ISER Working Paper Series

Job Polarization, Task Prices and the Distribution of Task Returns

Chiara Cavaglia
CEP and CVER, LSE

Ben Etheridge
University of Essex

No. 2017-09
June 2017

INSTITUTE FOR SOCIAL & ECONOMIC RESEARCH

University of Essex

E·S·R·C ECONOMIC & SOCIAL RESEARCH COUNCIL

www.iser.essex.ac.uk

Non-Technical Summary

The developed world has seen large shifts in the occupational composition over the last 30 years. To determine the causes of these shifts it is important to examine corresponding movements to occupational wages. Beyond observed wages it is important to decompose wage movements into more fundamental components. These components are, first, changes to the wage paid to each unit of effective labour (holding the quality of workers fixed) and, second, changes across occupational sectors in worker quality. This decomposition is especially important because if employment in a sector grows rapidly, it is likely that new workers will be of different average quality to the incumbents. Consequently, observed wages and unit wages will diverge.

In this paper, we use longitudinal data to track workers over time, and therefore to control for the changing quality component across sectors. We use this approach to provide new and coherent evidence for male workers from two large European countries: the UK and Germany. Specifically, we provide comparable cross-country evidence from the British Household Panel Survey and the German Socio-Economic Panel. As an organizing framework, we group workers into occupational sectors based on the nature of the job tasks performed. Within these groups, employment has grown fastest in ‘abstract’ (professional or knowledge) occupations and declined fastest in manual routine occupations.

We find that, in both countries, underlying unit wages have deviated substantially from observed average wages. Specifically, changes to underlying unit wages have been highly correlated with changes to employment share. On the other hand, the correlation between changes to employment share with observed wages is close to zero.

This finding has two important implications. First it implies that changes to the employment structure have been driven by changes to demand such as changes to technology. Second it illustrates the size of changes to the quality of workers across sectors over time. In particular, the average quality of workers in abstract jobs has declined substantially.

This study therefore raises important questions regarding the future of inequality. In particular, if demand for abstract work continues to grow, it is important to ascertain the capabilities (or “productivities”) of those who don’t currently carry out this type of work, but will in future. In the final section of the paper we begin to investigate this question by estimating inequality of productivities of *all* workers in *all* job sectors. This estimation is challenging because we often don’t see individuals work in more than one sector. We find that inequality is greatest in abstract work, and therefore that if employment in this type of work increases, inequality will grow.

Job Polarization, Task Prices and the Distribution of Task Returns*

Chiara Cavaglia,[†] Ben Etheridge[‡]

Abstract

We make two contributions to understanding the large shifts in occupational structure seen across developed countries. First, we estimate underlying prices on occupations, grouped by predominant task, using panel data from the UK and Germany. In both countries, price growth is positively associated with employment share growth. This pattern, which disappears with observed wages, is consistent with changes to labour demand, such as from technological changes. Second, we use the underlying Roy framework to further interpret these movements, by identifying the covariance structure of returns across tasks. The estimates show the importance of sorting based on productivity in abstract tasks.

JEL Classification: J20, J24, J31

Keywords: Job Polarization, Occupational Choice, Roy Model.

*We are grateful for discussions with and comments from Andrea Salvatori, Matias Cortes and Tom Crossley. Etheridge gratefully acknowledges funding from the ESRC Research Centre on Micro-social Change at the Institute for Social and Economic Research. All errors remain the responsibility of the authors.

[†]CEP and CVER, LSE. E-mail: c.cavaglia@lse.ac.uk

[‡]University of Essex. 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.

These results on prices, wages, and employment growth, are consistent with a Roy model of selection into occupational sectors. Our second contribution is to explore further the features of the Roy model that generate the patterns we see in the data. In particular we focus on the covariance structure of individual productivities across sectors. This covariance structure helps to predict a worker's expected productivity in the sectors away from which he selects. This covariance has been the focus of interest since Roy (1951) itself, and goes to the heart of debates on sorting in the labour market. In the context of job polarization, measuring this covariance is crucial to understanding the elasticity of supply across sectors, and the welfare effects of ongoing occupational change. In this paper we show how to identify this covariance structure using the measured sectoral prices and employment changes, among other features of the data. Our framework relates to that used by Autor and Handel (2013), who address this set of issues using cross-sectional data. Using panel data, we find consistent evidence on this covariance structure across both the countries we study. We use these results to interpret the observed patterns of wages and employment, and to characterize the nature of sorting in the labour market.

As is common in 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 tasks, and whether or not it is highly routine. In both countries, we see most prominently a rapid increase

[1] See, 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.

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. A corollary of these results is that the average quality of workers in abstract occupations has declined over time. Our results are robust to allocating jobs to tasks not by broad occupation but at a finer level of detail.[2]

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 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.[3]

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. Specifically we address the concerns about this approach raised by Gottschalk, Green, and Sand (2016). For example, in various specifications we control for heterogeneous human capital profiles, sector-specific job tenure profiles, and remove younger workers. 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.

The large changes in the occupational structure raise a second question: what happens to workers in declining sectors? This question has been addressed using cross-sectional data by, for example, Autor and Dorn (2013), and using panel data, by Cortes, Jaimovich, Nekarda, and Siu (2014).[4] To complement this empirical evidence on flows we take a more structural approach. As discussed above, we use the underlying Roy framework to identify key features of the joint distribution of returns across sectors. In addition to using the net flows across sectors and the estimated prices, we also show the importance of other features of the data. For example we show that the cross-sectional dispersion of wages of those who switch sectors is highly informative.

Our analysis of the distribution of returns has two key strengths. First we are able to identify the covariance structure without identifying all the parameters of the model. For example we can remain agnostic about whether different sectors have different amenities. Second, we use minimal parametric assumptions, about which we are explicit, and which we can easily adapt. In terms of results, we find a modest-to-strong correlation of returns across the sectors we study, in both countries. Moreover our results indicate the importance of heterogeneity in productivities

[2] In most of our analysis we use very broad occupational groups. Because we align these groups to predominant tasks, we call these groups 'sectors' or 'tasks' somewhat interchangeably.

[3] For 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.

[4] See also Autor and Dorn (2009) and Cortes (2016).

in abstract tasks: This heterogeneity is the largest component of persistent worker differences, and consequently plays a central role in occupational sorting.

Our paper relates most closely to three recent papers which estimate prices on tasks. First, as mentioned, Cortes (2016) uses panel data from the Panel Survey of Income Dynamics in the US. Two further papers use contrasting approaches using repeated cross-sections. Of these, Gottschalk, Green, and Sand (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 (2015) addresses selection into sectors using a control function approach. He uses data on ability scores in the US National Longitudinal Survey of Youth to control directly for otherwise unobservable characteristics. 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, changes to the education mix likely pose an even greater concern: in the UK at least, the proportion of new entrants to the labour market holding a university degree has increased dramatically. By tracking the same workers over time, the panel-data approach used here by-passes issues surrounding educational composition completely. We also advance on each of these papers by characterizing the nature of selection into each occupational sector.

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), who discuss Germany, while Beaudry, Green, and Sand (2016) address somewhat contrasting patterns in the US in the 2000s. Finally, our paper 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.

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 data. In particular it describes in detail the occupational classification used as well as showing patterns of wages and employment across the UK and Germany. Section 4 presents the main results on task prices. The identification and estimation of the covariances of returns is presented in section 5. The final section, 6, concludes. Extensive appendices provide further details and robustness for various features of our analysis. In particular we show how to modify our occupational classification to make it consistent with analyses from the US.

2 Framework

2.1 Econometric Framework

Our framework is based on a simple Roy (1951) 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 our model, the individual derives utility, and hence makes choices, from his own wages and the unobserved preference component; there are no savings and other family members are not explicitly modelled. 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 disamenities, 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. For our identification strategy to work, wages must depend on individual, time and sectoral characteristics in an additive way. Suppose that X_{it} is a set of observable characteristics capturing individual and time-specific factors. The return on these is given by δ_{jt} which may depend on sector-specific factors. We assume:

$$w_{ijt} = \delta_{jt} X_{it} + \theta_{jt} + \gamma_{ij} + v_{it} \tag{1}$$

where θ_{jt} and γ_{ij} are the primary objects of interest: θ_{jt} is the price of sector j at time t, and γ_{ij} is the (unobserved) ability of individual i in sector j. We separate θ_{jt} from $\delta_{jt} X_{it}$ in this expression precisely because of its importance. 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 $j^* = \text{argmax}_j \{u_{ijt}\}$ we can then define a set of binary indicators I_{ijt}, which capture sectoral choice as follows:

$$I_{ijt} = 1\,(j = j^*)$$

The observed wage, w_{it}, then takes the following form:

$$w_{it} = \sum_j I_{ijt} \left(\gamma_{ij} + \theta_{jt} + \delta_{jt} X_{it}\right) + v_{it} \tag{2}$$

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 strong assumptions on the functional form. Identifying parameters of interest in these types of environments is the subject of an established literature.[5]

In this paper, we use panel data for both the countries we study. Given panel data, and the assumed functional form on wages, it is possible to 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.[6] Intuitively, the parameters are identified in this approach from sector-specific wage growth. This type of approach raises some identification issues of its own, however, to which we return below.

