

Job Polarization and the Declining Quality of Knowledge Workers: Evidence from the UK and Germany*

Chiara Cavaglia[†], Ben Etheridge[‡]

July 1, 2020

Abstract

Across the developed world, employment has polarized clearly by occupation, but changes to the wage structure have been harder to interpret. We examine changes to the wage structure in the UK and Germany, two countries with apparently very different trends. Using panel data, we argue that changes to quality-adjusted prices for occupations grouped by predominant tasks correlate strongly with employment growth in both countries, consistently with task-based changes to labour demand. The gap between price and average wage changes is strongest in top (knowledge) jobs, implying that the average quality of these workers has declined over time. We obtain further direct evidence on changes to worker quality using rich data on individual characteristics.

JEL Classification: J20, J24, J31

Keywords: Job Polarization, Task Prices, Roy Model, Height, Cognitive Ability.

*This is a revised version of an earlier paper titled “Job Polarization, Task Prices and the Distribution of Task Returns”. We are grateful to the editor, Bernd Fitzenberger, and to two anonymous referees for suggestions that improved the paper. We are also grateful for discussions with participants at several seminars, as well as comments from Andrea Salvatori, Matias Cortes, Tom Crossley and Hans-Martin von Gaudecker. Etheridge gratefully acknowledges funding from the ESRC Research Centre on Micro-social Change at the Institute for Social and Economic Research, and from British Academy under grant DKB0400. Cavaglia gratefully acknowledges funding from the University of Essex ESRC Doctoral Training Centre. All errors remain the responsibility of the authors.

[†]Centre for Vocational Education Research, Centre for Economic Performance, London School of Economics, Houghton Street, London, WC2A 2AE. E-mail: c.cavaglia@lse.ac.uk

[‡]Corresponding Author. Address: Department of Economics, University of Essex, Wivenhoe Park, Colchester, UK, CO7 9AB. E-mail: bsethe@essex.ac.uk

1 Introduction

Across most developed economies, the occupational structure has shifted substantially over at least the last 30 years. This shift has typically seen employment decline in middle-earning occupations, and grow in jobs at the top of the wage distribution. Employment has also grown, to a lesser extent, in low-paying jobs, giving rise to a pattern of ‘polarization’.¹ This noticeable polarization has been attributed to a number of causes, such as changes to patterns of trade or to the task requirements in production. Testing the causes of polarization at even the most basic level requires examining the equilibrium movement across sectors of both employment and labour returns. Yet, although the patterns of employment changes are clear, the evidence on wages is harder to interpret. Exactly because of the large changes to employment, average wages observed across sectors are likely driven by composition effects. Intuitively, if employment in a sector grows, it is likely that new entrants are of different average quality to the incumbents. It is therefore important to look beyond average wages, and to identify the pure (selection- or composition-free) prices on labour. This distinction between average wages and the price on labour has long been made by labour economists. Within the polarization literature this distinction has recently been emphasized by [Acemoglu and Autor \(2011\)](#) and [Gottschalk, Green, and Sand \(2016\)](#) among others.

In this paper we make two contributions. First we estimate underlying prices of broad occupational groups using panel data for two important economies: the UK and Germany. In both countries we find that these prices deviate significantly from observed average wages. Importantly, in both countries, we find that price changes are noticeably positively associated with employment growth across sectors. Meanwhile if we compare employment changes not with prices but with changes to average wages, this association disappears. Overall, our evidence is consistent with occupational shifts being caused by changes in the demand for different types of labour, such as by changes to technology. Our evidence also therefore highlights the importance of identifying underlying prices when considering hypotheses about the labour market. Ours is the first paper to provide consistent evidence across countries on these task price movements.

A corollary of these results is that the average quality of workers across different sectors has changed substantially over time. Our second contribution, therefore, is to demonstrate these changes by providing additional evidence that goes beyond wage data. Specifically we examine two markers of individual productivity: cognitive ability scores and height. Both these markers are strongly correlated with labour market outcomes. While cognitive ability is possibly a closer and more obvious indicator of productive skills, height is extremely useful for the present application: it is highly stable within individuals over time; comparable across surveys, and, compared to cognitive scores, less sector-specific. Using both measures, and consistent with the evidence from wage data, we find evidence of changes in composition.

¹See, for example, [Goos, Manning, and Salomons \(2014\)](#), who document changes to the occupational structure across much of Western Europe since the early 1990s. See also the literature discussed below.

To the best of our knowledge, ours is the first paper to provide explicit evidence of changing sorting over time based on individual measures of worker quality in this context.²

In common with much of the literature, we group broad occupations according to their predominant task. Specifically we divide jobs into four categories, depending on whether or not the occupation is predominantly intensive in cognitive versus manual tasks along one axis, and whether or not it is highly routine along the other. In both countries, we see most prominently a rapid increase in cognitive and non-routine (‘abstract’) employment. We also see a decline in middle-earning routine-manual occupations among males. Correspondingly, in both countries we see a striking increase in the relative price on abstract labour, not matched by changes in observed wages. An implication of these results from wage data is that it is the average quality of workers in abstract occupations that has declined most strongly over time. Our results are robust to allocating jobs to tasks not by broad occupation but at a finer level of detail.³

To study wages, we use data from the British Household Panel Survey over 1991-2008, and from the German Socio-Economic Panel over 1985-2013. These periods are noteworthy because they feature particularly large changes in occupational composition. To estimate the task prices we use a standard wage model, in which a worker is endowed with differing productivities across sectors, and therefore makes occupational choices given aggregate prices and unobservable individual-specific returns. The model nests, and follows the logic of, a standard Roy framework. The model also allows for frictions, such as costs of switching occupations or jobs. As is standard, by using the panel data structure we net out the time-invariant unobservable components and provide estimates of the prices that are free of selection effects. Our approach follows that used by [Cortes \(2016\)](#) to study task prices in the US, as well as other studies in related settings.⁴

Our analysis of price changes has three main strengths. First, we go beyond [Cortes \(2016\)](#) in providing robust support for the empirical evidence, by introducing extensive controls for observable characteristics. Second, we extend the analysis to look at prices at a finer level of occupational aggregation. We find that patterns of employment and prices at this finer detail can be easily reconciled. Finally, by using a consistent occupational classification, we provide coherent evidence across countries. As is shown later, the UK and Germany experienced job polarization to differing degrees: our evidence provides a unified and coherent explanation of the movements of task prices and employment across both countries combined.

²In a different context [Ashworth, Hotz, Maurel, and Ransom \(2017\)](#) show, using the NLSY, that the average cognitive score of college graduates has declined over time relative to non-graduates.

³We use the broad occupational groups in most of our analysis. Because we align these groups to predominant tasks, we call these groups ‘sectors’ or ‘tasks’ somewhat interchangeably.

⁴For example [Combes, Duranton, and Gobillon \(2008\)](#) estimate city wage premia, taking into account sorting across locations, while [Solon, Barsky, and Parker \(1994\)](#) estimate labour prices over the business cycle, taking into account the non-random selection into unemployment.

We supplement our results with direct evidence of changes in worker quality in both countries across sectors. For the UK we augment our data with especially rich information on personal characteristics available in around 2010 from the Understanding Society survey. Despite the rich information available in both countries, data limitations require us to adapt our framework somewhat. To use the data to their fullest, we focus our analysis not on occupation but on education. Our favoured analysis shows that the average height of males with a degree has declined relative to those without a degree. To support this analysis, we show that in both countries degree holding has been a stable proxy for working in the abstract sector. Overall, we infer from our results that the quality of ‘knowledge workers’ has declined over time.

Using the UK’s rich data, we validate this result in several ways. First we use the cognitive ability test scores, also collected in 2010. In a related way to heights, we argue that the average cognitive score of degree holders has declined, relative to non-degree holders, over time. However, although cognitive ability is arguably a more direct measure of worker skills than height, it is sector-specific, less stable over time and less easy to compare across surveys. For these reason, and because appropriate comparable data for Germany are sparse, we use the analysis of cognitive ability as supportive evidence. Second, we combine our main data source on heights, which was collected around 2005 and 2010, with an additional set of measures taken in 1991, from the UK’s National Child Development Survey, to provide a repeated cross-section. Our results using education are confirmed. Most importantly, we show that height patterns are similar when we define sectors by occupation, although the sample sizes used here are smaller, and consequently estimates are less precise. Overall, we provide evidence of changes in quality compositions that is strong and robust.

Our paper relates most closely to recent papers which estimate prices on tasks using panel data. Besides [Cortes \(2016\)](#), [Yamaguchi \(2018\)](#) also uses data from the PSID and controls for unobserved skills by conditioning on full labour market histories. In ongoing work, [Böhm, von Gaudecker, and Schran \(2019\)](#) use administrative panel data from Germany, and build a rich structural framework with which they identify not only task price movements but also life-cycle skill formation. They identify both components simultaneously by partitioning their long panel and exploiting time variation in the underlying changes to labour demand. Relative to their work, our work provides cross-country evidence and explores individual characteristics. In terms of substantive results, it is reassuring that [Böhm et al.](#)’s findings on task price movements are similar to ours and also show a substantial divergence from observed average wages.

In addition to the papers using panel data, further papers employ contrasting approaches using repeated cross-sections. Of these, [Gottschalk et al. \(2016\)](#) use data on new entrants to the labour market from the US Current Population Survey. They estimate bounds on prices, taking into account selection effects, by trimming the observed wage distributions. Second, [Böhm \(2020\)](#) addresses selection into sectors using a ‘propensity index’ method. He uses data on ability scores in the US National Longitudinal Survey of Youth to model occupational choice. The main problem with both approaches is that it is difficult to

control successfully for changing characteristics of successive cohorts. In particular, not only has the employment structure in the US shifted, but so has the composition of educational attainment. In our context, and as discussed in both countries, degree holding and occupational attainment have increased in parallel. By tracking the same workers over time, the panel-data approach used here by-passes issues surrounding educational composition completely.⁵

More generally, our paper relates to the large literature on job polarization, recently discussed in [Autor \(2015\)](#). Of particular relevance is [Salvatori \(2015\)](#), who examines the UK, [Dustmann, Ludsteck, and Schönberg \(2009\)](#) and [Antonczyk, Fitzenberger, and Leuschner \(2009\)](#), who discuss Germany, while [Beaudry, Green, and Sand \(2016\)](#) address somewhat contrasting patterns in the US in the 2000s. Our paper also relates to important literatures on the evolution of inequality in the UK, and particularly in Germany. In addition to [Dustmann et al. \(2009\)](#), [Card, Heining, and Kline \(2013\)](#) emphasize the increased sorting of workers in Germany across establishments. Finally, our paper relates to the literature on labour market returns to individual characteristics. Perhaps most notable are [Heckman, Stixrud, and Urzua \(2006\)](#) on cognitive and other abilities, and [Case and Paxson \(2008\)](#) and [Persico et al. \(2004\)](#) on height.

The rest of the paper is organized as follows. Section 2 describes the wage model and places it in the context of wider theoretical debates. Section 3 describes the labour market data and shows patterns of wages and employment across the UK and Germany. Section 4 presents the main results on task prices. Section 5 shows results using individual measures of quality before the final section, 6, concludes. Extensive appendices provide further details, results, tests and robustness exercises for various features of our analysis. In particular in appendix A.3 we show some comparable results for the US.

2 Framework

2.1 Econometric Framework for Price Estimation

A number of frameworks have been developed to capture RBTC-type effects explicitly, including the workhorse model by [Acemoglu and Autor \(2011\)](#). The most relevant one for the current application features a continuum of skill types assigned to a set of occupational groups or tasks that is discrete and finite. Such a framework is used by [Cortes \(2016\)](#), in turn based on the model of [Jung and Mercenier \(2012\)](#) and on [Gibbons et al. \(2005\)](#). In this model, individuals sort into the finite occupational groups based on their comparative advantage.⁶ Our framework is similar. It is based on a simple [Roy \(1951\)](#)

⁵More recently, [Taber and Roys \(2019\)](#) combine both cross-section and panel data for the U.S. Additionally, another strand of literature departs from the Roy framework and considers both tasks and skills in a continuum. [Yamaguchi \(2012\)](#) and [Lindenlaub \(2017\)](#) are structural papers in this line.

⁶See also [Cozzi and Impulliti \(2016\)](#) who examine the effect of exposure to foreign technological competition.

model of selection into occupational sectors. Let w_{ijt} be the (log) wage for individual i in sector j at time t . The utility u_{ijt} derived from working in sector j and earning w_{ijt} is given by

$$u_{ijt} = w_{ijt} + \varepsilon_{ijt}$$

where ε_{ijt} is an idiosyncratic ‘shock’ affecting preferences for sectors. In the present context, the preference component, ε_{ijt} , can be modelled with a rich structure: it need not have zero mean and need not be uncorrelated across time. We might allow the shock to have non-zero mean, for example, if a sector provides amenities, or we might let the shock depend on previous sectoral choices to allow for non-pecuniary switching costs.

We next place further structure on wages. In the following exposition we propose a wage model which identifies prices in a transparent way. The model can be made more general, but our specification matches the empirical implementation we describe in section 2.2 and present in section 4.

Wages depend on individual, time and sectoral characteristics in an additive way. Suppose that X_{it} and W_{ijt} are both sets of observable characteristics capturing individual and time-specific factors, such as education and region of work. We additionally allow W_{ijt} to capture factors that vary by sector, such as, with restrictions that are discussed below, work tenure. Importantly, and as also discussed below, W_{ijt} cannot contain factors that are collinear with age or time. The returns on X_{it} and W_{ijt} are δ_t and β_j , which are time-specific and sector-specific respectively. We assume:

$$w_{ijt} = \underbrace{\delta_t X_{it} + \beta_j W_{ijt}}_{\text{observable factors}} + \underbrace{\theta_{jt} + \gamma_{ij} + v_{it}}_{\text{unobservables}} \quad (1)$$

where the set of coefficients $\Theta \equiv \{\theta_{jt}\}$, capturing the prices of sector j at time t , is the primary object of interest. γ_{ij} is the (unobserved) ability of individual i in sector j . Finally, v_{it} is an idiosyncratic shock to wages, which might also include measurement error. Importantly, and as discussed further below, v_{it} is common across sectoral choices, and is orthogonal to ε_{ijt} .

Given this structure, the individual chooses sector simply to maximize utility. Letting $j_{it}^* = \operatorname{argmax}_j \{u_{ijt}\}$ we can then define a set of binary indicators I_{ijt} , which capture sectoral choice as follows:

$$I_{ijt} \equiv I_{it}(j) \equiv \mathbf{1}(j = j_{it}^*)$$

where $\mathbf{1}(x)$ is an indicator function equaling 1 if the argument x takes the value ‘true’, and 0 otherwise. The observed wage, w_{it} , then takes the following form:

$$w_{it} = \sum_j I_{ijt} (\gamma_{ij} + \theta_{jt} + \delta_t X_{it} + \beta_j W_{ijt}) + v_{it} \quad (2)$$

In a regression framework, the problem for the econometrician is that the sectoral choice variables, I_{ijt} , on the right hand side of equation 2, are endogenously determined. In particular, the sectoral choices depend on the unobserved skills γ_{ij} . This problem is particularly acute when data are available in the cross-section only. The econometrician typically either needs to find an instrument for sectoral choice, or to make stronger assumptions on functional form. Identifying parameters of interest in these types of environments is the subject of an established literature.⁷

In this paper, we use panel data for both the countries we study. Given panel data we can control for selection effects by differencing out the unobserved component, γ_{ij} . We therefore use a fixed-effects panel estimator to take account of the endogeneity arising from selection. Specifically, all coefficients of interest are identified by running the regression using fixed effects at the sector-individual level. This type of approach is a common way to achieve identification in a model with selection.⁸ Intuitively, the parameters are identified in this approach from sector-specific wage growth.

Identification of the model accordingly depends on two standard assumptions of fixed-effect regressions. First, the wage residuals, v_{it} , must satisfy ‘strict exogeneity’. A formal statement in this context is as follows: let I_i be the set of sectoral choices, $\{I_{ijt}\}$, across time and sectors, and let X_i and W_i be defined similarly, then $\mathbb{E}(v_{it}|X_i, W_i, I_i, \Theta) = 0$. Therefore sectoral choice must be uncorrelated with wage residuals and depend purely on the unobservables, γ_{ij} and ε_{ijt} , the observables, X_{it} and W_{ijt} and the time effects θ_{jt} . The strict exogeneity assumption therefore also implies that the residuals v_{it} must be independent of the preference shocks ε_{ijt} .

Second, identification depends on the unobservable factor γ_{ij} being fixed over time. This second condition relates to a more subtle requirement: Identifying the coefficients of interest, θ_{jt} , separately from other effects requires a somewhat restrictive structure on both the unobservables γ_{ij} and observables, X_i and W_i . Even then the prices θ_{jt} are identified in the current approach only as changes with respect to a base year and a base sector. In the results reported in section 4, therefore, prices are usually interpreted as changes with respect to those in the year 1991 and with respect to the change to the price on routine manual work. As we also discuss below, we group occupations into sectors according to their predominant task. For this reason, we use the terms ‘sector’ and ‘task’ from now on somewhat interchangeably.

2.2 Further Discussion of Identification and Empirical Implementation

The set-up presented in equation (2) entails two main identification challenges. First, as discussed, the strict exogeneity restriction requires that individuals do not move sectors because of idiosyncratic sector-specific shocks. This restriction might be problematic given that the current framework is based on

⁷See Heckman and Honore (1990), and Dahl (2002) as classic references on estimation of Roy models using cross-sectional data.

⁸In addition to Cortes (2016), who looks at sectoral prices in the US, this approach has been used in a related contexts, by, for example, Combes et al. (2008) who estimate the evolution of city wage premia.

self-selection, and that individual sector-specific shocks might indeed occur. If strict exogeneity fails then those that stay in a given occupation are increasingly positively selected according to these shocks. If the shocks in one sector are systematically larger than the shocks in others, then wage growth would in turn be systematically higher for this sector and estimates of sectoral price growth would be biased upwards.

What matters is the quantitative relevance of this type of self-selection. Unfortunately, idiosyncratic sector-specific shocks are difficult to test for directly. An indirect test is to examine wage changes for those who switch occupations. As discussed by [Card et al. \(2013\)](#), a sufficient condition for strict exogeneity is ‘exogenous mobility’. Under exogenous mobility, and in [Card et al.’s](#) framework, the average wage change of switchers from sector j to j' should be the negative of the change for switchers in the opposite direction, from j' to j . Exogenous mobility is therefore testable. In contrast to [Card et al.](#), our framework is complicated by the fact that workers’ occupational choices are only assumed invariant to wage *shocks*: workers still self-select according to occupational prices, which fluctuate over time. However, such price fluctuations should have a small effect on average wage changes at high frequencies, such as over 2 or 3 years. For this reason, we examine [Card et al.’s](#) exogenous mobility test in appendix [A.2](#), and find evidence that is generally supportive.⁹

Perhaps the main challenge to our identification, however, is to capture pure time effects on task prices aside from effects on prices coming from age and tenure. This problem is reflected in the restrictive structure of equation (2). The general problem is that within each individual-sector spell, and from one wave of the panel to the next, time, age and tenure grow collinearly.¹⁰ This general problem manifests itself in several more specific issues raised by the literature, and which we address in turn. Our overall approach to eliminating the confounding effects of age and tenure is to include in the regressions rich sets of controls. These controls go beyond those included in the analysis of [Cortes \(2016\)](#).

The first specific identification challenge, discussed by [Gottschalk et al. \(2016\)](#), is to capture age profiles in human capital accumulation that are heterogeneous by sector. Unfortunately, and without some restrictions, complete sector-specific age profiles cannot be separately identified from the time effects of interest. This is because, as discussed above, sectoral prices are ultimately identified from sector-specific wage growth. We mitigate the concern in several ways. First, we control for heterogeneity in wage profiles by including in the set of covariates X_{it} an interaction of a quartic polynomial in age with a full set of education dummies, capturing seven levels of educational attainment. This interaction is included to pick up heterogeneity by worker skill level. Second, we also include in the set of covariates W_{ijt} a quartic polynomial in job tenure. In effect we here interact job tenure with task, allowing job-specific wage growth to vary by sector.

⁹An alternative assessment is provided by [Böhm et al. \(2019\)](#). Using their exclusion restriction of zero price changes in a base period, and using lagged sectoral choice as an instrument, they can identify the component of wage growth due to idiosyncratic shocks. In terms of sectoral price estimates they find that ignoring this factor has quantitatively small effects.

¹⁰The problem can be considered the panel data equivalent to the classic time-age-cohort identification problem.

Despite this, any returns to experience that differ by sector, and that are picked up by neither heterogeneous job tenure effects nor educational profiles, will potentially confound our estimates. This issue is addressed by [Böhm et al. \(2019\)](#) who also estimate occupational prices over time in Germany, similarly to the present paper, by using a similar data structure to us. They are able to estimate heterogeneous returns to experience by using a longer time period than we do and by exploiting apparent variation in prices over time. By assuming that prices are constant in an earlier period, they obtain an exclusion restriction with which both the returns to experience and prices in the later period can be identified together. They do indeed find heterogeneity in returns, although it should be mentioned that they do not otherwise adopt the rich specification that we do. With this in mind, and third, therefore, we further address this issue by estimating our prices on a narrower group of older workers. The idea here is that by excluding young workers from the sample we remove those for whom heterogeneous human capital accumulation is greatest. This approach is in the spirit of looking at ‘flat spot wages’, as in [Heckman, Lochner, and Taber \(1998\)](#).

A further, related identification problem is also discussed by [Gottschalk et al.](#), who suggest that differences in wage growth across sectors might be affected by differences in the nature of wage-tenure contracts. We control for this in our analysis first by including the interaction of task with job tenure, as above. We also capture heterogeneity in wage profiles caused by explicit contracting factors, by including an interaction of age with trade union status.¹¹ This interaction allows for workers whose pay is governed by union or collective bargaining to have wage profiles that are, for example, initially higher, but grow less quickly.

Finally, as emphasized by [Kambourov and Manovskii \(2009\)](#), wages are determined not only by tenure in the job, but also by occupational tenure itself.¹² When considering this factor, it should first be emphasized that the model can easily capture any occupational tenure effects that are common across sectors. Intuitively, common effects do not pick up any sector-specific differences in wage growth. Second, occupational tenure can be additionally included in W_{ijt} , and hence allowed to be sector specific, given further restrictions. The restriction in, for example [Kambourov and Manovskii \(2009\)](#) and in [Cortes \(2016\)](#), is that occupation-specific human capital is lost completely whenever an occupation spell finishes. In this way, occupational tenure is defined to not change collinearly with age, and identification is achieved. We discuss these issues further in section 4, although we relegate a full analysis to appendix A.3.

We summarize our empirical implementation as follows. We estimate a version of equation (2) by running regressions of wages on a set of individual-task fixed effects together with task-year fixed effects. We augment these regressions with rich sets of controls. These controls are, in all specifications, region of residence, marital status, a quartic polynomial in job tenure, a quartic in age, and interactions of age with

¹¹Although these controls may not capture all differences in implicit contracting across sectors, they should capture the more explicit contracting differences caused by national pay bargaining.

¹²A related point is that wages decline faster the further individuals move away from their initial occupation. See [Gathmann and Schoenberg \(2010\)](#).

education. We further include, in various specifications, interactions of job tenure with task, age with union status, task with trade union status and education with time. We also vary the set of ages in our sample frame. When computing standard errors we allow for arbitrary serial correlation in the residuals by clustering at the individual level.

2.3 Further Discussion of Sorting

The framework used by [Cortes \(2016\)](#) predicts that when technology causes a decline in demand for routine tasks, the relative price of the other tasks/occupations increases, as does their employment share.¹³ The observed wage effects and employment growth therefore depends on sorting, the precise nature of which depends on the distribution of latent skills, γ_{ij} .

