	30
B	Data appendix	39
B.1	Occupational categorization measures	39
B.2	Alternative occupational categorization	42
C	Alternative estimation approaches	43
C.1	Discussion of alternative estimation approaches of age, period, and cohort effects	43
C.2	Results	44

1 Introduction

Following the Great Recession, the UK labor market has been characterized by worsening economic opportunities for young workers. This development came on the back of a longer-term shift in the occupational composition of employment away from routine-intensive occupations towards low-paying manual and high-paying abstract jobs, contributing to rising wage inequality. In the last two decades, both trends have been particularly marked for young females without a college education. However, the link between these two trends and its potential impact over the course of the full careers of today’s young female workers is yet to be quantitatively assessed. In this paper, we provide a simple characterization of wage dispersion for non-college females in the UK between 2001 and 2019 through a life-cycle framework. We investigate the relative importance of employment shifts *across* occupations and changes in expected wages *within* occupations for the widening wage gap between young and prime-age non-college women. We also examine whether aggregate or cohort-specific factors underpin wage dispersion over time, quantifying their compounded effect on lifetime earnings for different cohorts.

The motivation for our analysis comes from the widening wage gap between young and older non-college females over the past two decades. Using the quarterly Labour Force Survey (LFS) for the UK, Figure 1a) plots the wage premium of prime-age (35 to 49 years) over young workers (21 to 34 years) by education (non-college vs. college educated) and gender (female vs. male) from 1998 to 2019.¹ Non-college females experienced by far the largest rise in the inter-age wage premium (approximately 16 percent) of the four groups. Notably, the widening of the gap predates the 2008 Great Recession. Moreover, Figure 1b) shows that the rising premium originated primarily from a persistent decline in the average hourly wage of the young since the mid-2000s, while the average wages of prime-age and “old” workers (aged 50-64) continued to rise and recovered from the Great Recession.²

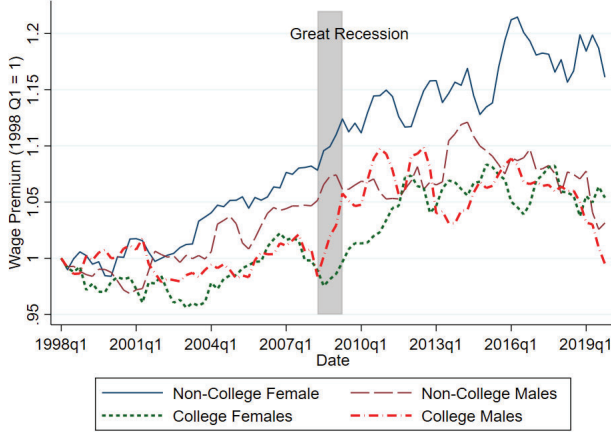
In line with the long-term aggregate “job polarization” trends identified by Goos and Manning (2007), the period 2001-2019 also featured large occupational shifts for non-college

¹We denominate the age range 21-34 as the “young”, which we compare to the “prime-age” (35-49) and the “old” (50- 64).

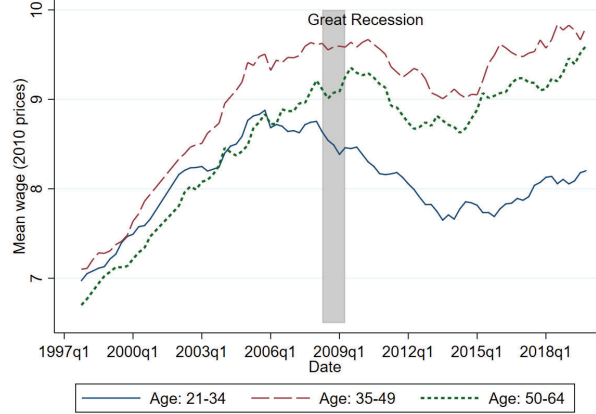
²As we show in the Online Appendix, the Great Recession also widened the unemployment gap between young and older worker. However, by 2016 the differential had reverted to its pre-recession level. Meanwhile, labor force participation rates continued along their respective long-run trends.

Figure 1: Increasing wage inequality between young and prime-age workers: 1998-2019.

(a) Wage premium of prime-age workers over young workers by education and gender



(b) Mean hourly wage by age group for non-college females



Notes: The left panel reports the ratio of the average wage of prime-age workers to that of young workers, normalized to one in 1998q1. In the right panel, wages are in GBP, chained at 2010 prices through the UK’s Consumer Price Index. The grey area represents the quarters of the Great Recession (2008q2-2009q2).
Sources: LFS and authors’ calculations.

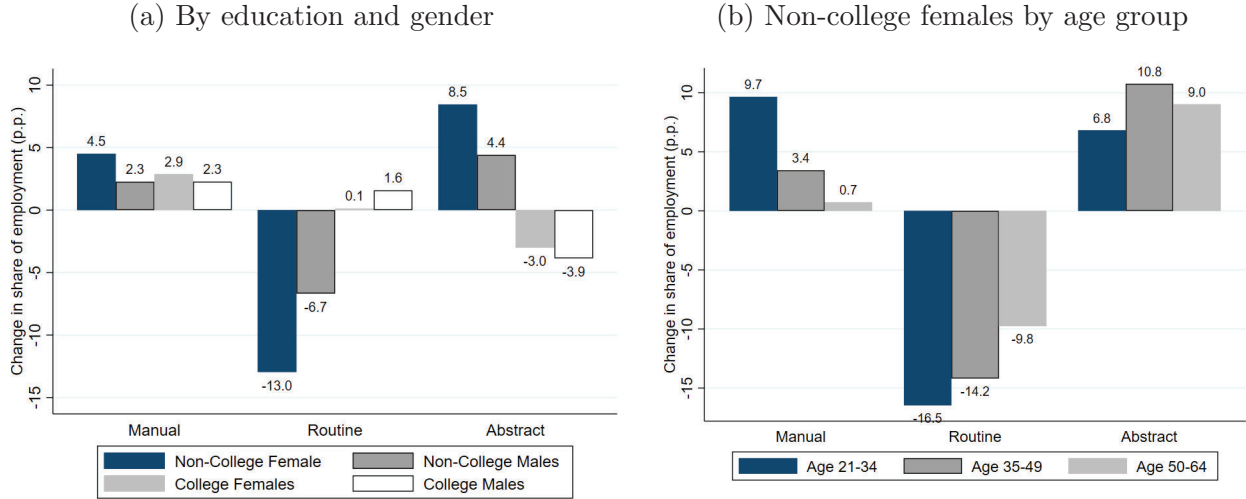
women.³ Following the literature on job polarization, starting with Autor et al. (2003), we divide occupations in three broad groups defined by the nature of the main tasks involved: manual, routine, and abstract.⁴ Importantly, the three groups are broadly characterized by an ascending order of average wages from manual to abstract. As shown in Figure 2), low-skill females experienced a marked shift of employment away from routine occupations (-13 p.p.) towards both manual (+8.5 p.p.) and abstract jobs (+4.5 p.p.). For other groups, this pattern was either less pronounced (non-college males) or absent (college workers). Moreover, the occupational shift for non-college females was heterogeneous across age groups. Figure 2b) shows that young workers experienced the steepest decline in routine jobs and the largest rise in manual jobs. For older groups the drop in routine jobs was smaller and reallocation was concentrated in abstract jobs.

For women in particular, occupational changes over the life-cycle are an important driver of wage dynamics. Goldin et al. (2017) show that women’s slower progression into high-earning industries accounts for the widening of the gender wage gap in later stages of workers’ careers. Through the lens of the life-cycle, two channels could underpin the above trends. First, the wage

³As explained in Section 3 the shorter time period compared to the wage series is due to the availability of a consistent occupational classification.

⁴More details on the classification of jobs are provided in Section 4.

Figure 2: Changes in employment shares of occupation types: 2001-2019



Notes: Each bar reports the change in the share (percentage points) of a given occupation type between 2001 and 2019 for the respective worker group by education and gender (left plot) or age group among non-college females (right plot).

Sources: LFS and authors' calculations.

gap between young and older workers can widen if recent labor market entrants do not face the same opportunities for career progression as previous generations. Younger generations may thus be permanently trapped into lower-paying occupations. However, a second channel may be at play. *Ceteris paribus*, the young may also receive lower wages than previous generations across all jobs even with similar chances of switching occupations over their careers.

We thus employ a life-cycle framework to study the importance of these two channels and examine the nexus between the shifting occupational composition and inter-age earnings inequality. Using the LFS, we estimate an age-period-cohort model of the propensity to be employed in each occupation type and the expected wage therein to disentangle inter-generational differences from aggregate trends and the life-cycle component. We conduct a series of counterfactual exercises to quantify the contribution of each component to (i) the widening wage gap between young and older female workers between 2001 and 2019 and (ii) the changes from one cohort to another in total expected earnings over workers' life-cycle.

We find that uneven patterns of job polarization across cohorts played a quantitatively limited role in the worsening labor market prospects of young non-college females. Overall, cohort effects capture a general worsening in labor market prospects for those born since 1980 through both channels mentioned above. First, post-1980 cohorts are more likely to be employed

in low-paying manual jobs and less likely to hold abstract positions compared to the 1970s generation. Second, they also receive significantly lower wages within all occupations. However, the latter channel is quantitatively larger, suggesting that the main driver of lower wages for young non-college females is falling returns to labor rather than uneven occupational shifts. Counterfactual exercises show that cohort effects contribute to an 8.6% decline in the average wage of young non-college females between 2001 and 2019 while lowering the wages of the prime-age by just 1.4% and raising that of older workers by 2.7%. These worsening labor market prospects captured by cohort effects account for approximately half the increase in wage inequality between young and prime-age non-college females over the past two decades. Furthermore, we estimate that non-college females born in 1990-1994 can expect approximately 10% lower lifetime earnings relative to those born in the 1970s.

We consider three potential trends driving these findings, for which we provide descriptive evidence. First, younger non-college female workers have moved increasingly away from routine-intensive industries towards low-paying manual-intensive ones, while sectoral reallocation has been more uniform for older workers. Second, young workers today are less likely to receive training on the job than in the past, while older workers increasingly receive training. Finally, the steep rise in university attendance by women over the past 20 years may have entailed a negative selection within the low-education pool based on unobservable skills.

The rest of this paper is structured as follows. Section 2 relates our work to the literature. Section 3 describes the data and the occupational categorization we apply. Section 4 presents descriptive evidence on changes in the occupational composition of employment of non-college females and discusses differences across cohorts. Section 5 outlines the empirical strategy. Section 6 presents the estimation of age, period, and cohort effects in employment propensities and wages. In Section 7, we carry out a set of counterfactual exercises to quantify the relevance of the cohort effects for earnings inequality across ages and for workers' expected lifetime income. Section 8 discusses potential channels driving the lower economic prospects for younger cohorts that are reflected in the estimated cohort effects. Section 9 concludes.

2 Relation to the literature

A large literature examines job polarization and heterogeneous shifts in the occupational composition of employment (see Cortes et al., 2020; Bussolo et al., 2018; Autor and Dorn, 2009, for evidence on the US and other countries). Our focus on non-college females is motivated by the unique characteristics of this group of workers. Numerous studies highlight the role of industrial and occupational composition of the economy in explaining gender differences in labor market outcomes (Goldin, 2006; Olivetti and Petrongolo, 2014, 2016; Ngai and Petrongolo, 2017). However, most of the improvement in labor market prospects for females in advanced economies in recent decades has come from growth in abstract-intensive service jobs, typically requiring an advanced set of cognitive and interpersonal skills (Cortes et al., 2021). Hence, it is not clear *a priori* that low-skill females stand to gain from the shift from routine jobs into abstract ones.⁵ In fact, in the UK, although female labor force participation has increased and the gap between average male and female wages converged over the last decades, non-college females exhibit the lowest wages in the labor force with little growth over the life-cycle (Costa Dias et al., 2021). We contribute to this research by showing that occupational patterns for low-skill females have evolved asymmetrically for younger and older generations. This asymmetry was also accompanied by a more general worsening of labor market outcomes for younger cohorts.

Our paper is also related to studies of occupational change in the UK (Goos and Manning, 2007; Salvatori, 2018; Montresor, 2019). While remaining agnostic on the causes of occupational shifts, our contribution is to quantify their role in widening the wage gap between younger and older non-college female workers over the past 20 years. We do so through the lens of the worker’s life-cycle, similar to Cristini et al. (2017) and Antonczyk et al. (2018). Rather than interpreting each age group as simply a cross-section, the life-cycle approach provides insights into the shifting occupational landscape for individual workers’ prospects. This framework can shed light on the drivers of inequality and project the impact of each cohorts’ expected earnings over their entire (future) careers.

The widening income inequality between young and older workers in recent decades has been the subject of much debate (see Cribb, 2019, for recent evidence on the UK). Several papers

⁵For instance, Brussevich et al. (2019) document that women, and especially those with lower levels of education, tend to work in jobs that are more routine-intensive in many OECD countries.

document the disproportionately adverse impact of the 2008 Great Recession on labor market outcomes for young workers (Elsby et al., 2010; Boeri et al., 2016; Hur, 2018). By framing the analysis from a cohort perspective, we show that part of the slowdown in earnings for young female workers preceded the recession. To some extent, the slowdown is thus related to long-run trends in occupational shifts and wages within occupations. In this regard, although focusing on a different group of workers, our finding is aligned with Beaudry et al. (2014), who document that the employment share of abstract occupations in the US stagnated for subsequent cohorts of college-educated workers since 2000.

3 The data

We use data from the UK Labour Force Survey (LFS), available at quarterly frequency from the mid-1990s. Our main sample includes employed workers aged 21 to 65 over the period 2001:Q2-2019:Q4, for which a consistent occupational classification is available.⁶ We focus our analysis on females with no qualifications above secondary education (A-levels or similar). Henceforth, we use the word “workers” to refer exclusively to non-college female workers unless otherwise specified.

The data is comprised of repeated cross-sections that are used to construct synthetic cohorts of workers.⁷ The provided observation weights make the sample representative of the UK population. For wages, our variable of interest is a worker’s gross hourly pay, as pre-computed in the LFS dataset. The second data source comprises the categorization of occupations based on their “task content” as proposed by Cortes et al. (2020). We merge these categorical definitions of occupations with occupational classifications for the UK. While Cortes et al. (2020) classify occupations along two orthogonal dimensions, routine vs. non-routine and manual vs. abstract, we combine the two routine groups. This yields three categories: abstract, manual, and routine.⁸ Although this grouping may appear generic, it captures well the information con-

⁶Although the LFS contains occupation and wage information since 1997, the switch from the 1990 Standard Occupational Classification (SOC1990) to the updated SOC2000 in 2001 creates a discontinuity in the analysis. Therefore, the analysis considers the period 2001:Q2-2019:Q4 for which the occupations are classified using SOC2000.