By using the panel model approach, we identify the task prices semi-parametrically. In particular, identification requires no assumptions on the distributions of the unobservables: γ_{ij}, v_{it} and ε_{ijt}. In our approach, identification depends on the standard assumptions in fixed-effect regressions. In this regard, first we assume strict exogeneity of

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

[6] In 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.

the wage residuals, ν_{it}. Formally, if I_i is the set of sectoral choices, $\{I_{ijt}\}$, across time and sectors, then we require that $\mathbb{E}(\nu_{it}|X_i, I_i) = 0$. Therefore sectoral choice must be uncorrelated with wage residuals and depend purely on the unobservables, γ_{ij}, and ε_{ijt}, observables, X_{ijt}, and the time effects θ_{jt}. The strict exogeneity assumption also implies that the residuals ν_{it} must be independent of the preference shocks ε_{ijt}. We provide a test of strict exogeneity in appendix A.

Second, we assume that (unobserved) individual-level sector-specific skills, captured by γ_{ij}, are fixed over time. We therefore require that unobserved human capital factors are fixed and any changes to human capital are captured by the observables. We model human capital accumulation in the application by allowing for a rich set of covariates. As discussed below, these might capture tenure effects, in addition to, say, education effects and pecuniary switching costs.

An advantage of the approach used here is that we follow a given set of workers over time. Other approaches to identifying changes to task prices typically use repeated cross-sections. In particular, Gottschalk, Green, and Sand (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. According to the logic of the Roy model, if employment in the abstract sector grows, for example, it is likely that new entrants are of below-median quality. Gottschalk, Green, and Sand therefore obtain an upper bound on the growth of the (selection-free) abstract task price by trimming the wage distribution at the bottom by enough to keep total employment at the size of the base year. A problem with this approach is that it requires holding constant the distribution of abilities in successive cohorts. This assumption is problematic given that the composition of educational attainment in the US has also changed. In our context, changes to the education mix is likely to be even more of a concern: in the UK at least, the proportion of new entrants to the labour market holding a university degree has increased dramatically. By tracking the same workers over time, the panel-data approach used here by-passes issues surrounding educational composition completely.

A second approach is pursued by Böhm (2015), who controls for quality directly using ability scores in the US NLSY. He addresses selection into tasks using a control function approach. His analysis suffers from the standard problem with such approaches: namely that it is difficult to find instruments that determine sectoral choice without determining wages.

Aside from the task prices, θ_{jt}, we also focus on identifying the distribution of unobservable skills, γ_{ij}, within and across individuals. Identifying this distribution is important for considering how those in routine manual jobs coped with the apparent decline in demand for their labour. Of course, we observe the distribution of γ_{ij} only for those selected into sector j. And we observe the joint distribution of γ_{ij} only for those who work in multiple sectors. Identifying features of the unconditional joint distribution is therefore challenging. We turn attention to estimating these distributions in section 5.

2.2 Empirical Implementation

As discussed, we estimate equation 2 using fixed effects at the sector-spell level for each individual. This estimator controls for the time invariant component, γ_{ij}, by de-meaning the wage within each individual-sector spell. The task price θ_{jt} is captured by interacting dummies for occupational groups and years. In fact, these prices are identified only as changes with respect to a base year and a base sector. In the results reported below, therefore, these coefficients should be interpreted as the changes to the price over time, with respect to that in the year 1991 and with respect to the change to the price on routine manual work.

In the estimation we include a rich set of controls. These are: a quartic polynomial in age, interacted with a full set of education dummies (capturing seven levels of educational attainment), region of residence, marital status, and the year dummies. In some specifications, we also control for interactions of task with education, of task with job tenure, and for the interaction of each of age and task with trade union status. When computing standard errors we allow for arbitrary serial correlation in the residuals by clustering at the individual level.

Perhaps the main challenge to our identification is being able to capture pure time effects on task prices aside from age effects and tenure effects. The problem is that within each individual-sector spell, and from one wave of the panel to the next, time, age and tenure grow collinearly.[7] Our identification problem becomes more pronounced given that the trends in occupational employment are monotonic over time. Therefore, it is difficult, at least in the UK, to argue for time effects based on conspicuous time-series variation. More generally, this identification problem is discussed, for example, by Gottschalk et al. (2016) who suggest that differences in wage growth across sectors might be affected by differences in the nature of wage-tenure contracts. In our analysis we confront these issues by including in our regressions the rich set of controls discussed above. These controls go beyond what is included in the analysis of Cortes (2016). In particular, by including interactions of a polynomial in age with a full set of education dummies, we aim to capture heterogeneous age profiles by skill type. Moreover, to capture the heterogeneity in age profiles caused by explicit contracting differences, we include an interaction of age with trade union status.[8]

Similarly, we address concerns about heterogeneous tenure effects by using available information on job tenure. Specifically, and as mentioned above, we interact a quartic polynomial in job tenure with the occupational groups. Of course, wages are determined not only by tenure in the job, but also by occupational tenure itself, as emphasized by Kambourov and Manovskii (2009). Similarly, wages decline faster the further individuals move away from their initial occupation, as shown by Gathmann and Schoenberg (2010). It should be emphasized throughout that our framework easily incorporates occupational tenure effects as long as they are common across sectors. However, to capture tenure profiles that vary by sector, we consider controlling for job tenure rather than occupational tenure directly in many ways more attractive. In particular, specific occupational tenure effects can be identified separately from time effects only by imposing quite strong restrictions. The restriction, for example, in Cortes (2016) is that occupation-specific human capital is lost completely whenever the occupation spell finishes. In contrast, we can introduce heterogeneous job-tenure effects cleanly within our framework. These job-tenure effects are naturally lost completely whenever the individual switches jobs, assuming that they never return.[9] Finally, and as a practical matter, 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 wave of the survey. 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

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

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

[9] Notice 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, even though job- and occupational tenure often switch together.

therefore left-censored.[10] Nevertheless, we perform an analysis controlling for heterogeneous occupational tenure profiles in our additional robustness exercise, shown in appendix A.

As a final point, it is worth adding that our framework is able to control for differences in returns to tasks across educational groups. In particular, we might expect that the return to working in abstract jobs is greater for the higher educated than the low educated. This fact would explain why educational attainment and sectoral choice are positively correlated. These differences in returns can be assessed by including an interaction between education and sector. Of course, because this interaction is constant within an individual-task spell, it is differenced out by the fixed effect estimator, and so is rendered irrelevant in the main regressions. The interaction can, however, be picked up when estimating by OLS.[11]

2.3 Theoretical Background

The main motivation for estimating task prices is to provide evidence on wider debates about the evolution of the wage structure. It is therefore important to root our framework and approach in the wider theoretical literature.

First, it is important to emphasize that the underlying sectoral prices that we measure are themselves equilibrium outcomes, capturing the effects of changes in demand for tasks and supply of effective labour. While our approach controls for sorting into occupational groups, it is silent on the factors that drive task prices themselves. Nevertheless, as discussed in the introduction, our results indicate that the changes in occupational structure are demand led. For example, and given that the largest employment declines are seen in routine manual work, our results are broadly consistent with explanations based on 'routine-biased technical change' (RBTC).

A number of frameworks have been developed to capture RBTC-type effects explicitly. A workhorse model is given by Acemoglu and Autor (2011). The most suitable framework 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. This model predicts that when technology causes a decline in demand for routine tasks, then the relative price of the other tasks/occupations increases, as does their employment share.[12] The employment growth therefore depends on sorting, the precise nature of which depends on the distribution of latent skills, γ_{ij}.

The second aim of our work is therefore to estimate features of the distribution of these latent productivities. The framework used to estimate task prices here allows for the correlation of skills across sectors to be unspecified. The framework therefore nests two benchmark models of this distribution, discussed by Gottschalk et al. (2016) and originally classified by Willis (1986). First, if latent skills are perfectly correlated across sectors, then this gives rise to a 'hierarchical ability' model. Second and alternatively, if latent skills are uncorrelated across sectors, then

[10] We have the the date at which the individual started the "current" job, but he may have had a previous job in the same occupational group beforehand, which is not accounted for in the survey. In this case, the derived occupational tenure would be shorter than the true tenure.

[11] It is worth considering whether these heterogeneous returns to task across educational groups change with time or over the life cycle. To estimate this requires introducing a triple interaction of education, task and, say, time. We did experiment with these triple interaction models, but they are typically not robust even with large datasets: these interactions soon lead to hundreds of dummies in the estimating equation, which are estimated imprecisely. Investigating changes in the return to task by education is, however, an interesting question for future research.

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

an 'independent productivity shocks' model is implied. The hierarchical ability model implies that skills can by captured by a single measure in one dimension and is typically used in the theoretical papers discussed above.

As a preview of the analysis in section 5, it is worth commenting on how these simple frameworks relate to basic features of the data. A simple version of the hierarchical ability model, for example, states that those with high skills have a comparative advantage in sectors that are most productive. This model is therefore consistent with the data only if employment is perfectly segmented into different sectors in equilibrium, with the highest skilled sorting into the top sector, and the lowest skilled sorting into the bottom. In fact this model is falsified by the data immediately, because residual wages overlap substantially across sectors. Indeed, as shown later in section 3, even though the occupational groups are clearly ordered by average wage within each country, these average wages for the bottom three tasks are close. The overlap in residual wages therefore implies a model featuring some lower correlation of productivities and where workers receive differing returns in each sector: some individuals receive good wages in otherwise low-paying sectors because their skills there are particularly high.