The theoretical models, such as that presented in [Jung and Mercenier \(2012\)](#), typically incorporate variation in skills along a single dimension. In these models, employment in equilibrium is perfectly segmented into different sectors, with the highest skilled sorting into the top sector, and the lowest skilled sorting into the bottom. In terms of implications for the observed average sectoral wages, this distribution of skills results in relatively little attenuation, because when workers reallocate from one sector to the next, they move from, say, the top of a lower-ranking sector to the bottom of a higher-ranking sector, and the average quality of both sectors declines roughly in unison. In contrast, the econometric framework used to estimate task prices here allows for the distribution of γ_{ij} to be potentially richer. For example, it allows for latent skills to be potentially uncorrelated across sectors.¹⁴ In this case, those who are sucked into the top sector may come from anywhere in the distribution of the other sectors. Accordingly, attenuation of average wages can be much greater.

There are good reasons to believe that the correlation of skills across sectors is positive but far from perfect. An obvious reason is that residual wages overlap substantially across sectors. This is true even though, as shown later in section 3, the occupational sectors are clearly ordered by average wage within each country. The overlap in residual wages therefore implies that some individuals receive good wages in otherwise low-paying sectors because their skills there are particularly high. In a different vein, data on abilities, such as that from the US-based NLSY typically shows a positive but imperfect correlation across sector-specific skills ([Böhm, 2020](#)).

We conclude this section by extending the wage model to provide a framework for our later results on sorting. Given a positive but imperfect correlation on underlying skills we can rewrite the sector-specific

¹³See also [Cozzi and Impulliti \(2016\)](#) who examine the effect of exposure to foreign technological competition.

¹⁴This is the ‘independent productivity shocks’ model in the language of [Gottschalk et al. \(2016\)](#) and of [Willis \(1986\)](#). The case when skills are unidimensional is captured by the ‘hierarchical ability’ model.

wage equation (1) as follows:

$$w_{ijt} = \underbrace{\dots}_{\text{observable factors}} + \theta_{jt} + \lambda_j \zeta_i + \xi_{ij} + v_{it} \quad (3)$$

where we have glossed over the observable factors. The new terms are a ‘universal’ skill type, ζ_i , the sector-specific payoff to this skill, λ_j , and other sector-specific skills, ξ_{ij} , which are now mean zero, orthogonal to ζ_i , though not necessarily orthogonal to each other. In short, we have factorized γ_{ij} as $\lambda_j \zeta_i + \xi_{ij}$. Now suppose that sectors are ordered in terms of their average wage, such that $\mathbb{E}_i(w_{ijt} | j_{it}^* = 1) < \dots < \mathbb{E}_i(w_{ijt} | j_{it}^* = N)$. This ordering can be generated by imposing $\lambda_N > \dots > \lambda_1$ and $\theta_{1t} > \dots > \theta_{Nt} \forall t$. Intuitively, individuals with low ζ_i on average sort into a low-ranking sector because of the ranking of the θ_{jt} terms, while individuals with high ζ_i on average sort into the higher ranking sectors. Changes in wages and in sorting over time are driven by changes in task prices θ_{jt} , while task-specific returns λ_j are assumed constant over time.

We can use equation (3) to unify the analyses in sections 4 and 5. In section 4 we estimate the task prices θ_{jt} . The price estimates and the observed average wage changes are consistent with strong sorting based on ζ_i and ξ_{ij} . We examine this sorting further in section 5, using data on individual characteristics. The characteristic we focus on in particular is height. As we discuss later, height has been shown to be correlated strongly with outcomes across all tasks and related to all underlying abilities. We therefore consider height as a good proxy for ζ_i , the latent unidimensional factor.

3 Labour Market Data

Our wage analysis uses publicly-available micro datasets. From Germany we use the German Socio-Economic Panel (SOEP). From the UK we use the British Household Panel Survey (BHPS). The BHPS and SOEP provide internationally-comparable labour market data covering similar long time spans, and including comparable and extensive covariates.

3.1 Labour Market Variables

At least since Autor et al. (2003) it has been recognized that changes in employment structure are strongly associated with changes in the demand for job tasks. For this reason we proceed by grouping occupations at the 1-digit classification level according to the type of their predominant task, using the approach in Acemoglu and Autor (2011). This approach is standard in the literature on assessing the extent of routine-biased technical change.

For the basic assignment of workers to occupations, we use the 1988 International Standard Classification of Occupations (ISCO-88). This classification relates closely to most national systems, particularly to the UK 1990 Standard Occupational Classification (SOC90). We use the ISCO classification, rather than relying on national systems, for comparability of results. For robustness, however, we repeat the main analysis using the SOC90 and report results in appendix B.

Following [Acemoglu and Autor \(2011\)](#), we then construct four broad task groups by merging the occupational categories of ISCO-88. We follow a similar categorisation to that which applies to US data.¹⁵ The first group includes all the non-routine cognitive occupations: these are ‘Legislators’, ‘Professionals’, and ‘Technicians and associate professionals’. We term this group “abstract”. The second group clusters occupations which are both routine and manual. These are ‘Craft and related trade workers’, and ‘Plant and machine operators and assemblers’. The third group includes routine cognitive occupations: ‘Clerks’, and ‘Sales’. The fourth group includes mainly non-routine, manual occupations: ‘Service workers’, ‘Skilled agricultural and fishery labourers’ and ‘Elementary occupations’.¹⁶

We explored alternative approaches extensively. Specifically, we also mapped the occupational data to occupational classifications in the US Department of Labor’s Dictionary of Occupational Titles (DOT) and its successor, the Occupational Information Network (O*Net). We then grouped occupations using the task information in these surveys more directly. These alternative approaches, and the resulting task price estimates are discussed in appendix B, where we also discuss the German Klassifizierung der Berufe (KldB). Of course, all these approaches are broad and provide a connection with job tasks that is fairly loose. For example, since only predominant tasks are used to categorize workers, we neglect the fact that jobs typically require tasks of many types, the prices on each of which may be changing. Similarly we neglect the fact that the task composition within sectors has changed over time ([Spitz-Oener, 2006](#)), and so the return to workers skills evolve for this reason too. Nevertheless, our approach is standard in the literature. Additionally, as described in section 4, we also go beyond the simple grouping described here to examine occupational groupings at the 1-digit level.

As is well-known, a complication when using the occupational status variable is measurement error. In both surveys, a non-negligible number of individuals appear to switch occupation one year, only to revert immediately to their original occupation the next. For this reason we derived a ‘smoothed’ measure of occupation, explained in appendix B. In fact, when estimating task prices, both smoothed and raw

¹⁵Note that we implicitly assume that the task content of each occupation is similar in the UK, in Germany, and in the US.

¹⁶ISCO-88 does not have a major group for Sales occupations, differently from the national classifications based on the Standard Occupational Classification. Part of sales occupations are with service occupations. Others are spread in the other major groups. However, according to [Acemoglu and Autor \(2011\)](#), whereas service workers do a non-routine manual job, salesmen involve mainly routine cognitive tasks. Therefore, we create an extra occupational category for Sales occupations. By confronting the UK 1990 SOC and the US 2010 SOC, we classify workers as salesmen if their ISCO-88 code is one of the following: 3415, 3416, 3417, 9113, 5220, 5230. Notice as well that recent analyses, based on the classification of [Dorn \(2009\)](#), exclude agricultural workers (see [Goos et al., 2014](#); [Autor and Dorn, 2013](#); [Cortes, 2016](#)). In our sample, we only consider employed individuals, and therefore, the percentage of workers in agriculture is low. We keep the individuals in this occupational category, although excluding them does not affect the estimates of the task prices, as we show in Appendix B.

measures provide similar results (see also appendix B). We therefore mainly use the raw data throughout the paper. However, we use the smoothed measure when examining occupational mobility explicitly, such as for the strict exogeneity test presented in appendix A.2.

Our wage measure is constructed as the log of current gross monthly earnings divided by weekly working hours and multiplied by 12/52. The earnings measure captures usual pay received by employees in their current main job before tax and other deductions. Similarly, the measure of hours worked includes overtime. In SOEP, this variable is bounded at 80 hours per week. We therefore bound it similarly in the BHPS to ensure comparability of results.¹⁷ Finally, wages are deflated by the local 2010 consumer price indices.

3.2 Construction of the Labour Market Samples

The UK sample consists of the original BHPS survey running from 1991 to 2008, and excludes subsequent booster samples which were designed to increase representation from Wales, Scotland, Northern Ireland, and households with low incomes. Similarly the German sample consists of the original West German SOEP survey running from 1984 onwards and excludes the booster samples added subsequently. Specifically, and consistently with the existing literature, we exclude East Germany because its wage structure is very different from the West (see, for example, [Dustmann, Ludsteck, and Schönberg, 2009](#)).¹⁸

For wages, we consider men in their prime age, between 25 and 60, unless otherwise stated. We focus on prime-aged males for the usual reason that we aim to estimate market prices using a group that is as homogeneous as possible in terms of labour market attachment. The lower bound on age of 25 removes labour market entrants, many of whom do casual or part-time jobs. We set the upper bound at 60 as those working after this age are also increasingly selected. Additionally, because we rely on the panel dimension, individuals are only included if information is provided about their earnings in at least 5 waves.¹⁹

For this wage sample, we exclude self-employed or workers in the armed forces. Because we consider hourly wages and we do not want to add extra selection criteria, we include both full- and part-timers.²⁰ Finally, we exclude observations with missing values of occupation or labour force status, number of

¹⁷Only 0.65% of the observations concerning employees reported a total number of hours larger than 80.

¹⁸See [Wagner et al. \(2007\)](#) for further details on the scope and evolution of SOEP, as well as on the differences and similarities with BHPS and PSID.

¹⁹We do not trim the hourly wage upwards or downwards. One consideration here is that the minimum wage in the UK was first introduced only for some occupations and only in the late-1990s. The consequence of our choice is to keep some non-credibly low wages, for example lower than 2 GBP/EUR. Therefore, we checked whether excluding the lower tail of the distribution affects our main results. The estimated occupational prices in the main analysis are consistent with results on a sample of workers excluding the bottom 1% and the top 0.25%.

²⁰Only 2.5% and 2.7% of the prime-age sample in Britain and Germany respectively work more than 0 but less than 30 hours per week.


worked hours and years of job tenure (if employed) or completed level of education. As the appendix table A.1 indicates, the resulting wage samples consist of 24,364 and 36,918 observations from the UK and Germany, coming from 2,285 and 2,739 individuals respectively. The table also reports further summary statistics.

The statistics for total employment are constructed using a wider sampling frame consisting of all males in the labour market aged 16-64. Of course, when computing market prices and quantities, various levels of aggregation may be of interest. Here we choose a wider group of males for our quantity concept, but in fact all our results look very similar when using the narrower frame. Finally, in this paper, we mostly ignore females. In fact, market prices for females might differ from those for males, if males and females are imperfect substitutes in the labour market. Of course, estimating labour market prices properly is harder for females because of further issues of labour-market selection. To provide a rough check, however, we estimate prices and quantities for females simply, and without correcting for selection, in appendix A.4, where we find results that are remarkably consistent with those for males.

3.3 Trends in Employment and Wages

Before turning to the main results, we show overall trends in employment and wages across the two countries. Figure 1 plots the fraction of workers in each occupational sector over time. As discussed above, we classify the occupations into four broad groups, according to the type of their predominant task.²¹

Figure 1: Employment by Occupational Sector



Notes: Based on a sample of 16 to 64 y-o employees. ‘Abstract’ stands for abstract task. ‘R Man’ indicates routine manual. ‘NR Man’ indicates non-routine manual. ‘R Cog’ indicates routine cognitive.

²¹The figures in Appendix B indicate that the patterns identified here below are robust across occupational classifications.

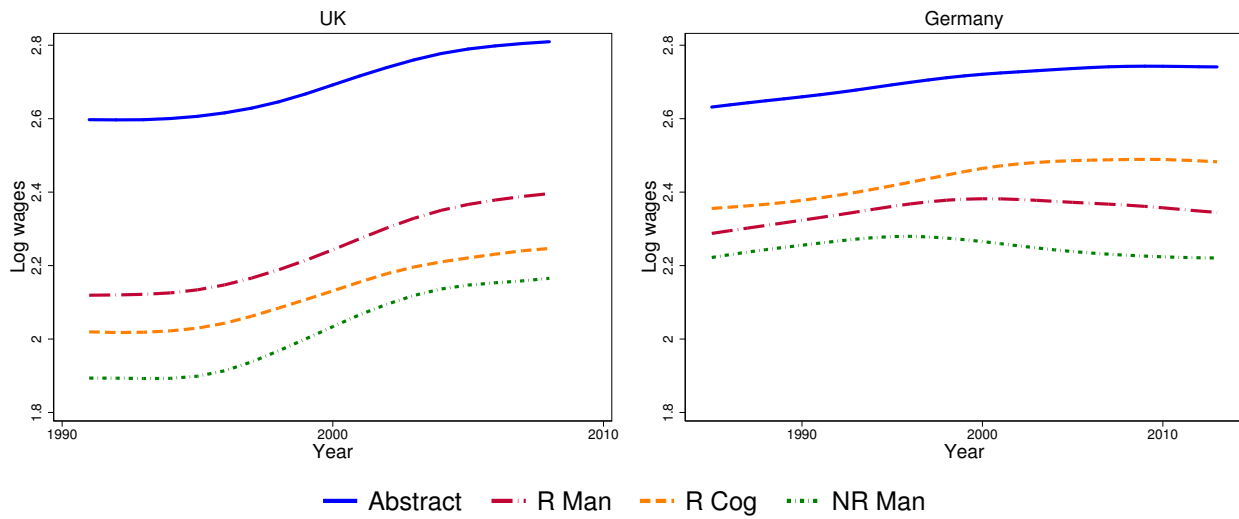
The figure shows that, in both countries, males belong mainly to two categories of jobs. In the UK, shown in the left-hand panel, abstract and routine manual workers account for 70% of workers, averaged over all years. As for the trend over time, the most striking feature is the large increase in the fraction of abstract workers, partially offset by the decrease in the percentage of routine manual workers. In 1991, these two sectors each represented around 35% of the sample. Over 1991 – 2008 the abstract share then grew relative to routine manual by around 15 percentage points. For Germany, shown in the right panel, the patterns are even more pronounced. On average over 75% of men in the sample work in an abstract or in a routine manual job. Similarly to the UK, the relative shift away from routine manual jobs to abstract was large, amounting to around 20 percentage points over 1991-2008, or over 25 percentage points over 1985-2013. In this country, however, routine manual workers decreased more than in the UK in absolute terms: over 10 percentage points versus over 5 percentage points for the same period.

Turning to wages, figure 2 plots the raw median wage for each of the four main occupational groups for UK (left panel) and German (right panel) workers. In both countries workers in abstract occupations earn by far the most, and workers in non-routine manual occupations earn the least. Nevertheless the two countries display some noteworthy differences. First, wage growth was higher in the UK than in Germany over the sample periods as a whole, and particularly when we compare 1991 to 2008. For example, over this comparable period, median routine manual wages in the UK grew by around 28%, but in Germany, by only 10%. Second, the figure indicates that whereas routine manual workers on average earn more than routine cognitive workers in the UK, in Germany the average wage order is reversed. However, remember that here we are displaying the median of raw wages. As we show later in section 4, when we condition on covariates, wages in the two routine groups line up similarly in both countries, and, if anything, workers in routine manual are comparatively better paid in Germany than in the UK.

Most strikingly, perhaps, the figure shows differences in the trends in inequality across countries, and between occupational sectors. Sectoral wages in Germany diverged, and inequality increased. The increase in inequality in Germany more generally has been the subject of a prominent literature (Dustmann, Ludsteck, and Schönberg, 2009 and Card, Heining, and Kline, 2013). In the UK, on the other hand, the occupational wage structure was roughly flat. The fact that employment in abstract occupations grew so fast while relative average wages were flat conflicts, *prima facie*, with an explanation of occupational shifts based on changes to the demand for labour. One of the main aims of this paper is to explain these patterns taking into account shifts in employment composition.

As a final piece of descriptive discussion we link both employment changes and wage changes to initial 1991 wages at the 1-digit ISCO88 level. The results are shown in figure 3. It shows observed changes alongside quadratic smoothed plots, following Goos and Manning (2007). For employment shares, on the left-hand side, the figure shows a remarkable similarity in patterns across countries, indicative of job polarization. In both countries, abstract occupations (ISCO88 categories 1, 2 and 3) experienced the highest employment growth and are also characterized by the highest wages. And in both countries,

Figure 2: Median Log Wage by Occupational Sector



Notes: Based on a sample of 16 to 64 y-o employees. ‘Abstract’ stands for abstract task. ‘R Man’ indicates routine manual. ‘NR Man’ indicates non-routine manual. ‘R Cog’ indicates routine cognitive. The vertical axis shows logs of hourly wages in GBP for both countries.

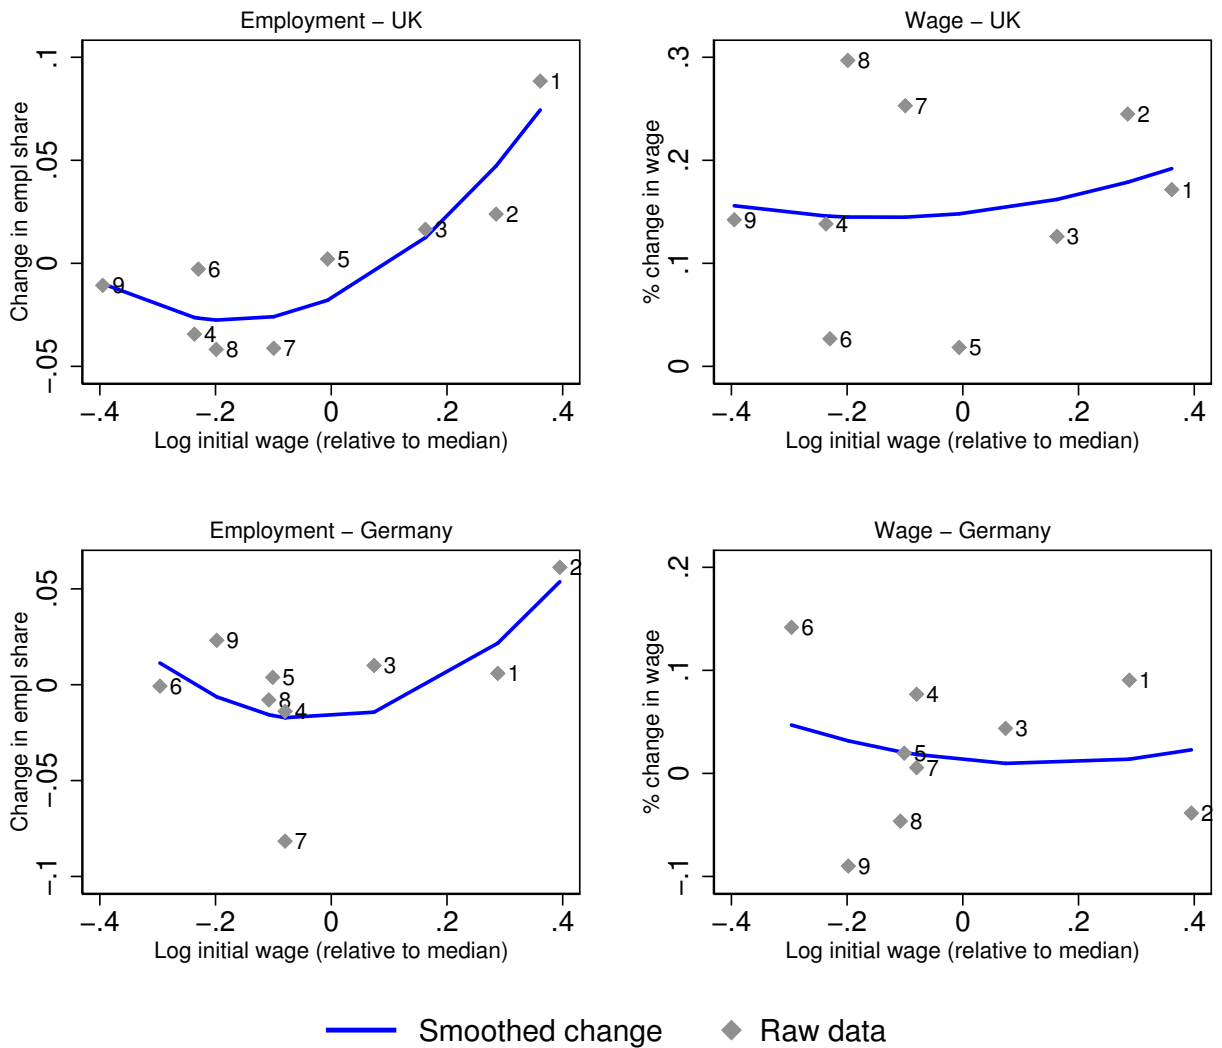
occupations in the middle of the wage distribution experienced mild declines in employment share, while occupations at the bottom of the wage distribution were roughly static. For wages, the evidence is much less clear; the smoothed profiles indicate mild polarization, but to a much lower degree. Moreover, whereas the raw scatter points show a clear pattern for employment, the pattern for wages is harder to discern. In the next section, we compare the change in employment share not with initial occupational wages or wages changes but with changes in occupational price, and argue that the evidence on prices across both countries is more strikingly consistent with changes to demand for labour.

4 Estimates of Task Prices

We now show the estimates of the sectoral/task prices, using the fixed-effect regressions. In this section, we discuss the prices of all occupational sectors, but we focus mainly on the abstract price versus that for routine manual work. This is because, in both countries, these sectors still comprise the bulk of employment. Moreover, the relative shift in employment for these groups has been particularly large: routine manual has seen the largest declines in employment share, and the abstract sector the largest gains. This focus also links closely to the ensuing analysis in section 5 where we examine ‘knowledge’ workers classified, with greater emphasis, by degree-holding status.

Figure 4 shows the main results for both the UK and Germany, split into two columns of panels. The right-hand panels show the estimates for prices on each task relative to routine manual and indexed to 0 in 1991. These panels show raw estimates by year as well as smoothed profiles generated using a 3-period

Figure 3: Employment and Wage Changes by 1-digit Occupation and by Initial Average Wage



Notes: Left-hand side shows changes in employment share over 1991-2008. Right-hand side shows changes in mean log wages over the same period. These units are consistent with those used in figure 4. Smoothed profiles are predicted values from a regression of the outcome measure on log wage and wage squared in the initial year, as in [Goos and Manning \(2007\)](#). The label represent the 9 categories of the 1-digit ISCO-88: 1 ‘Legislators’, 2 ‘Professionals’, 3 ‘Technicians and associate professionals’, 4 ‘Clerks’, 5 ‘Service and sale workers’, 6 ‘Skilled agricultural and fishery labourers’, 7 ‘Craft and related trade workers’, 8 ‘Plant and machine operators and assemblers’, 9 ‘Elementary occupations’.