⁷See Deaton and Paxson (1994) and Attanasio (1998) for a discussion of synthetic cohort analysis.

⁸We use a cross-walk between the SOC2000 and the International Labour Organization’s ISCO08 to merge the UK data with the occupational categorization and task-content data, which are based on US classifications. See the Online Appendix for details.

tained in alternative measures of task content, such as the Routine Task Index (RTI) developed by Autor and Dorn (2013) (see the Online Appendix).⁹

4 Occupational shifts by cohorts

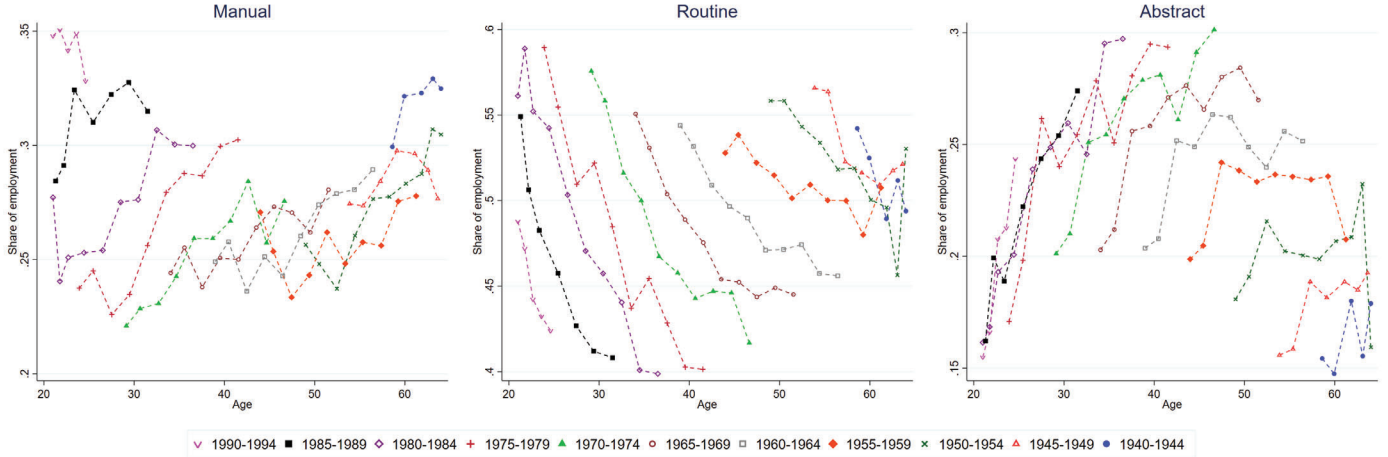
As shown above, low-skill females experienced a large widening of the wage gap between young and older workers in the past two decades concurrently with the long-term phenomenon of job polarization. However, this reallocation across occupations has occurred unevenly across age groups, with young workers shifting predominantly toward low-paying manual jobs and older workers toward high-paying abstract work.

While suggesting a possible connection between these developments, the stylized facts reported in Figures 1 and 2 do not provide a quantitative assessment of the importance of occupational shifts for the widening wage gap across ages. Furthermore, the raw differences in wages and occupational shifts do not partial out possible inherent differences in labor market prospects for different cohorts of labor market entrants which may persist through workers' entire careers. Since occupations are closely associated with wages, inter-cohort differences have implications for earnings inequality. In what follows, we thus introduce cohorts as an additional dimension of analysis.

Figure 3 presents synthetic cohort plots of the employment share in each occupation. We divide workers into 5-year cohorts based on birth year. For each cohort, we plot the share in each occupation and the average age in a given year (taking averages across 2-year windows to reduce noise). Given the time range of the data, the earliest cohort is comprised of workers born between 1940 and 1944 (i.e., Cohort 1940), while the last cohort includes those born between 1990 and 1994 (i.e., Cohort 1990). Each line pertains to a different cohort. The variation over age for a given cohort is informative of the overall life-cycle profile of a variable. At the same time, changes over cohorts for a given age shed light on aggregate changes in the economy over time, which could affect all cohorts equally. These plots reveal several patterns. Overall, the life-cycle profile is concave and hump-shaped for abstract jobs but convex for routine and

⁹A further caveat, as pointed out by Bhalotra and Fernandez (2018), is that task-based indices and categorizations constructed on US occupations may not be fully representative of the task content of occupations in another country. However, for a relatively similar economy like the UK the risk of significant differences in the nature of the occupations is limited. For a deeper analysis of the numerous combinations of skills involved in different occupations in the UK see Nedelkoska and Quintini (2018).

Figure 3: Cohort plots for females without a college education: 2001-2019



Notes: Each marker reports the share of employment in a given occupation, for a given cohort of females without a college degree at a given age. Markers report averages across two years of the LFS, from 2001-2002 until 2017-2018. Not all cohorts are present in all years as the plot is restricted to the age range 21-64.
Sources: LFS and authors’ calculations.

manual jobs. The figures suggest that workers progress into higher-paying occupations over the course of their careers.¹⁰ There are also significant differences across cohorts in the likelihood of being in manual and abstract jobs. While the share of females in abstract occupations has risen over time for workers aged 40 and above in 2019, it has remained unchanged for earlier ages. On the other hand, young workers in later cohorts are more likely to be employed in manual jobs compared to earlier cohorts. One interpretation of this finding is that the reallocation away from routine jobs has induced a process of “upgrading” for older generations but “downgrading” for younger females just entering the workforce, which may persist throughout their full careers.

5 Empirical strategy

This section describes our empirical strategy. We start from an intuitive representation of the average wage of workers within each age group \mathbb{A} , $W_t^{\mathbb{A}}$, as a function of the occupational distribution within the group and the average wages therein:

$$W_t^{\mathbb{A}} = \sum_c s_t^{\mathbb{A}c} \sum_o p_t^{co} w_t^{co}, \quad (1)$$

¹⁰While the importance of occupational mobility for wage growth is not new to the literature (see Kambourov and Manovskii, 2009b; Groes et al., 2015), the contribution of switches across types of occupations has not been focus of these studies.

where c is the index for cohorts and o is the index for occupation type. The share of cohort c within the total employed population in \mathbb{A} , $s_t^{\mathbb{A}c}$, captures the turnover of the cohorts over time within the age groups. Their employment propensity in occupation type o is p_t^{co} and their average wage therein is w_t^{co} .

We assume that both the likelihood of employment in each occupation type and the expected wage are functions of the worker’s cohort and age, and of aggregate developments over time. In what follows, we discuss how we estimate p_t^{co} and w_t^{co} as functions of these factors. We then use the results and the workers’ shares $s_t^{\mathbb{A}c}$, taken from the LFS, to decompose the shift in the wage gap between young and older non-college females over 2001-2019. Finally, we propose a similar approach to make projections of each cohort’s expected earnings over their full working lives.

Estimating age, period, and cohort effects

First, we estimate probabilities of employment in each occupation and expected wages therein as a function of age, period, and cohort effects. Our approach is similar to that of MaCurdy (1995), and Kambourov and Manovskii (2009a). The main deviation comes from first estimating a linear probability model for employment in each occupation and then a set of occupation-specific wage equations. We control for cohorts using a set of dummy variables for 5-year windows of birth years, with those born between 1940 and 1944 as the reference cohort. We model age effects through a cubic polynomial, and year effects through one-year dummies.

For each occupation, the dependent variable equals one if worker i is employed in that occupation and zero otherwise.¹¹ The baseline specification of this linear probability model for a worker i , from cohort c , at time t is

$$\textit{Employed in } \textit{Occ.}_{ict} = \alpha_0 + \alpha_1 \textit{age}_{ict} + \alpha_2 \textit{age}_{ict}^2 + \alpha_3 \textit{age}_{ict}^3 + \gamma_t + \delta_c + \varepsilon_{ict} \quad (2)$$

where $\textit{Occ.}$ is routine, abstract, or manual occupation, γ_t is the yearly period effect, and δ_c is the cohort effect.¹² The equation is estimated through Ordinary Least Squares (OLS) using heteroskedasticity-robust standard errors.

¹¹We exclude the unemployed and those out of the labor force. The sample size is 922,952.

¹²We also included in the specification a set of dummies for quarter of the year.

Using the subset of LFS workers who report their labor earnings, we use the same specification as in (2) to decompose (log) wages within each occupation into age, period, and cohort effects:¹³

$$\text{Log}(w_{ict}) = \alpha_0 + \alpha_1 \text{age}_{ict} + \alpha_2 \text{age}_{ict}^2 + \alpha_3 \text{age}_{ict}^3 + \gamma_t + \delta_c + \varepsilon_{ict} \quad (3)$$

Counterfactual exercises: looking back and looking ahead

Considered individually, each regression in (2) and (3) can uncover heterogeneity across cohorts and time for a given occupation type. In practice, both cohort and period effects represent vertical shifts in the age polynomials. Cohort effects capture permanent differences across generations of workers in the propensity to be in each occupation and the average wage therein. Year effects encompass both cyclical fluctuations and long-term trends specific to the occupation type and wages that affect all workers equally regardless of age or year of birth. However, a joint framework is needed to draw comprehensive conclusions on inequality across age and generations. We thus carry out a set of counterfactual exercises to gauge the quantitative relevance of cohort and year effects for inter-age earnings inequality among low-skill females. We first examine the drivers of wage growth over the period 2001-2019. Using the same tools, we then take a forward-looking approach to project each cohort’s average lifetime earnings.

Decomposing wage growth between 2001 and 2019

The predicted share of workers in each occupation \hat{p}_t^{Aco} and the expected wage \hat{w}_t^{Aco} are computed using the estimated OLS coefficients from equations (2) and (3), using the average age of workers from cohort c in year t , age_t^c , together with the appropriate cohort and year dummies. For clarity of comparison with the counterfactual, it is useful to explicitly express them as functions of age, cohort, and year effects: $\hat{p}_t^{\text{Aco}} = p(\text{age}_t^c, \text{cohort}^c, \text{year}_t)$, and $\hat{w}_t^{\text{Aco}} = w(\text{age}_t^c, \text{cohort}^c, \text{year}_t)$.

There are several drivers of variation in W_t^{A} over time and across age groups. To examine their contribution, we construct counterfactual series for W_t^{A} . First, the composition of the population, represented by the vector of s_t^{Aco} s, changes as new cohorts enter the labor force. Second,

¹³The sample sizes are 64,525 for manual, 129,853 for routine, and 54,652 for abstract jobs. Because only a subset of workers report information on earnings, these three samples do not add up to the sample size from the employment regression.

aggregate developments may affect different age groups heterogeneously due to their different exposure. For instance, a fall in wages within routine jobs would affect disproportionately age groups that have a higher share of routine occupations.

Based on these considerations, the counterfactual series are as follows:

a . Only cohort effects

$$W_t^{\Delta, COHORT} = \sum_c s_{2001}^{\Delta c} \sum_o p(\text{age}_{2001}^c, \widetilde{\text{cohort}}_t^c, \text{year}_{2001}) w(\text{age}_{2001}^c, \widetilde{\text{cohort}}_t^c, \text{year}_{2001}) \quad (4)$$

b . Only year effects

$$W_t^{\Delta, YEAR} = \sum_c s_{2001}^{\Delta c} \sum_o p(\text{age}_{2001}^c, \text{cohort}^c, \text{year}_t) w(\text{age}_{2001}^c, \text{cohort}^c, \text{year}_t) \quad (5)$$

where $\widetilde{\text{cohort}}_t^c$ is a counterfactual cohort dummy constructed to allow for cohort turnover while all other components remain fixed to their 2001 levels.¹⁴

Counterfactuals and life-cycle projections

The second set of counterfactual exercises compute each cohort's expected lifetime earnings and assess the contribution of year and cohort effects. While the previous exercise focused on understanding the drivers of wage growth over 2001-2019, the spirit of this analysis is forward-looking. Making assumptions regarding occupational shifts and wage growth pre-2001 and post-2019, the lifetime wage of workers in cohort c can be computed as follows:

$$Y^c = \sum_{t=t_0^c}^{T^c} \beta^t \sum_o \widehat{p}_t^{co} \widehat{w}_t^{co} \quad (6)$$

where \widehat{p}_t^{co} and \widehat{w}_t^{co} are as defined above, $0 < \beta < 1$ is a discount rate, and the index t encompasses the cohort-specific working years from t_0^c to T^c . In computing \widehat{p}_t^{co} and \widehat{w}_t^{co} , we apply the 2001 year effects for all preceding years and the 2019 effect for all subsequent years, thus assuming the change to be permanent. For each cohort and worker group, we compute Y^c and two counterfactuals in which cohort and year effects are excluded, respectively. For the

¹⁴For example, if in 2001 the 20 year olds are from Cohort 1970, in 2011 they are from Cohort 1980. However, we assume that in 2011 Cohort 1980 will have the same population share that Cohort 1970 had in 2001.

cohort exercise, we set all cohort effects equal to Cohort 1940. For the year counterfactual, we set all year effects equal to their 2001 value. We assume that each year is discounted at the rate of 2.5% (i.e., $\beta = 0.975$).

Discussion of age, period, and cohort effects

It is worth providing further discussion of what each term in the equations (2) and (3) represents, particularly for the former, which is less common in the literature. The age polynomials capture the employment share in each occupation over a worker’s age that is common to all cohorts and constant over time. They reflect workers’ natural progression into different occupations as they grow older. The year and cohort effects represent upward or downward shifts in the age profiles that are common to all workers in a given year or to workers belonging to the same cohort, respectively, relative to the reference year (2001) and the reference cohort (1940). In other words, year effects trace the aggregate shifts owing to changes in the economy over time rather than the generational turnover of workers. In the context of occupations, year effects may reflect shifts in industrial composition of the economy as well as the decline of specific occupations that are replaced by technology (e.g., blue-collar and clerical occupations). Cohort effects capture aggregate shifts in occupational composition driven by inherent “unobserved” differences between successive generations of workers that persist throughout a generation’s entire career. In order to be generation-specific, these effects likely capture features of the workers’ human capital formation prior to beginning their career (e.g., changes in the educational system), or the conditions in which they entered the labor force (e.g., the scarring effect of graduating during a recession).

It is important to note that equations (2) and (3) do not contain interaction terms between period effects and the age polynomials. In other words, the equations do not aim to capture heterogeneity across age groups in the temporary impact of aggregate events. As a result, this specification does not capture the potentially heterogeneous impact of business cycle fluctuations on workers of different ages, and only captures the average change over time in the employment propensity of each occupation.¹⁵ This choice is motivated by our interest in understanding job polarization as a medium-to-long-run phenomenon.