3 Data

3.1 The BHPS and GSOEP Datasets

For the UK, we base our analysis on the British Household Panel Survey (BHPS). The BHPS began in 1991 with a representative sample of about $5,500$ households and $10,300$ individuals. The longitudinal dimension allows us to investigate the patterns of occupational change and its impact on household income over almost twenty years, until 2008. The survey follows all the adult members of a given household over time, even after joining a new household. We use the baseline survey and ignore the later booster samples from Scotland, Wales and Northern Ireland and of low-income households.

For Germany, we use the German Socio-Economic Panel (SOEP), a longitudinal survey currently comprising approximately $11,000$ private households. The survey provides similar information to the BHPS and the US Panel Study of Income Dynamics (PSID).[13] The first wave was collected in West Germany in 1984 and consisted of about 6,000 households and 12,000 individual respondents. Progressively, additional samples were introduced to provide adequate coverage to specific groups, such as immigrants and high-income earners. In 1990, East German states were also added to the survey. Similarly to the case of the UK, we only consider the original sample. Specifically, and consistently with the existing literature, we decide to exclude East Germany because its wage structure is very different from the West (see, for example, Dustmann, Ludsteck, and Schönberg, 2009).

3.2 Variable Definitions

Our analysis centres around groupings of occupations that correspond to predominant job tasks. This approach is standard in the literature on assessing, for example, the extent of routine-biased technical change. Our analysis is more focussed on estimating underlying prices of different types of labour than on directly testing hypotheses for the shifts in employment and wage structures. Therefore, our occupational groupings are somewhat arbitrary.

[13] Wagner et al. (2007) provide further details on the scope and evolution of GSOEP, as well as on the differences and similarities with BHPS and PSID.

Nevertheless, in order to inform on these wider debates, and because some occupations characterized by specific tasks have undergone noticeable changes, we use this task-based approach.[14]

We assign workers to occupations, or tasks, using the 1988 International Standard Classification of Occupations (ISCO-88), approved by the ILO Governing Body in 1988. This international classification relates closely to some national systems, particularly to the British 1990 Standard Occupational Classification (SOC90).[15] We use the ISCO classification, rather than relying on national systems, for comparability of results. For robustness, however, we repeat the main analysis using national classifications, and report these in appendix B.[16]

We follow a similar categorisation to that which Acemoglu and Autor (2011) apply to US data. We construct four broad groups by merging the occupational categories of ISCO-88 rather than imputing task data to these categories.[17] Specifically, the four groups classify the occupations according to the type of their predominant task. 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'.[18]

The task approach represents an organizing framework to investigate the role of technological change on the structure of wages and employment. However, as discussed in Autor and Handel (2013) the mapping from the occupational classification to the predominant task of a given occupation is somewhat arbitrary. A more precise approach may be to impute task requirements from the the US Department of Labor's Dictionary of Occupational Titles (DOT) and its successor, the Occupational Information Network (O*Net).[19] An advantage of this method with respect to the one adopted here is that the predominant task is attributed to each occupational title, allowing variation within occupation major groups.

Occupations in DOT and O*Net are expressed in terms of US Social Occupational Classifications. Therefore, to be able to use this we would need several steps to convert the ISCO-88 classification. Given that at each conversion some precision is lost, as there is no unambiguous correspondence across classifications, we decided to

[14] A related strand of the literature classify workers by skill or qualification level, rather than by occupation or task. For example, Micheals et al. (2014) found that the development of technology increased the demand of college graduates at the expense of those with a mid-level qualification in the US, Japan and nine European countries.

[15] There may be some differences between the task classification based on the national Standard Occupational Classifications and that based on ISCO-88. Each ISCO category is based on the complexity and the specialization of the skills required to fulfil the tasks and duties. In practice, however, for the UK, the task classifications seem very similar. For example, for the UK the correlation between the task variable based on ISCO-88 and the one based on the 1990 Standard Occupational Classification is 0.84.

[16] The German Klassifizierung der Berufe (KldB) is structured differently, as explained in Appendix B. Therefore, we focus on applying the SOC classification to the GSOEP as well.

[17] Note further that this approach implicitly assumes that the task content of each occupation is similar in Europe and in the US.

[18] ISCO-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. Therefore, the percentage of workers in these occupations is low. We keep the individuals in this occupational category, even though excluding them would not affect the estimates of the task prices, as we show in the Appendix.

[19] See for example Autor and Dorn, 2013 and Goos et al. (2009).

adopt Acemoglu and Autor's approach.[20] Nonetheless, as a robustness check and to allow a more straightforward comparability of our study with those for the US, we have also estimated the task prices using the DOT- and O*NET-based task classifications. The results, as well as detailed instructions on how we attribute DOT- and O*NET measure indicators to each ISCO-88 occupation, are reported in the appendix.

A possible issue with the occupational status variable is measurement error. In the BHPS, for example, a substantial number of individuals appear to switch occupation one year, only to revert immediately to their original occupation the next. Additionally, there are cases when the occupation changes from one year to the other even though respondents state that there has been no change in terms of their job or position, within the same firm or elsewhere. This phenomenon is less common after 2006, when dependent interviewing was introduced.[21] To give an idea of the magnitude of the problem, we compute the mobility rate, as the share of workers with a different occupation group with respect to the previous year and we plot the rate over time. The appendix figure B.1 illustrates that the mobility rate drops from just below 20% in 2005, the last year of independent interviewing, to around 7% in 2008. A similar issue arises with GSOEP. The GSOEP survey alternates years with full and partial interviewing. Here, average occupational mobility is lower in the years based on partial surveys, as indicated in figure B.2.[22] The mobility rate changes from around 14% in the years of full interviews to less than 4% in years with partial interviews. To address this issue, we construct a corrected version of occupational group, using a slightly different procedure for the two data sets. The full procedure is explained in the appendix. This procedure is not perfect: specifically, the risk is to eliminate some true changes in occupation. For this reason, when estimating the task prices we use the uncorrected measure, and we report results using the corrected measure in the appendix. On the other hand, the analysis of distributional features in section 5 relies more on information about year-to-year switchers. For that analysis it is appropriate to isolate those for whom we are sure that occupational switches are genuine, and so we use the corrected measure.

We construct the natural logarithm of hourly wages in the local currency. This wage measure is constructed as current gross monthly earnings divided by weekly working hours, 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 GSOEP, this variable is bounded at 80 hours per week. We therefore bound it similarly in the BHPS to ensure comparability of results.[23] Finally, wages are deflated by the 2010 Consumer Price Index.

3.3 Construction of the Samples

For wages, we consider men in their prime age, between 25 and 60. We focus on males for the usual reason that selection into the labour market itself poses fewer econometric challenges than for females. The lower bound on

[20] Acemoglu and Autor (2011) compare the categorisation done using the occupational categories of the US Standard Occupational Classification and the classification obtained by attributing each occupation a task measure using the US Department of Labor's Dictionary of Occupational Titles (DOT). The authors conclude that the pattern of task intensity across the occupations for the two measures is comparable.

[21] Surveys 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.

[22] In 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 GSOEP.

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

age of 25 removes labour market entrants, many of whom do casual or part-time jobs at the beginning of their career or while in education. We set the upper bound at 60 as those working after this age are also increasingly highly selected. For example, older workers might change from their job to a lighter occupation instead of officially retiring. Additionally, in order to be able to consider dynamics, individuals are only included in the sample if they have provided information about their earnings in at least 5 waves.

For the wage data, 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. Including part-time workers is unlikely to affect the main results because only 2.5% and 2.7% of the prime-age sample in Britain and Germany respectively work less than 30 hours per week.[24] Finally, we exclude observations with missing values of occupation or labour force status, number of worked hours and years of job tenure (if employed) or completed level of education. As the appendix table A.1 indicates, the resulting 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.

3.4 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 the four different occupational groups over time. It also displays the fraction of working-age men in the labour market but out of work. As discussed above, we classify the occupations into four broad groups, according to the type of their predominant task.[25]

Figure 1: Employment Patterns by Occupational Sector

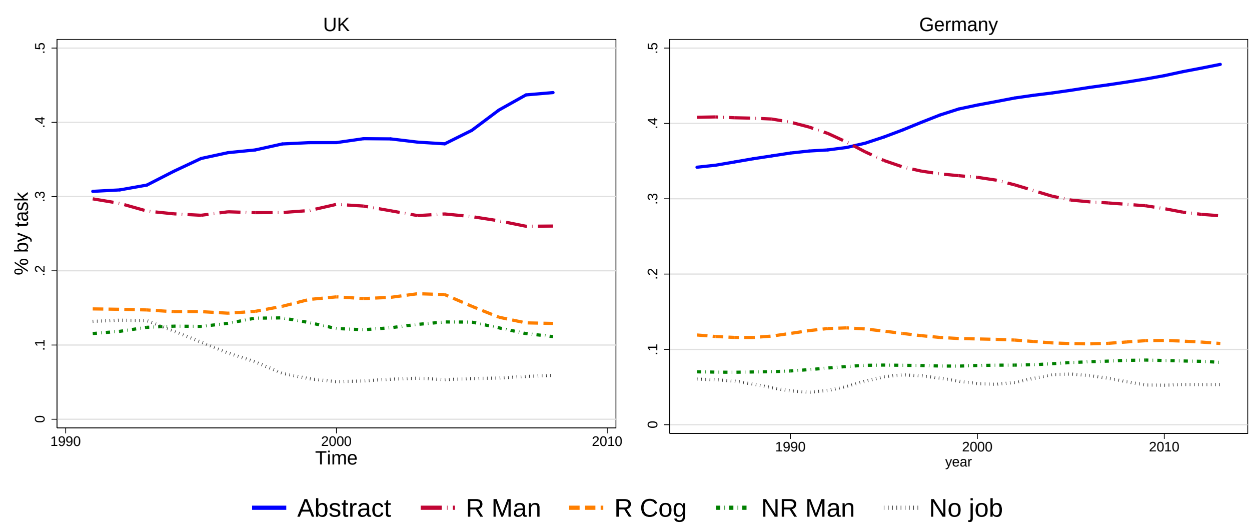

Notes: Sample includes all men aged between 16 and 64 and in the labour market. 'Abstract' stands for abstract task. 'R Man' indicates routine manual. 'NR Man' indicates non-routine manual. 'R Cog' indicates routine cognitive.