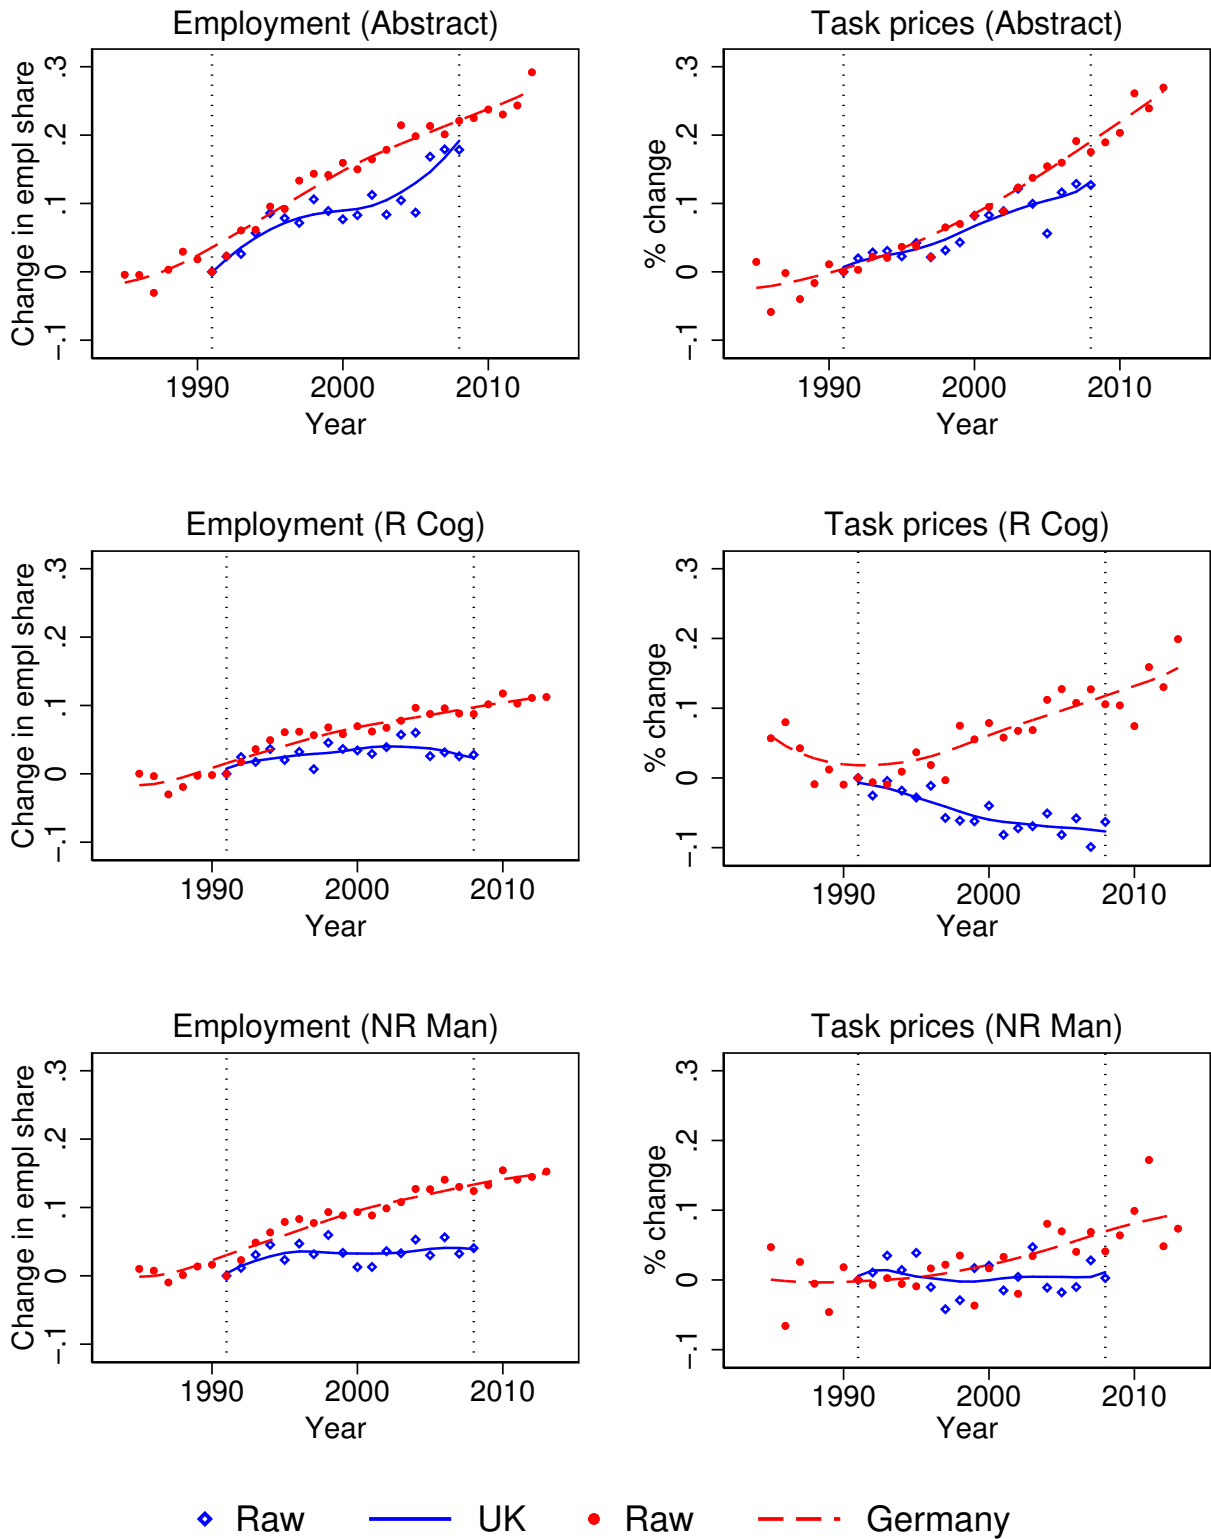
moving average. The left-hand panels show percentage growth in the employment of males in each task relative to routine manual, again from 1991. The vertical lines in all panels indicate the start and end years of the BHPS.

The top row of figure 4 shows that, compared to routine manual, task prices for abstract jobs increased markedly in both countries. In the UK, the price increased relatively by around 13% by 2008, or a little over a half a percent per year. The relative price change in Germany was even larger, reaching around 18% by 2008. In Germany, moreover, the relative task price continued to increase after 2008, after the global financial crisis, at roughly the same pace. Although not indicated here, these price estimates are significantly different from base for both countries at the 1% level from the mid-2000s onwards. Linking these to changes in labour supplied, the top row left panel shows that changes in employment look markedly similar. Employment in abstract occupations in the UK grew relatively by 18 percentage points over 1991-2008. Again it also shows that relative employment growth was even stronger in Germany. There, it grew by around 22 percentage points over 1991 to 2008, again continuing after the global financial crisis.

As discussed in section 3, statistics for employment are computed using a broader population than is used to estimate the prices. In particular, and as discussed in section 3, here we use all males aged 16 – 64, to provide a more comprehensive measure of labour supplied. More broadly, it is worth remembering that female employment also saw a large shift towards abstract occupations; even larger, in fact than for males. In Germany, in particular, whereas both genders witnessed an increase in abstract occupations, women experienced a much larger decrease in routine occupations than men (Black and Spitz-Oener, 2010). We show equivalent estimates for prices and for employment changes for females in appendix A.4. Overall, and in conclusion therefore, it seems that there was a strong increase in demand for abstract occupations in both countries. Moreover, it seems, this increase was stronger in Germany than in the UK, in line with the extra increase in the occupational price shown in the top right-hand panel.

The bottom two rows of figure 4 show the evolution of prices and employment over time for both countries and for the remaining two sectors: routine cognitive and non-routine manual. Again we show both prices and quantities relative to the routine manual task. In both countries, these sectors employ, or have employed, a large fraction of women, but their share of male employment has always been smaller. In terms of prices, the non-routine manual task shows little systematic difference from the base category. Its price appears to have grown a little in Germany since 2000, although the estimate is rarely significant. The prices on the routine cognitive sector, however, are noticeably different across countries. The point estimate is negative in the UK, though it is only mildly significant (not shown). In Germany, on the other hand, the routine cognitive price has risen significantly compared to routine manual. It is first worth pointing out that both routine cognitive and non-routine manual employment grew more in Germany relative to routine manual, than in the UK. It is also worth pointing out that these price changes can still be reconciled within the Roy framework. With multiple occupations, then the relationship between employ-

Figure 4: Price and Employment Changes by Task: Relative to Routine Manual



Notes: The figure shows changes in employment and in task price of workers in each sector with respect to manual workers. Task prices estimates come from models with controls for region, marital status, year dummies, quartics in age and in job tenure, and an interaction between the age polynomial and education, as well as individual-sector fixed effects. Estimates correspond to those in column 2 in Table 1. Employment change is computed for 16 to 64 y-o men. 'Abstract' stands for abstract task. 'NR Man' indicates non-routine manual. 'R Cog' indicates routine cognitive.

ment changes and price changes can be complex. For example, if routine cognitive jobs in Germany are a close substitute to abstract jobs in terms of latent skill requirements, an increase in price can be consistent with little change in employment. This pattern would happen if the price on the abstract task grows faster. This is because, even though the equilibrium price on both occupations increases, individuals move to the close substitute instead. Nevertheless, the overall impression from figure 4 is that employment changes do match price changes closely in both countries.

As discussed in section 2, we need to check that our results do not depend on confounding factors, such as heterogeneous tenure profiles. We therefore perform our analysis on alternative specifications. Results are shown in table 1. The results from the benchmark model, and shown in the figure above, are summarized in the second column. It shows the relative growth in prices for each sector between 1991 and 2008 only, coming from the interaction of task and year dummies, and omitting estimates from intervening years. The prices for the UK are given in the top panel; those for Germany in the bottom panel. Here we show results for Germany for 1991-2008 only, for direct comparison with the UK.²² Corresponding with the figures above, the column shows that the price on the abstract task grew by 13% relative to routine manual occupations over 1991 to 2008 in the UK and around 18% in Germany. The results for routine cognitive and for non-routine manual also correspond to those shown in figure 4.

Before looking at alternative specifications we show results from raw OLS regressions in the first column. These OLS regressions use the same specification as the benchmark fixed-effect model. Most importantly, and as discussed in section 2, we control for heterogeneous wage profiles over age by including interactions of a polynomial in age with a full set of education dummies. The OLS regressions also pick up the level of average sectoral wages in the base year, 1991. The results show that abstract jobs have always paid substantially more than other sectors, in both the UK and Germany, even conditional on other observable characteristics. They also imply that, at least in 1991, routine manual jobs were comparatively better paid in Germany than in the UK. Most importantly, the first two columns show markedly different results for the growth in sectoral prices and wages. For example, the OLS results for the UK imply that average wages in the abstract sector did not grow relative to routine manual. The results from fixed effects, on the other hand, which address selection into each sector based on unobservable characteristics, show that the growth in the relative price on abstract occupations was pronounced. For Germany, the OLS results show that average wages in abstract jobs did grow, in terms of magnitude, relative to routine manual. Yet still, the coefficients are not statistically significant. Moreover, the second column shows that growth in prices was even larger. These columns therefore highlight the difference between growth in average wages, which includes changes to the average quality of workers in each sector, and changes to pure sectoral prices, which capture the price paid to an effective unit of labour supplied.

The first two columns of table 1 also show results for the other sectors, still relative to routine manual. In both other sectors, and in both countries, sectoral prices grew faster relative to routine manual jobs than

²²Note that the coefficients corresponding to the other years are plotted in figure 4.

Table 1: Changes in Task Price 1991-2008: Additional Specifications

	OLS (1)	Benchmark FE (2)	Educ. Returns (3)	Job Tenure (4)	TU (5)	35-60 y-o (6)	FE fewer controls (7)
UK							
1991 Abstract	0.259*** (0.03)						
1991 R Cog	0.055* (0.03)						
1991 NR Man	-0.033 (0.04)						
Δ Abstract	0.010 (0.04)	0.126*** (0.04)	0.161*** (0.04)	0.126*** (0.04)	0.124*** (0.04)	0.078* (0.05)	0.157*** (0.03)
Δ R Cog	-0.146*** (0.05)	-0.060 (0.04)	-0.047 (0.04)	-0.075 (0.05)	-0.059 (0.04)	-0.126** (0.06)	-0.058 (0.04)
Δ NR Man	-0.079 (0.05)	0.002 (0.05)	0.001 (0.05)	-0.014 (0.05)	-0.004 (0.05)	-0.043 (0.06)	0.014 (0.05)
Constant	2.169*** (0.03)	2.484*** (0.06)	2.491*** (0.06)	2.485*** (0.06)	2.463*** (0.06)	2.352*** (0.06)	2.568*** (0.05)
Observations	24364	24364	24364	24364	24364	16955	24364
Adjusted R^2	0.376	0.247	0.246	0.247	0.250	0.145	0.240
Germany							
1991 Abstract	0.157*** (0.02)						
1991 R Cog	0.028 (0.02)						
1991 NR Man	-0.071*** (0.03)						
Δ Abstract	0.042 (0.03)	0.175*** (0.03)	0.156*** (0.03)	0.203*** (0.04)	0.181*** (0.03)	0.110** (0.05)	0.224*** (0.03)
Δ R Cog	0.062 (0.05)	0.106*** (0.04)	0.108*** (0.04)	0.072 (0.06)	0.108*** (0.04)	-0.007 (0.07)	0.121*** (0.04)
Δ NR Man	-0.046 (0.05)	0.041 (0.04)	0.031 (0.04)	0.061 (0.08)	0.053 (0.04)	0.038 (0.13)	0.056 (0.05)
Constant	2.674*** (0.03)	2.767*** (0.07)	2.767*** (0.06)	2.764*** (0.06)	2.766*** (0.07)	2.790*** (0.08)	2.867*** (0.06)
Observations	36918	36918	36918	36918	36032	27124	36918
Adjusted R^2	0.309	0.144	0.142	0.144	0.146	0.058	0.133
Controls:							
Baseline controls	x	x	x	x	x	x	
Education*Year			x				
Job tenure*Task				x			
Union*(Task,Age)					x		

Notes: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Standard errors clustered at individual level in parenthesis. For models (1) to (6), controls include region, quartics in age and in job tenure, interaction between age and education, marital status, year dummies. In (1), the benchmark model is estimated with OLS; in (2), with panel fixed effects; in (3), an interaction between education and year is added, and the interaction of education with age is dropped; in (4), model (2) is augmented with the interaction between occupation and job tenure, and between year and job tenure; in (5), model (2) is augmented with a dummy for trade union (or collective bargaining) association along with its interactions with occupation and with age; in (6), model (4) is estimated on a group of mature workers, from 35 to 60; (7) is based on model (2) but dropping job tenure, education and the interaction of education with age.

did average wages. However, for both these sectors, and in both countries the difference between wage and price changes is smaller than for abstract jobs, and the standard errors are slightly larger.

A potentially confounding explanation for our results is changes in returns to education. This is because occupational choice is correlated with educational status, and so apparent changes in returns to the former may be explained by changes in returns to the latter. *A priori*, this factor is unlikely to be important in the UK at least, because most analyses show that the return to education was flat over the period (Blundell, Green, and Jin, 2018). Nevertheless, we take account of these changing returns by including in the regressions interactions of time dummies with a full set of education dummies. To do this, we have to remove the interaction of education with age. The relevant controls are therefore the interaction of education with year and the polynomial in age. Results are shown in table 1 in the third column, and are very similar to those for the benchmark regressions.

Also discussed in section 2, our estimates are potentially improved by controlling flexibly for tenure effects. For example, and as discussed by Gottschalk et al. (2016), if wages in routine manual jobs have a flat tenure profile, because of implicit contracting considerations, then its estimated task price growth may be biased downwards, or the relative price growth of the other tasks may be biased upwards. We therefore control for this factor by including an interaction of sector with a quartic polynomial in job tenure. The estimates are shown in the fourth column. For the UK, these estimates are very similar to the benchmark. For Germany, it is noticeable that the estimated growth in the abstract price is pushed up. This suggests that tenure profiles, if anything, flatten wages in the abstract sector in Germany, and that workers in this sector receive larger wage growth when switching employers.

The fifth column of table 1 shows results when we control for contracting effects in a different way, by controlling for trade union status. Our union status variable here is designed to pick up workers whose pay is governed by collective bargaining. For the UK, we use the measure contained in the BHPS on the presence of union or staff association at the workplace. For Germany we use public information on coverage of collective agreements, which we map into SOEP at the 2-digit industry level.²³ Here we interact the union status of the worker separately with sector and with age, to pick up heterogeneous effects. The results are identical to the benchmark regression for the UK, and very similar for Germany.

Next, we check the robustness of results by restricting the sample to a group of mature workers aged between 35 and 60. This sample removes completely those in early career and restricts to those in mid- and late- career. We do this to address concerns about heterogeneous and unobserved human capital effects, particularly in early career. This sample is chosen to respond specifically to the criticisms raised by Gottschalk et al. (2016), discussed above and in section 2. We therefore also include the controls for

²³In more detail for Germany, we use the tariff archives of WSI at www.tarifvertrag.de (Statistisches Taschenbuch Tarifpolitik). We use the information about the collective agreements in place in historic years. We then use the 2-digit NACE indicator to identify the industry in SOEP data. We create a dummy equal to 1 if that industry was covered by a collective agreement.

sector-specific job-tenure effects. The results are shown in the sixth column. It shows that the increase in task price on abstract jobs remains robust, even though the point estimates are slightly lower than in the benchmark regressions.

Finally, we show the effects of relaxing the controls used in our benchmark model. This final specification, shown in column 7, differs from the benchmark by excluding job tenure, education and the interaction of age with education, bringing it closer to the benchmark in [Cortes \(2016\)](#).²⁴ The results, shown in the final column, show that the estimated coefficients are higher than in the benchmark, and noticeably so in Germany. In an additional exercise (not shown) we relaxed the age restriction on our sample to include all working-age males as in [Cortes](#), with little effect on estimates. Additional results from all the regressions featured in table 1 are presented in appendix sections A.3 and A.4. In appendix B we show results when using different occupational classifications.

It is worth comparing our results against the logic of the bounds on task prices derived by [Gottschalk et al. \(2016\)](#). They derive an upper bound on the increase in price on the abstract task by computing statistics after trimming the bottom of the abstract wage distribution. The idea is that this bound captures the extreme case of selection according to a Roy model. In their case they compute medians, but here we compute means, using the regression function. Focussing on this sector, and using the UK as an example, employment in abstract jobs increased by around 13 percentage points, or a third of the 2008 total. Therefore we obtain a quick estimate of the upper bound by comparing the raw mean in the abstract sector with the mean obtained from trimming a third of wages from the bottom. We do this for the abstract sector in 2008 using residual wages after regressions on age and education to take account of observable factors. We find that the trimmed mean is 25 log points higher than the raw mean. We can compute an implied selection component from our estimates in table 1 by comparing the fixed-effect estimates in the second column with those in the first column. These imply a selection effect of around 11 log points, comfortably within this upper bound. The fact that the point estimate is well below the upper bound implies that selection into abstract jobs is mostly, but not always, below the average wage.


We build on our results further by examining effects at a finer level of occupational detail. Specifically, we aggregate closer to the 1-digit level: we break down the abstract sector into legislators and managers, professionals and associates. Similarly we break down routine manual into machine operatives and crafts. We do not break down the occupations going into routine cognitive and non-routine manual groups, because these are already small and the resulting estimates are imprecise.

The results of this analysis, for both the UK and Germany, are shown in figure 5. For both countries we use changes in employment share and changes in sectoral prices over 1991 to 2008. The figure shows that, for both countries, the task price increase and employment growth are strongly correlated. For example, within the abstract sector, and in the UK, the largest employment growth has been for

²⁴This model differs from that in [Cortes \(2016\)](#) by excluding the control for union status, by including controls for age, and by sampling a narrower age group.

legislators (managers), who have also seen the price on their tasks rise the fastest. In the interpretation of the Roy model, this implies that it is for legislators that the average quality of employees has fallen the most as this group has become increasingly negatively selected. For completeness, and to confirm this intuition, figure 6 shows the flat relationship between employment share growth and growth in average wages. These average wages are actually the coefficients from an OLS regression on sector and detailed covariates. They therefore capture composition changes based on observables.

Figure 5: Changes in Price and Employment: Finer Occupational Grouping




Notes: Employment change is computed on a sample of 16 to 64 y-o men. Categories of the same shape and colour belong to the same task group. The size of the markers is proportional to the 1991 employment share. Task prices are estimated relative to crafts. The straight line is the coefficient of a regression of changes in task prices on changes in employment shares over 1991-2008.

We finish this section by discussing some additional analyses that we show in the appendix. We now consider occupational-specific human capital accumulation, along the lines of the discussion in section 2.2. It is worth re-emphasizing here that we have already allowed for heterogeneous tenure effects extensively, although we have focussed not on occupational tenure but instead on job tenure. We have done this for two reasons. First, as discussed, specific occupational tenure effects can be identified separately from time effects only by imposing quite strong restrictions. In contrast, heterogeneous job-tenure effects can be specified cleanly: these job-tenure effects are naturally lost completely whenever the individual switches jobs, assuming that they never return.²⁵ Second, and more practically, job tenure is recorded in our data as a direct survey question. The question is retrospective, therefore we have information about job tenure also for jobs that have started before the first waves. Measures of occupational tenure used in Cortes (2016) or Kambourov and Manovskii (2009), on the other hand, need to be computed by hand from the sample observations. The occupational tenure of the first job of each individual is therefore

²⁵Notice that job tenure and occupational tenure are far from co-linear. Several authors, such as Groes, Kircher, and Manovskii (2015), have documented that individuals often change occupation within the same firm, and switch jobs without switching occupations. We also see this in our data. Nevertheless there remains a sizeable correlation between job and occupational switches.

Figure 6: Changes in Average Wage and Employment: Finer Occupational Grouping



Notes: Employment change is computed on a sample of 16 to 64 y-o men. Categories of the same shape and colour belong to the same task group. The size of the markers is proportional to the 1991 employment share. Average wage changes are computed from the coefficients of an OLS regression (equivalently to the model in the first column of table 1). They are estimated relative to crafts.

left-censored. This being said, we perform an analysis controlling for heterogeneous occupational tenure profiles in this way in a robustness exercise, shown in appendix A.3, where we find results that are very similar to our benchmark.

Finally, we compare our task price estimates to the results in Cortes (2016) for the U.S. Results are shown in appendix A.3 in figure A.4, where we analyse the PSID data used in Cortes, but apply our sampling framework and estimation specification. The figure shows the changes in prices are remarkably consistent across countries. A slight difference for the U.S. is comparatively higher growth in the price on non-routine manual work. Overall, employment changes in the U.S. (not shown) were also in line with changes in the UK and Germany over the period; as discussed by Beaudry et al. (2016) employment in routine jobs fell drastically compared to the abstract sector. Again a slight difference for the U.S. was the relatively larger growth in manual work, particularly after 2000, when there was an overall dampening of demand in the labour market.

5 Direct Evidence on Worker Quality

In section 4 we argued that the price on occupational sectors correlates strongly with employment share growth. Moreover, both price and employment growth were shown to be highest for the abstract sector. A corollary of this finding is that the quality of workers in abstract occupations has declined over time. However, the approach taken relies on identifying restrictions that are difficult to test using wage data alone. To strengthen the argument, therefore, in this section we provide more direct evidence of changing

worker quality across broad groups. We use evidence on two measures of quality: test scores for cognitive ability and, with greater emphasis, height.²⁶

Height is of great use in assessing worker quality. It has long been documented that height is associated strongly with labour earnings. In the analysis of [Case and Paxson \(2008\)](#), an extra inch (2.5 cm) of height for men is associated with an increase in earnings of around 2 percent. Height is associated not only with physical capabilities, but is also correlated with social and intellectual capabilities ([Persico, Postlewaite, and Silverman, 2004](#) and [Lundborg, Nystedt, and Rooth, 2014](#)). In line with the discussion in section 2, therefore, we interpret height as a universal measure of quality. Equally importantly, adult height is plausibly exogenous to choices of education and occupation that we examine here.²⁷ In terms of identification, therefore, and as we discuss, height offers substantial advantages over cognitive ability as a measure. For these reasons, height is our preferred metric of quality in the present analysis.