¹⁵In this set-up, identifying true cohort effects requires observing each cohort through a long enough period for the temporary effect of a specific recession or recovery to wane.

To overcome the collinearity issues in estimating all three sets of effects, we impose some parametric assumptions. First, the cohorts are comprised of 5-year intervals, so that they are not equal to year minus age. This imposes the implicit restriction that workers in adjacent birth years are equally affected within each 5-year window. Second, we model age through a cubic polynomial rather than a set of dummy variables, thus restricting somewhat the shape of the resulting life-cycle paths.

While the above explanation of the age, period, and cohort effects is intuitive, a large literature has shown the challenges of estimating these effects separately, given their perfect collinearity, and the non-trivial implications of seemingly innocuous parameter restrictions (see Fannon and Nielsen, 2019, for a recent discussion).¹⁶ Identification is thus conditional on imposing enough zero restrictions on the values the coefficients, assuming a parametric structure, or recovering a given “transformation” of the effects. While numerous such approaches have been proposed, the tenability of given a method intrinsically depends on the context of the empirical question itself. For this reason, in the Online Appendix we discuss alternative approaches to our baseline specification, which we undertake to check the robustness of our results.

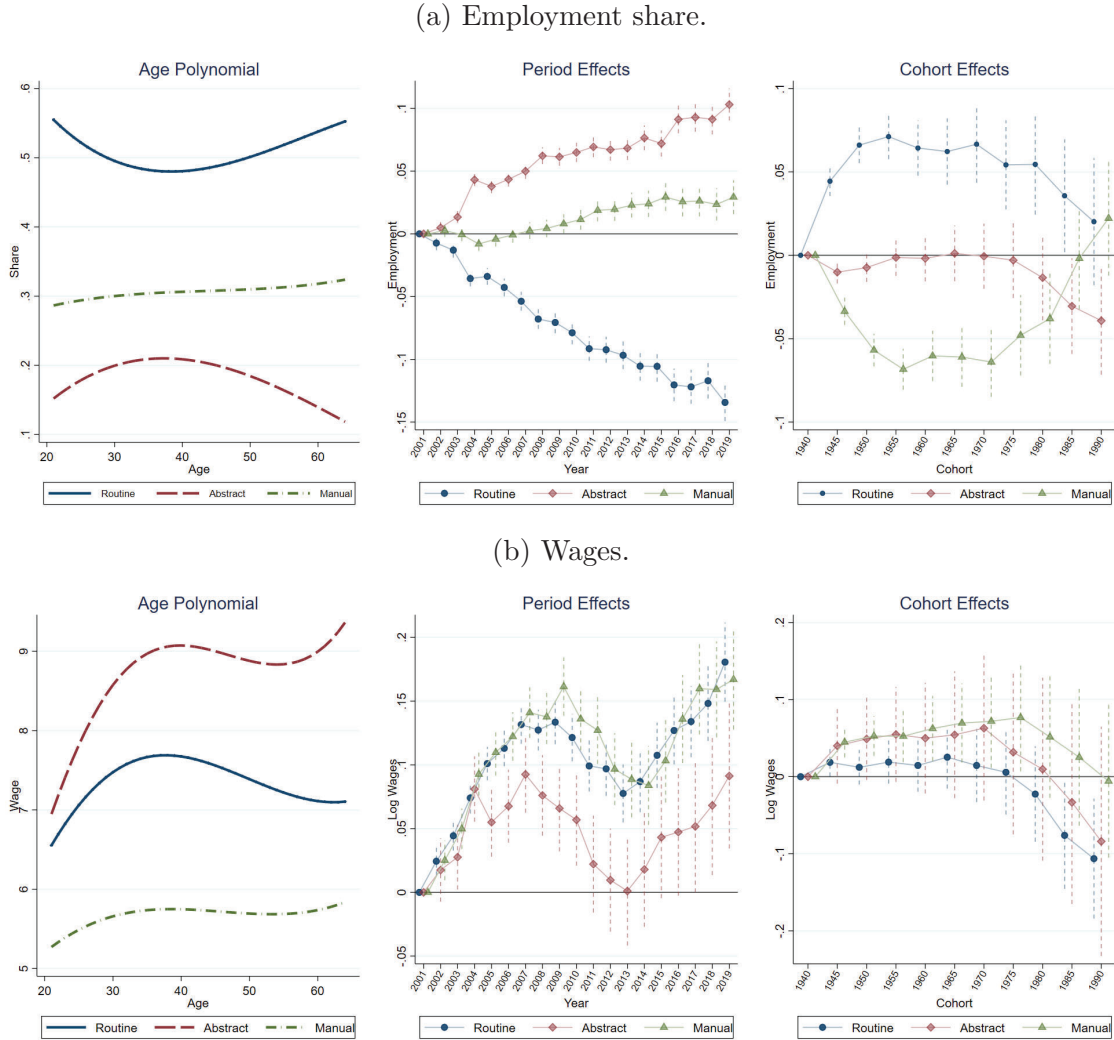
6 Results for employment propensities and wages

The estimation results of (2) and (3) are presented in Figure 4. Each subplot in the figure reports the estimated age polynomials, period, or cohort effects of employment propensity (top panel) or wages (bottom panel) in the three occupation types.

Employment shares The age polynomials trace a mild U-shape for employment in routine jobs, with the lowest point reached just before age 40, followed by a gradual rise. Overall, the polynomial remains in the 0.5-0.6 interval for most of the age range, implying that, abstracting from aggregate developments over time and cohort turnover, routine jobs are the most prevalent for non-college females. This path is mirrored by the hump shape of abstract employment, although with a much lower average share over the life-cycle. The share of manual jobs remains almost flat around 0.3. The period effects reflect a fall in routine jobs of around 13 percentage

¹⁶More specifically, as Fannon and Nielsen (2019) note, the unidentifiability only relates to the linear component of the effects, while the higher-order components are separately identifiable.

Figure 4: Age polynomials, period effects, and cohort effects.



Notes: The top panels report the estimated age polynomials, the yearly period effects, and the 5-year cohort effects of employment shares in each occupation type as in Eq. (2). The bottom panels report the same figures for the expected wages within each occupation as in Eq. (3). For period and cohort effects, the standard errors reported comprise the 95% confidence interval. Age polynomials for wages are reported as the exponential of the estimate polynomial for log wages.

Sources: LFS and authors' calculations.

points between 2001 and 2019, compensated by a 10 percentage point rise in abstract jobs and a small rise in manual jobs. These year effects thus capture a large aggregate shift for all non-college women. The cohort effects provide a perspective of heterogeneity across cohorts along this aggregate trend. They trace non-linear paths, implying that the largest differences in occupational composition occur between the most recent cohorts and those born in the 1960s and 1970s. For instance, while the “Baby Boomer” cohorts were progressively less likely to enter a manual job compared to previous cohorts, the trend reversed when “Generation X” (born

between 1965 to 1979) entered the labor force. The youngest “Millennial” females (born 1990-1994) without a college degree are almost 10 percentage points more likely to be in a manual occupation compared to workers born in the mid-1960s. Combined with the year effects, this reversal implies that, while on the whole abstract jobs have risen for low-skill females (via period effects), this trend was milder for younger generations, who experienced a sharper decline in routine jobs mirrored by a larger rise in manual jobs.

Wages The age polynomials for wages suggest that, on average, abstract jobs pay more than routine jobs, while manual ones are the lowest-paying at every age. Moreover, abstract jobs show a steeper wage progression from early ages to workers’ prime years compared to routine jobs, while manual jobs show virtually no progression throughout one’s career. Combined with the changing employment propensities discussed above, these polynomials suggest that a significant component of earnings growth over a worker’s career comes from switching from manual jobs to higher-paying routine and abstract occupations. The year effects show a common business cycle fluctuation for all occupations, with positive wage growth until 2007 and steep decline during the Great Recession. Routine and manual jobs also exhibit a positive trend throughout the period so that in 2019 wages are more than 15% higher than in 2001. Abstract jobs only show moderate growth throughout the period, with the 2019 coefficient suggesting an increase in wages of only 9% relative to 2001.

The cohort effects suggest that non-college females born since the late 1970’s have experienced a pronounced fall in the permanent component of their real wages across all occupation types. For instance, those born in 1990-1994 on average receive wages that are 14% lower than those of workers born in 1965-1969 within all occupations.¹⁷

¹⁷Although standard errors are larger compared to the employment regressions, many coefficients are statistically significant at the 95% level. The confidence intervals for cohort effects tend to increase for younger cohorts. This is due to the choice of Cohort 1940 as the reference group. This cohort has a relatively smaller sample size, as it was only followed for a few years after 2001. The relative wages of other cohorts, especially when also small in sample size, are therefore noisily estimated compared to the reference group. Choosing a later cohort as the reference group, e.g. 1970, which is present through the whole period and hence has a larger sample size, would lower the standard errors. Additionally, in the Online Appendix we report the p-values of statistical tests of the null hypothesis that each coefficient is equal to that of Cohort 1970. For cohorts 1980 and later, the null hypothesis of equality with the coefficients for Cohort 1970 is rejected in almost all regressions at the 10% significance level or lower.

Sensitivity analysis

In the Online Appendix we present the results from a series of robustness checks for the estimations of (2) and (3) along multiple dimensions.

Controls, changes in sample First, we use an alternative approach to categorize occupations into manual, routine, and abstract. Second, we change the sample to include foreign-born workers or exclude self-employed workers. Finally, we add a set of controls for individual worker characteristics that may affect employment propensity and wages. These are marital status, number of children, and region of residence for both equations, and part-time, self-employment, and contract type for (3) only. The results hold in all cases, with only minor quantitative changes.

Expanding the time sample For the main results we choose the period 2001:Q2-2019:Q4 because the LFS provides employment information using a consistent occupational classification, the SOC 2000, over this period. In order to extend the task-based categorization of occupations backward in time we use a set of crosswalks, as described in the Online Appendix, that bridge between the SOC 2000, the SOC 1990 (for the years 1993-2000), and the preceding KOS. We extend the sample back until 1986 using the quarterly files from 1993 to 2001 and the yearly datasets for the preceding years. This extension allows us to add two preceding cohorts to the sample: Cohorts 1930 and 1935. However, the bridging across classifications creates small but visible breaks in the aggregate share of workers in each occupation type.

Since the LFS does not contain wage data before 1997, we do not estimate the wage equation for this check. Incorporating the extra cohorts shows some accentuated inter-cohort differences. The hump-shaped paths of routine and manual job propensities are even more pronounced when comparing the 1970s cohorts to those born in 1930.

Quantile regressions For wages we inspect whether the age polynomials and the year and cohort effects vary across the wage distribution. To this end, we use a quantile regression approach similar to Gosling et al. (2000) and Antonczyk et al. (2018) to estimate the parameters at the 10th, 25th, 50th, 75th, and 90th percentiles. The results for period and cohort effects hold broadly across the wage distribution within each occupation type.

Alternative estimation methods We further show the robustness of the results to alternative estimation methods. These are the maximum entropy method by Browning et al. (2013), the intrinsic estimator of Yang et al. (2008), and the use of the relative price of IT equipment and productivity as proxies for period effects in employment and wages, respectively. We leave a detailed discussion of the alternative methods to the Online Appendix.

7 Decomposition exercises

In this section, we present the results of the exercises outlined in Section 5. We first examine the drivers of the rising wage gap between young and older female workers over 2001-2019. Subsequently, we use the life-cycle framework to study the cumulative impact of cohort and year effects on workers' expected career earnings.

Table 1 compares the empirical change in the wage gap from the LFS with that implied by the estimated regression models, combined as in (1). The last two columns of Table 1 show that the regressions explain 50 % and 60% of the increases in the wage gap between young and prime-age and old workers, respectively, that are observed in the LFS. The various rows of the table decompose the changes in wages between 2001 and 2019 into cohort and period effects, further breaking them down into the contribution of employment shares in each occupation type and wages therein. The rightmost columns of the table show that cohort effects contribute the most to the widening gap across age groups that is explained by the regression model: 77% (0.56 out of 0.73) for the young-prime-age gap and 85% for the young-old gap (0.89 out of 1.04). Furthermore, most of this contribution is driven by cohort effects for wages *within* each occupation, while the remainder (between a quarter and a fifth) is driven by the shares of employment in each occupation.

Figure 5 provides a graphical understanding of the results from Table 1. The solid black line plots the average wages from the LFS for each age group, while the grey circles plot the predicted wage computed in (1) at annual frequency. While the best-fit prediction closely tracks the empirical series in general, the main deviations occur during the years following the Great Recession, where wages of young females are overpredicted and that of older workers underpredicted.¹⁸

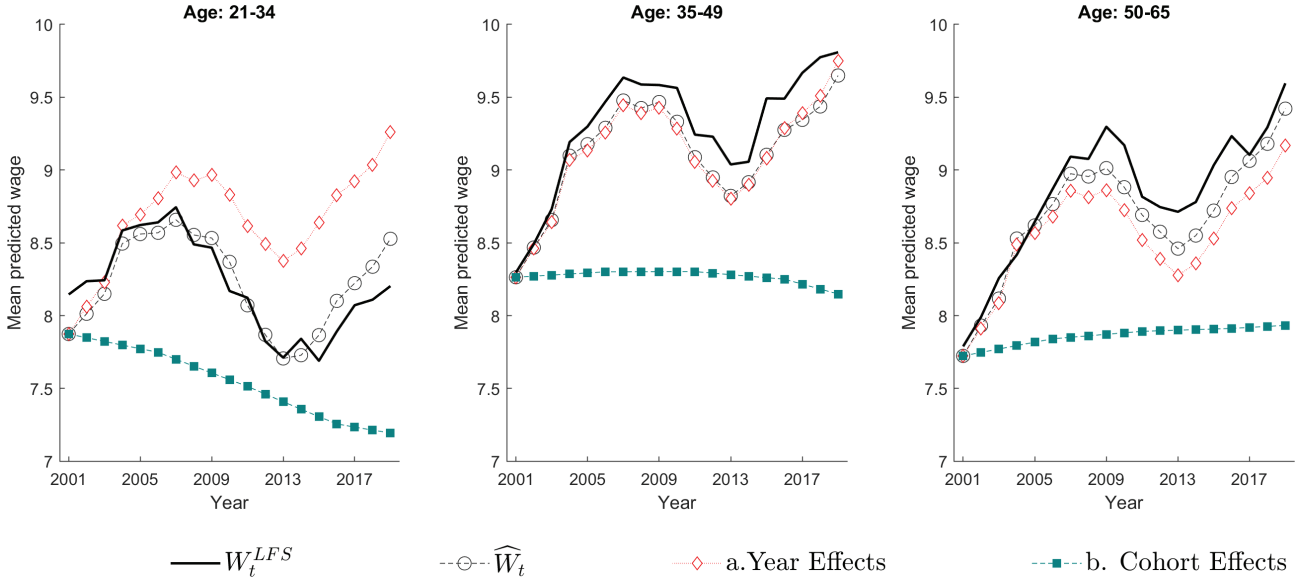
¹⁸This implies that the cyclical component of wages is larger for younger workers and has resulted in a more