[24]Similarly, we do not bound the hourly wage upwards or downwards. It would be difficult to identify the threshold that allows a fair comparison across occupations, without depriving each occupation of its specificity. This is especially a problem since the minimum wage was first introduced only for some occupations and only in the mid-1990s. The consequence 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 would affect the main results of this analysis. 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%.

[25]The 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 75% 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, abstract and routine manual workers each represented 30% of the sample. Over 1991 – 2008 the abstract share then grew relative to routine manual by 18 percentage points. For Germany, shown in the right panel, the patterns are even more pronounced. On average, 77% 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 22 percentage points over 1991-2008, or 29 percentage points over 1985-2013. In this country, however, routine manual workers decreased more than in the UK: around 10 percentage points versus less than 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.

Figure 2: Median Log Wage by Sector

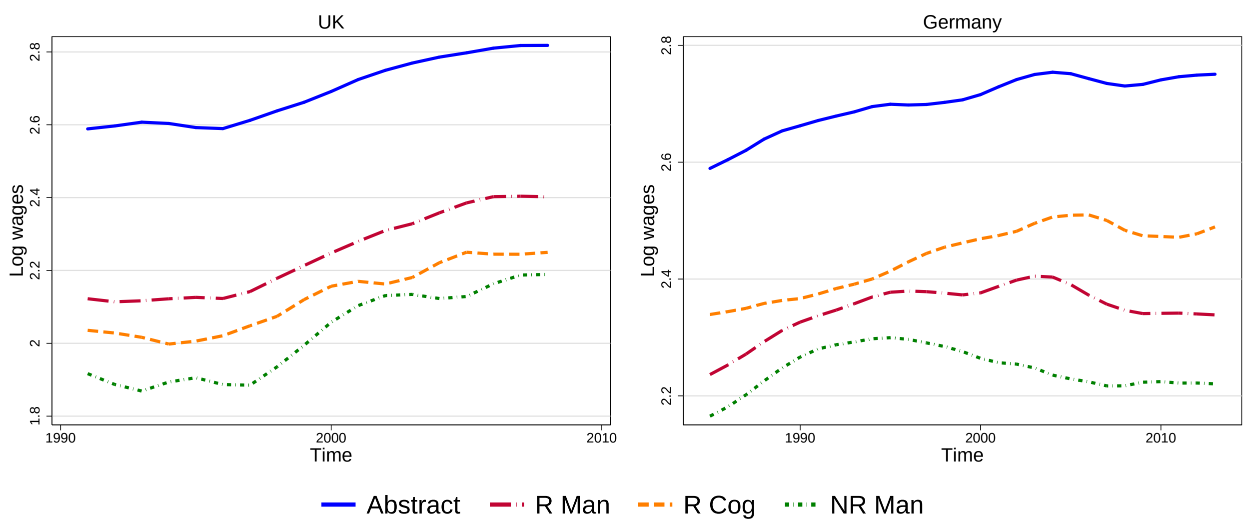

Notes: Based on a sample of 16 to 64 y-o employees. 'Abstract' stands for abstract task. 'R Man' indicates routine manual. 'NR Man' indicates non-routine manual. 'R Cog' indicates routine cognitive.

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 wages were roughly flat. The fact that employment in abstract occupations grew so fast while relative average wage 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 employment changes to initial wages at the 1-digit ISCO88 level. The results are shown in figure 3. It shows the predicted values from a regression of the change in employment share on a quadratic function of the initial log wage, following Goos and Manning (2007). The figure shows a remarkable similarity in patterns across countries, which is 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, 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. In the next section, we compare the change in employment share not with initial occupational wages but with changes in occupational price, and argue that the evidence on prices across both countries is also strikingly consistent with changes to demand for labour.

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

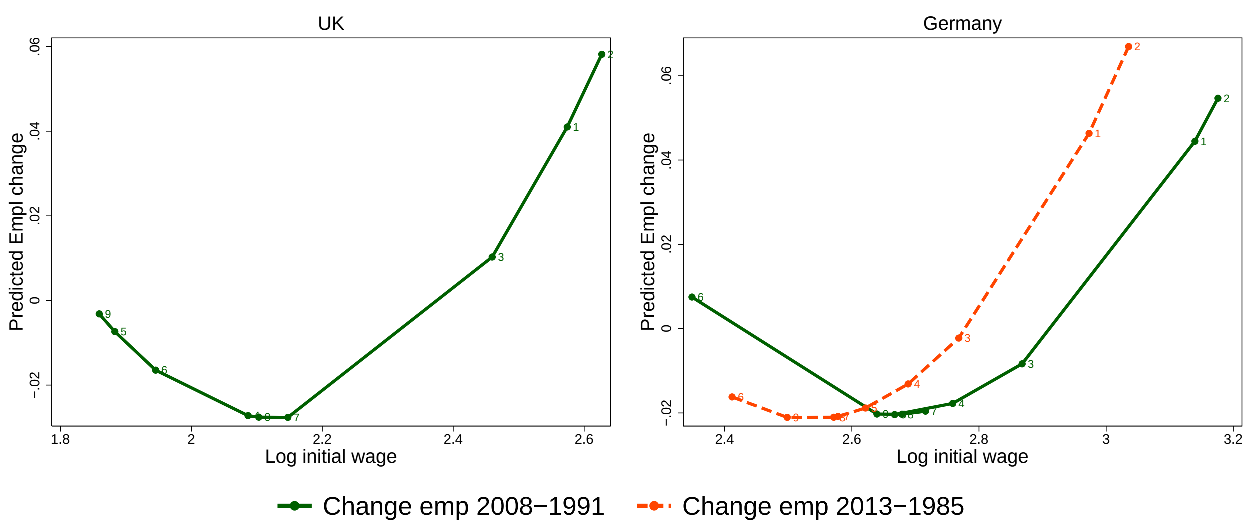

Notes: Figure shows predicted values from a regression of the change in employment share on log wage and wage squared in the initial year, as in Goos and Manning (2007). The label represent the 9 categories of the 1-digit ISCO-88: 1 'Legislators', 2 'Professionals', 3 'Technicians and associate professionals', 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'

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.

Figure 4 shows the main results for both the UK and Germany, split into two panels. The right hand panel shows estimates for the price on the abstract task relative to the routine manual task, indexed to 0 in 1991, and smoothed using a 3-period moving average. The figure shows that, compared to routine manual, task prices for abstract jobs increased markedly. In the UK, the price increased relatively by around 13% by 2008, or a little over a half a percent per year. The panel also shows the relative price change for Germany, together with vertical lines to indicate the start and end of the BHPS sample, for easy comparison across countries. The relative price change in Germany was even larger than in the UK, 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. The stars on the price estimates indicate that, when significant, changes from the base are significant at the 1% level.

Figure 4: Price and Employment Changes for Abstract Relative to Routine Manual

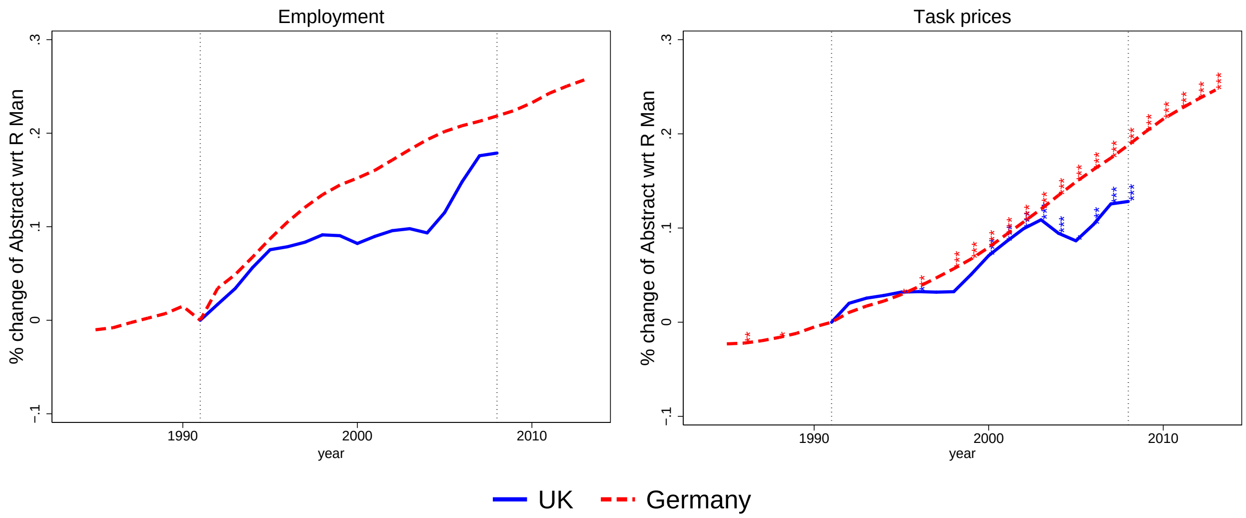

Notes: The figure illustrates changes in employment and in task prices of abstract workers with respect to manual workers. Employment change is computed for 16 to 64 y-o men. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

One main focus of this paper is to link the changes to task prices with shifts in employment. With this in mind, figure 4 also shows growth in the employment share in abstract occupations relative to routine manual occupations in the left hand panel. Again it shows results for both countries, indexed to 0 in 1991, and also smoothed using the 3-period moving average. The panel shows, for example, that 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.