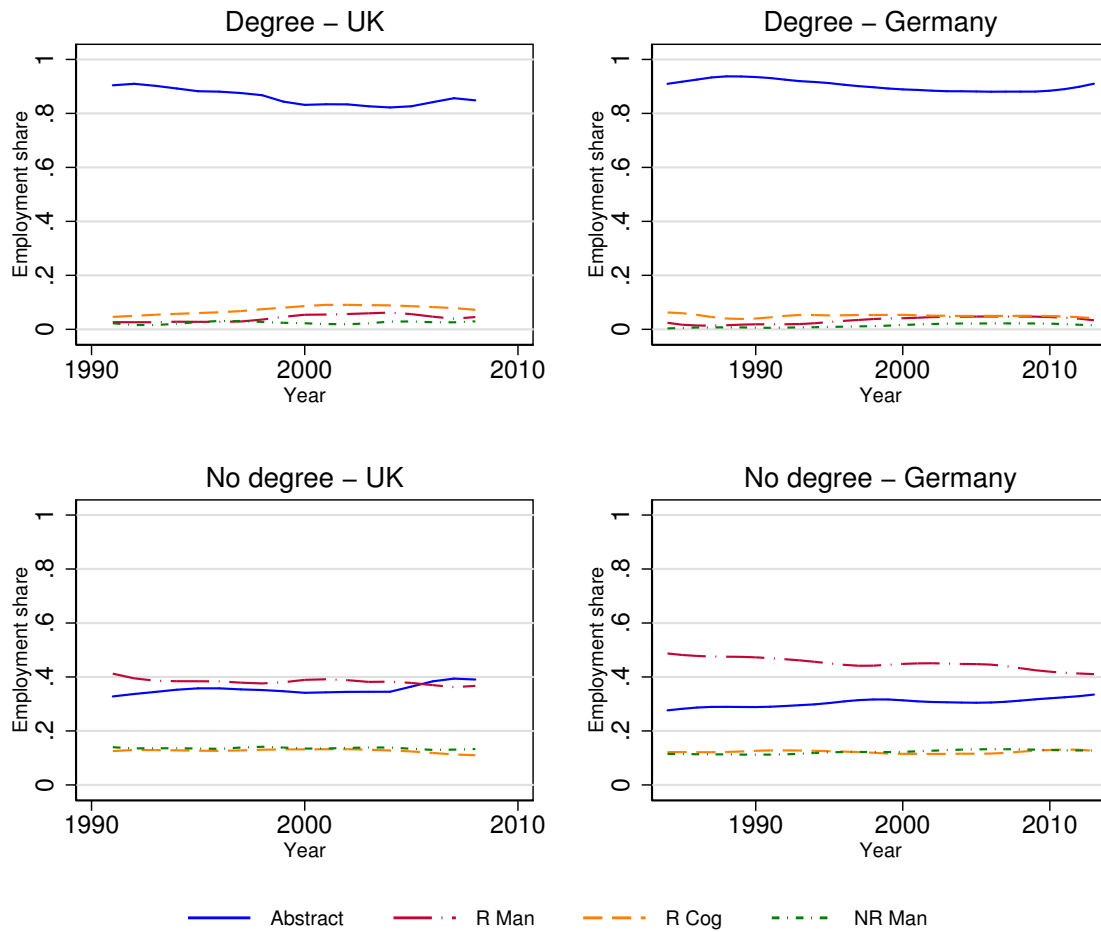
Even with the measurements we have, assessing worker quality across sectors is challenging. The ideal analysis would provide an aggregate measure of quality across the different occupational sectors both at the beginning of our period (around 1990) and at the end (around 2010). Unfortunately this analysis places almost unachievable requirements on the data, because very few measures of quality are available from before around 2000. For this reason we favour other approaches, using measures of both quality and sectoral choice that are themselves fixed over working life. Our most favoured analysis exploits height data and compares, not occupational sectors, but educational attainment. Specifically, we focus on degree holders compared to non-degree holders. In this way we can exploit the data on characteristics in their entirety, and particularly the unusually rich data from the UK. For the UK in particular, we find convincing and robust evidence that the composition of heights has changed over time. We find similar evidence for Germany, although the data there are more sparse.

By pursuing this analysis we are effectively using degree holding as a proxy for working in the abstract sector. This approach is valid if the increase in degree holders is driven by the increase in demand for abstract work, and the change in educational attainment rate is a movement along a supply curve. Of course, we recognize that our results may be explained by differential selection patterns into education rather than into occupations. Specifically, if the supply of degree holders has increased even faster, and this group has accordingly moved down the occupational ladder on average, then our results would be overstated. To assess this issue in more detail, we show occupational choice by educational attainment over time for both the UK and Germany in figure 7. For the UK the proportion of male degree holders

²⁶In theoretical terms, recall from section 2 that I_{iat} is the indicator for selection into the ‘knowledge’ sector, here defined as performing abstract work. Here, we investigate the change over time in $\mathbb{E}_i(\zeta_i | I_{iat} = 1) - \mathbb{E}_i(\zeta_i | I_{iat} = 0)$, claiming that changes to this selection term are driven by movements to θ_{jt} . It should be noted that, when the unobservables follow a Gaussian distribution, then changes to this selection ‘difference-in-characteristics’ term are theoretically determined. This selection difference declines as the smaller sector expands, and grows as the smaller sector shrinks. In our application the smaller sector is expanding, so we should expect a convergence of average characteristics. However, when unobservables are not Gaussian then this theoretical prediction is not so clear cut. Ultimately what matters for the present paper is whether the selection difference did indeed decline, which is purely an empirical question.

²⁷We discuss potential endogeneity as this section develops.

Figure 7: Occupational Sector by Educational Status for Males



Notes: Employment is for males aged between 25 and 64 working full time. Educational qualifications are classified according to the International Standard Classification of Education (ISCED) reported in the surveys.

working in the abstract sector declined initially, but then rose towards the end of the sample period such that the beginning-to-end fall was less than 5 percentage points. For Germany, the proportion is almost completely flat. For both countries, the figure indicates that non-degree holders (bottom row) have been increasingly likely to work in abstract jobs, but again that this change has been moderate. It seems, therefore, that the correlation between degree holding and abstract work has weakened slightly, but overall that the former has been a stable proxy for the latter.

As we discuss below, we can in principle develop this analysis further by bringing in additional data sources and looking at height measures from earlier periods. By doing this we examine occupational attainment directly at comparable ages. As also discussed, we can examine a more direct measure of worker quality in cognitive ability. All these additional analyses, however, bring their own problems. We therefore use these as robustness exercises, discussing them later in this section and in appendix A.5.

5.1 Evidence from Heights

Data

The data used in this section for Germany come, as before, from the SOEP. For the UK we augment data from the BHPS with data from its successor, Understanding Society (UKHLS).²⁸ The UKHLS is an ongoing household panel survey, originating in 2009. Relative to the BHPS, the SOEP and similar surveys around the world, the UKHLS has some unrivaled and relevant strengths. These strengths are its focus on biomarker data and its unusually large sample size. The collective samples of the UKHLS consist of approximately 40,000 households, far larger than the BHPS or SOEP, or say, the PSID.

For both the UK and Germany, we use self-reported height measurements, capturing height without shoes. In both countries, the accuracy of the response is assessed by the interviewer. In the UK the interviewer also prompts a reassessment if the respondent claims to be less than 150cm or taller than 200cm (or the equivalent in feet and inches). In Germany, where measurements are collected in multiple survey waves, height reports that are over 10cm different from at least two other reports are automatically deleted and replaced with imputations. As we discuss later, the UKHLS also provides clinical height measurements obtained from nurse visits. We don't use these in our main results because the number of measurements is fewer. However, as we discuss below, these measurements do offer a check on final results. They also offer a check on the raw data: Overall self-assessed and nurse-measured height show a correlation across the UK sample as a whole of over 0.95.

For the UK, the self-reports were requested from all respondents in UKHLS wave 1, taken in 2009-10, as well as in BHPS waves 14 and 16, taken in 2004 and 2006. In total this mode yielded over 47,000 height measurements in 2009 and over 15,000 in the BHPS waves from 2004 and 2006. For Germany, height self-reports have been collected in the SOEP every two years since 2002. Across the whole SOEP, and up to 2012, there are over 129,000 valid height measurements, from over 40,000 distinct individuals.

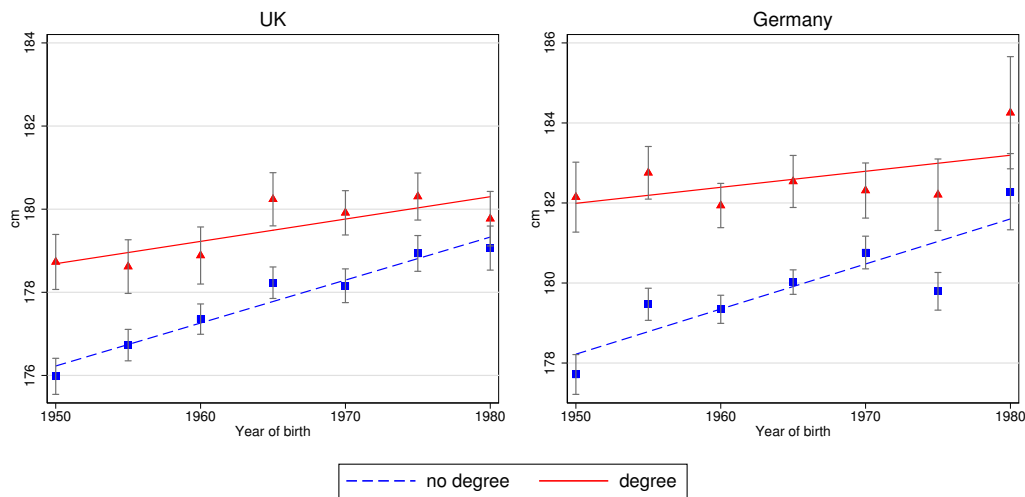
For the UK analysis, we use the main UKHLS sample (the general population sample) together with the main BHPS sample. UKHLS also includes a large sampling frame targeted at ethnic minorities, which we exclude. For Germany, to improve sample size, we combine the original SOEP sample with the subsequent booster samples, excluding the 1990 East German sample. Given that our chosen samples are not completely nationally representative, in all analyses we use appropriate survey weights. After selecting on the originating samples, on age, and using only males, we use over 11,700 self-reported heights for the UK. For Germany we use over 25,000 observations, but these come from a little under 7,000 individuals. Fuller information on sampling and sample sizes is given in the overview of descriptive statistics, provided in appendix table [A.5](#).

²⁸The abbreviation UKHLS comes from Understanding Society's alternative name, the UK Household Longitudinal Survey.

Results

To motivate the results, figure 8 shows the key piece of evidence in this section in graphical form. It shows mean height by education group and by 5-year birth cohort since 1950 across both the UK and Germany. The figure shows averages by cohort bin marked as dots, alongside 95 percent confidence intervals, and alongside linear regression lines. The bin marked as 1950, for example, captures those born between 1950 and 1954. The averages here are obtained after controlling for ethnicity, using race dummies for the UK and parental country of origin for Germany. For presentation purposes, we also adjust for age and use estimated maximal height. We do this by adding 0.1cm to measured height for every year of age above 40, as suggested by [Merrill and Richardson \(2009\)](#). Here we use the original samples only to minimize the need for further adjustment. The figure shows clearly that average heights have converged across groups: in both countries the difference between groups has declined by around 50 percent from the oldest cohort to the youngest. It should be noted that the overall increases in height are consistent with the evidence in [Hatton \(2013\)](#) on trends shown over the longer-term across Europe.

Figure 8: Mean Height by Cohort and by Educational Attainment



Notes: Data based on the BHPS and UKLHS original samples for the UK, and on the 1984 West German and migrant samples for Germany using self-reported height. Sample is 25-60 year-old males, born between 1950 and 1984, split into 5-year bins according to birth date. Height is adjusted by ethnicity. The vertical bars show 95% confidence intervals.

We now investigate the convergence of heights more formally, presenting the results in table 2. The table shows regressions from the UK in the left-hand panel, and from Germany on the right. All regressions use males aged strictly between 25 and 60, use sample weights, feature a dummy for holding a degree interacted with year of birth, and then include various controls. In addition to birth year we also include a quadratic in age. This age effect is identified because of the collection of measurements across multiple years. For the UK, we report robust standard errors. For Germany, where we usually have multiple reports for each individual, we include all observations and cluster standard errors at the individual level.

Table 2: Height by Cohort and by Educational Attainment

	Panel A. UK				Panel B. Germany					
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
Degree	1.911*** (0.195)	2.122*** (0.284)	2.058*** (0.192)	2.231*** (0.280)	1.912*** (0.182)	1.864*** (0.328)	2.030*** (0.473)	1.736*** (0.327)	1.918*** (0.473)	1.642*** (0.315)
Degree × Year of birth	-0.040** (0.016)		-0.034** (0.016)		-0.031** (0.015)	-0.056* (0.031)		-0.057* (0.031)		-0.050* (0.030)
Degree × (Gen X-BB)		-0.556 (0.368)		-0.436 (0.361)			-0.490 (0.637)		-0.554 (0.635)	
Degree × (Gen Y-BB)		-0.977** (0.429)		-0.872** (0.422)			-1.087 (0.850)		-1.002 (0.841)	
n	11790	11790	11779	11779	11779	6,857	6,857	6,850	6,850	6,850
Observations	11790	11790	11779	11779	11779	25439	25439	25403	25403	25403
Adjusted R ²	0.028	0.028	0.056	0.055	0.057	0.070	0.068	0.101	0.098	0.101
Controls:										
Age, Region, Native Born	x	x	x	x	x	x	x	x	x	x
Ethnicity			x	x	x			x	x	x
No extreme heights					x					x

Notes: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$ Sampling weights used in all regressions. Robust standard errors in Panel A, in parentheses. Standard errors clustered at individual level in Panel B, in parentheses. Gen Y born after 1975. Gen X born 1961-1975. BB, or Baby Boomers, born 1950-1960. Panel A: Coefficient on degree in first row refers to male born in 1960.

The first columns of each panel (columns 1 and 6) show the raw interaction of education with birth year, controlling only for age, region and for being native born, either in the UK or Germany. The first row here shows the predicted average difference in heights between education groups for males born in 1960. The second row shows the interaction of education with year of birth. This interaction term is strong; according to the point estimates heights converged by 1.2cm over 30 years for the UK, and 1.7cm over 30 years for Germany. As with all results reported here, the estimates are more precise for the UK than for Germany, because, as indicated previously, the sample sizes for the UK are larger. The number of observations and number of distinct individuals used are shown in the bottom rows.

In columns (2) and (7) we show an interaction not with a linear year-of-birth term, but instead with three cohort dummies. Specifically we split the sample into those born before 1960 (the ‘Baby Boomers’, now shown as the reference group in the first row), those born between 1961 and 1975 (‘Generation X’) and those born after 1975 (‘Generation Y’). These columns show a steady convergence in heights for the later groups, which is again significant for the UK, but not for Germany. In columns (3) and (8) we control also for ethnicity. As [Case and Paxson \(2008\)](#) discuss, this control is important, because there are noticeable differences in height by ethnicity which do not correspond with other measures of ability and do not correspond with labour earnings. In the UK data, ethnicity can be captured by dummies for race. In SOEP, where direct questions about race are precluded, we use the parental country of origin. In the regressions, the interaction term now goes down slightly for the UK, indicating that the ethnicity of degree holders has changed somewhat there, with relevant composition effects. Nevertheless the interaction term is still significant and quantitatively significant. It implies a convergence of heights in the UK of around 0.9cm over 30 years. In columns (4) and (9) we use the interaction with cohort dummies. In columns (5) and (10), we windsorize the top and bottom 2.5 percent of heights, to show that the results are not driven by outliers. We don’t show the full cohort dummies here, but the results follow the patterns shown from previous columns.

To conclude the discussion of these findings, we discuss results when we include full interactions of degree holding with both birth year and age separately, using the multiple years of the BHPS and UKHLS for the UK, and SOEP for Germany. Although both interactions are identified theoretically, this identification is weak because the number of observations from, for example, the BHPS is relatively small, the gap in time between measurements in both countries is short. As suspected, the standard errors on the main interaction (birth year with degree holding) increase substantially. In the UK for example, they go up from around 0.015 to 0.07. The estimates are therefore too imprecise to use, and so we do not report them.²⁹

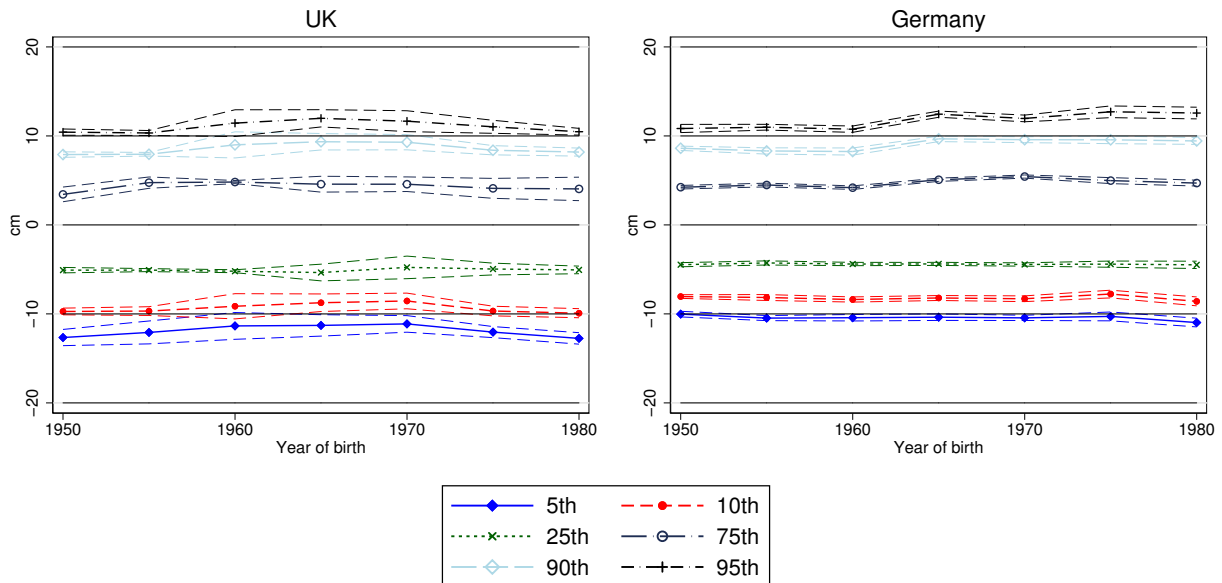
To probe this evidence further it is worth considering two other possible explanations for the effects shown. First, it is possible that age profiles of adult height are not common across socio-economic

²⁹Incidentally, the point estimates in a variety of regressions are substantially negative, indicating strong convergence. In the richest specification, and for the UK, the p-value on the hypothesis of no convergence is 0.19.

status. In fact, [Huang et al. \(2013\)](#) find evidence that those with lower education in China shrink faster than the highly educated. This phenomenon would generate the pattern shown in figure 8 which comes mainly from cross-sectional variation. However, the populations studied here are very different to the Chinese population studied by [Huang et al. \(2013\)](#), who spent much of their adult life in poverty. Indeed, in a related study of old age shrinking for the UK using the English Longitudinal Study of Ageing, [Fernihough and McGovern \(2015\)](#) find no significantly different changes for degree holders versus non-degree holders. In fact their point estimates suggest that degree holder shrink *faster*, consistent with the notion that taller individuals shrink faster in general.³⁰

Second, and in a related way, it is possible that inequality in maximally-attained heights has been declining over time. This declining inequality would also generate a convergence on average between groups. We examine this possibility in figure 9, where we show quantile plots of self-reported height by birth cohort, pooling all individuals. The quantiles are shown as absolute differences from the cohort median and shown with 95 confidence intervals. The figure shows that the distribution as a whole has not become noticeably compressed. In fact, when comparing a raw standard deviation of heights, the data indicate that inequality in heights has *increased* in both countries over time. Finally, given that the oldest men we examine here are aged below 60, it is unlikely that survivor bias plays a role.

Figure 9: Quantile Differences from Median Height by Cohort



Notes: Figure shows various quantiles of male heights as absolute differences in cm from median. Sample is 25-60 year-old males, born between 1950 and 1984, split into 5-year bins according to birth date. Data based on the BHPS and UKLHS original samples for the UK, and on the 1984 West German and migrant samples for Germany using self-reported height. Dashed lines show 95% confidence interval on quantile difference obtained using a bootstrap with 1000 repetitions.

Finally, we describe briefly how we may push our analysis here further. As discussed when introducing

³⁰It should be noted that females show different patterns of height decline, and in fact shrink much faster than men. See [Sorkin et al. \(1999\)](#).

this section, an ideal analysis would incorporate measures of worker quality from the beginning of our sample period (around 1990), and would examine not education, but occupation choice at fixed ages. For the UK at least, we can extend the analysis in this way, by bringing in the rich data from the country's cohort studies, and specifically the National Child Development Survey (NCDS). The NCDS provides data on children born in 1958, with their heights measured at age 33 in 1991. We can therefore compare the heights for this cohort by occupational sector, with those at comparable ages in the UKHLS around 20 years later. We discuss results for these data fully in appendix A.5, and find them to be consistent with what is already shown for the UK. We relegate this extra analysis to the appendix mainly because the sample sizes considered are much smaller, and consequently estimates are less precise.

5.2 Evidence on Cognitive Scores

To finish this section we show some available evidence from cognitive scores. We do this because, although the analysis using height is overall stronger, cognitive ability is more closely linked to productivity, and is therefore more commonly used in economic analyses in general. For cognitive ability, the data from UKHLS are very rich. In comparison, the usable sample sizes in SOEP are less than a tenth as large. In this subsection we therefore drop Germany and examine the UK only.

Data

The data come from a specific module in wave 3 of UKHLS carried out in 2011-12.³¹ The module follows much of the design used by the US Health and Retirement Survey, and subsequently employed elsewhere. It includes, for example, tasks on word recall, and completing numerical/mathematical tasks. The module was given to all adult respondents who provided full interviews.³² The total number of responses is over 35,000. The number of usable responses for men aged between 25 and 60 on, for example, the immediate word recall test is 11,504.

We use scores on 4 tests: completing number series; immediate word recall, delayed word recall, and verbal fluency. Each of these questions is designed to assess a different facet of cognitive functioning. For example the word recall test assesses episodic memory, whereas the number series assesses fluid

³¹Even within country, comparing measures of cognitive ability across surveys is not as straight forward as comparing heights, because even slight changes in test procedure can potentially make comparison impossible. (See [Altonji et al., 2012](#) for an analysis combining cognitive assessments in the NLSY79 with the NLSY97. See also [Attanasio et al., 2020](#) for an analysis comparing socio-emotional characteristics across two cohort studies in the UK.) For this reason, here we use data from UKHLS alone.

³²The vast majority of interviews were conducted face-to-face, although some telephone interviews were also performed. Questions were translated and asked in the respondent's native language. Again, we use the main UKHLS sample excluding the ethnic minority boost. This includes original BHPS members carrying on into the UKHLS frame.

reasoning.³³ A complete description of the data can be found in [McFall \(2013\)](#). As we discuss shortly, we also form principal components of the survey measures to compute a measure of cognitive function that is uni-dimensional.³⁴

Results

Just as for height our analysis is based on comparing across cohorts and across educational attainment. We show results from regressions of various ability measures on year-of-birth interacted with degree holding. To control for curvature in the age profile of test scores we also control for a quadratic in age, although now the first-order term is not separately identified from year of birth, given the cross-sectional nature of the data.

We present the results in [table 3](#). The table shows six separate regressions using the various measures. The first four regressions show the individual types of cognitive functioning. As discussed, these metrics are themselves derived from a battery of assessment questions. The first row shows, as expected, that degree holders exhibit substantially higher test scores than non-degree holders. The second row shows the coefficient on the interaction of education and age. It is negative for all the measures, though only significantly so for verbal reasoning and immediate word recall. These negative point estimates indicate convergence for later birth cohorts. Finally, the bottom two rows show results from a non-linear interaction of cohort with degree holding. We do this by grouping males into the same three cohort groups as above: Generation Y (born after 1975), Generation X, (born 1961-1975), and Baby Boomers, (born 1950-1960). The results are consistent with those in the top panel.

We synthesize these findings on the right-hand side by using combined indices of these scores. We do this by using the first principal component from a factor analysis. The approach of taking a single component is in line with theories of general cognitive ability, which posits a strong single underlying factor to cognitive functioning ([Carroll, 1993](#)).³⁵ This first component corresponds strongly with all the measures, though it loads most on the number series. As the table shows, there is a significant convergence for the high education group with the low education group. Quantitatively, the results indicate that, over 30 years, the gap in test scores between the two groups has declined by over 20 percent. This relative decline is less than for height, but still quantitatively non-negligible. To show that this result is not driven by any

³³In more detail, the number series test consists of a number of logic-style puzzles and as designed to assess fluid reasoning. The immediate word recall task involves reading a list of 10 words to the respondent and asking them to repeat as many of the words as possible. In the delayed word recall task, the respondent is asked to repeat the words after a separate test has been administered. Both of these test episodic memory. The verbal fluency test requires respondents to name as many instances of a category (such as animals) as possible. It requires mental flexibility, organization and abstraction.