Table 1: Change in Wage Gap of Young Workers Relative to Prime-Age and Old Workers: Actual Values and Counterfactuals

	Mean Wages									Δ Gaps	
	W^{21-34}			W^{35-49}			W^{50-64}			W^{35-49}	W^{50-64}
	2001	2019	% Δ	2001	2019	% Δ	2001	2019	% Δ	$-W^{21-34}$	$-W^{21-34}$
LFS	8.15	8.20	0.7%	8.30	9.81	18.2%	7.79	9.60	23.2%	1.45	1.75
APC Model	7.87	8.53	8.3%	8.26	9.65	16.8%	7.72	9.42	22.0%	0.73	1.04
Cohort Effects		7.20	-8.6%		8.15	-1.4%		7.93	2.7%	0.56	0.89
Shares		7.74	-1.7%		8.24	-0.3%		7.81	1.1%	0.11	0.22
Wages		7.31	-7.2%		8.17	-1.2%		7.86	1.7%	0.47	0.70
Period Effects		9.26	17.6%		9.75	18.0%		9.17	18.7%	0.10	0.06
Shares		8.00	1.6%		8.44	2.1%		7.91	2.4%	0.05	0.06
Wages		9.22	17.1%		9.66	16.9%		9.06	17.3%	0.05	-0.01

Note: The first row reports the average wage of the age groups 21-34, 35-49, and 50-64 in 2001 and 2019 from the LFS, and the percentage change between the two years and the change in the wage gaps between the first group and the other two. The second row reports the respective values predicted by the estimated set of age-period-cohort regressions. The other lines report the predicted values explained by cohort and period effects in employment shares and wages. Values are in GBP deflated to 2010 prices with the UKCPI, and percentage changes when specified.

Figure 5: Decomposition of average wage by age group into period and cohort effects.



Notes: The solid black line reports the average wage of each age group in the LFS at yearly frequency. The black circles report the predicted wages computed from the set of age-period-cohort regressions, as explained in the text. The red diamond markers report counterfactual wage series where only year effects vary over time. The green squares report the counterfactual series where only cohort effects vary over time.

Sources: LFS and authors' calculations.

The remaining lines in Figure 5 plot counterfactual series, decomposing the predicted average wage as described in Section 5. The contribution of the year effects (red diamonds),

severe wage downturn after 2008. That is not captured by the regressions in (2) and (3), where the year effects are not interacted with age. This also suggests that our estimation does not confound differences in cyclical behavior (which drive transitory fluctuations in wage inequality) with cohort effects (which are permanent).

which capture cyclical fluctuation in wages, is similar across age groups. The cohort effects (green squares) drive diverging trends across the three groups, with young non-college females experiencing a fall in average wages, prime-age workers' wages remaining almost unchanged, and those of old workers experiencing a mild rise. These diverging paths reflect the widening wage gaps driven by cohort effects reported in Table 1.

Figure 6 presents the results for the lifetime earnings exercise. In the left panel, for a given cohort, a positive blue bar indicates that the cohort effects (relative to Cohort 1940) contribute positively to lifetime earnings. The figure shows that cohort effects have contributed positively to workers' lifetime earnings over successive generations until Cohort 1980. That is, younger cohorts of low-skill females have higher earnings prospects over the course of their entire careers. However, this contribution has been non-monotonic. Since Cohort 1965, each successive cohort exhibits a lower effect compared to the preceding one, with the effect turning negative for those born since 1985.

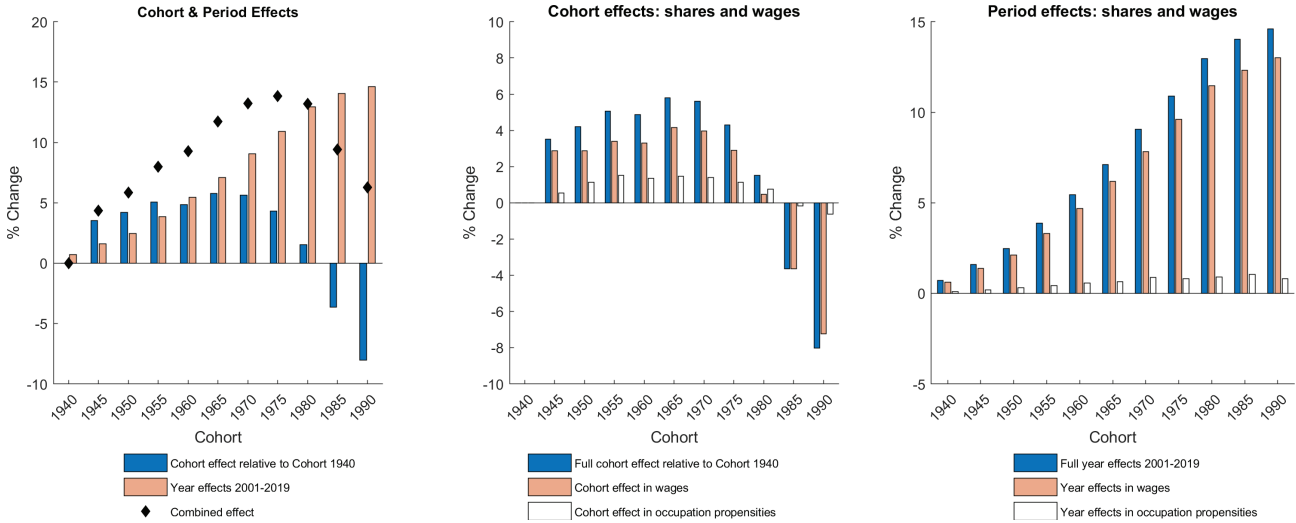
A similar intuition applies to the larger positive contribution of period effects, which are monotonically increasing between 2001 and 2019. Assuming that the 2019 effect persists in future years implies that the most recent cohorts benefit from the aggregate growth in wages throughout the rest of their careers.¹⁹

Finally, adding together the cohort and period effects, the black diamonds report their combined impact on lifetime earnings. These show that the negative cohort effects for female workers born after 1975 mitigate substantially the aggregate growth in earnings coming from the period effects.

As cohort and period effects enter into both \widehat{p}_t^{co} and \widehat{w}_t^{co} , we can also disentangle the employment propensity and wage channels by computing counterfactual lifetime earnings that exclude only cohort or period effects in either occupational shares or wages. The center plot of Figure 6 shows that the contribution of cohort effects on earnings is mostly driven by those of wages within each occupation. However, the likelihood of being employed in different occupations still accounts for about one third of the total positive contribution until Cohort 1975, reflecting the fall in the propensity of employment in manual jobs. Importantly, effects on wages within occupations almost entirely drive the fall in lifetime earnings between Cohort 1975 and

¹⁹Assumptions about wage growth in future years are also crucial for this result. Alternative scenarios that assume positive wage growth after 2019 would mitigate the total contribution of year effects for younger cohorts.

Figure 6: Contribution of cohort effects (relative to Cohort 1940) and year effects (2001-2019) to the present value of lifetime earnings by cohort.



Notes: The blue bars represent the % contribution of the cohort effect (relative to Cohort 1940) for each cohort's average lifetime earnings. The orange bars represent the contribution of the year effects (relative to a fixed 2001 effect) for each cohort's average lifetime earnings.

Sources: LFS and authors' calculations.

Cohort 1990. This implies that, among the post-1975 cohorts, the rise in lower-paid manual jobs is quantitatively less important than the fall in expected wages in all occupations. In other words, as wages fall across the occupational spectrum, progressing into other occupations over the career has a less significant effect on lifetime earnings.

The right panel of Figure 6 reports the same exercise for the period effects. Once again, the effects on wages is found to be the main driver. The aggregate shift towards abstract and manual jobs, on the other hand accounts for about a tenth of the increased lifetime earnings prospects for each cohort.

8 Possible drivers of cohort effects

In this section, we discuss other trends over the past 20 years that may be linked to the cohort effects we observed for recent generations of non-college females.

First, the large increase in university attendance, especially by females, has spurred a debate about the possible impact of sorting along unobserved worker skills on the wages of college and non-college workers. As the cost of (returns to) a university education falls (rise), more individuals who in the past would not have attended university now do so. A sorting-

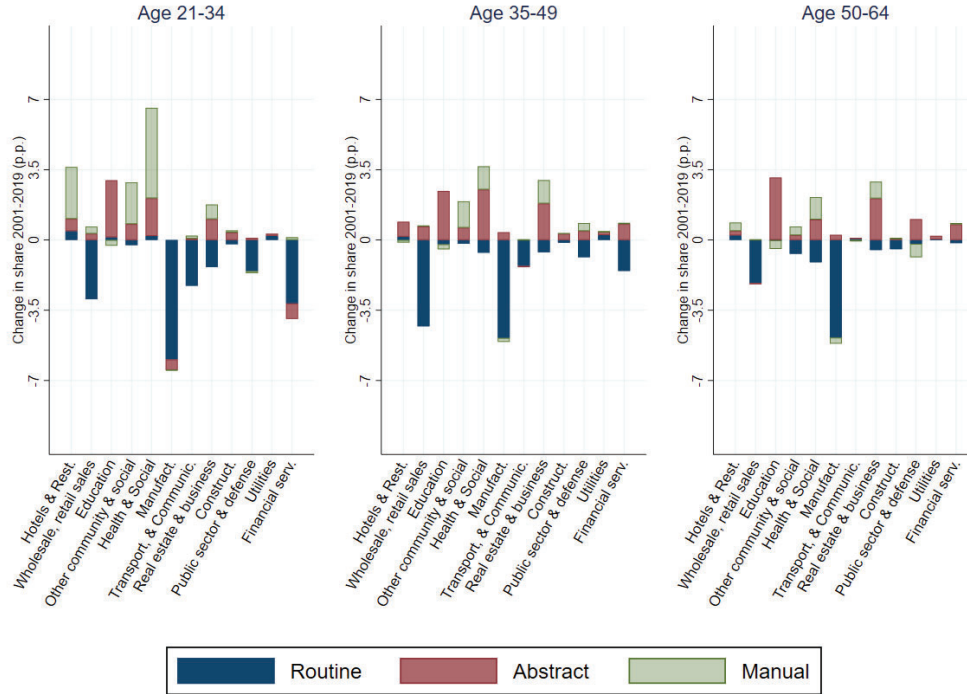
type model with unobservable skills would predict that these “marginal graduates” belong to the higher end of the unobserved skill distribution, thus entailing a fall in the average wages of the non- college group. However, empirical support for this hypothesis remains mixed. While Belfield et al. (2018) and Cavaglia and Etheridge (2020) find evidence suggesting sorting dynamics connected to the rise in university attendance in the UK over the past decades, Blundell et al. (2016b) do not find strong evidence when testing the predictions of a sorting model between the mid-1990s and 2014.

Second, routine-replacing technical change could have a differential impact on labor reallocation across sectors. Grigoli et al. (2020) find that automation of tasks where labor is easily substitutable by capital weighed on the labor force participation rates of prime-age female workers for a large sample of European countries. For the UK, the data also suggests that reallocation of low-skill workers across industries took place unevenly for young and older non-college women. Figure 7 reports the change in the employment share by industry across age groups, with industries sorted based on their median wage in 2001 from the lowest to the highest. The figure shows that for prime-age and older non-college females, the decline in routine jobs was concentrated in sectors with relatively low wages, such as trade and manufacturing. For younger workers, routine jobs declined across a larger number of industries, some of which are characterized by relatively higher median wages. Furthermore, the reallocation of employment for young female workers has been concentrated in relatively low-wage industries such as health and elderly care, driven in part by population ageing, while it has been more widespread across sectors for older workers. OECD (2020) reports that non-college women in the mid-2010s were more likely to work in low-skill occupations than in the mid-1990s even compared to similarly educated men.

Third, changes in family formation and numerous reforms to the tax and welfare system that took place in 1990s and 2000s, with major increases to in-work benefits, or tax credits could be a potential explanatory factor.²⁰ These reforms took place over the life-cycle of individuals

²⁰The 1990s-2000s period saw numerous reforms to the specific parameters determining entitlement to benefits and tax liabilities. The most significant was the sequence of reforms to the benefits of families with children that occurred between the fall 1999 and April 2002, which introduced the means-tested Working Families Tax Credit (WFTC) and changed the Income Support (IS) benefits for low-income families. The WFTC reform substantially increased the maximum benefit award both directly and through increases in support for childcare. The subsidy for working mothers involved a 25 percent rise (in constant wage levels) in the maximum award for single mothers of one child and a drop in the withdrawal rate from 70 to 55 percent.

Figure 7: Change in employment shares of manual, routine, and abstract occupations by industry and age group between 2001 and 2019



Notes: Each plot reports the change in the employment share of a given industry for non-college women, by occupation type. Industries are sorted in ascending order based on the median hourly wage in 2001. Industries are classified using the UK SIC 1992 classification. Agriculture, fishing, mining, work in international organizations are excluded.

Sources: LFS and authors' calculations.

of different cohorts at different ages, differentially affecting their returns to work. In the last three decades, the age profile of childbearing in the UK has shifted to the right, with fewer children and later births in women's lives. At the same time, the share of single mothers is noticeably high in the UK (ranking 4th in the OECD), particularly among non-college females. Using data for the period 1991-2008, Blundell et al. (2021), show that women with less than high school qualifications tend to have at least one child by age 23 compared to 4 percent of university-educated women. The dip in employment rates of women that typically occurs in the middle of their working lives due to childbearing happens earlier and is more pronounced for the lower educated. The same period is accompanied with an increase in part-time hours, particularly for those with high school qualifications and less. Blundell et al. (2016a) show that the increase in tax credits in the late 1990s, while inducing many non-college mothers into work, did not affect their wages and employment in the long term because their design encouraged

part-time work. This effect was particularly pronounced for single mothers. Part-time workers are generally paid less than their full-time counterpart, invest less in training, and have limited scope to transition to full-time permanent employment (OECD, 2020), resulting in lower overall earnings progression.