Statistics for employment are computed using a broader population than is used to estimate the prices. In particular, here we use all males aged $16-64$, to provide a more comprehensive measure of labour supplied. Strictly speaking, we should measure all effective labour provided across the whole economy, including from females. 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 non-routine cognitive occupations, women experienced a much larger decrease in routine occupations than men (Black and Spitz-Oener, 2010). 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 correspondence with the extra increase in the occupational price shown in the right hand panel.[26]

[26]It is interesting to note that the data show a slowdown in the employment shift towards abstract occupations in the UK in the middle of the sample period. However, there is reason to think this is due to sampling variation. Data from the much larger UK Labour Force Survey

For completeness, figure 5 shows the evolution of prices over time for both countries and for all sectors. The additional sectors are routine cognitive and non-routine manual alongside the abstract group already shown. Again we show these prices 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 small. 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. 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, as shown in figure 1. It is also worth pointing out that these price changes can still be reconciled within the Roy framework, as discussed by Böhm (2015). With multiple occupations, then the relationship between employment 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 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, as we show later, employment changes in fact match price changes closely in both countries.

Figure 5: Task Prices Relative to Routine Manual

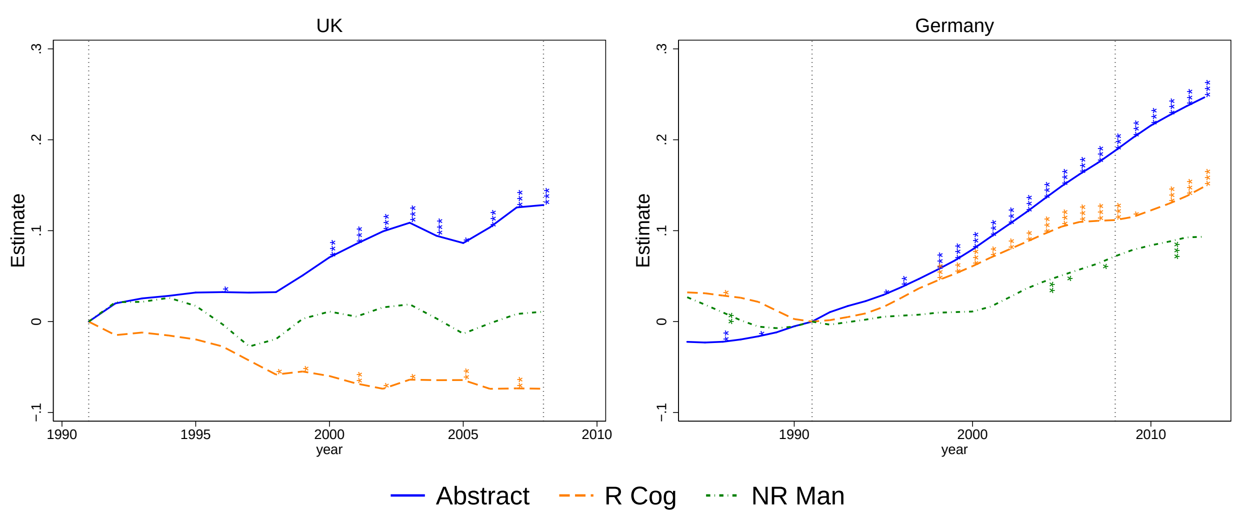

Notes: Based on a sample of 25 to 60 y-o men. 'Abstract' stands for abstract task. 'NR Man' indicates non-routine manual, and 'R Cog' indicates routine cognitive. The vertical dashed line indicates the common period in the two data sets, from 1991 to 2008. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

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 figures above, are summarized in the second column. It shows the relative growth in prices for each sector in 2008 only, 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 only for 1991-2008 for direct comparison with the UK. Corresponding with the figures above,

discussed in, for example, Carrillo-Tudela et al. (2016), show that the shift away from routine occupations was monotonic and constant. We therefore do not treat this slowdown as a significant feature of the data.

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 5.

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 show 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 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, 2016). 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

Table 1: Changes in task price by occupation: different specifications

	OLS		Benchmark FE		Educ. Returns		Job Tenure		TU		35-60 y-o	
UK												
1991 Abstract	0.259***	(0.03)										
1991 R Cog	0.055*	(0.03)										
1991 NR Man	-0.033	(0.04)										
2008-1991 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)
2008-1991 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)
2008-1991 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)
Constant	2.169***	(0.03)	2.484***	(0.06)	2.491***	(0.06)	2.485***	(0.06)	2.518***	(0.06)	2.352***	(0.06)
Observations	24364		24364		24364		24364		24364		16955	
Germany												
1991 Abstract	0.157***	(0.02)										
1991 R Cog	0.028	(0.02)										
1991 NR Man	-0.071***	(0.03)										
2008-1991 Abstract	0.042	(0.03)	0.175***	(0.03)	0.156***	(0.03)	0.203***	(0.04)	0.150***	(0.04)	0.110**	(0.05)
2008-1991 R Cog	0.062	(0.05)	0.106***	(0.04)	0.108***	(0.04)	0.072	(0.06)	0.106**	(0.05)	-0.007	(0.07)
2008-1991 NR Man	-0.046	(0.05)	0.041	(0.04)	0.031	(0.04)	0.061	(0.08)	0.094*	(0.05)	0.038	(0.13)
Constant	2.674***	(0.03)	2.767***	(0.07)	2.767***	(0.06)	2.764***	(0.06)	2.670***	(0.09)	2.790***	(0.08)
Observations	36918		36918		36918		36918		21791		27124	
Adjusted R^2	0.309		0.144		0.142		0.144		0.101		0.058	
Controls:												
Baseline controls	x		x		x		x		x		x	
Education∗Year					x							
Job tenure∗Task							x				x	
Union membership∗(Task,Age)									x			

Notes: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$ Standard errors clustered at individual level in parenthesis. For all models, 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; in 4, model 2 is augmented with the interaction between occupation and job tenure, and between year and job tenure; in 5, model 2 includes a dummy for trade union association and interactions with occupation and with and age; in 6, model 3 is estimated on a group of mature workers, from 35 to 60 trade union status. Here we interact the union status of the worker separately with sector and with age, to pick up heterogeneous effects.[27] The results are identical to the benchmark regression for the UK, and very similar for Germany.[28]

Finally, 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 sector-specific job-tenure effects. The results are shown in the last 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. More complete results from all these regressions are presented in appendix A. In appendix B we show results when using different occupational classifications.

It is worth comparing the results as a whole 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 compared to the routine manual task by computing statistics after trimming the top of the distribution of routine manual wages and the bottom of the abstract wage distribution. In their case they compute medians, but we compute means. Using the UK as an example, employment in the abstract occupation 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 occupation with the mean obtained from trimming a third of wages from the bottom. We do this for the abstract occupation 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 the bounds. 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. Note that we do not directly address selection effects on the other sectors in this simple exercise.

We conclude this section 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 6. 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. In particular, for example, within the abstract sector, the largest employment growth has been for legislators (managers), who have also seen the price

[27] We recognize that controlling for union status in this way is an imperfect way of picking up contracting effects caused by centralized pay negotiation. In many occupations, and many jobs, pay may be bargained centrally for all workers regardless of the union status of the individual. Wages are determined this way in the UK in the higher education industry, for example. A more thorough approach would therefore be to identify occupations at a fine level for which pay is organized collectively. Nevertheless, our results are indicative of those likely to come from a more thorough treatment.

[28] It is worth noticing that information about trade union membership is not available for all waves in GSOEP. Therefore, it is calculated on a smaller sample. When we estimate the benchmark model on the same reduced sample the coefficients are almost identical. For example, the estimate of 2008-1991 Abstract would be 0.154.

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 7 shows the flat relationship between employment share growth and growth in average wages.