³⁴The cognitive ability module also includes tests on numeric ability and subtraction. We excluded these because the response distributions are non-gaussian and, in fact, double peaked. Hence, they do not fit into the factor analysis so well. Including these, however, has little effect on our results.

³⁵The first principal component explains 51 percent of the variance of the scores, and is the only component with an eigenvalue greater than 1. We therefore focus on this first component in line with the Kaiser criterion for factor selection.

one response alone, we then compute a principal component which *excludes* the most salient component: immediate word recall.³⁶ We see a convergence over time that is lower, but still significant at the 5 percent level. Finally, examining the cohort interactions with education status on the bottom rows, the results indicate that the main action is between the earliest cohort and the latter two. Overall, these results are consistent with those from [Ashworth et al. \(2017\)](#), who, using the two waves of the NLSY, also show that the ability scores of US degree holders has declined over time relative to non-degree holders.

To finish this section, it is worth providing context and briefly relating the analysis to that for heights shown above. Most importantly, and as discussed extensively, cognitive function is less plausibly exogenous than height to education and to labour market outcomes. However, the literature has found little evidence of systematic differences in cognitive declines across education groups. The evidence that most closely relates to the present study comes from [Rusmaully et al. \(2017\)](#), using the large scale UK Whitehall II cohort study. They find that, if anything, those with more education see faster declines, presumably because their baseline level of functioning is so much higher. If this is the case, it implies that the results here are attenuated, and the true cohort effect is in fact even larger.

Table 3: Cognitive Ability by Cohort and by Educational Attainment in the UK

	Number series	Verbal fluency	Immediate Word recall (IWR)	Delayed Word recall	1st Principal component	1st Principal component, excl. IWR
Panel A						
Degree	0.642*** (0.026)	0.450*** (0.028)	0.557*** (0.026)	0.478*** (0.026)	1.031*** (0.037)	0.886*** (0.032)
Degree × Year of birth	-0.025 (0.020)	-0.041* (0.021)	-0.046** (0.021)	-0.030 (0.020)	-0.075** (0.029)	-0.059** (0.025)
Observations	10215	10587	10469	10678	10047	10185
Adjusted R^2	0.080	0.039	0.080	0.082	0.122	0.112
Panel B						
Degree	0.654*** (0.039)	0.473*** (0.041)	0.625*** (0.038)	0.520*** (0.038)	1.120*** (0.055)	0.935*** (0.048)
Degree × (Gen X-BB)	-0.047 (0.053)	-0.047 (0.056)	-0.171*** (0.052)	-0.091* (0.052)	-0.213*** (0.076)	-0.118* (0.066)
Degree × (Gen Y-BB)	-0.045 (0.060)	-0.115* (0.062)	-0.110* (0.060)	-0.095 (0.059)	-0.201** (0.087)	-0.156** (0.074)
Observations	10215	10587	10469	10678	10047	10185
Adjusted R^2	0.080	0.039	0.081	0.082	0.122	0.113

Notes: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$ Robust standard errors in parentheses. Gen Y born after 1975.

Gen X born 1961-1975. BB, or Baby Boomers, born 1945-1960.

Coefficient on degree in A refers to male born in 1960. In all models, a quadratic polynomial in age is included.

³⁶Again the analysis implies the importance of a single underlying factor, which this time explains 59 percent of the variance of the scores. Again this factor loads slightly more on number series.

6 Conclusions

Across most developed economies, the occupational structure has shifted substantially over at least the last 30 years. This shift has followed a pattern of job polarization, in which employment in routine jobs in the middle of the wage distribution has declined, and employment in non-routine and, in particular, abstract jobs at either end of the wage distribution has increased. Identifying the causes of this polarization requires examining the equilibrium movement of both employment and wages across occupations. In this paper we estimate prices of broad occupational groups taking into account composition effects. We do so using panel data for two large developed economies: the UK and Germany. We find a noticeable increase in the price on top-earning sectors that correlates strongly with employment growth. Overall, our estimates are consistent with occupational shifts being caused by changes in the demand for different types of labour, such as by changes to technology.

Our paper therefore contributes to the literature on polarization by rationalizing movements in wage data. While the US labour market of the 1990s showed strong polarization in wages (Acemoglu and Autor, 2011), the evidence from other countries has been far from clear-cut (see Salvatori, 2015 for the UK, and Dustmann, Ludsteck, and Schönberg, 2009 for Germany). Our paper shows the importance when explaining the wage structure of accounting for composition effects. As such it complements related recent papers which similarly explore this factor (Gottschalk, Green, and Sand, 2016, Cortes, 2016). In particular Böhm, von Gaudecker, and Schran (2019) also find that composition effects obscure a strong increase in prices on abstract work in Germany.

One strength of our paper is that it provides consistent and coherent evidence across two important economies. Indeed we find that occupational change in Germany was more pronounced than in the UK. Accordingly, we find that the estimated price on abstract jobs at the top of the distribution grew faster in Germany too. Similarly, when looking at a finer level of occupational aggregation we find that the evidence on price growth and employment growth lines up across countries.

We provide additional evidence supporting the underlying mechanism using information on individual characteristics from both countries. Specifically we use two high quality proxies for individual productivity across jobs: height and, for the UK, cognitive ability. To use the data to its fullest we also focus more on degree holding as a correlate of working in abstract jobs. We find robust evidence of composition effects due to sorting over time, and a decline in the relative quality of these knowledge workers.

References

- ACEMOGLU, D. AND D. AUTOR (2011): “Skills, Tasks and Technologies: Implications for Employment and Earnings,” in *Handbook of Labor Economics*, ed. by D. C. Ashenfelter and Orley, Elsevier, vol. 4, Part B, 1043–1171.
- ALTONJI, J. G., P. BHARADWAJ, AND F. LANGE (2012): “Changes in the characteristics of American youth: Implications for adult outcomes,” *Journal of Labor Economics*, 30, 783–828.
- ANTONCZYK, D., B. FITZENBERGER, AND U. LEUSCHNER (2009): “Can a task-based approach explain the recent changes in the German wage structure?” *Jahrbücher für Nationalökonomie und Statistik*, 229, 214–238.
- ASHWORTH, J., V. J. HOTZ, A. MAUREL, AND T. RANSOM (2017): “Changes across cohorts in wage returns to schooling and early work experiences,” Tech. rep., National Bureau of Economic Research.
- ATTANASIO, O., R. BLUNDELL, G. CONTI, AND G. MASON (2020): “Inequality in socio-emotional skills: A cross-cohort comparison,” *Journal of Public Economics*, 104171.
- AUTOR, D. H. (2015): “Why Are There Still So Many Jobs? The History and Future of Workplace Automation,” *The Journal of Economic Perspectives*, 29, 3–30.
- AUTOR, D. H. AND D. DORN (2013): “The Growth of Low-Skill Service Jobs and the Polarization of the US Labor Market,” *The American Economic Review*, 103, 1553–1597.
- AUTOR, D. H., F. LEVY, AND R. J. MURNANE (2003): “The Skill Content of Recent Technological Change: An Empirical Exploration,” *The Quarterly Journal of Economics*, 118, 1279–1333.
- BEAUDRY, P., D. GREEN, AND B. SAND (2016): “The great reversal in the demand for skill and cognitive tasks,” *Journal of Labour Economics*, 34, S199–S247.
- BEERTEN, R., L. RAINFORD, AND A. JONES (2001): “Changing to Standard Occupational Classification (SOC) 2000: dual coding on the Labour Force Survey.” Tech. rep., Labour Market Division, Office for National Statistics.
- BLACK, S. E. AND A. SPITZ-OENER (2010): “Explaining Women’s Success: Technological Change and the Skill Content of Women’s Work,” *Review of Economics and Statistics*, 92, 187–194.
- BLUNDELL, R., D. GREEN, AND W. JIN (2018): “The UK education expansion and technological change,” *IFS Working Paper*.
- BÖHM, M. (2020): “The price of polarization: Estimating task prices under routine-biased technical change,” *Quantitative Economics*.

- BÖHM, M., H.-M. VON GAUDECKER, AND F. SCHRAN (2019): “Occupation Growth, Skill Prices, and Wage Inequality,” .
- CARD, D., J. HEINING, AND P. KLINE (2013): “Workplace heterogeneity and the rise of West German wage inequality,” *Quarterly Journal of Economics*, 128, 967–1015.
- CARROLL, J. B. (1993): *Human cognitive abilities: A survey of factor-analytic studies*, Cambridge University Press.
- CASE, A. AND C. PAXSON (2008): “Stature and status: Height, ability, and labor market outcomes,” *Journal of Political Economy*, 116.
- COMBES, P., G. DURANTON, AND L. GOBILLON (2008): “Spatial wage disparities: Sorting matters!” *Journal of Urban Economics*, 63, 723–742.
- CORTES, G. M. (2016): “Where have the middle-wage workers gone? a study of polarization using panel data,” *Journal of Labour Economics*, 34, 63–105.
- COZZI, G. AND G. IMPULLITI (2016): “Globalization and Wage Polarization,” *Review of Economics and Statistics*, 98, 984–1000.
- DAHL, G. (2002): “Mobility and the return to education: Testing a Roy model with multiple markets,” *Econometrica*, 70, 2367–2420.
- DORN, D. (2009): “Essays on Inequality, Spatial Interaction, and the Demand for Skills,” Tech. rep., PhD dissertation 3613, University of St. Gallen.
- DUSTMANN, C., J. LUDSTECK, AND U. SCHÖNBERG (2009): “Revisiting the German wage structure,” *Quarterly Journal of Economics*, 124, 843–881.
- FERNIHOUGH, A. AND M. E. MCGOVERN (2015): “Physical stature decline and the health status of the elderly population in England,” *Economics & Human Biology*, 16, 30–44.
- FORTIN, N. M. (2006): “Higher-Education Policies and the College Wage Premium: Cross-State Evidence from the 1990s,” *The American Economic Review*, 96, 959–987.
- GATHMANN, C. AND U. SCHOENBERG (2010): “How General is Specific Human Capital? A Task-Based Approach,” *Journal of Labor Economics*, 28, 1–49.
- GIBBONS, R., L. F. KATZ, T. LEMIEUX, AND D. PARENT (2005): “Comparative Advantage, Learning, and Sectoral Wage Determination,” *Journal of Labor Economics*, 23, 681–724.
- GOOS, M. AND A. MANNING (2007): “Lousy and lovely jobs: The rising polarization of work in Britain,” *The review of economics and statistics*, 89, 118–133.

- GOOS, M., A. MANNING, AND A. SALOMONS (2009): “Job polarization in Europe,” *The American Economic Review*, 58–63.
- (2014): “Explaining Job Polarization: Routine-Biased Technological Change and Offshoring,” *American Economic Review*, 104, 2509–26.
- GOTTSCHALK, P., D. GREEN, AND B. SAND (2016): “Taking Selection to Task: Bounds on Trends in Occupational Task Prices for the US, 1984–2013,” *mimeo*.
- GROES, F., P. KIRCHER, AND I. MANOVSKII (2015): “The U-Shapes of Occupational Mobility,” *The Review of Economic Studies*, 82, 659–692.
- HATTON, T. J. (2013): “How have Europeans grown so tall?” *Oxford Economic Papers*, 66, 349–372.
- HECKMAN, J. AND B. HONORE (1990): “The empirical content of the Roy model,” *Econometrica: Journal of the Econometric Society*, 58, 1121–1149.
- HECKMAN, J. J., L. LOCHNER, AND C. TABER (1998): “Explaining rising wage inequality: Explorations with a dynamic general equilibrium model of labor earnings with heterogeneous agents,” *Review of economic dynamics*, 1, 1–58.
- HECKMAN, J. J., J. STIXRUD, AND S. URZUA (2006): “The effects of cognitive and noncognitive abilities on labor market outcomes and social behavior,” *Journal of Labor economics*, 24, 411–482.
- HUANG, W., X. LEI, G. RIDDER, J. STRAUSS, AND Y. ZHAO (2013): “Health, height, height shrinkage, and SES at older ages: evidence from China,” *American Economic Journal: Applied Economics*, 5, 86–121.
- JUNG, J. AND J. MERCENIER (2012): “Routinization-Biased Technical Change and Globalization: A Theoretical Exploration of Labor Market Polarization,” *Working Paper, RWTH Aachen University*.
- KAMBOUROV, G. AND I. MANOVSKII (2009): “Occupational Specificity of Human Capital,” *International Economic Review*, 50, 63–115.
- KAMPELMANN, S. AND F. RYCX (2011): “Task-Biased Changes of Employment and Remuneration: The Case of Occupations,” Tech. rep., Social Science Research Network, Rochester, NY.
- LINDENLAUB, I. (2017): “Sorting multidimensional types: Theory and application,” *The Review of Economic Studies*, 84, 718–789.
- LONGHI, S. AND M. BRYNIN (2009): “Occupational Change in Britain and Germany,” Tech. rep., Social Science Research Network, Rochester, NY.
- LUNDBORG, P., P. NYSTEDT, AND D.-O. ROTH (2014): “Height and earnings: The role of cognitive and noncognitive skills,” *Journal of Human Resources*, 49, 141–166.

- MC FALL, S. (2013): “Understanding Society: UK household longitudinal study: Cognitive ability measures,” *Institute for Social and Economic Research, University of Essex*.
- MERRILL, R. AND J. RICHARDSON (2009): “Validity of Self-Reported Height, Weight, and Body Mass Index: Findings From the National Health and Nutrition Examination Survey, 2001-2006,” *Preventing Chronic Disease*, 6.
- PERSICO, N., A. POSTLEWAITE, AND D. SILVERMAN (2004): “The Effect of Adolescent Experience on Labor Market Outcomes: The Case of Height,” *Journal of Political Economy*, 112.
- ROY, A. (1951): “Some thoughts on the distribution of earnings,” *Oxford economic papers*, 3, 135–146.
- RUSMAULLY, J., A. DUGRAVOT, J.-P. MOATTI, M. G. MARMOT, A. ELBAZ, M. KIVIMAKI, S. SABIA, AND A. SINGH-MANOUX (2017): “Contribution of cognitive performance and cognitive decline to associations between socioeconomic factors and dementia: A cohort study,” *PLoS medicine*, 14, e1002334.
- SALVATORI, A. (2015): “The Anatomy of Job Polarisation in the UK,” Tech. rep., IZA Discussion Papers.
- SOLON, G., R. BARSKY, AND J. PARKER (1994): “Measuring the cyclicity of real wages: how important is composition bias,” *Quarterly Journal of Economics*, 109, 1–25.
- SORKIN, J. D., D. C. MULLER, AND R. ANDRES (1999): “Longitudinal change in height of men and women: implications for interpretation of the body mass index: the Baltimore Longitudinal Study of Aging,” *American journal of epidemiology*, 150, 969–977.
- SPITZ-OENER, A. (2006): “Technical change, job tasks, and rising educational demands: Looking outside the wage structure,” *Journal of labor economics*, 24, 235–270.
- TABER, C. R. AND N. A. ROYS (2019): “Skill Prices, Occupations, and Changes in the Wage Structure for Low Skilled Men,” Tech. rep., National Bureau of Economic Research.
- WAGNER, G. G., J. R. FRICK, AND J. SCHUPP (2007): “The German Socio-Economic Panel study (SOEP)-evolution, scope and enhancements,” *Journal of Applied Social Science Studies*, 127, 139–169.
- WILLIS, R. (1986): “Wage determinants: A survey and reinterpretation of human capital earnings functions,” in *Handbook of Labor Economics*, vol. 1, 525–602.
- YAMAGUCHI, S. (2012): “Tasks and heterogeneous human capital,” *Journal of Labor Economics*, 30, 1–53.
- (2018): “Changes in returns to task-specific skills and gender wage gap,” *Journal of Human Resources*, 53, 32–70.

Our approach to assessing the role of idiosyncratic shocks is to examine wage growth around occupational switches, along the lines of [Card et al. \(2013\)](#), who emphasize that exogenous mobility is a sufficient condition for exogeneity in the wage equation more broadly. As discussed in [section 2](#), the idea is that if occupational switches are exogenous then we should observe wage ‘symmetry’: apart from marginal changes in the sectoral prices, the average wage change for switches from sector j to sector j' should be the negative of the average wage change for switches from j' to j . In contrast, if switches are driven by, say, idiosyncratic wage offers then we should observe wage growth in both directions.

We implement this test by performing an event-study of wages in a 3-year window around switches between any of our 4 main sectors. The full results of average log wages are displayed in [table A.2](#), which also includes average wage progression for those remaining in a sector for 3 years continuously. In this table, year t captures the first observation in the new sector, year $t - 1$ is the last year in the originating sector. For a visual interpretation the results for those switching into and out of the abstract sector are given in [figure A.1](#), which displays the important switches most starkly. Both the table and graph display average log wages, unadjusted by any controls. As such it should be noted that the strict exogeneity test may be slightly confounded by systematic sectoral price changes. However, any effect is likely small at the frequencies shown: the price on the abstract sector grew by around 0.6% pa relative to manual routine according to our central estimates.

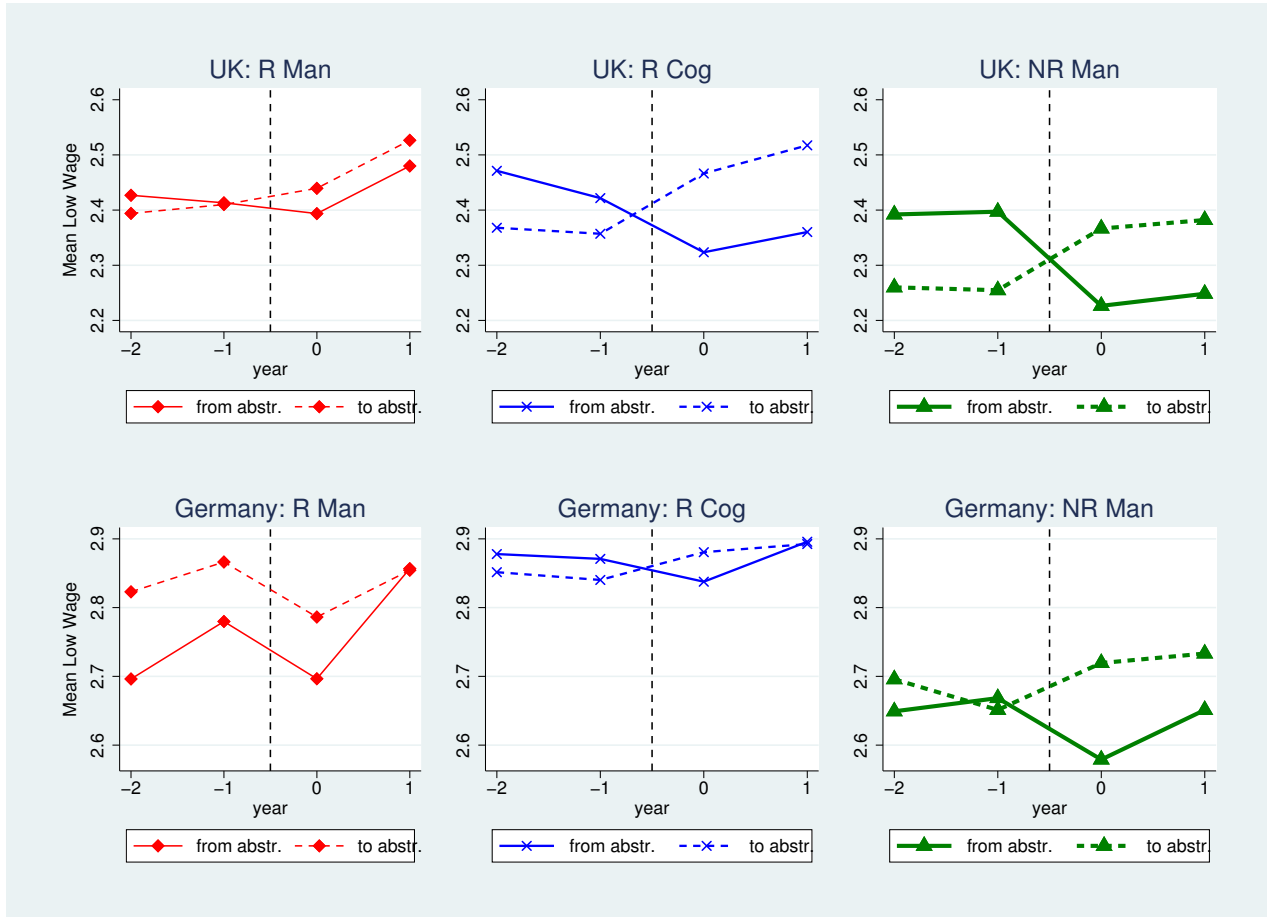
For both countries, [figure A.1](#) shows that switches into and out of both cognitive routine and manual non-routine are both consistent with exogenous mobility: the positive effect of switching into abstract is matched by the negative effect of switching into these lower-ranked sectors. The only switches which provide a different pattern are those into and out of manual routine. In the UK, it appears that wages for switchers in both directions grow in year $t + 1$. In Germany, the evidence is more mixed; wages grow in year $t + 1$ but decline in year t , just after the switch occurs. There is therefore some evidence that switches between these sectors are slightly more complex than in other cases. However, [table A.2](#) shows further that it is only this combination of sectors that departs from wage symmetry.

A.3 Further Information on Task Price Estimates

The following subsection provides further details on the estimation of the task prices on different specifications. For completeness, the [tables A.3](#) and [A.4](#) present additional coefficients of the specifications discussed in [section 4](#) and presented in [table 1](#).

[Figures A.2](#) and [A.3](#) plot the task prices over time for the UK and Germany, respectively. For each country, the first four graphs correspond to the coefficients associated to the model specified in [column 2](#) to [4](#) of [table 1](#).

Figure A.1: Mean Wages of Occupational Switches to/from Abstract Work



Notes: Figure shows mean log wages for those switching into and out of abstract work, pooled across years. Year = 0 indicates the first observation in the new occupation. Sample includes all workers switching between time $t = -1$ and $t = 0$ irrespective of their occupation in time $t = -2$ or $t = 1$. For fuller results see table A.2. Occupation measure obtained after correcting for ‘noisy switches’, as discussed in section B.1. See also section 3 for further discussion of sample selection.