Fourth, as in other OECD countries, the UK experienced an increasing share of the workforce entering into non-standard work during the period under study, exemplified by the rising percentage of workers who are self-employed or on flexible and temporary contracts that have no minimum number of work hours provisions (zero-hours contracts). Temporary and zero-hours contracts are more common in low-paying, low-skilled sectors, and for part-time work, and usually given to younger workers. For instance, more than a third of workers on zero-hours contracts in the UK were under the age of 24 in 2017 and almost half of all zero-hours contracts were in low-skilled occupations (OECD, 2017).

Finally, the declining bargaining power of workers could be a potential candidate. Young, low-skilled women could, for instance, have a lower bargaining power compared to their employer as they are less mobile, particularly single mothers. Coudin et al. (2018) show that women's income falls after childbirth compared to men because women tend to self-select into firms offering flexible employment near their living place, granting them a lower wage bargaining power. This could potentially reinforce a low-wage trap for single parents.

9 Conclusion

This paper uses a life-cycle framework to examine the nexus between occupational shifts and inter-age wage inequality for non-college female workers in the UK over 2001-2019. The life-cycle approach provides an intuitive lens to understand the drivers of structural change and derive quantitative estimates of its impact on the labor market prospects of distinct cohorts.

The results unveil a marked worsening of labor compensation for young non-college women but suggest that job polarization played only a limited role. Estimating age, period, and cohort effects for the probability of employment in routine, abstract, and manual occupations, and the expected wages therein, we find marked changes in the labor market prospects of different generations. Cohorts born between 1940 and 1970 saw a progressive fall in the likelihood of being employed in low-paying manual occupations and a rising share of routine jobs. However,

this trend reversed with those born since 1980, with a steep rise in manual employment and a fall in routine occupations. Aggregate job polarization, captured by period effects, further accelerated the decline in routine jobs. Cohort effects from the wage equations point to a slowdown in wages for non-college female workers born since 1980 in all types of occupations. Occupational shifts contribute to a significant portion of the growth in lifetime earnings between Cohort 1940 and Cohort 1970, but their role is very modest for recent cohorts. Quantitatively, the within-occupation decline in wages for recent cohorts explains most of the aggregate wage gap between young and older non-college workers since 2001.

Further research is needed to better understand the drivers of the widening inter-age wage gaps. Structural models could provide insights on the existence and quantitative relevance of competing channels such as the negative selection among the non-college educated labor force, the role of labor market training, flexible employment contracts, and family policies. From a policy perspective, our results also raise questions about the long-term consequences: today's young non-college workers face lower lifetime earnings and will reach retirement with lower savings while also facing longer life expectancies. Furthermore, their limited wage and occupational progression could exacerbate gender inequality and stunt aggregate productivity.

References

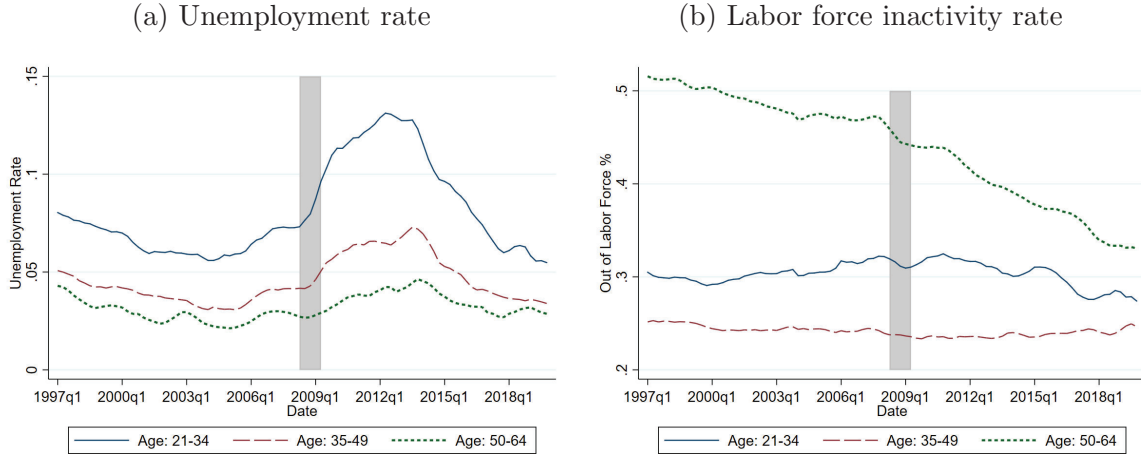
- ACEMOGLU, D. AND D. AUTOR (2011): “Skills, Tasks and Technologies: Implications for Employment and Earnings,” in *Handbook of Labor Economics*, ed. by O. Ashenfleter and D. Card, Elsevier, vol. 4b, 1043–1171.
- ANTONCZYK, D., T. DELEIRE, AND B. FITZENBERGER (2018): “Polarization and Rising Wage Inequality: Comparing the U.S. and Germany,” *Econometrics*, 6, 1–33.
- ATTANASIO, O. (1998): “Cohort Analysis of Saving Behavior by U.S. Households,” *Journal of Human Resources*, 33, 575–609.
- AUTOR, D. AND D. DORN (2009): “This Job Is ”Getting Old”: Measuring Changes in Job Opportunities Using Occupational Age Structure,” *American Economic Review*, 99, 45–51.
- AUTOR, D. H. AND D. DORN (2013): “The Growth of Low-Skill Service Jobs and the Polarization of the U.S. Labor Market,” *American Economic Review*, 103, 1553–1597.
- AUTOR, D. H., F. LEVY, AND R. MURNANE (2003): “The Skill Content of Recent Technological Change: An Empirical Exploration,” *The Quarterly Journal of Economics*, 118, 1279–1333.
- BEAUDRY, P., D. A. GREEN, AND B. M. SAND (2014): “The Declining Fortunes of the Young Since 2000,” *American Economic Review: Papers & Proceedings*, 104, 381–386.
- BELFIELD, C., J. BRITTON, F. BUSCHA, L. DEARDEN, M. DICKSON, L. VAN DER ERVE, L. SIBIETA, A. VIGNOLES, I. WALKER, AND Y. ZHU (2018): “The Impact of Undergraduate Degrees on Early-Career Earnings,” Institute of Fiscal Studies Research Report.
- BHALOTRA, S. AND M. FERNANDEZ (2018): “Women’s Labor Force Participation and the Distribution of the Gender Wage Gap,” *IZA Working Paper*, 11640.
- BLUNDELL, R., M. COSTA-DIAS, D. GOLL, AND C. MEGHIR (2021): “Wages, Experience, and Training of Women over the Life Cycle,” *Journal of Labor Economics*, 39, S275–S315.
- BLUNDELL, R., M. COSTA DIAS, C. MEGHIR, AND J. SHAW (2016a): “Female Labor Supply, Human Capital, and Welfare Reform,” *Econometrica*, 84, 1705–1753.
- BLUNDELL, R., D. A. GREEN, AND W. JIN (2016b): “The UK Wage Premium Puzzle: How did a Large Increase in University Graduates Leave the Education Premium Unchanged?” *IFS Working Paper*.
- BOERI, T., P. GARIBALDI, AND E. MOEN (2016): “A Clash of Generations? Increase in Retirement Age and Labor Demand for Youth,” *CEPR Working Paper*, 11422.
- BROWNING, M., I. CRAWFORD, AND M. KNOEF (2013): “The age-period cohort problem: set identification and point identification,” *Cemmap Working Paper*.
- BRUSSEVICH, M., E. DABLA-NORRIS, AND S. KHALID (2019): “Is Technology Widening the Gender Gap? Automation and the Future of Female Employment,” *IMF Working Paper*, 19/91.

- BUSSOLO, M., I. TORRE, AND H. J. WINKLER (2018): “Does Job Polarization Explain the Rise in Earnings Inequality? Evidence from Europe (English),” *Policy Research working paper*, no. WPS 8652, World Bank.
- CAVAGLIA, C. AND B. ETHERIDGE (2020): “Job polarization and the declining quality of knowledge workers: Evidence from the UK and Germany,” *Labour Economics*, 66, 101884.
- CORTES, G. M., N. JAIMOVIC, C. J. NEKARDA, AND H. SIU (2021): “The Growing Importance of Social Tasks in High-Paying Occupations: Implications for Sorting,” Working Paper.
- CORTES, G. M., N. JAIMOVICH, C. J. NEKARDA, AND H. E. SIU (2020): “The dynamics of disappearing routine jobs: A flows approach,” *Labour Economics*, 65, 101823.
- COSTA DIAS, M., R. JOYCE, AND F. PARODI (2021): “The gender pay gap in the UK: children and experience in work,” *Oxford Review of Economic Policy*, 36, 855–881.
- COUDIN, E., S. MAILLARD, AND M. TÔ (2018): “Family, Firms and the Gender Wage Gap in France,” Working Papers 2018-09, Center for Research in Economics and Statistics.
- CRIBB, J. (2019): “Intergenerational Differences in Income and Wealth: Evidence from Britain,” *Fiscal Studies*, 40, 275–299.
- CRISTINI, A., A. GERACI, AND J. MUELLBAUER (2017): “Sifting through the ASHE: Job Polarisation and Earnings Inequality in the UK, 1975-2015,” *INET Oxford Working Paper*, No. 2018-05.
- DEATON, A. AND C. PAXSON (1994): “Intertemporal Choice and Inequality,” *Journal of Political Economy*, 102, 437–467.
- DICKERSON, A., R. WILSON, G. KIK, AND D. DHILLON (2012): “Developing Occupational Skills Profiles for the UK: A Feasibility Study,” Evidence Report No. 44 for the UK Commission for Employment and Skills.
- ELSBY, M., B. HOBIJN, AND A. SAHIN (2010): “The Labor Market in the Great Recession,” *Brookings Papers on Economic Activity*, 41, 1–69.
- FAGERENG, A., C. GOTTLIEB, AND L. GUIISO (2017): “Asset Market Participation and Portfolio Choice over the Life-Cycle,” *Journal of Finance*, 72, 705–750.
- FANNON, Z. AND B. NIELSEN (2019): “Age-Period-Cohort Models,” in *Oxford Research Encyclopedia of Economics and Finance*, Oxford University Press.
- GOLDIN, C. (2006): “The Quiet Revolution that Transformed Women’s Employment, Education, and Family,” *American Economic Review*, 96, 1–21.
- GOLDIN, C., S. P. KERR, C. OLIVETTI, AND E. BARTH (2017): “The Expanding Gender Earnings Gap: Evidence from the LEHD-2000 Census,” *American Economic Review*, 107, 110–14.

- GOOS, M. AND A. MANNING (2007): “Lousy and Lovely Jobs: The Rising Polarization of Work in Britain,” *Review of Economics and Statistics*, 89, 118–133.
- GOSLING, A., S. MACHIN, AND C. MEGHIR (2000): “The Changing Distribution of Male Wages in the U.K.” *Review of Economic Studies*, 67, 635–666.
- GRIGOLI, F., Z. KOCZAN, AND P. TOPALOVA (2020): “Automation and labor force participation in advanced economies: Macro and micro evidence,” *European Economic Review*, 126.
- GROES, F., P. KIRCHER, AND I. MANOVSKII (2015): “The U-Shape of Occupational Mobility,” *Review of Economic Studies*, 82, 659–692.
- HUR, S. (2018): “The Lost Generation of the Great Recession,” *Review of Economic Dynamics*, 30, 179–202.
- KAMBOUROV, G. AND I. MANOVSKII (2009a): “Accounting for the Changing Life-Cycle Profile of Earnings,” Working Paper.
- (2009b): “Occupational Mobility and Wage Inequality,” *Review of Economic Studies*, 50, 731–759.
- MACURDY, T. AND MROZ, T. (1995): “Measuring Macroeconomic Shifts in Wages from Cohort Specifications,” Working Paper.
- MONTRESOR, G. (2019): “Job Polarization and Labour Supply Changes in the UK,” *Labour Economics*, 59, 187–203.
- NEDELKOSKA, L. AND G. QUINTINI (2018): “Automation, Skills Use and Training,” *OECD Social, Employment and Migration Working Papers*, 202.
- NGAI, R. AND B. PETRONGOLO (2017): “Gender Gaps and the Rise of the Service Economy,” *American Economic Journal: Macroeconomics*, 9, 1–44.
- OECD (2017): “OECD Economic Surveys: United Kingdom 2017,” Organization for Economic Cooperation and Development.
- (2020): “OECD Employment Outlook 2020: Worker Security and the COVID-19 Crisis,” Organization for Economic Cooperation and Development.
- OLIVETTI, C. AND B. PETRONGOLO (2014): “Gender gaps across countries and skills: Demand, supply and the industry structure,” *Review of Economic Dynamics*, 17, 842–859.
- (2016): “The Evolution of Gender Gaps in Industrialized Countries,” *Annual Review of Economics*, 8, 405–434.
- SALVATORI, A. (2018): “The Anatomy of Job Polarisation in the UK,” *Journal for Labour Market Research*, 52.
- YANG, Y., S. SCHULHOFER-WOHL, W. J. FU, AND K. C. LAND (2008): “The Intrinsic Estimator for Age-Period-Cohort Analysis: What It Is and How to Use It,” *American Journal of Sociology*, 113, 1697–1736.

A Additional Tables and Figures

Figure A.1: Unemployment rate and labor force inactivity rate by age group for non-college females in the UK: 1997-2018



Notes: The left panel plots the unemployment rate by age group, while the right panel plots the Out-of-Labor-Force rate for the same groups. The grey shaded area encompasses the dates of the Great Recession.

Sources: LFS and authors' calculations.