Figure 6: Task Prices and Changes in Employment at the 1-Digit Level

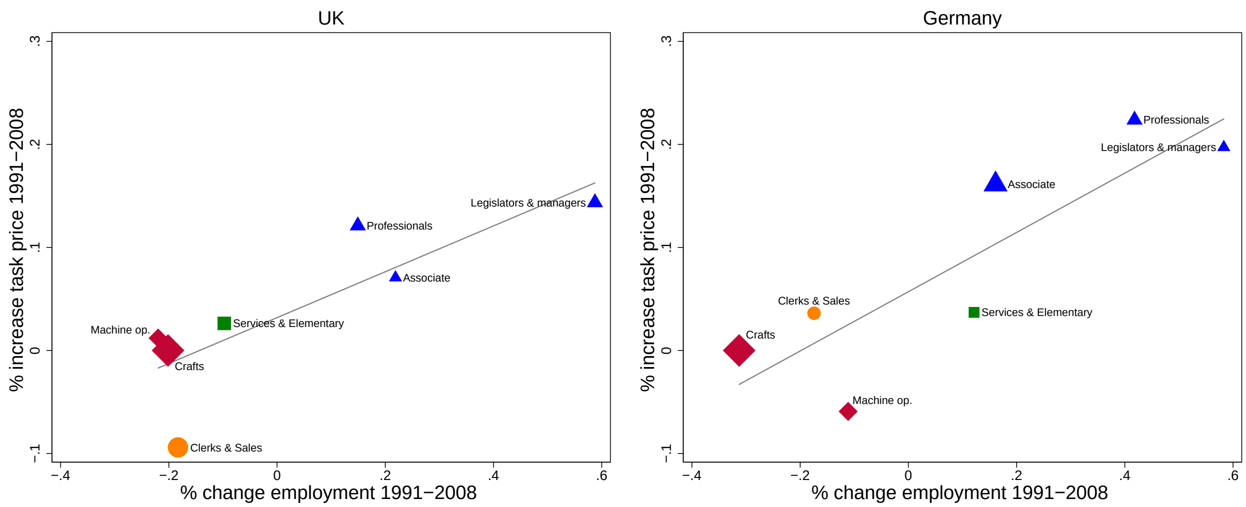

Notes: Employment change is computed on a sample of 16 to 64 y-o men. Categories of the same shape and colour belong to the same task group. Bigger shapes indicate larger 1991 employment shares. Task prices are estimated relative to crafts. The straight line is the coefficient of a regression of changes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

5 Estimating the Covariance Structure of Sectoral Returns

Section 4 showed how to estimate task prices using the fixed effect regressions. As discussed, we are also interested in identifying the distribution of unobserved returns, γ_{ij}, within, and across individuals. In this section we identify and estimate a specific feature of the joint distribution, the covariance structure, using unexploited features of the panel data.

Identifying features of this joint distribution allows us to interpret the movements in prices, wages and employment shown in section 4. Our analysis therefore complements, for example, Autor and Dorn (2013) who provide a detailed empirical assessment of outcomes for routine workers in the US.[29] In this sense, our analysis provides a structural interpretation of the long-run supply curve across tasks.

Moreover, identifying the covariance structure also allows us to interpret to what extent those who are productive in one sector are likely to be productive in another. This is important for assessing the welfare consequences of ongoing changes to the employment structure: for example, if those in disappearing occupations have low productivity in all other occupations then the welfare consequences of occupational change are particularly severe.

[29] See also Cortes (2016) for a panel analysis of routine workers in the US using the PSID and Cortes et al. (2014) who use the CPS. See also Autor and Dorn (2009).

Figure 7: Average Wages and Changes in Employment at the 1-Digit Level:

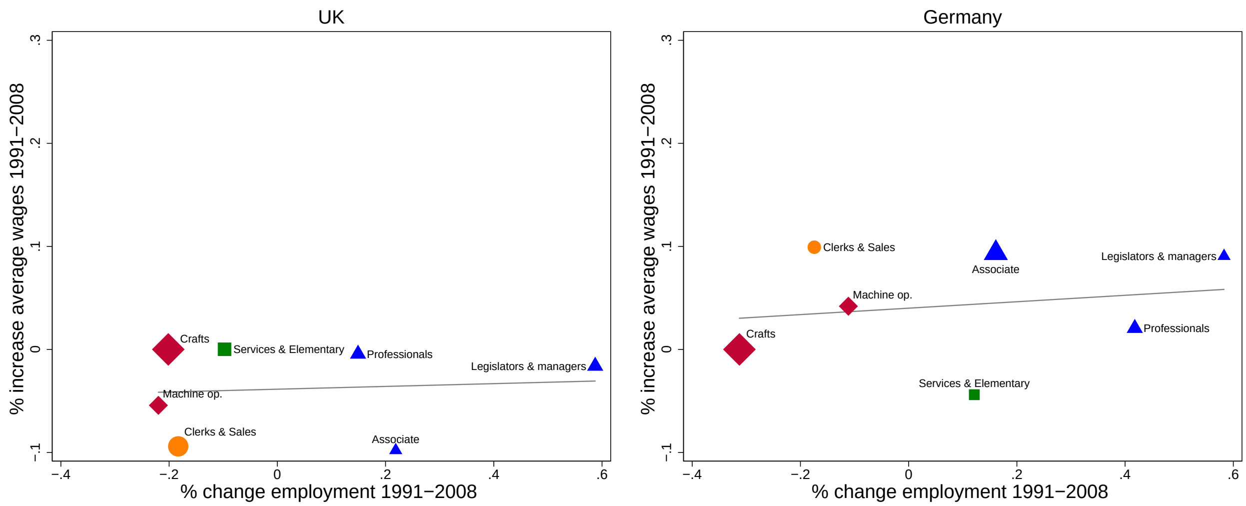

Notes: Categories of the same shape and colour belong to the same task group. Bigger shapes indicate larger 1991 employment shares. Average wage changes are computed from the coefficients of an OLS regression (similar to the model in the first column of table 1). They are estimated relative to crafts.

Identifying the distribution of returns across occupations has been of interest at least since Roy (1951)'s original article itself. Identification of any Roy model is challenging. As discussed, for example, by French and Taber (2011), any strategy using only the cross-sectional assignment of workers to sectors either requires an exclusion restriction or is highly dependent on functional form assumptions. However, given panel data and sufficient variation in prices, Heckman and Honore (1990) show that model parameters can be identified non-parametrically. Here we use the spirit of this result to identify the covariance structure of returns in an intuitive way. We show how the covariance structure maps into the elasticity of employment growth with respect to sectoral price. We also show how the covariance structure maps into the dispersion of wages for those who switch sector. Our approach requires minimal parametric assumptions, about which we are explicit. A strength of our analysis is that we show how to identify the covariance structure, while remaining agnostic about other parameters: the intercept of underlying returns and sector-specific disamenities. By ignoring these incidental parameters, our identification argument is transparent and intuitive.

Our analysis relates to that in Autor and Handel (2013), who address a similar set of issues by introducing a general framework distinguishing between occupations and tasks. In their framework, workers have different capabilities in various tasks, while different occupations offer different task returns. The framework used in our paper can be thought of as a special case of Autor and Handel's framework in which occupational groups and tasks are linked by a one-to-one mapping. As such, our individual-specific fixed effects γ_{ij} capture both individual-specific capabilities and the sector-specific return to those capabilities. In the current paper, we usually refer to this component as an individual-specific return, while recognizing that it captures factors that also reflect capabilities or skills.

In this section we aggregate sectors further to 2 broad groups: abstract, and the other occupations combined. We do this for two reasons. First, using only 2 sectors makes the discussion of identification simpler and more intuitive. Second, in both countries, the largest relative change in both price and employment is between these broad groups. In the UK, for example, the price between routine manual and non-routine manual changed little. Unfortunately, by aggregating in this way, we cannot directly interpret outcomes for routine workers. Nevertheless, the argument developed easily translates into 3 or more sectors. With 3 sectors, however, the calculations become highly involved and lengthy, for little extra insight.

As we shall see, we find a modest-to-strong correlation of returns across sectors in both countries. When we re-cast the returns in terms of orthogonal components, we find that the 'pure' abstract component is the largest source of persistent worker differences. This result indicates that heterogeneity in productivities in abstract tasks plays a key role in labour market sorting. On the other hand, there is little heterogeneity in abilities specifically in other tasks. A stylized interpretation of our results is that workers sort into abstract tasks if and only if they have particularly high abstract skills: workers sort away from abstract tasks because of low abstract skills and not because of particularly high skills in, say, routine or manual tasks.

5.1 Identification

We derive three sets of moment conditions. These are: the price elasticity of employment changes; the variance of wages for those who switch sector; and the cross-sectional variance of wages observed in each sector.

Before deriving these conditions we restate in more detail, and refine, the basic framework. Compared to the framework addressed in section 2, we assume away transitory shocks to non-pecuniary sector-specific returns, so that individuals switch sector purely because of their individual-specific (fixed) skills and because of changes to the aggregate price. The effect of shocks is discussed below and addressed in appendix C. They do not substantively affect the analysis. As discussed, we restrict the environment to 2 sectors, which we call 'abstract' and 'non-abstract', or 'other'. In this environment, using the same notation as in section 2, individual i selects into abstract at time t if and only if

$$\delta_a X_{it} + \theta_{at} + \gamma_{ia} + \varepsilon_a > \delta_o X_{it} + \theta_o + \gamma_{io}$$

where ε_a is a mean non-pecuniary preference for the abstract sector, which here we assume is constant across individuals and time. The non-abstract group is denoted by the 'o' subscript, for 'other'.

We need consider only changes in the *relative* price on abstract work. Therefore the non-abstract price θ_o is held constant, and, in fact, could be normalized to zero. We also assume for convenience that returns on observable characteristics other than sectoral choice are constant over time.

For the following formulae it is useful to define $\eta_i = \gamma_{ia} - \gamma_{io}$ as individual i's extra return in the abstract sector. Moreover, we define η_i^* to be the standardized version of this variable. Then the individual selects into the abstract sector if and only if

$$\begin{aligned} & \eta_i > (\delta_o - \delta_a) X_{it} + \theta_o - \theta_{at} - \varepsilon_a \\ \iff & \eta_i^* > \beta_{Xt} \end{aligned}$$

where β_{Xt} indicates the (standardized) selection threshold at time t, which depends on observable characteristics. For example, the threshold may be lower for highly educated individuals if their return in the abstract sector is, on average, particularly high.