Table A.2: Mean Log Wages Before and After Occupational Change

Origin/destination Occupation	Number of Observations (1)	Mean Log Wage of Switchers									
		$t^* - 2$ (2)		$t^* - 1$ (3)		t^* (4)		$t^* + 1$ (5)		3-year Change (6)	
<u>UK</u>											
Abs to Abs	9768	2.72	(0.005)	2.74	(0.005)	2.78	(0.005)	2.80	(0.005)	0.08	(0.004)
Abs to RMa	138	2.43	(0.04)	2.41	(0.03)	2.39	(0.04)	2.48	(0.03)	0.05	(0.04)
Abs to RCo	148	2.47	(0.05)	2.42	(0.04)	2.32	(0.04)	2.36	(0.04)	-0.11	(0.05)
Abs to NRM	50	2.39	(0.07)	2.40	(0.08)	2.23	(0.07)	2.25	(0.08)	-0.14	(0.09)
RMa to Abs	189	2.39	(0.03)	2.41	(0.03)	2.44	(0.03)	2.53	(0.03)	0.13	(0.03)
RMa to RMa	6369	2.29	(0.005)	2.30	(0.005)	2.32	(0.004)	2.34	(0.005)	0.05	(0.005)
RMa to RCo	93	2.25	(0.04)	2.17	(0.04)	2.16	(0.04)	2.21	(0.04)	-0.04	(0.04)
RMa to NRM	95	2.07	(0.04)	2.08	(0.05)	2.01	(0.04)	2.07	(0.03)	0.00	(0.05)
RCo to Abs	237	2.37	(0.03)	2.36	(0.03)	2.47	(0.02)	2.52	(0.03)	0.15	(0.03)
RCo to RMa	60	2.15	(0.05)	2.07	(0.06)	2.14	(0.05)	2.17	(0.057)	0.01	(0.06)
RCo to RCo	2765	2.27	(0.008)	2.28	(0.008)	2.30	(0.007)	2.34	(0.008)	0.07	(0.008)
RCo to NRM	42	2.03	(0.07)	1.93	(0.08)	1.98	(0.05)	2.05	(0.05)	0.02	(0.07)
NRM to Abs	71	2.26	(0.08)	2.25	(0.07)	2.37	(0.06)	2.38	(0.05)	0.12	(0.07)
NRM to RMa	104	2.01	(0.05)	1.97	(0.04)	2.06	(0.04)	2.08	(0.04)	0.07	(0.05)
NRM to RCo	44	2.04	(0.07)	2.09	(0.08)	2.07	(0.05)	2.16	(0.04)	0.12	(0.08)
NRM to NRM	2287	2.23	(0.010)	2.23	(0.010)	2.26	(0.009)	2.29	(0.010)	0.05	(0.009)
<u>Germany</u>											
Abs to Abs	14568	3.08	(0.004)	3.09	(0.004)	3.11	(0.003)	3.13	(0.004)	0.05	(0.003)
Abs to RMa	82	2.70	(0.07)	2.78	(0.08)	2.70	(0.05)	2.86	(0.08)	0.16	(0.08)
Abs to RCo	126	2.88	(0.05)	2.87	(0.05)	2.84	(0.04)	2.90	(0.06)	0.02	(0.07)
Abs to NRM	29	2.65	(0.08)	2.67	(0.12)	2.58	(0.08)	2.65	(0.08)	0.00	(0.11)
RMa to Abs	180	2.82	(0.04)	2.87	(0.04)	2.79	(0.03)	2.85	(0.03)	0.03	(0.04)
RMa to RMa	11956	2.78	(0.003)	2.78	(0.003)	2.79	(0.003)	2.81	(0.003)	0.03	(0.004)
RMa to RCo	71	2.73	(0.03)	2.72	(0.07)	2.63	(0.03)	2.68	(0.03)	-0.04	(0.04)
RMa to NRM	104	2.61	(0.05)	2.54	(0.06)	2.49	(0.05)	2.56	(0.06)	-0.05	(0.08)
RCo to Abs	183	2.85	(0.04)	2.84	(0.03)	2.88	(0.04)	2.89	(0.03)	0.04	(0.04)
RCo to RMa	46	2.59	(0.05)	2.56	(0.07)	2.56	(0.06)	2.58	(0.07)	-0.01	(0.08)
RCo to RCo	3954	2.88	(0.006)	2.88	(0.006)	2.90	(0.006)	2.92	(0.006)	0.04	(0.006)
RCo to NRM	23	2.47	(0.08)	2.48	(0.07)	2.56	(0.13)	2.55	(0.07)	0.08	(0.08)
NRM to Abs	47	2.70	(0.06)	2.65	(0.11)	2.72	(0.09)	2.73	(0.04)	0.04	(0.06)
NRM to RMa	66	2.59	(0.09)	2.57	(0.05)	2.46	(0.08)	2.57	(0.04)	-0.02	(0.10)
NRM to RCo	24	2.38	(0.09)	2.55	(0.15)	2.64	(0.09)	2.68	(0.13)	0.30	(0.14)
NRM to NRM	2720	2.70	(0.007)	2.70	(0.007)	2.70	(0.006)	2.73	(0.006)	0.03	(0.007)

Notes: Table displays mean log wages for those switching between occupations, pooled across years. Standard errors on wages in parentheses. Abs stands for abstract. RMa stands for routine manual. RCo stands for routine cognitive. NRM stands for non-routine manual. Time = t^* indicates the first observation in the new occupation. Sample includes all workers switching between time $t = t^* - 1$ and $t = t^*$ irrespective of their occupation in time $t^* - 2$ or $t^* + 1$. Occupation measure obtained after correcting for 'noisy switches', as discussed in section B.1. See also section 3 for further discussion of sample selection.

Table A.3: Changes in Task Price in the UK: Controls in the Additional Specifications

	OLS	Benchmark FE	Educ. Returns	Job Tenure	TU	35-60 y-o
Primary ed	-0.195*** (0.03)	-0.007 (0.05)	0.015 (0.05)	-0.009 (0.05)	- (0.05)	0.046 (0.06)
Secondary ed.	-0.069 (0.04)	-0.147** (0.06)	-0.029 (0.05)	-0.146** (0.06)	-0.133** (0.06)	-0.039 (0.06)
High secondary	0.064** (0.03)	0.060* (0.04)	0.108*** (0.04)	0.060* (0.04)	0.058* (0.04)	0.092** (0.04)
High vocational	0.124*** (0.03)	0.084*** (0.03)	0.081** (0.04)	0.084*** (0.03)	0.085*** (0.03)	0.114*** (0.04)
First degree	0.288*** (0.04)	0.118* (0.06)	0.148** (0.07)	0.116* (0.06)	0.116* (0.07)	0.086 (0.07)
Higher degree	0.370*** (0.05)	0.201*** (0.07)	0.276*** (0.10)	0.200*** (0.07)	0.206*** (0.07)	0.166 (0.10)
Job Tenure	0.022*** (0.00)	0.006** (0.00)	0.007** (0.00)	0.006 (0.00)	0.005* (0.00)	0.005 (0.00)
Job Tenure2	-0.002*** (0.00)	- (0.00)	- (0.00)	- (0.00)	- (0.00)	- (0.00)
1991 Abstract × Job Tenure				-		0.002 (0.00)
task=2 × Job Tenure						
1991 R Cog × Job Tenure				0.003 (0.00)		0.004 (0.00)
1991 NR Man × Job Tenure				0.003 (0.00)		0.005* (0.00)
Union in workplace					0.073*** (0.01)	
1991 Abstract × Union in workplace					-0.027 (0.02)	
1991 R Cog × Union in workplace					-0.040 (0.03)	
1991 NR Man × Union in workplace					0.001 (0.03)	
Union in workplace × age					-0.001 (0.00)	
Constant	2.169*** (0.03)	2.484*** (0.06)	2.491*** (0.06)	2.485*** (0.06)	2.463*** (0.06)	2.352*** (0.06)
Observations	24364	24364	24364	24364	24364	16955
Adjusted R ²	0.376	0.247	0.246	0.247	0.250	0.145

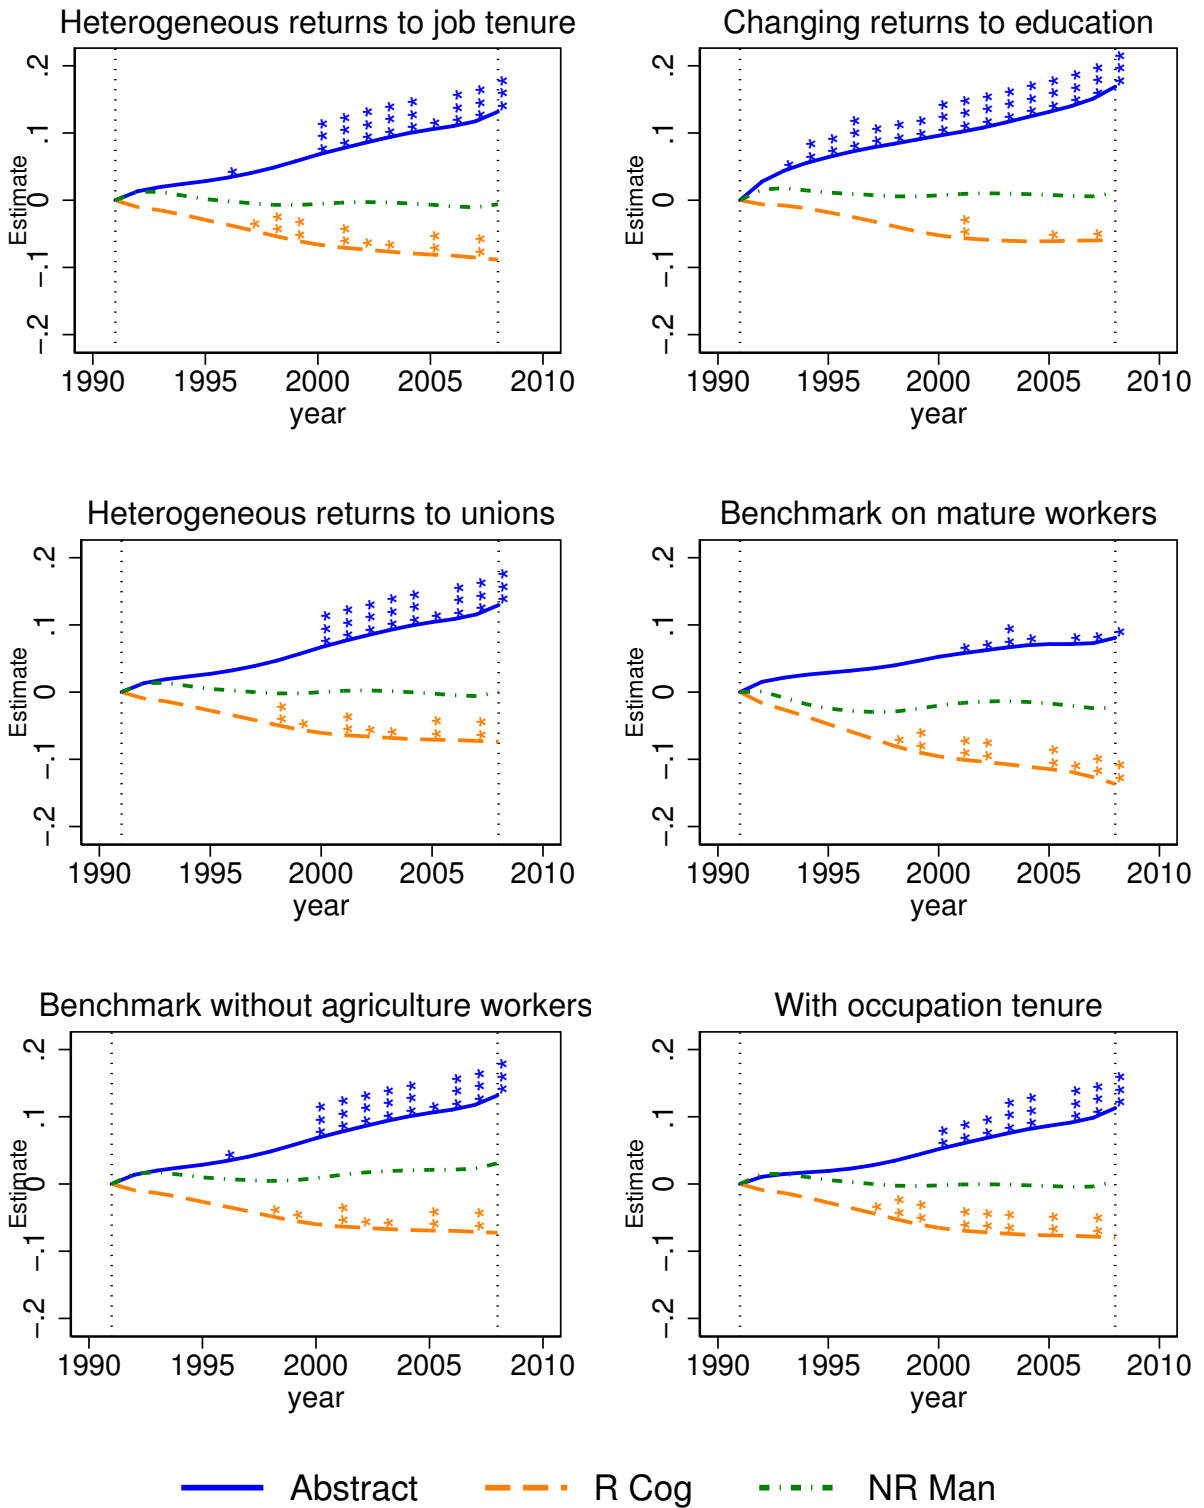
Notes: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$ Standard errors clustered at individual level in parenthesis. For all models, controls include region, quartic in age and in job tenure, education, interaction between age and education, education and year, marital status, year fixed effects. In 1, the benchmark model is estimated with OLS; in 2, with panel fixed effects; in 3, an interaction between education and year is added; in 4, model 2 is augmented with the interaction between sector and job tenure, and between year and job tenure; in 5, model 2 includes a dummy for trade union association and its interaction with sector and age; in 6, model 3 is estimated on a group of mature workers, from 35 to 60; (7) is based on model (2) but dropping job tenure, education and the interaction of education with age. The omitted category for education is mid secondary.

Table A.4: Changes in Task Price in Germany: Controls in the Additional Specifications

	OLS	Benchmark FE	Educ. Returns	Job Tenure	TU	35-60 y-o
FE fewer controls						
Primary ed	-0.044** (0.02)	-0.012 (0.02)	-0.010 (0.03)	-0.011 (0.02)	-0.009 (0.02)	-0.019 (0.02)
vocational + Abi	0.092** (0.04)	0.099*** (0.03)	0.049 (0.05)	0.101*** (0.03)	0.102*** (0.03)	0.009 (0.03)
High vocational	0.108*** (0.02)	0.073*** (0.02)	0.058** (0.03)	0.072*** (0.02)	0.076*** (0.02)	0.030 (0.03)
Degree or more	0.307*** (0.02)	0.182*** (0.03)	0.091** (0.04)	0.182*** (0.03)	0.183*** (0.03)	0.023 (0.05)
Job Tenure	0.032*** (0.00)	0.013*** (0.00)	0.016*** (0.00)	0.014*** (0.00)	0.014*** (0.00)	0.010*** (0.00)
Job Tenure2	-0.002*** (0.00)	-0.001*** (0.00)	-0.001*** (0.00)	-0.001*** (0.00)	-0.001*** (0.00)	- (0.00)
1991 Abstract × Job Tenure				-0.002 (0.00)		-0.003* (0.00)
task=2 × Job Tenure						
1991 R Cog × Job Tenure				0.003 (0.00)		0.003 (0.00)
1991 NR Man × Job Tenure				-0.002 (0.00)		-0.004 (0.01)
Collective bargaining					0.012 (0.01)	
1991 Abstract × Collective bargaining					-0.009 (0.02)	
1991 R Cog × Collective bargaining					-0.028 (0.03)	
1991 NR Man × Collective bargaining					0.012 (0.03)	
Collective bargaining × age					0.001 (0.00)	
Constant	2.674*** (0.03)	2.767*** (0.07)	2.767*** (0.06)	2.764*** (0.06)	2.766*** (0.07)	2.790*** (0.08)
Observations	36918	36918	36918	36918	36032	27124

Notes: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$ Standard errors clustered at individual level in parenthesis. For all models, controls include region, quartic in age and in job tenure, education, interaction between age and education, education and year, marital status, year fixed effects. In 1, the benchmark model is estimated with OLS; in 2, with panel fixed effects; in 3, an interaction between education and year is added; in 4, model 2 is augmented with the interaction between sector and job tenure, and between year and job tenure; in 5, model 2 includes a dummy for trade union association and its interaction with sector and age; in 6, model 3 is estimated on a group of mature workers, from 35 to 60; (7) is based on model (2) but dropping job tenure, education and the interaction of education with age. The omitted category for education is mid secondary.

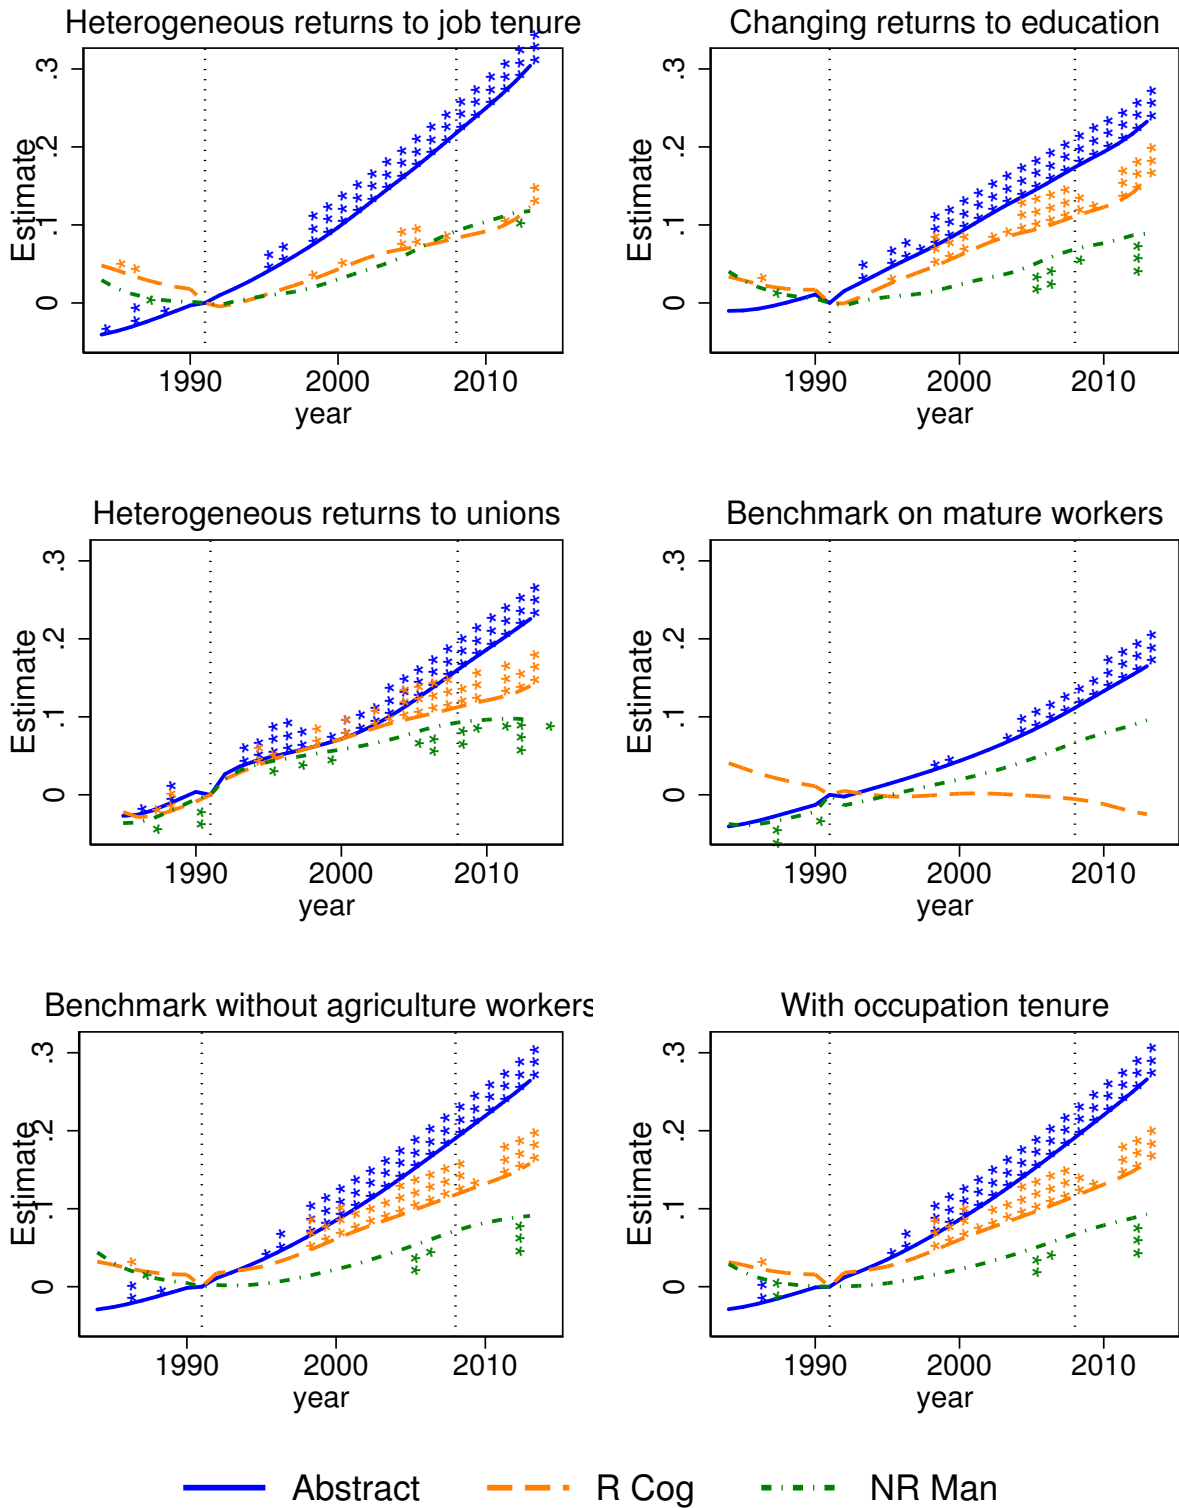
Figure A.2: Task prices in the UK: Supplementary Specifications



Picture is based on a sample of 25 to 60 y-o men.

'Abstract' stands for abstract task. 'NR Man' indicates non-routine manual, and 'R Cog' indicates routine cognitive. The vertical dashed line indicates the common period in the two data sets, from 1991 to 2008.

Figure A.3: Task prices in Germany: Supplementary Specifications



Picture is based on a sample of 25 to 60 y-o men.

'Abstract' stands for abstract task. 'NR Man' indicates non-routine manual, and 'R Cog' indicates routine cognitive. The vertical dashed line indicates the common period in the two data sets, from 1991 to 2008.

In the last row, we add two additional robustness checks. The left graph plots the coefficients estimated on a subsample of workers, after we exclude agriculture and fishery workers. Finally, the right panel in the last row is based on a similar model as the one in the fourth column of table 1. Here, however, we consider an indicator of occupational tenure instead of job tenure. The indicator of occupational tenure follows [Kambourov and Manovskii \(2009\)](#) and [Cortes \(2016\)](#). Tenure results from the number of consecutive years we observe an individual in the same broad occupational group. The years of tenure cumulate until we observe a switch. An occupation switch occurs if an individual's current sector is different from the group of the previous year, without conditioning on an employer or job switch. Notice that to account for respondents not being interviewed in certain waves but reappearing later, when there is a gap year we replace the missing information about the sector with the sector of the previous year, if it is the same as the sector in the following year. Notice as well that if a worker is unemployed for a given period, the years of tenure before the unemployment spell are not added to the years following the spell, even if the individual works in the same broad occupational group. Despite the differences with the indicator of job tenure, however, the figure shows that the results are consistent with those in the fourth column of table 1. Overall, as the figures show, the results are consistent with those in the main analysis.