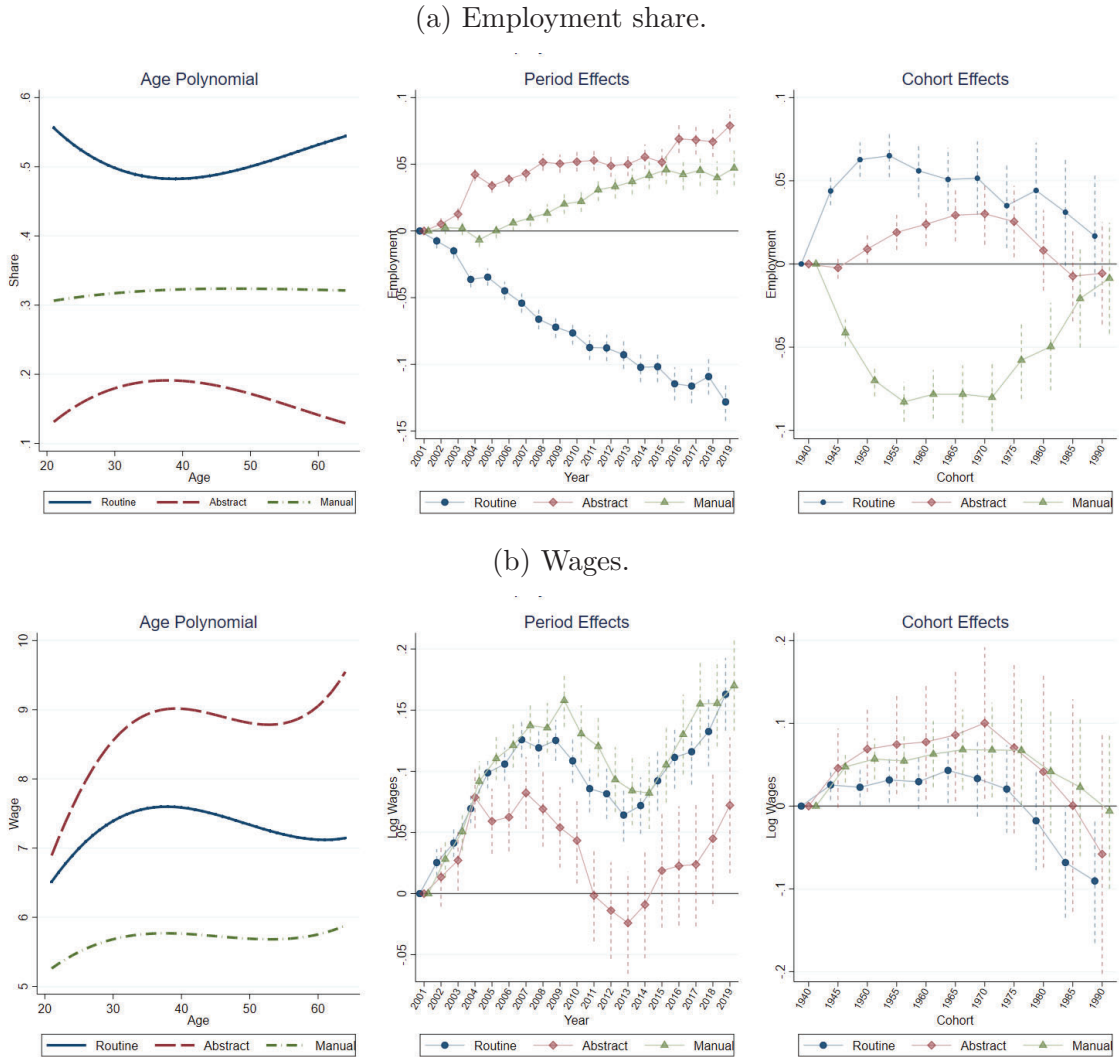
Table A.1: P-values from two-sided tests for equality of the coefficients of each cohort and Cohort 1970.

		Cohort										
		1940	1945	1950	1955	1960	1965	1970	1975	1980	1985	1990
Employment	Manual	0.00	0.00	0.30	0.45	0.32	0.20	-	0.00	0.00	0.00	0.00
	Routine	0.00	0.02	0.95	0.47	0.58	0.09	-	0.00	0.01	0.00	0.00
	Abstract	0.97	0.21	0.30	0.87	0.72	0.43	-	0.67	0.00	0.00	0.00
Log Wages	Manual	0.02	0.28	0.36	0.22	0.40	0.75	-	0.62	0.10	0.01	0.00
	Routine	0.51	0.86	0.88	0.74	0.99	0.07	-	0.15	0.00	0.00	0.00
	Abstract	0.11	0.45	0.62	0.63	0.37	0.37	-	0.02	0.00	0.00	0.00

Note. The table reports the p-values from two-sided F-tests of the null hypothesis that the coefficient for the respective cohort is equal to that of Cohort 1970. The first three rows report the tests for the regressions of employment in each occupation type, as specified in (2). The last three rows report the tests for the regressions of log wages in each occupation type, as specified in (3).

A.1 Sensitivity analysis of baseline regression

Figure A.2: Age polynomials, period effects, and cohort effects including immigrants in the sample.

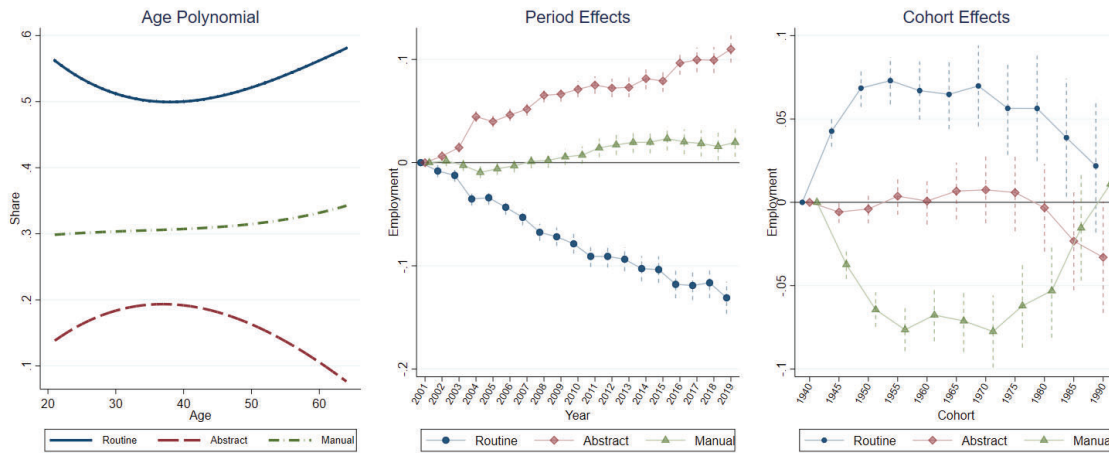


Notes: The top panels report the estimated age polynomials, the yearly period effects, and the 5-year cohort effects of employment shares in each occupation type as in Eq. (2). The bottom panels report the same figures for the expected wages within each occupation as in Eq. (3). For period and cohort effects, the vertical bars represent the 95 percent confidence interval. Age polynomials for wages are reported as the exponential of the estimate polynomial for log wages.

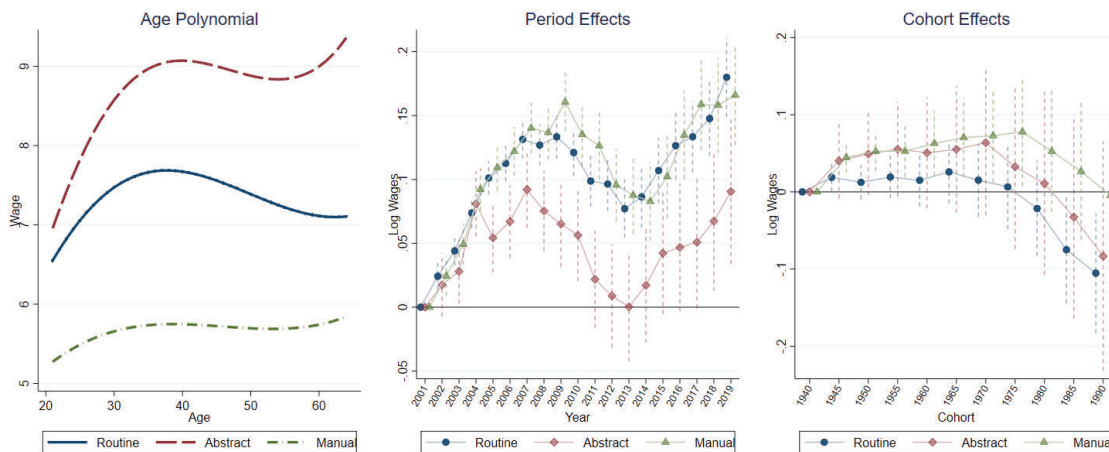
Sources: LFS and authors' calculations.

Figure A.3: Age polynomials, period effects, and cohort effects excluding self-employed workers from the sample

(a) Employment share.



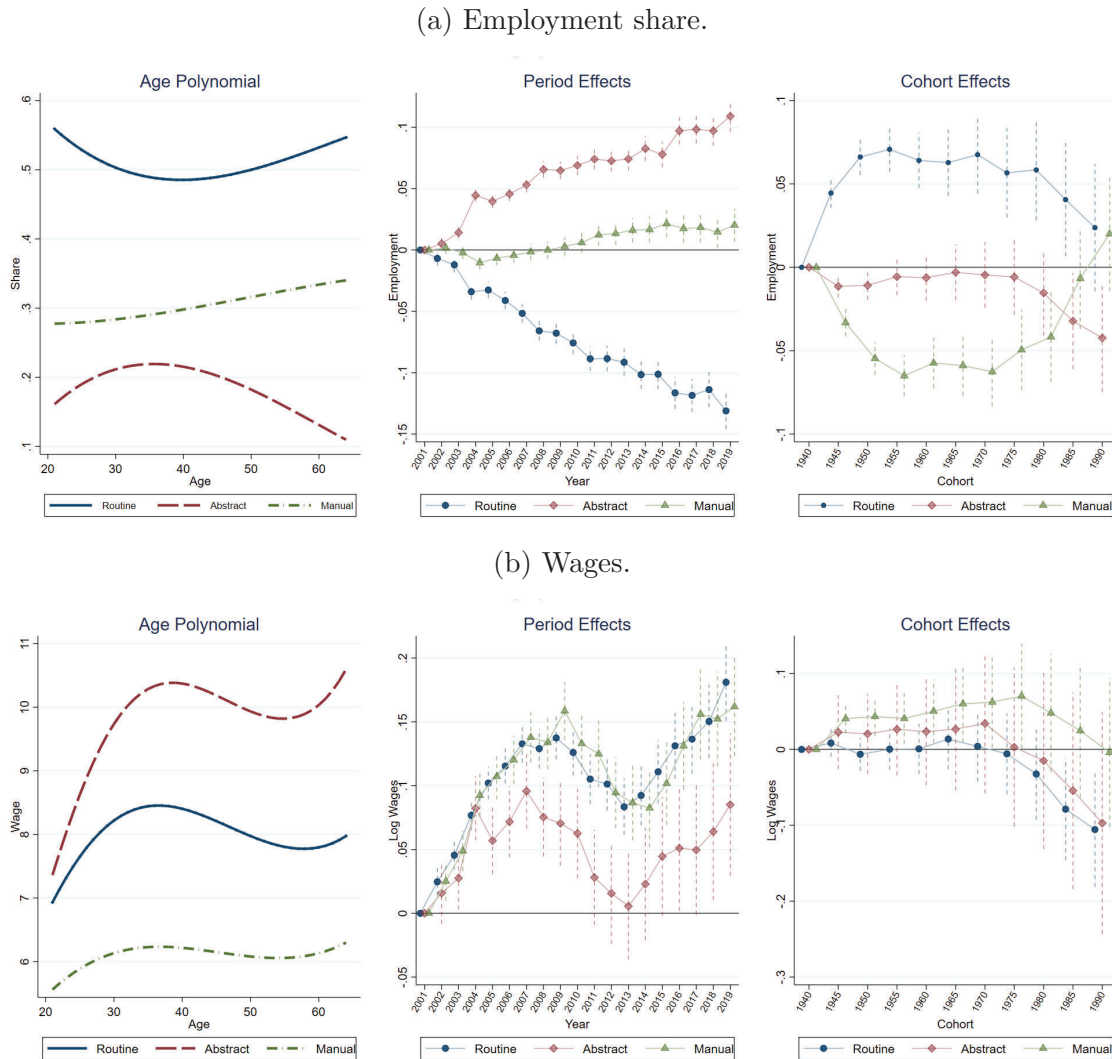
(b) Wages.



Notes: The top panels report the estimated age polynomials, the yearly period effects, and the 5-year cohort effects of employment shares in each occupation type as in Eq. (2). The bottom panels report the same figures for the expected wages within each occupation as in Eq. (3). For period and cohort effects, the vertical bars represent the 95 percent confidence interval. Age polynomials for wages are reported as the exponential of the estimate polynomial for log wages.

Sources: LFS and authors' calculations.

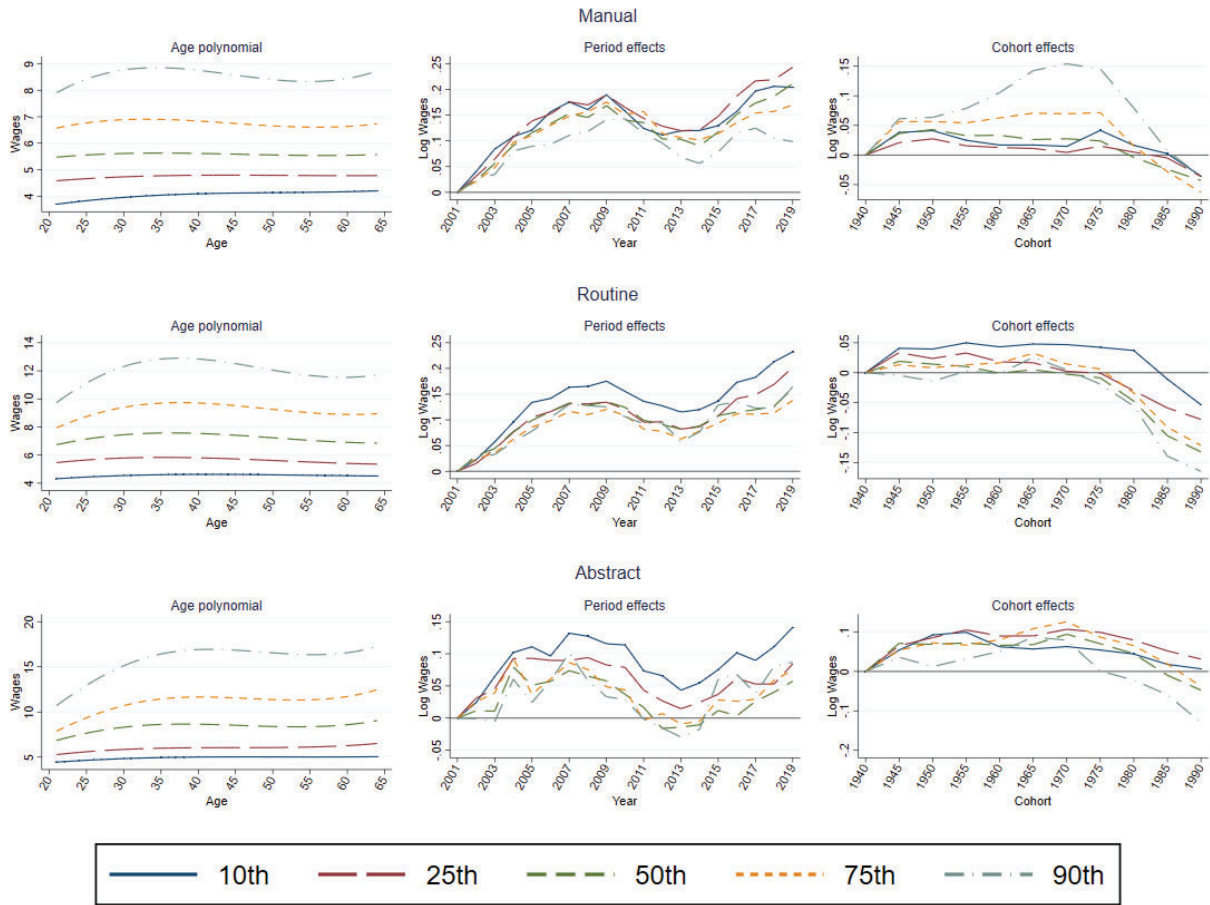
Figure A.4: Age effects, period effects, and cohort effects estimated adding individual controls



Notes: The top panels report the estimated age polynomials, the yearly period effects, and the 5-year cohort effects of employment shares in each occupation type as in Eq. (2). The bottom panels report the same figures for the expected wages within each occupation as in Eq. (3). For period and cohort effects, the vertical bars represent the 95 percent confidence interval. Age polynomials for wages are reported as the exponential of the estimate polynomial for log wages.