Employment Elasticity

We first show how parameters of the model map into the elasticity of employment changes with respect to price. Omitting i subscripts, the proportion of individuals working in the abstract sector at time t is given by

$$E_{at} = \sum_{X \in \{X_t\}} \omega_t(X)\left(1 - F_{\eta^*}(\beta_{Xt})\right)$$

where $F()$ is the cumulative distribution function of η_i^*. Here we consider observable characteristics, X_t, to form a discrete set. The proportion of the population taking each characteristic is given by $\omega_t(X)$.

We can relate the change in the employment share to the change in sectoral prices as follows:

$$dE_a = -\sum_X \omega(X) f_{\eta^*}(\beta_X) \frac{d\beta_X}{d\theta_a} d\theta_a$$

where $f_{\eta^*}()$ is the probability density function of η^*. We can therefore further express the change in employment share between time t and t' as follows:

$$\frac{\Delta E_{at'}}{\Delta \theta_{at'}} \approx -\sum_X \omega(X) f_{\eta^*}\left(\tilde{\beta}_X\right) \frac{d\tilde{\beta}_X}{d\theta_a} \tag{3}$$

for any $\tilde{\beta}_X \in (\beta_{Xt}, \beta_{Xt'})$. Expression 3 is the moment condition we take to the data.

In expression 3, we observe $\Delta\theta_a$ as the relative price change estimated earlier, which doesn't depend on any features of the distribution. Moving to the right hand side, first notice that $\omega(X)$ can be taken straight from data. This leaves $f_{\eta^*}\left(\tilde{\beta}_X\right)$ and $\frac{d\tilde{\beta}_X}{d\theta_a}$.

In fact, for all β, given the standardization, we have

$$\frac{d\beta}{d\theta_a} = -\frac{1}{\sqrt{\sigma_a^2 + \sigma_o^2 - 2\rho\sigma_a\sigma_o}}$$

where σ_j^2 is the variance of returns in sector j and ρ is the correlation coefficient on sectoral returns. These are the parameters to be estimated. Finally, to take the moment condition to the data, we need to compute the probability density function, $f_{\eta^*}(\beta_X)$, which is unobserved. However, the cdf $F_{\eta^*}(\beta_X)$ is observed simply as the proportion of individuals working in the non-abstract sector, by observable characteristics. Therefore given information on $F_{\eta^*}()$, or any parametric choice, we can extract the value of the unobserved density at the selection threshold via $\beta_X = F_{\eta^*}^{-1}(E_{X,o})$. Recall that η^* has mean zero and unit variance by construction.

Before moving on, a few features of expression 3 should be pointed out. First, notice that the condition depends explicitly neither on non-pecuniary returns in each sector, nor on mean skills, nor on the level of prices. These factors are captured in β.

Second, it is worth considering the relationship between the parameters and the size of the theoretical moment. In fact, conditional on all other parameters and on the selection threshold, the relationship between the correlation coefficient, ρ, and the elasticity of employment growth is monotonically increasing: a higher correlation coefficient implies a higher elasticity. This monotonicity implies that identification of the correlation using this moment condition is clean. The relationship between the variances and the theoretical moment is a little more complex. It depends on some key relationships. For example, as long as $\sigma_j > \rho\sigma_k$ for $k \neq j$ then an increase in σ_j causes the employment elasticity to decrease. If we consider ρ in terms of these variances, then the relationship is clear: if, for example $\rho < \min\left(\frac{\sigma_a}{\sigma_o}, \frac{\sigma_o}{\sigma_a}\right)$ then the employment elasticity is decreasing in both the variances. In practice, this is the relevant range for ρ.

Finally, it is worth noting that the current framework abstracts from frictions in the labour market, such as occupational switching costs. As shown by Kambourov and Manovskii (2009), Gathmann and Schoenberg (2010) and Cortes and Gallipoli (2017), among others, these costs are large. Of course, we rationalize the framework used here as being more suitable for understanding movements over the long-run. It may be, however, that switching costs are sufficiently large that they do not 'wash out' over the standard length of a career. Unfortunately, allowing for labour market frictions is beyond the scope of the present analysis.[30]

Variance of Wages for Switchers

Here we show that the dispersion in wages for these switchers can similarly be expressed in terms of our key parameters only.

To proceed, as standard we write $\gamma_{ij}^* = \rho_{j\eta}\eta_i^* + \xi_{ij}$ where γ_{ij}^* is the standardized version of γ_{ij} for $j \in \{a, o\}$, $\rho_{j\eta}$ is the correlation of γ_{ij} and η_i, and ξ_{ij} is a mean-zero random variable that is uncorrelated with η_i. Henceforth, we assume that ξ_{ij} is homoskedastic conditional on η_i^*. This assumption is guaranteed, for example, by the Gaussian distribution. Then we consider those individuals who switch from non-abstract to abstract, given a marginal change in prices. We obtain:

$$\begin{aligned} Var_t\left(\gamma_{ia} \middle| \text{ switch to abstract}, X\right) &= \sigma_a^2\left[\rho_{a\eta}\left[Var\left(\eta_i^* \middle| \beta_{Xt-1} > \eta_i^* > \beta_{Xt}\right)\right] + Var\left(\xi_{ia}\right)\right] \\ &\approx \sigma_a^2\left[Var\left(\xi_{ia}\right)\right] \end{aligned} \tag{4}$$

where $Var(\xi_{ia}) = 1 - \rho_{a\eta}^2$ and $\rho_{a\eta}$ depends purely on σ_a, σ_o and ρ.[31] We can observe the left hand side of equation 4 as the dispersion in returns for year-to-year switchers. Equation 4 therefore gives an additional moment restriction in terms of key parameters. Notice that it does not depend on particular parametric assumptions, other than the homoskedasticity of ξ_{ia}.

Again, before moving on, a couple of features of expression 4 should be pointed out. First, we should consider the empirical relevance of equation 4. In particular, our model implies that the gross flow of workers across sectors is very small, while a large literature shows that a larger volume of workers switch sectors in any year, and in all directions.[32] We can defend our framework, however, by arguing that those who switch sector must be approximately on the margin. This is true even in a model with, say, sizeable search frictions. Accordingly, the moment condition in equation 4 captures distributional features of those on this margin, who can come from a diverse pool. To explore this point further, in appendix C we consider an extension of the model with shocks. These shocks, which can be interpreted either as wage, or as preference shocks, can generate a realistic volume of sectoral switching. The appendix shows that the moment condition holds approximately, even in this richer framework.

[30] Although we have to abstract from frictions in this section, we can and do include them when estimating the prices in section 4. See also appendix A.

[31] $\rho_{a\eta}$ is the correlation between γ_a and $\gamma_a - \gamma_o$. Therefore $\rho_{a\eta} = \frac{\sigma_a - \rho\sigma_o}{\sqrt{(\sigma_a^2 + \sigma_o^2 - 2\rho\sigma_a\sigma_o)}}$

[32] See, for example, Carrillo-Tudela et al. (2016) for the UK.

Second, we can test the theoretical approach by examining related moments for switchers. For example consider:

$$Var_t\left(\gamma_{io}\right| \text{switch to abstract})$$

which is the variance of returns in the source sector, non-abstract, for those who switch to abstract. This moment and that given in equation 4 are theoretically identical. Similarly, we observe in the data $Var_t\left(\gamma_{io}\right| \text{switch to non-abstract})$ and $Var_t\left(\gamma_{ia}\right| \text{switch to non-abstract})$. Therefore, we can use the 4 moments as a useful test of the theoretical approach, by comparing their size. Notice as well, that as the moments observed in the data are theoretically equivalent, they do not provide any extra information for identification.

Variance of Observed Sectoral Wages

For identification, we need at least one additional moment condition. We consider the variance of latent returns, conditional on working in each sector at time t. The moment for abstract workers is given by:

$$Var_t\left(\gamma_{ia}\right| \text{work in abstract}, X) = \sigma_a^2\left[\rho_{a\eta}\left[Var\left(\eta_i^*|\eta_i^* > \beta_{Xt}\right)\right] + Var\left(\xi_{ia}\right)\right] \tag{5}$$

The expression for non-abstract is similar. The moments on the left hand side of this equation are estimable, while, the expressions on the right hand side are similar to those in equation 4. The right hand side here differs by containing the extra term $Var\left(\eta_i^*|\eta_i^* > \beta_{Xt}\right)$. This conditional variance depends on the choice of the functional form for the cdf of η_i^*. Therefore, the moment condition in expression 5 is more dependent on functional form assumptions than the variance for switchers. This may be problematic because the conditional variance can either tend to zero or increase without bound as β tends to infinity, depending on the shape of the distribution of η_i^* in the tails. However, we can check the size of the moments empirically. In fact, as we see later, the empirical moments imply the conditional variance is small. As a final point, notice that this moment condition theoretically varies over time. However, in practice, the time-series variation is small; so, in the empirical implementation, we average it over the sample period.