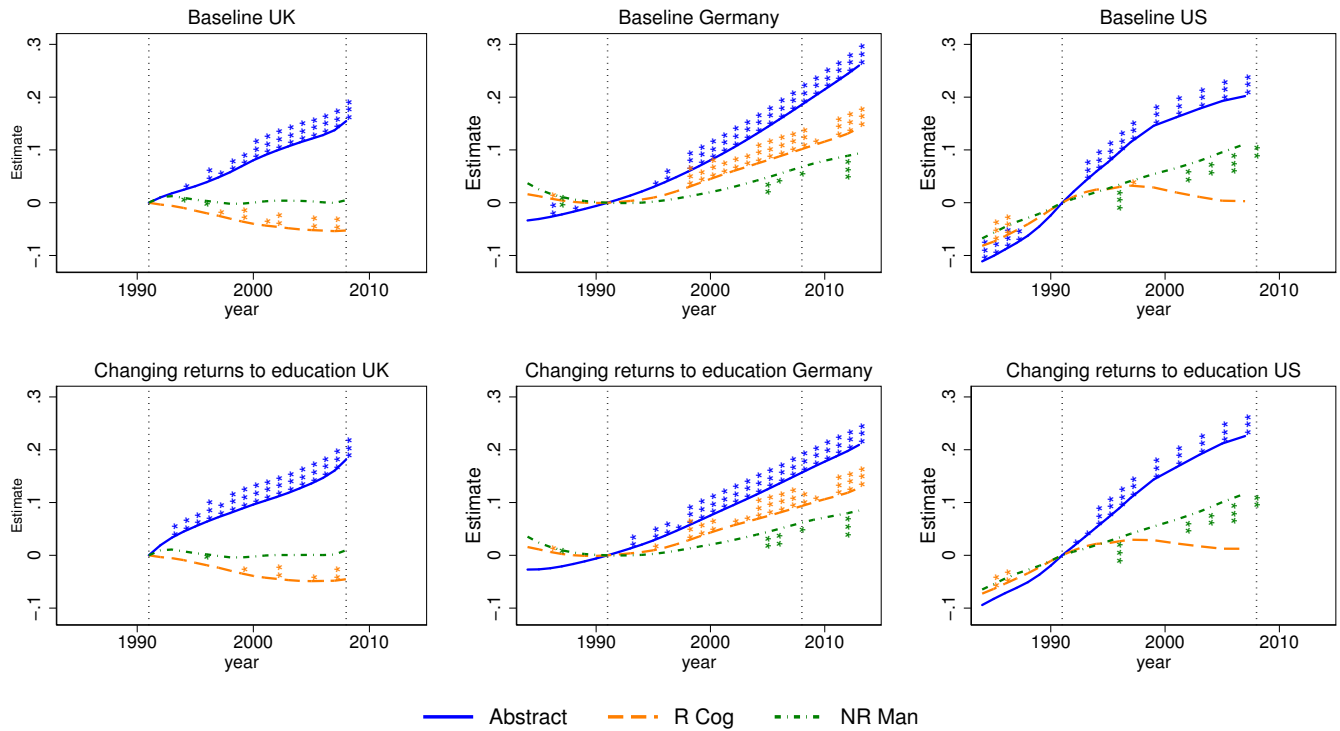
Figure A.4 shows task price estimates for the UK and Germany, alongside those from the U.S. For the U.S. we take the data and sample used by [Cortes \(2016\)](#), using the Panel Study of Income Dynamics (PSID). For consistency with our analysis we split the 'routine' occupational sector used by [Cortes](#) into our component sectors: routine manual and routine cognitive. Again for consistency with our analysis, we also use routine manual as the base occupation, where [Cortes](#) uses for his base category non-routine manual.

We show task prices for two specifications, our benchmark model and the specification with changing returns to education. This latter specification is also discussed in section 4 with results shown in table 1. This latter specification is interesting because patterns of returns to education appear to have differed by country: returns to education in the UK since the early 1990s have been flat ([Blundell et al., 2018](#)), while the U.S. saw increases in returns, albeit at a slowing rate over time ([Fortin, 2006](#)). In fact in this specification it is for Germany that the estimated price increase for abstract compared to routine manual is most reduced compared to the benchmark. For the U.S. and the UK the implied price increase is more than in the benchmark specification, though only marginally so.

A.4 Price and Employment Estimates for Females and from the UKHLS

Figure A.5 shows task price estimates and employment changes for females. The results are very consistent with those for males. Figure A.6 shows prices and employment changes for males using data from the UKHLS up to 2015. It shows that prices in the UK after 2008 were fairly flat across sectors. It also shows that employment in all sectors grew strongly relative to manual routine. These results are *prima*

Figure A.4: Task Prices: Comparison with the U.S.



Notes: Prices for U.S. based on PSID data from Cortes (2016), restricted to men aged 25-60. The occupational classification is defined to be the same as in the rest of this paper. Similarly is true for estimation specifications.

facie at odds with the argument presented in the rest of the paper. However, they can be explained by a model in which certain sectors, and particularly manufacturing, were strongly hit by negative finance shocks after 2008. According to such a hypothesis, employment in manual routine jobs declined steeply even if the price on labour held up. Notice further that the price on all other sectors held up relative to each other.

A.5 Further Background and Results on Height

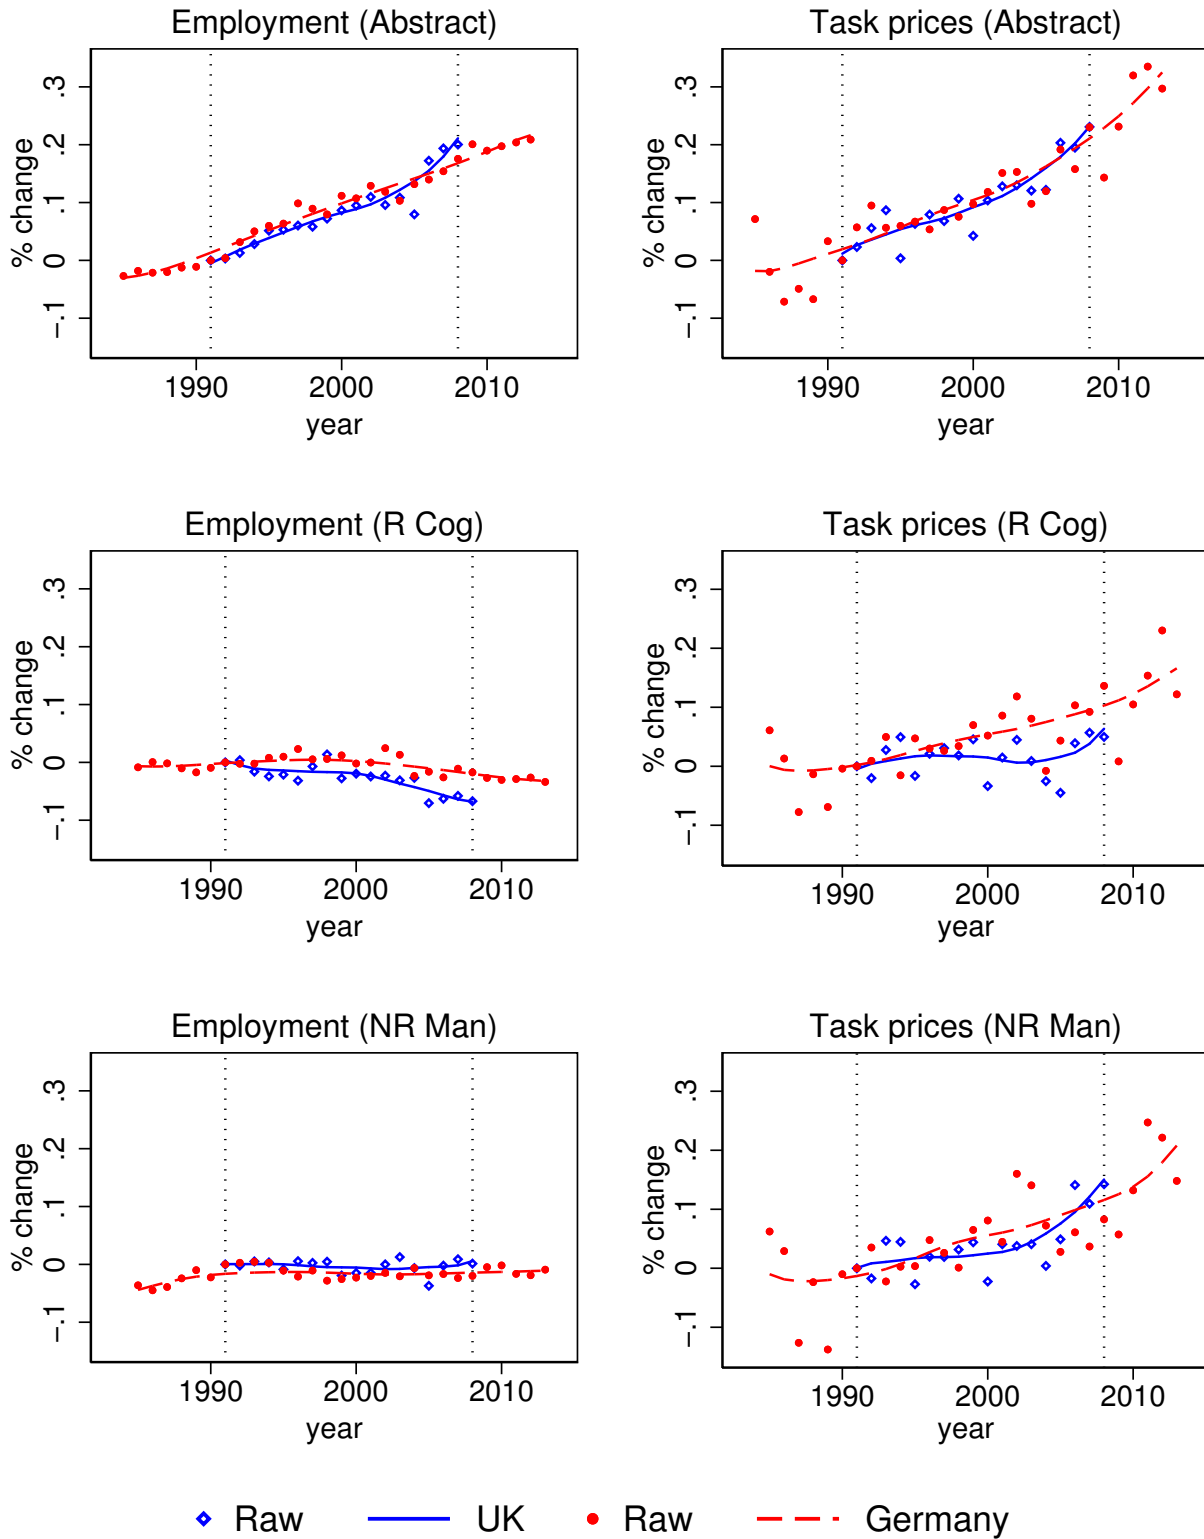
A.5.1 Summary statistics of height measurements

Table A.5 shows the mean and standard deviation of self-reported height (in cm) at each step of the sample selection.

A.5.2 Results from analysis of National Child Development Survey

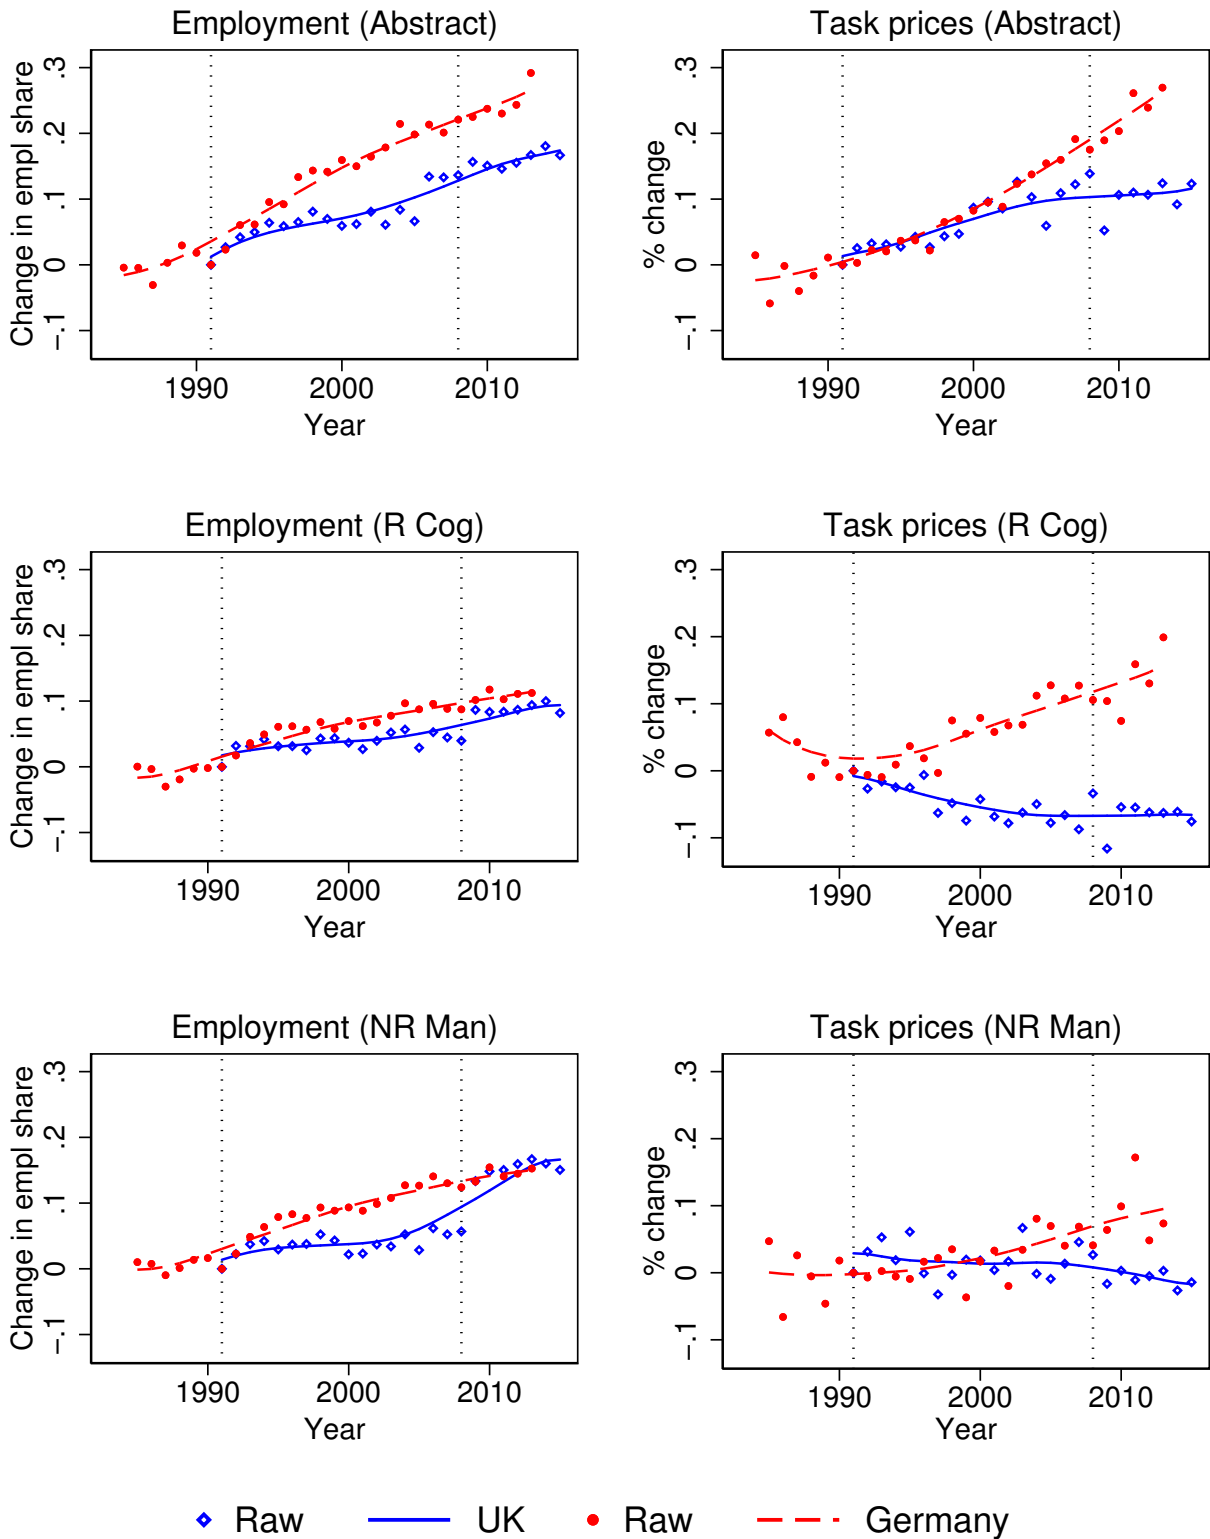
Following the discussion in section 5, in this appendix we discuss evidence in this direction using the UK's National Child Development Survey (NCDS).

Figure A.5: Price and Employment Changes by Task for Females



Notes: The figure illustrates changes in employment and in task prices of workers with respect to manual workers. Employment change is computed for 16 to 64 y-o women.

Figure A.6: Price and Employment Changes by Task in the UK: Data from UKHLS



Notes: The figure illustrates changes in employment and in task prices of workers with respect to manual workers. This figure extends the analysis for the UK by including data from the UK Household Longitudinal Survey. Employment change is computed for 16 to 64 y-o men.

Table A.5: Valid height measurements and sample selection

	Total observations	Number of individuals	Height in cm (mean)	Height in cm (sd)
UK				
All individuals	76,097	63,399	168.32	10.44
All males	33,742	28,038	175.75	9.01
Original BHPS and USoc samples	24,109	20,915	176.42	8.63
25-60 year-olds	14,616	12,609	177.07	8.63
Born 1950 or later	13,909	12,128	177.22	8.56
Non-missing variables	13,520	11,790	177.25	8.49
Germany				
All individuals	129,158	40,874	171.40	9.32
All males	58,352	16,354	177.96	7.28
Exclude 1990 East German sample	49,874	14,290	178.01	7.31
25-60 year-olds	30,042	9,260	179.05	7.09
Born 1950 or later	26,936	7,820	179.34	7.06
Non-missing variables	25,439	6,857	179.34	7.05

The NCDS is a UK-based cohort survey, which continues to track around 17,000 individuals born in March 1958. Relevant to present purposes, sweep 5 of the NCDS, carried out in 1991, contains information on individuals' education, occupation and height at age around 33. This survey yields 5,501 male height observations from a total sample size of 5,634. The attraction of using these data is that we can combine them with data on males born in the late 1970s and early 1980s in UKHLS and make direct cross-cohort comparisons on height and on occupation directly. The downside of this analysis is that the obtained sample sizes are far smaller than those used in section 5.

In contrast to the analysis in section 5, which relied on height self reports, the data in the NCDS are obtained using clinical measures from nurse visits. Fortunately, the UKHLS also contains clinical height measures which can be compared directly.³⁷ However, in contrast to the data on self-reports, the sample size obtained from clinical measures is quite small. Therefore, in this appendix we show results using both types of measurement.

We now perform an analysis similar to that presented in table 2. We pool data from the NCDS obtained in 1991 with data from a later cohort measured in the UKHLS in 2010. In order to provide a decent sample size, we use the 'Generation Y' cohort from the UKHLS and pool males born between 1976 and 1984. To be specific on this UKHLS sample, therefore, we observe self-reported heights on individuals at some point between ages 27 and 35, nurse-assessed heights sometime between ages 26 and 34, and education and occupation attained in individuals' early 30s in various survey waves. In order to include the younger members of the cohort and to maximize the sample size we include all observations of occupation between ages 31 and 33. In all results we therefore cluster standard errors at the individual level to allow for these multiple measurements. We interpret the NCDS cohort as representing the Baby Boomers, although it should be remembered this cohort is close to the threshold with Generation X.

³⁷We replicated figure 8 using the nurse-assessed measure. Similarly to figure 8, it shows strong convergence of heights over time. Although not showed here in the interest of space, it is available upon request.

Table A.6 presents the results of four regressions provided with a 2×2 structure. We alternate in the regressions self-reported height from UKHLS (to maximize sample size) and nurse-measured height in UKHLS (for best comparison with NCDS). We also alternate choices of grouping: in one set of regressions we use education as in the main analysis; the other set of regressions uses occupation.

Table A.6: Height by Cohort using NCDS and UKHLS

	(1)	(2)	(3)	(4)
	S	N	S	N
Degree	1.751*** (0.271)	1.751*** (0.271)		
Degree*(Gen Y-BB)	-0.935* (0.527)	-0.469 (0.647)		
Abstract			1.526*** (0.189)	1.526*** (0.189)
Abstract*(Gen Y-BB)			-0.355 (0.488)	-0.409 (0.589)
Observations	7962	6907	7585	6690
Controls:				
Region, UK Born	x	x	x	x
Ethnicity	x	x	x	x

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Robust standard errors on the left panel.

Clustered standard errors on the right panel. S indicates that the height for the UKHLS is self-reported; N, measured during nurse visits. BB is the NCDS cohort at 33 years of age. Gen Y is the cohort from UKHLS+, born after 1975, with occupation observed between 31 and 33 years of age.

Overall, the evidence from combining NCDS with UKHLS is similar to that presented in the main analysis in section 5, though less precise, because the sample sizes are smaller. The first row shows the difference in height between degree and non-degree holders for those born in 1958. The difference lines up closely with the equivalent cohorts using the cross-sectional evidence from UKHLS only. The second row shows the interaction term with the cohort dummies. The first column shows results using self-reported heights from UKHLS, which again are similar to those from main table 2. The second column shows results when using nurse measurements. The interaction term is now somewhat smaller but is statistically indistinguishable from the first column, and the sample size is smaller. In fact the standard error is sufficiently large that we can barely rule out complete convergence over time at the 5% level.

We turn now to occupation, where, as discussed above, we compare those working in the abstract sector to all the other sectors combined. Again we see a similar difference in heights for those born in 1958. The bottom row of this panel then shows the interaction term which shows the convergence in heights across cohorts. For occupation, we no longer see a significant convergence for the later cohorts. However, the point estimates are similar to the other evidence presented, and, again, a 95% confidence interval is

sufficiently wide that we can barely rule out complete convergence across cohorts.

B Appendix to Section 4: Robustness to Alternative Treatments of the Data

In this section, we check whether the results in the main analysis are robust to alternative ways of classifying occupations.

B.1 Corrected Measure of Occupation

As mentioned in Section 3, occupational status may suffer from measurement error, the extent of which appears to differ according to how the interviews are conducted. In the UK, measurement error seemed to decline after 2006, when dependent interviewing was introduced.³⁸ In Germany, the SOEP survey alternates years with full and partial interviewing, with its own effect on measurement accuracy.³⁹ In this section we show in more detail the effect of this variation in interviewing technique. First, though, we describe the method we use to derive a consistent and ‘corrected’ measure of occupational status within each survey.

With BHPS, we correct the occupational status variable using a variable indicating the first year of the current job. Specifically, each respondent is asked the following question: *“What was the date you started working in your present position? If you have been promoted or changed grades, please give me the date of that change. Otherwise please give me the date when you started doing the job you are doing now for your present employer.”* First, we manually correct for the few inconsistencies in this job-start variable. Second, we use the year when the respondent has changed job to correct the ISCO-88 code. Specifically, we revert the occupational code to that of the previous year if reported occupation changes but the respondent reports no job changes. In addition, we correct for 1-year ‘transitory’ switches in occupation when the job start measure is missing; for example, if a respondent reports working as a fireman in 1999, 2000, 2002 and 2003, and a teacher in 2001, we recode the occupational code of the year 2001.

With SOEP, we use a similar procedure. We combine a variable indicating the year when the respondent started at the current employer with an indicator for whether the respondent’s job situation has changed since the previous interview. First, we check whether there are inconsistencies in the variable indicating

³⁸Surveys based on ‘dependent interviews’ update the respondent’s occupational code only if the respondent reports a change in job or position. With ‘independent interviews’, the occupational information is gathered from all workers every wave. Prior to 2006, independent interviewing was used and respondents were always asked about their occupation.


³⁹In years with a partial survey only new respondents or employed individuals who changed jobs are asked about their occupation. The years with partial survey are 1985, 1986, 1987, 1988, 1990, 1992, 1994, 1996, 1999, 2001, 2003, 2005, 2006, 2008, 2010 and 2012. This question is always asked in the waves with full interviews. See, for example, Longhi and Brynin (2009) for a discussion on the BHPS and SOEP.

employer’s tenure. Then, we correct the ISCO-88 occupational codes according to the combination of these two variables, just as for the BHPS.

Figures B.1 and B.2 plot the mobility rate on the uncorrected and on the corrected task measure. The mobility rate is computed as the share of workers who change their task from that in the previous year with respect to the overall sample of workers. The left-hand plots of figure clearly shows that, for both surveys, the mobility rate is much higher in the years when the respondents are asked about their precise occupation. As mentioned in the main section, this has happened in the BHPS since 2006, when dependent interviewing was introduced. In the SOEP, this occurs in the years with full interviews.

Figure B.3 reports the occupational wage premia estimated on the baseline model (2nd column of table 1) and based on the corrected version of the occupational group. The figure shows that the results are almost identical to those in the main analysis.