Sources: LFS and authors' calculations.

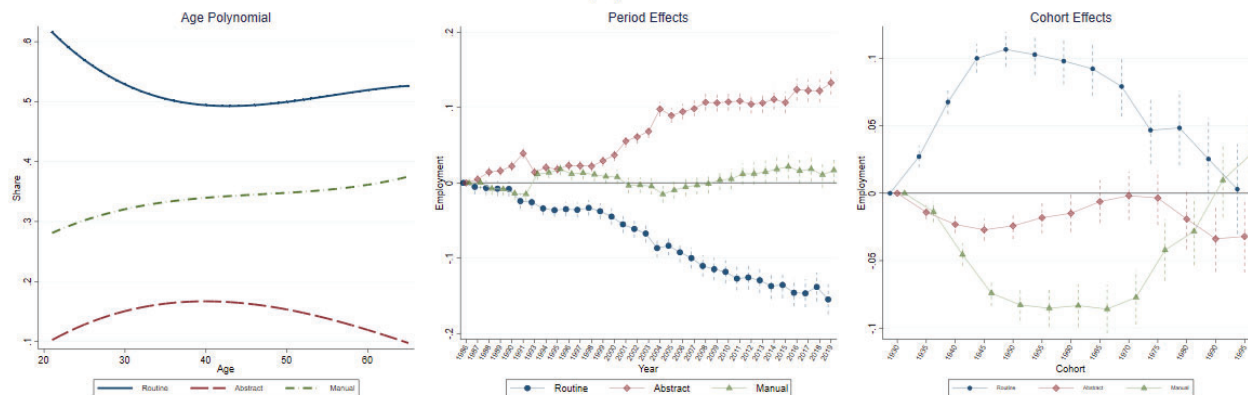
Figure A.5: Age effects, period effects, and cohort effects for wages estimated using quantile regressions at the 10th, 25th, 50th, 75th, and 90th percentiles



Notes: Each row reports the age polynomials, the yearly period effects, and the 5-year cohort effects of wages in each occupation type as in Eq. (3) estimated using a quantile regression. The lines report the estimated coefficients at the 10th, 25th, 50th, 75th, and 90th percentiles. Age polynomials are reported as the exponential of the estimate polynomial for log wages.

Sources: LFS and authors' calculations.

Figure A.6: Age effects, period effects, and cohort effects for employment share extending the sample backwards until 1986

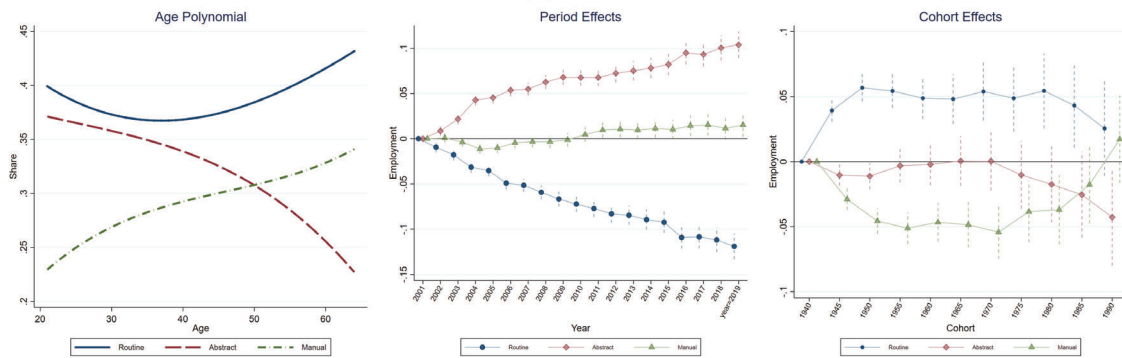


Notes: The panels report the estimated age polynomials, the yearly period effects, and the 5-year cohort effects of employment shares in each occupation type as in Eq. (2). For period and cohort effects, the standard errors reported comprise the 95 percent confidence interval.

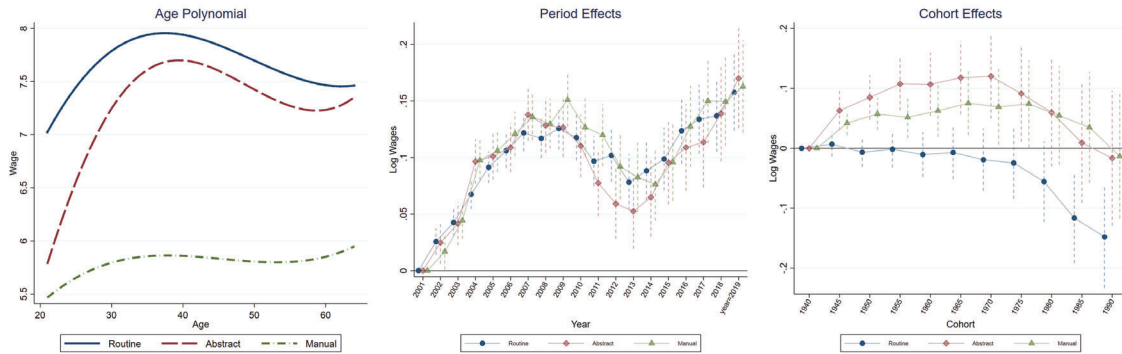
Sources: LFS and authors' calculations.

Figure A.7: Age polynomials, period effects, and cohort effects using the alternative occupational categorization

(a) Employment share.



(b) Wages.

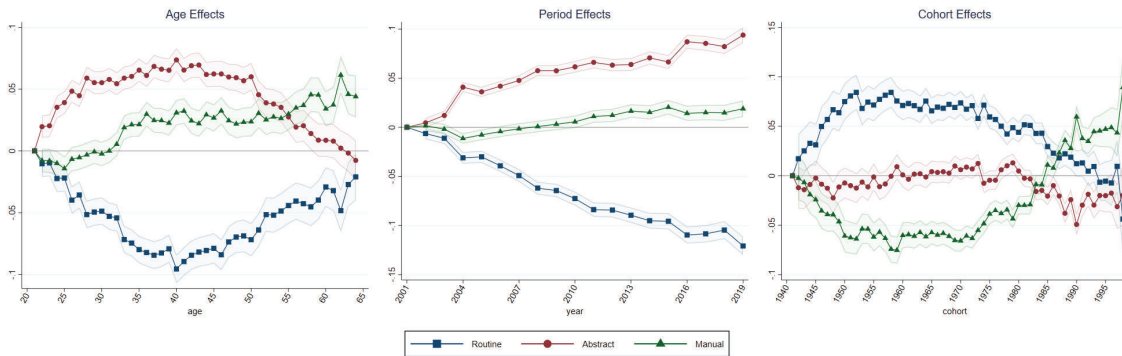


Notes: The top panels report the estimated age polynomials, the yearly period effects, and the 5-year cohort effects of employment shares in each occupation type as in Eq. (2). The bottom panels report the same figures for the expected wages within each occupation as in Eq. (3). For period and cohort effects, the standard errors reported comprise the 95 percent confidence interval. Age polynomials for wages are reported as the exponential of the estimate polynomial for log wages.

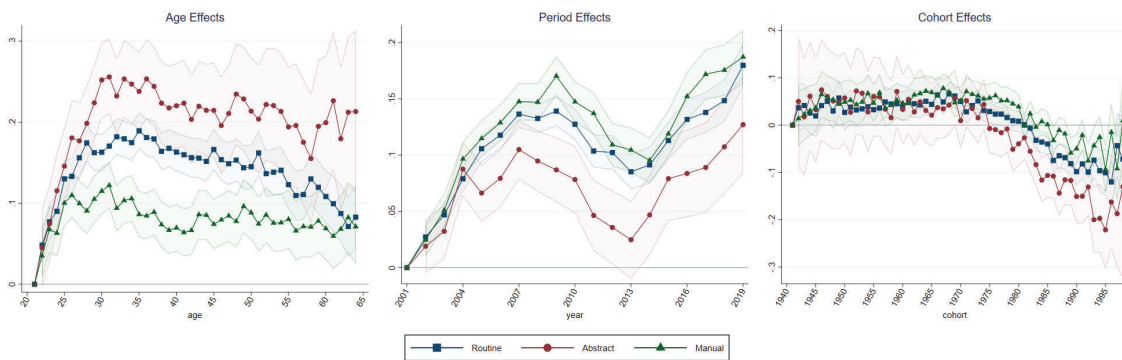
Sources: LFS and authors' calculations.

Figure A.8: Age effects, period effects, and cohort effects using the maximum entropy estimation method

(a) Employment share.



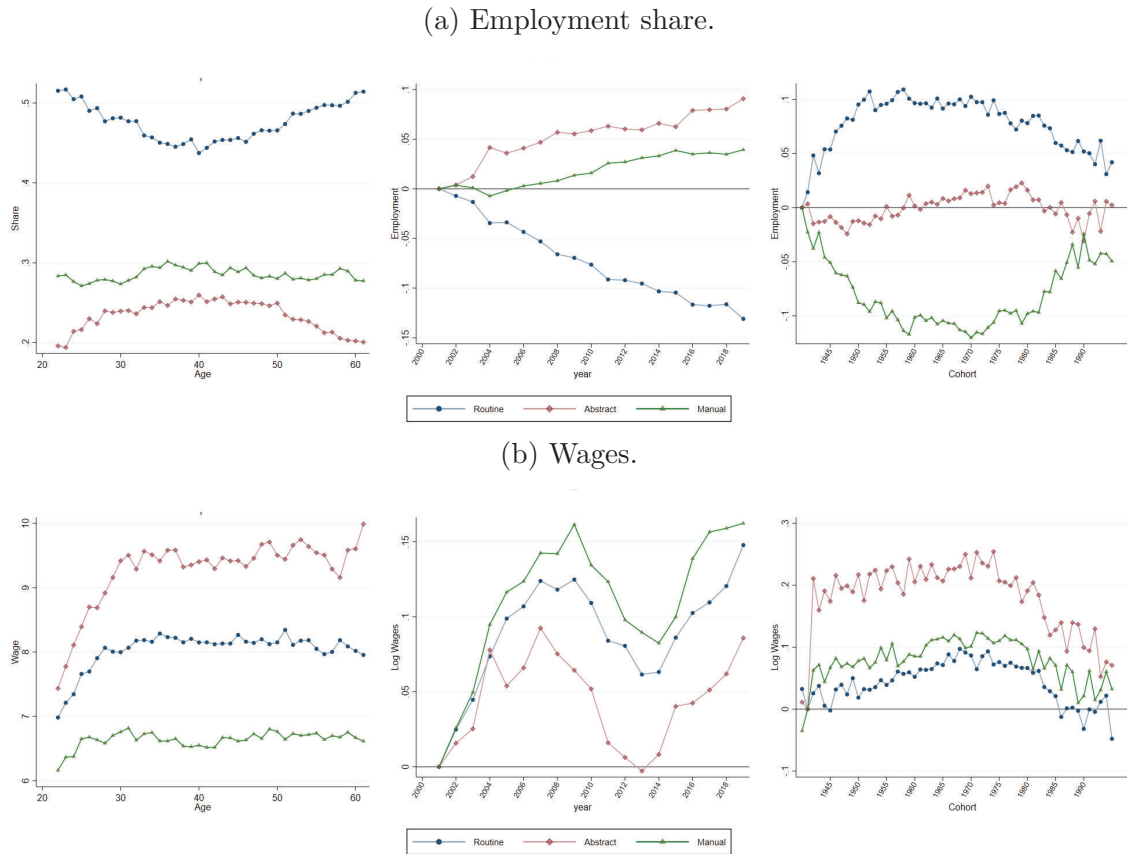
(b) Wages.



Notes: The top panels report the age, period, and 1-year cohort effects of employment shares in each occupation type as in Eq. (2) estimated using the “maximum entropy” method. The bottom panels report the same figures for the expected wages within each occupation as in Eq. (3). For period and cohort effects, the standard errors reported comprise the 95 percent confidence interval.

Sources: LFS and authors’ calculations.

Figure A.9: Age effects, period effects, and cohort effects using the intrinsic estimator method

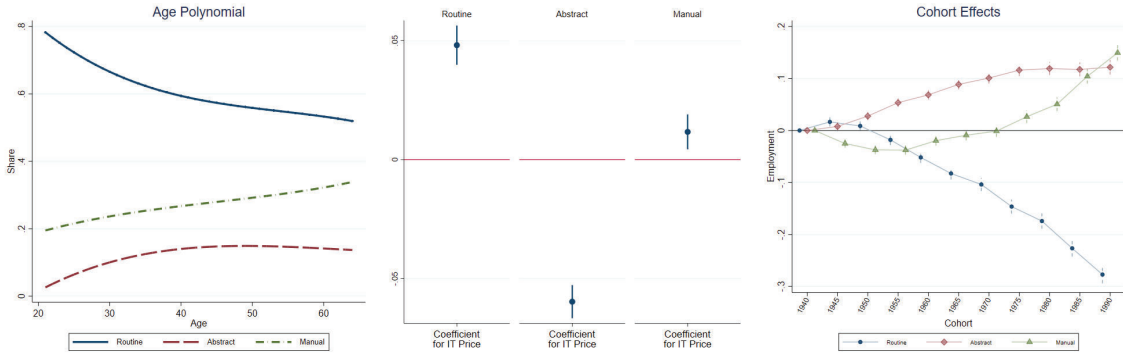


Notes: The top panels report the age, period, and 1-year cohort effects of employment shares in each occupation type as in Eq. (2) estimated using the “intrinsic estimator” method. The bottom panels report the same figures for the expected wages within each occupation as in Eq. (3). Age effects include the constant term.

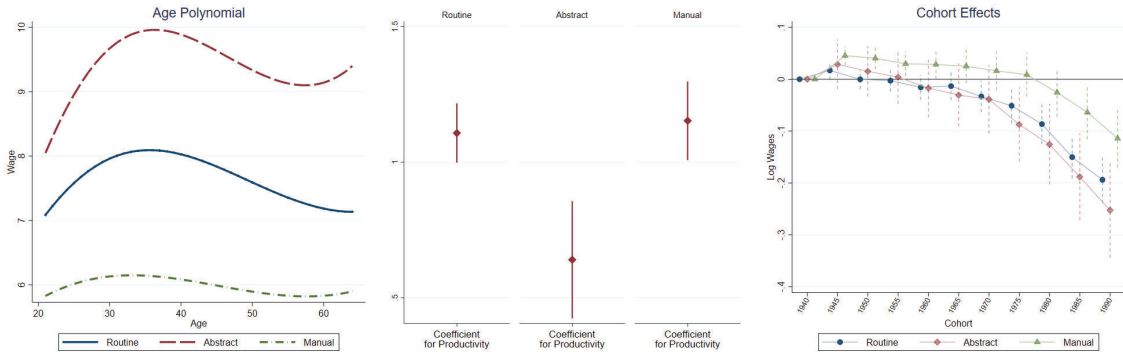
Sources: LFS and authors’ calculations.

Figure A.10: Age polynomials, cohort effects, and coefficients on the proxies for year effects

(a) Employment share.



(b) Wages.



Notes: The top panels report the age polynomials, the coefficients of the variable proxying for period effects, and the 5-year cohort effects of employment shares in each occupation type as in Eq. (2). The bottom panels report the same figures for the expected wages within each occupation as in Eq. (3). Age polynomials for wages are reported as the exponential of the estimate polynomial for log wages. Period effects in employment shares are proxied using the price of IT equipment relative to the price of all capital goods. The coefficients across the three regressions are restricted to sum up to zero. Period effects in wages are proxied using labor productivity. For period proxy and cohort effects, the standard errors reported comprise the 95 percent confidence interval

Sources: LFS and authors' calculations.

B Data appendix

B.1 Occupational categorization measures

This section describes the categorization of occupations into three types: manual, routine, and abstract. As we base the definition on the original categorization of Cortes et al. (2020) for US occupations, a key step entailed translating this categorization to the UK occupational classifications. The quarterly LFS uses three different occupation classifications from 1994 to 2019: the SOC 1990 (until 2001), the SOC 2000 from 2001 to 2011), and SOC 2010. Although (partial) crosswalks between these three classifications are available, changes between the two were substantial and can create large discontinuities at the aggregate level in the share of workers in each occupation type. However, all LFS waves using the SOC2010 also provide a double-coded version of the worker’s occupation variable using the SOC2000. We thus work with the period 2001-2019, using the double-coded SOC2000 variable after 2011.

The original task content measures are based on US occupation classifications (either US SOC or US Census). The Cortes et al. (2020) classification is based on the US census classification OCC 2000. To map it into the UKSOC 2000, we use the following crosswalks: OCC 2000 → US SOC 2000 → US SOC 2010 → ISCO 08 → UK SOC 2010 → UK SOC 2000. For the US crosswalks, when multiple occupations were combined into a single ISCO-08 code, we used simple averages.

We also made ad-hoc adjustments to some individual correspondences. In these cases, we used the non-unique occupational mapping provided by Dickerson et al. (2012) using the word-based matching software CASCOT.²¹ The authors used the software to match descriptions of UK jobs with those of US jobs from the O*NET database. The adjustments made are:

- UK SOC 2010 9232 was originally paired with ISCO 08 9613, which in turn was originally paired with US SOC 373019, which has no information on O*NET. We thus paired it with 474051 based on CASCOT information from Dickerson et al. (2012).
- UK SOC 2010 3561 and 3565 were originally paired with ISCO 3359 which was originally paired with US SOC 452011 (agricultural inspectors), but the original UK SOC occupations resemble more civil sector clerks. We thus repaired them with the ISCO 08

²¹We thank Prof. Rob Wilson from the Warwick Institute for Employment Research for sharing this mapping with us.

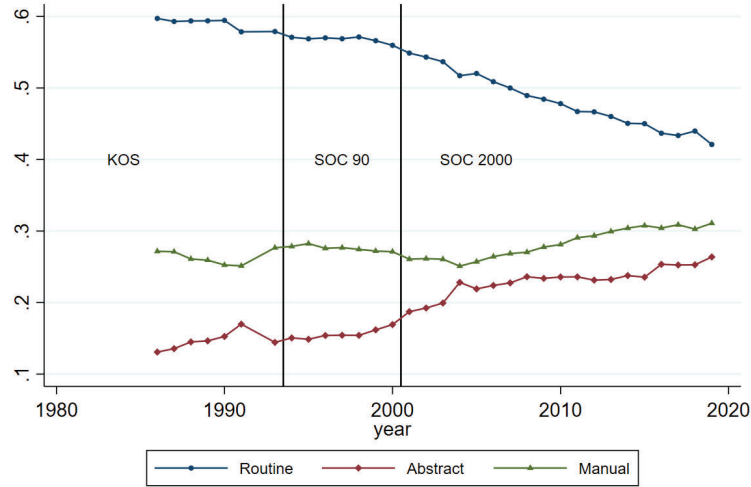
occupations corresponding to the US SOC 2010 codes that are matched with them in the CASCOT-based mapping from Dickerson et al. (2012).

As explained in the paper, we adjust the original classification by Cortes et al. (2020) by grouping together routine-cognitive and routine-manual occupations into a single routine group. Our manual and abstract categories correspond to the non-routine-manual and non-routine-cognitive categories, respectively.

In the robustness checks, when using the LFS waves containing the SOC 1990 we used the following mapping: ISCO 08 \rightarrow ISCO 88 \rightarrow ISCO 88 (COM) \rightarrow SOC 1990. For the crosswalk between the SOC 1990 and the KOS we use the LFS 1991 dataset, in which workers' occupations are double-coded using both classifications. The mapping between the two occupations is not one-to-one. Therefore, we first merge the categorization of the SOC 1990 with the dataset. Using worker's weights, we then compute the KOS-based categorization through a modal-based matching. For instance, if 80 percent of the workers in a given KOS occupation are also matched with a SOC 1990 occupation that is classified as routine, we classify that KOS occupation as routine. These extensive cross-walks imply some degree of approximation compared to the original task content dataset and a discontinuity in the aggregate statistics using the LFS. However, the essence of the information is very similar.

Figure B.1 plots the share of routine, abstract, and manual jobs for females without a post-secondary education from 1986 to 2019.

Figure B.1: Shares of routine, abstract, and manual occupation for non-college females: 1986-2019



Notes: The figures plot the shares of routine, abstract, and manual jobs from 1981 onwards. Series breaks due to changes in the occupational classifications are marked by vertical lines.
Source: LFS and authors’ calculations.

We also check the alignment between our occupational categories and the RTI index constructed by Autor et al. (2003) based on the US Census OCC1990 classification. We followed the mapping OCC 1990 → OCC 2000 → US SOC 2000 → US SOC 2010 → ISCO 08 → UK SOC 2010 → UKSOC2000. The RTI is calculated using the task-content measures produced through the 1977 US Dictionary of Occupational Titles. The formula is $RTI = \text{Log}(\text{Routine Task}) - \text{Log}(\text{Abstract Task}) - \text{Log}(\text{Manual Task})$. Table B.1 shows that our categorization is closely aligned with routinization measured through the RTI.

Table B.1: Summary Statistics of Routine Task Index by Occupation Type in 2001

	Males				Females			
	Mean	Median	25th pct.	75th pct.	Mean	Median	25th pct.	75th pct.
Routine	0.72	0.56	-0.31	1.61	2.35	2.51	0.81	3.34
Abstract	-0.03	-0.45	-0.65	0.53	-0.12	-0.46	-1.15	0.67
Manual	0.06	0.09	-0.3	0.67	0.45	0.32	-0.27	0.96

Source: LFS and authors’ calculations.

Note: Each line reports the mean, median, 25th and 75th percentiles of the distribution of the RTI across all occupations in each category, separately for males and female.

B.2 Alternative occupational categorization

As a robustness check, we classify occupations into the manual, routine, and abstract categories following Bhalotra and Fernandez (2018). In particular, we merge the UK SOC 2000 occupation with task content measures computed by Autor and Dorn (2013) from the Dictionary of Occupational Titles 1977, as explained in the previous section.

First, we calculate the percentile corresponding to each occupation’s value in the distribution of each task component—routine, manual, and abstract. Because we do this for each of the task components, every occupation classification has three different percentiles. This procedure informs us of where each occupation stands relative to the others regarding the content of each task. We then assign each occupation to the category corresponding to its highest percentile. For instance, if an occupation is in the 90th percentile for manual, 70th for routine, and 23rd for abstract, we categorize it as a manual occupation. Because percentiles range from 0 to 100, they are comparable across categories.

Table B.2 shows the cross-distribution of non-college female workers in the two occupational groupings for the first and last year of the sample. Overall, the two methods are consistent, and the majority of workers remain in the same category.

Table B.2: Cross-Distribution of Baseline and Alternative Categorizations for 2001 and 2018 by Worker Groups

		Alternative					
		2001			2019		
		Manual	Routine	Abstract	Manual	Routine	Abstract
Baseline	Manual	83.5	10.9	5.7	82.6	10.2	7.2
	Routine	6.0	68.7	25.3	6.4	62.6	31.0
	Abstract	2.8	5.7	91.4	3.9	5.0	91.0
	All	25.6	41.8	32.6	29.2	32.4	38.5

Figure A.7 reports the estimated age polynomials, period effects, and cohort effects from Eq. (2) and (3) using the alternative occupation categorization. Results from the baseline specification are broadly robust to this categorization. Age polynomials for the employment shares are somewhat different with regards to their linear trends, but maintain the same shape in terms of their non-linear parts (i.e., concavity or convexity).

C Alternative estimation approaches

In this section, we briefly discuss alternative estimation methods for age, period, and cohort effects. We consider the merits, underlying assumptions, and requirements of each method and whether, in our view, they are appropriate for the empirical question investigated in our paper. We present the estimation results for the approaches that provide viable alternatives.

C.1 Discussion of alternative estimation approaches of age, period, and cohort effects

Zero restrictions on the values of specific coefficients or their linear combinations can often be motivated on the basis of economic theory. For instance, in studying households' savings rate, Deaton and Paxson (1994) impose the condition that period effects must sum to zero, as they capture temporary business cycle fluctuations. While this assumption is common in economic problems, it is not applicable in our case since the year effects in our framework are expected to capture long-run trends in the structure of the economy. For the same reason, applying the restriction to cohort effects would not be tenable.²²

Browning et al. (2013) and Yang et al. (2008) propose two methods to side-step the need to impose parametric or zero restrictions. Browning et al. (2013) note that, when the value of the dependent variable is bounded, the structure of the problem imposes some constraints on the possible (unknown) values of the true coefficients of the age, cohort, and period factors (i.e., there is a bounded set of identifiable coefficients). Their “maximum entropy” method thus aims to estimate the coefficients as the expected values from the “most informative” (in Shannon’s entropy terms) probability distribution over possible values of the coefficients within the identifiable set. Yang et al. (2008) show that, although the matrix $X'X$ of cohort, age, and period effects is not full-rank, the structure of the problem makes the type of linear dependence among the variables in X clear. Hence, in their “implicit estimator” method all the effects can be extracted by first regressing the dependent variables on transformed variables created using the eigenvectors corresponding to the non-zero eigenvalues of $X'X$. We apply both methods as

²²An even more restrictive choice would be to fully exclude one set of effects, thus assuming that no change in occupational shares or wages occurs along a chosen dimension. However, F-tests of the restricted model without one of the effects against a saturated model can be used to assess whether the data rejects this restriction. In our case, we find that the exclusion of any of the effects is rejected in the regressions (2) and (3) for all three occupation types.

robustness checks for our baseline results.²³

A further option is to substitute one set of effects with a proxy variable that captures the same variation in the data. For instance, Fagereng et al. (2017) use stock market returns when an individual was in their “formative year” (age 18 to 25) to proxy for the cohort effects in the share of equities in their assets. In the context of our question, period effects allow for more intuitive proxies than cohort effects, as they are likely linked to measurable variables. Based on the literature linking job polarization to automation (Acemoglu and Autor, 2011), we use the price of IT equipment relative to the price of all capital stock as a proxy for aggregate time variation in the share of routine, abstract, and manual jobs.²⁴²⁵ Further, we use labor productivity as a proxy for period effects in hourly wages.²⁶

C.2 Results

Figure A.8 presents the results using the “maximum entropy” method of Browning et al. (2013), while Figure A.9 reports the results using the “intrinsic estimator” proposed by Yang et al. (2008). In both cases, the results are very similar. The use of single-year cohorts adds noise in the year-to-year changes across cohorts (for instance, for abstract job wages under the “intrinsic estimator”) but leave the overall features unchanged. The linear trends are mostly captured by the period effects, while the same nonlinearities in cohort and age effects remain.

Figure A.10 shows the result when using the relative price of IT equipment and labor productivity to proxy for year effects in employment shares and wages, respectively. The middle subplots report the estimated coefficients of these variables instead of the year effects. Most results from the baseline specification carry through, with only a substantial difference in the cohort effects of employment propensities, which now contain very pronounced linear trends. In particular, cohort effects seem to capture a very large fall in the probability of routine jobs from Cohort 1940 to Cohort 1990 of almost 30 percentage points. The non-linear components of

²³The requirement of a bounded dependent variable technically makes the “maximum entropy” unsuitable for the wage equation. While acknowledging the violation of this assumption, we show that our main results hold in practice by specifying a very large upper bound that is above all wage levels in the data.

²⁴The prices of IT equipment and capital in the UK over the period 2001-2017 are taken from the EU KLEMS 2019 database.

²⁵We estimate (2) simultaneously for all three occupation types using the Seemingly Unrelated Regression method and adding the restriction that the three coefficients on the proxy variables must add up to zero.

²⁶Annual labor productivity data is taken from the Office of National Statistics website. The coefficients of the proxy variable is left unrestricted.

cohort effects, however, remain similar, with cohort effects for routine jobs exhibiting a concave shape and those of manual jobs a convex one.



PUBLICATIONS

Are Low-Skill Women Being Left Behind? Labor Market Evidence from the UK
Working Paper No. WP/22/42