Figure B.1: Mobility Rate for Original vs Corrected Task Measures: BHPS



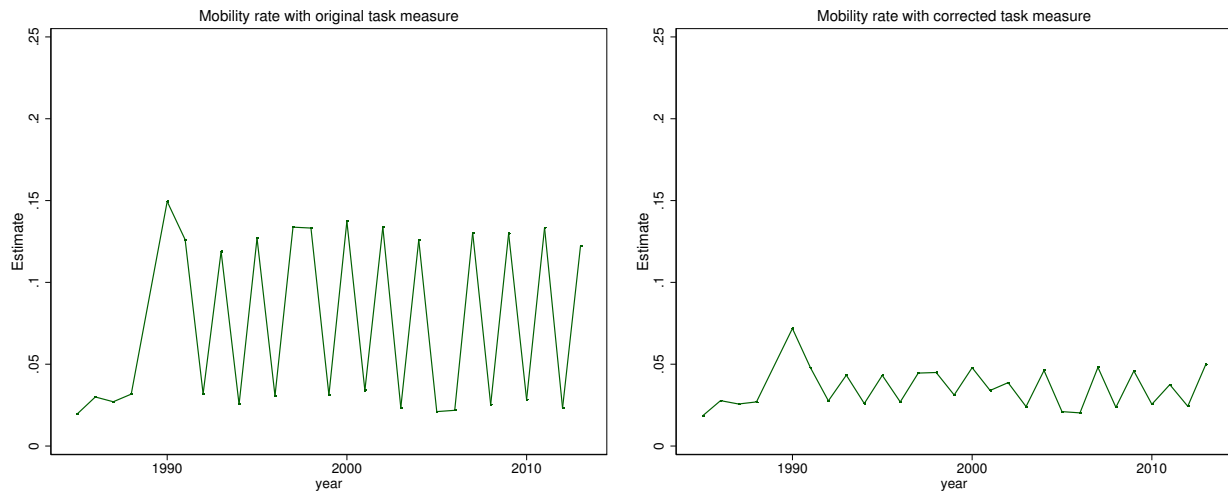
Notes: Mobility rate computed as the year-on-year switches between our four task groupings. Picture is based on a sample of 16 to 64 y-o men

B.2 Task Prices Computed on Alternative Classifications

B.2.1 National classifications

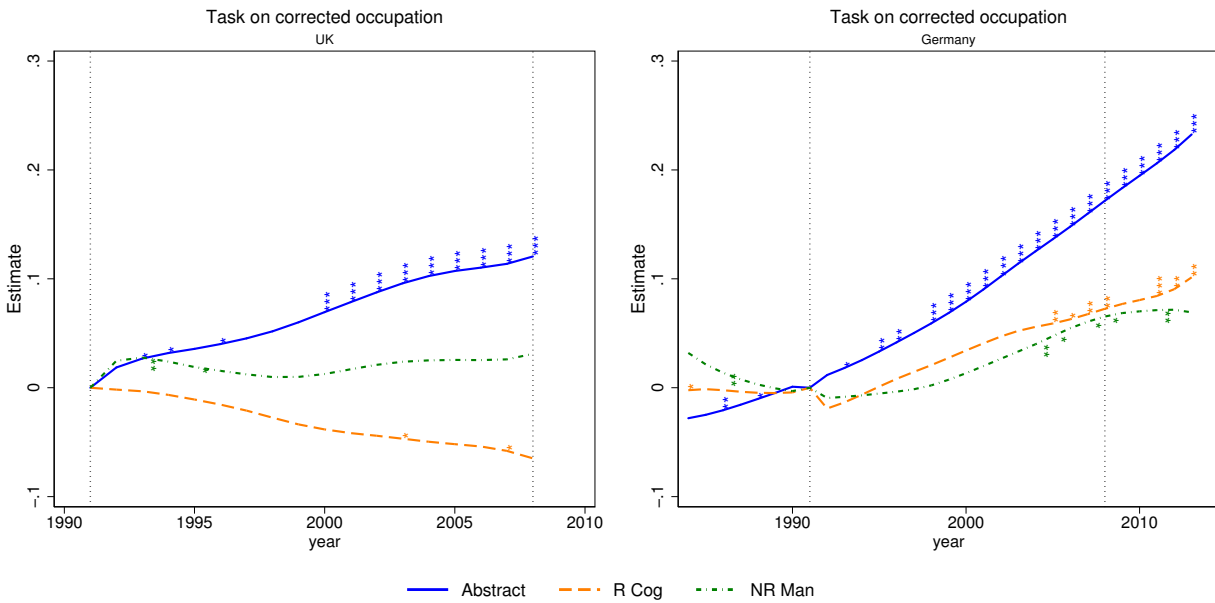
In section 3 we described the assignment of employment into occupations based on the ISCO-88 classification, which is designed for international comparability. In fact both the SOEP and BHPS dataset also contain classifications based on domestically-used classifications. The two classifications are explained below.

Figure B.2: Mobility Rate for Original vs Corrected Task Measures: SOEP



Notes: Mobility rate computed as the year-on-year switches between our four task groupings. Picture is based on a sample of 16 to 64 y-o men

Figure B.3: Task Prices: Corrected Task Measure



Picture is based on a sample of 25 to 60 y-o men.
 'Abstract' stands for abstract task. 'NR Man' indicates non-routine manual, and 'R Cog' indicates routine cognitive.
 The vertical dashed line indicates the common period in the two data sets, from 1991 to 2008.

The classification in SOEP was created by the German Federal Statistical Office in 1992 (KldB 92). Differently from ISCO-88 and the national classifications based on the Standard Occupation Classification (SOC), where the skills needed or tasks involved are the key to define the groupings, the six major groups of KldB are based upon industrial sector. Indeed all these groups include managerial occupations. This structure therefore makes it more difficult to derive a task categorization than the one based on SOC. Indeed, the majority of the relevant literature on Germany either uses additional information to group the KldB according to the task content of the occupation or selects other occupational classifications, such as ISCO-88.⁴⁰

The reference occupational classification in the BHPS is the Standard Occupation Classification (SOC). The members of the original BHPS sample are classified in terms of the 1990 Standard Occupational Classification (SOC90) for the whole period, even after the introduction of the 2000 Standard Occupational Classification. This is an advantage given that there is no perfect correspondence between the two classifications.⁴¹ The SOC90 is very similar to the US SOC classification used in [Acemoglu and Autor \(2011\)](#).

Given the shortcomings of KldB 92, we do not use this classification for our analysis. Instead, we check the robustness of our results on occupational groupings based on the UK SOC90 classification, for both the UK and Germany.⁴² We construct four broader groups by merging the nine occupational categories of the 1990 SOC, like we did with ISCO-88.⁴³ The first includes all the cognitive non-routine occupations: these are, ‘Managers and administrators’, ‘Professional occupations’, ‘Associate professional and technical occupations’. The second cluster is for the occupations involving manual and routine tasks, and comprises ‘Craft and related occupations, plant and machine operators’. The third group includes the jobs who are predominantly characterised by cognitive and routine tasks: ‘Clerical and secretarial occupations’, and ‘Sales occupations’. The fourth group consists in the occupations that involve mainly manual and non-routine tasks: ‘Personal and protective service occupations’, and ‘Other occupations’.

Figure [B.4](#) plots the share of workers in the four occupational groups over time, as in figure [1](#). When we compare the two figures, it is worth underlining that for the UK the patterns are almost identical. For

⁴⁰For example [Black and Spitz-Oener \(2010\)](#) combine the KldB and the “Qualification and Career Survey”. The grouping of [Goos et al. \(2009\)](#) is based on ISCO. An exception is [Kampelmann and Rycx \(2011\)](#), who use SOEP and some tailored questions in the survey to create task measures. However, this question was only collected for the employees in 1985, 1987, 1989, 1995 and 2001.

⁴¹This is well illustrated in the technical report of the Office for National Statistics ([Beerten et al., 2001](#)). The report is based on the Labour Force Survey (LFS) for the summer quarter of 2000. This LFS is dual coded to both the 1990 and the 2000 SOC to clarify the impact of the revision in the SOC. According to that report, only 70% of the occupations are in the same major groups using the summer 2000 Labour Force Survey Data. When considering the task content of each occupation, only 61% of the 2000 occupations can be unambiguously attributed a task following the 1990 classification. For 30% of the cases, there are two concurrent tasks. Three tasks correspond to the same 2000 occupational code in 8% of the occupations. Finally, around 1% of 2000 occupational codes could be expressed equally in terms of all four tasks of the SOC90. Obviously, the risk would be to observe a change in the task content of the occupation which is only due to the change in the used classification.

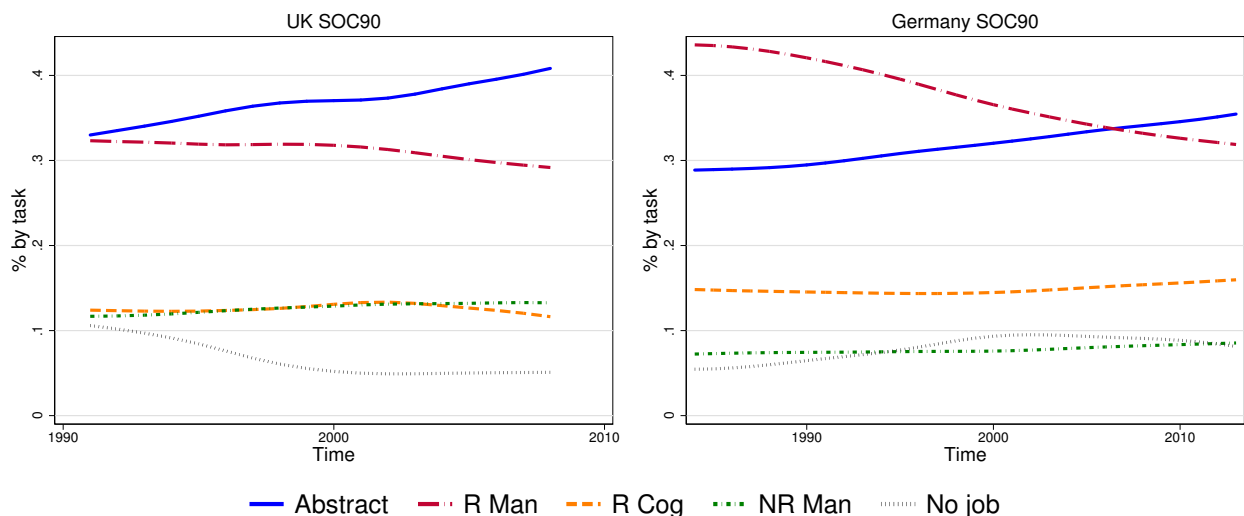
⁴²We derive the occupation for Germany in terms of SOC90 by using the correspondence between SOC90 and ISCO-88.

⁴³Example of another UK study that uses the same methodology is [Salvatori \(2015\)](#).

Germany, about 5% of workers who were classified as abstract according to ISCO-88 are here classified as routine cognitive.

Despite these differences, however, figure B.5 indicates that the occupational returns computed on the SOC90 classification are consistent with those based on the ISCO-88, in the main figure 4. The only difference is the higher returns for routine cognitive occupations in Germany, which may be attributed to the different categorization of routine and abstract workers, as explained above.

Figure B.4: Employment Patterns by Occupational Sector: SOC90 Classification



Picture is based on a sample of 16 to 64 y-o men

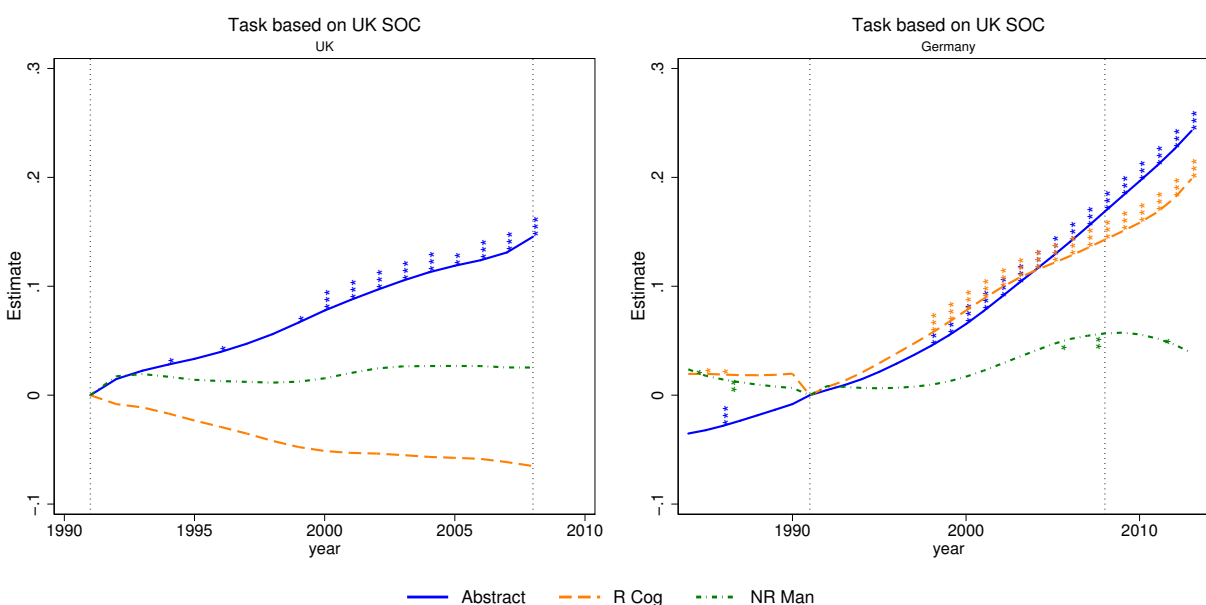
B.2.2 Measures of occupation based on the US occupation classifications

This section presents the estimates of task prices based on the US Department of Labor’s Dictionary of Occupational Titles (DOT) and its successor, the Occupational Information Network (O*Net). Scholars have not only used these classifications for US studies (such as Autor et al. 2003; Autor and Dorn 2013) but also for cross-country comparisons (such as Goos et al., 2009).

For comparison with the existing literature (such as Cortes, 2016), we classify occupations in three categories: abstract, routine and manual. In order to have a clearer idea of how this relates with our study, the first two panels in figure B.6 plot the task classification that we have used so far, but where we aggregate routine manual and routine cognitive activities. In addition, note that for this exercise we only use the subsample of workers for which we can provide a crosswalk with US classifications, which we describe now.⁴⁴

⁴⁴Notice that both in BHPS and in SOEP workers are sometimes assigned a more general 4-digit ISCO-88 code. This can end with a 0, often when the provided job description does not allow assignment at the 4-digit level. Because in the main analysis, we assign the main tasks on the basis of 1-digit occupational groups, we keep these observations in our sample. These codes, however, are not included in the available crosswalks. We exclude them to have the same sample used for the US classifications.

Figure B.5: Task Prices: SOC90 Classification



Picture is based on a sample of 25 to 60 y-o men.

‘Abstract’ stands for abstract task. ‘NR Man’ indicates non-routine manual, and ‘R Cog’ indicates routine cognitive. The vertical dashed line indicates the common period in the two data sets, from 1991 to 2008.

The crosswalk we construct aims at creating a correspondence path between the main tools used to categorise the main task content of occupations (DOT, O*NET) that are based on the US Standard Occupational Classification and different classification systems (Census Occupational Classifications and ISCO of various years).

Specifically, we proceed in different steps. First, we use a set of specific measure indicators from the DOT and the O*Net to classify occupations into task groups. These variables are associated to occupations based on the US Standard Occupational Classification. Therefore, the second part of the process consists in matching these occupations into the classifications that are present in the main datasets (such as Census Occupation Classification for US surveys and ISCO for European datasets). We will do this in several steps.

We use the datasets developed by [Autor et al. \(2003\)](#) that contain two versions of DOT (1977, 1991) associated to the 1990 US Standard Occupation Classification. As per the O*Net, we select the variables of the 2006 version that are associated to the 2000 US Standard Occupation Classification. Then, we convert the 1990 and 2000 Standard Classification Occupations into an harmonised version of the 1990 Census Occupation, developed by [Dorn \(2009\)](#).

For the European datasets we exploit the fact that most contain the ISCO classification. For both BHPS and SOEP, we use ISCO88, which is available for all the waves. We match the ISCO classification with the harmonized 1990 Census Occupation. We do this in subsequent steps. First, we convert ISCO-88 into

the 2000 US Census Occupation Classification. Then, we translate the 2000 US COC to the harmonised 1990 COC, using the crosswalks of [Autor and Dorn \(2013\)](#).

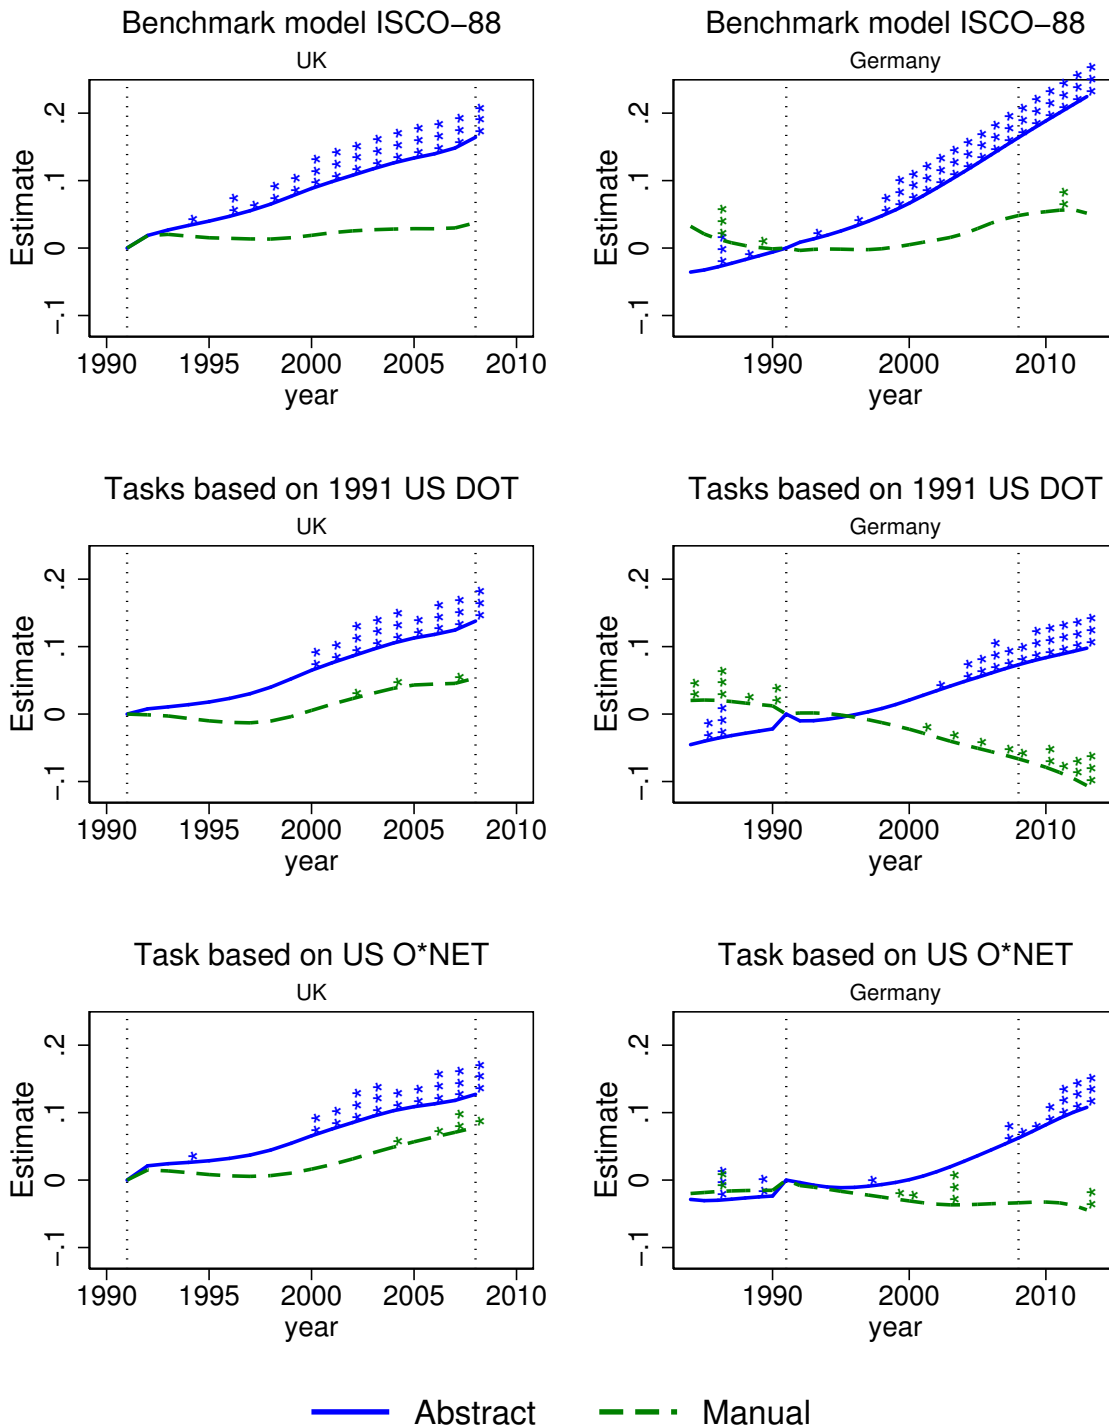
The steps used to categorize occupations into tasks using DOT and O*Net are similar. In particular, we first convert the ISCO-88 into the 2000 Census Occupation Codes (COC). We employ the crosswalks used in [Autor and Dorn \(2013\)](#) in order to convert the DOT (and O*Net) SOC occupations to the 2000 census occupations. Second, we use a set of specific measure indicators from the DOT and the O*Net to classify occupations into task groups. To construct the task indicators we use US employment shares as weights, as in [Goos et al. \(2009\)](#). Third, we compute the relative task share for each task and each occupation, as the ratio between the score for a task for a given occupation and the mean for that task across occupations. These shares indicate the task intensity of a given occupation. Finally, we attribute to each occupation the task with the highest share for that occupation. Notice that not all occupations have measures for task. When it is possible we substitute their values with the mean values of the closest groups.⁴⁵ Additionally, we attribute the ‘Legislators’ to the abstract group.

Although similar in their structure, the DOT and the O*Net have different indicators. As with the DOT, we follow [Autor and Dorn \(2013\)](#) whereas we select the same O*Net variables of [Acemoglu and Autor \(2011\)](#) to reproduce a classification based on three main groups.

The second and last rows of figure [B.6](#) report the task prices estimated on the DOT and on the O*Net, respectively. The results are consistent with the classification based on ISCO-88. It is interesting to notice, though, how the selected classification affects the results. The omitted category here is the group of occupations including routine manual and routine cognitive activities. Figure [4](#) in the main section indicates that in the UK the task prices of cognitive routine occupations declined compared to those of manual occupations. Indeed, here the task prices of manual occupations grow compared to the omitted category. The opposite is true for Germany. Indeed, figure [B.6](#) shows that the task prices of manual routine are flat or slightly declining with respect to those of the omitted category. For Germany, task price growth for abstract occupations is reduced once we aggregate routine manual and cognitive workers.

⁴⁵For example, the category ‘43-2021 Telephone Operators’ does not have these variables. However, we have these values for ‘43-2011 Switchboard Operators, Including Answering Service’ and ‘43-2099 Communications Equipment Operators, All Other’. Therefore we assign the mean values of the two latter to the former. Notice that excluding these categories does not affect the results as the number of occupational codes with missing values is low.

Figure B.6: Task Prices: US Occupational Classifications



Picture is based on a sample of 25 to 60 y-o men.

'Abstract' stands for abstract task. 'Manual' indicates non-routine manual, The omitted category, 'Routine' includes both manual and cognitive occupations.

The vertical dashed line indicates the common period in the two data sets, from 1991 to 2008.