Discussion

To aid intuition, we illustrate the identification strategy in figure 8, which shows plots in four quadrants, according to the classic working of the static Roy model. Each quadrant shows the individual-specific returns γ_{ia} and γ_{io} in the axes. The two upper plots show the assignment choices into each sector; first in a static setting in the left plot, and second (right plot), when the common price on abstract tasks increases. The two lower plots show the contours from gaussian distributions of latent returns. Here we emphasize how the theoretical moments are affected by the correlation coefficient. Accordingly, the left plot presents a positive correlation of returns, while the correlation in the right plot is negative. The distribution of those on the margin of switching is shown in dark red. In general, the distribution of the switchers depends on the underlying parameters in a non-linear way. However, the figure illustrates the identification argument in our particular case clearly. Holding fixed the mass of workers in each task, and the dispersion of underlying productivities, for example, there is a clear relationship between the correlation of returns, ρ, and our first two moment conditions. When latent returns are negatively correlated (right plot), switchers generally come from a narrow pool of talents, and the dispersion in their wages is low. When latent returns are

Figure 8: Identification of Covariance Structure with 2 Sectors

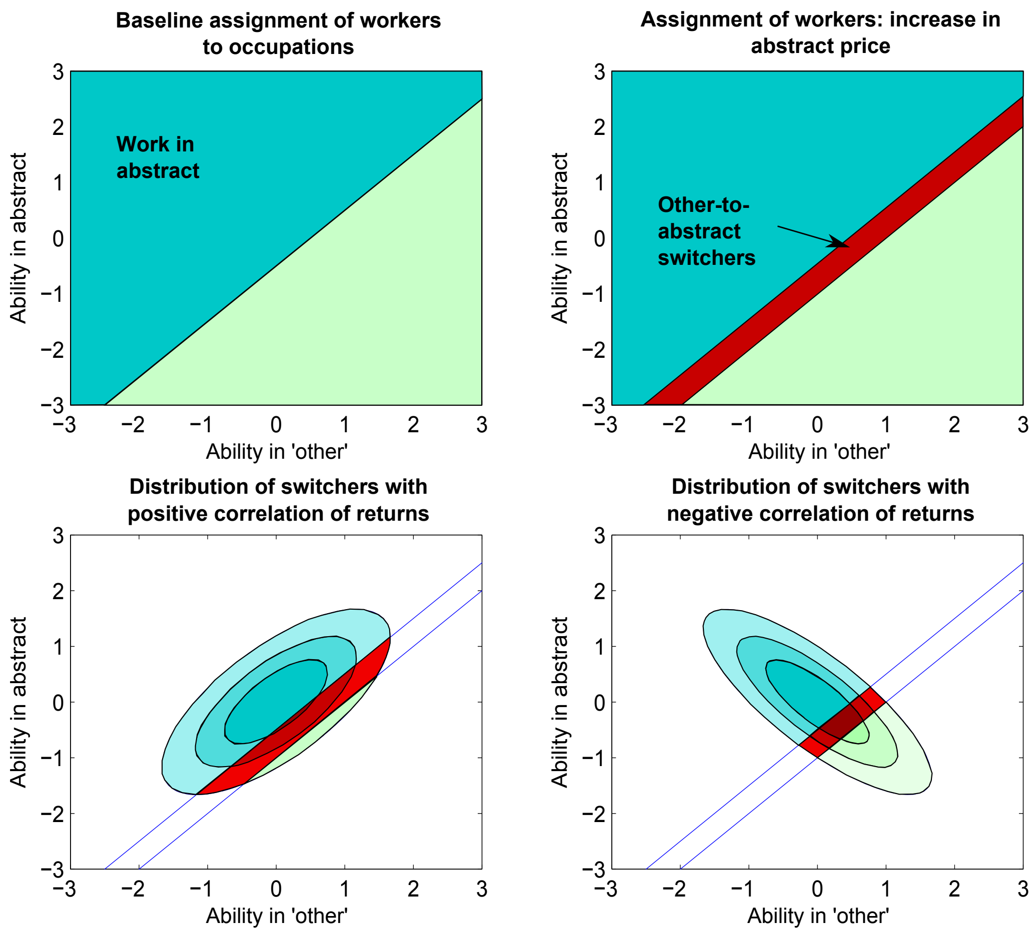

positively correlated (left plot), switchers generally come from a wider pool of talents and the dispersion in their wages is higher. Furthermore, a positive correlation of returns implies that the volume of switchers is larger, and so the employment elasticity is higher.

To summarize, we approach the data attempting to estimate 3 parameters with 7 moments, 4 of which are theoretically identical but which can be used to test the framework. Importantly, and to emphasize, none of these moment conditions depends on the location of the underlying returns, or the preference parameter ε_a, given information on the observed employment shares. We have therefore been able to identify the parameters of interest without identifying the full distribution. Moreover, we have been explicit about which functional form assumptions are needed and where.

5.2 Empirical Implementation

We estimate the framework on both the UK and Germany separately. To implement it, we assume that $\begin{pmatrix} \gamma_a \\ \gamma_o \end{pmatrix}$ follows a bivariate gaussian. This assumption ensures homoskedasticity of ξ_a and ξ_o and provides a functional form for η^*. To perform the empirical analysis we use the same sample as in section 4. We also have to take some additional steps with the data, which we now describe.

First we estimate the growth in employment across tasks holding fixed the skill composition of the work force. We do this because we are effectively tracing out the supply curve and so we wish to hold fixed supply factors other than the change in price. The main issue here, as discussed previously, is that educational attainment has increased for more recent cohorts, and high educational attainment likely increases the return to working in abstract jobs. Presumably, this change in attainment partly reflects an endogenous response to the increase in demand for abstract labour. However, as argued by Salvatori (2015) for the UK, for example, it also likely reflects more exogenous shifts in supply caused by the relaxation of rationing of university places. To address these issues we estimate the change in abstract employment over time holding fixed cohort factors by regressing sector on cohort and time dummies. We also remove those aged under 35. We do this because the share of workers in abstract increases noticeably in early career due to pure life-cycle factors.

Second, we run a first-stage OLS regression of wages on characteristics that are fixed at the individual-sector spell level: these include region and the interaction of education with task. We do this because these characteristics are differenced out in the fixed effect estimation. If we were to omit this first stage, then the return to education, for example, would be absorbed in γ_{ij} and generate a spurious positive correlation between returns across the different tasks. In practice this first stage regression has little effect on any of the estimates.

The full estimation procedure therefore is as follows. We compute the estimated residuals, $\hat{\gamma}_{ij}$, by running OLS on stable characteristics, then fixed effects on time-varying characteristics. We also compute aggregate employment share growth after stripping out age and cohort effects. We use the results of these steps to compute the moments indicated by equations 3, 4 and 5. We compute standard errors for all these moments using 1000 iterations of the bootstrap taking into account multiple steps in the procedure. It is worth emphasizing that here we use the corrected measure of occupation of work, described in section 3 and in appendix B, which drastically reduces the number of observed switches. Notice, however, that the results are very similar if we use the uncorrected measure in the raw data. For the estimation we use 1991 to 2008 for the UK. For Germany, we use a longer period, 1991 to 2013, because the shift towards abstract continued after 2008, and a longer sample period should provide more precise estimates.

We then estimate the covariance matrices using classical minimum distance. To weight the moments we take the diagonal of the optimal weighting matrix. Because the switching moments all provide theoretically the same information, we average these by dividing their elements in the weighting matrix by 4.

Table 2 shows the parameter estimates for both countries. Alongside the estimates we present 95% confidence intervals, obtained by bootstrap. These confidence intervals are not necessarily symmetric because of skewness in the distribution of the raw moments. The first two rows show the variance of returns in each main task. In both countries we find a wider dispersion in latent returns in the abstract task versus all the other tasks combined. The variance of returns in the non-abstract task is somewhat higher in the UK than in Germany. The bottom row shows the correlation across tasks. The point estimate here implies a moderate-to-strong correlation in both countries. The estimate is higher in the UK than in Germany, although the confidence intervals overlap sizeably. Notice that the confidence interval just about excludes 0, even in Germany.

To understand these estimates we need to examine the raw moments and their model fit. These are presented in table 3. We first examine the raw moments, which are presented alongside standard errors. The first row shows the employment share gradient in the abstract task. In the UK we estimate that each percentage increase in the price on abstract jobs increases employment by 0.3 percentage points. This corresponds to the estimated price increase

Table 2: Covariance Estimates: 2 Occupational Groups

		UK		Germany	
Variance returns in abstract	σ_a^2	0.16	[0.15,0.18]	0.14	[0.12,0.16]
Variance returns in non-abstract	σ_o^2	0.12	[0.11,0.13]	0.07	[0.06,0.08]
Correlation of returns	ρ	0.50	[0.37,0.65]	0.26	[0.00,0.54]

Notes: 95% confidence interval in square brackets, obtained from 1000 bootstrap repetitions.

of around 13% over the period and an increase in the share employed in abstract work of around 4 percentage points. In Germany the increase in the abstract employment share was larger, at around 10 percentage points and the employment share gradient is consequently higher.

Table 3: Empirical Moments and Fit: 2 Occupational Groups

		UK			Germany		
Description	Label	Data		Fit	Data		Fit
Grad. of emp. Change	$\Delta E_a/\Delta\theta_a$	0.334	(0.240)	0.816	0.569	(0.152)	0.729
Var of observed skills	$Var(\gamma_a\|\text{abstr})$	0.124	(0.006)	0.127	0.085	(0.011)	0.091
	$Var(\gamma_o\|\text{other})$	0.104	(0.004)	0.106	0.060	(0.003)	0.061
Var of skills for switchers	$Var(\gamma_a\|o\to a)$	0.108	(0.008)	0.100	0.115	(0.018)	0.057
	$Var(\gamma_a\|a\to o)$	0.119	(0.011)		0.128	(0.027)	
	$Var(\gamma_o\|o\to a)$	0.118	(0.009)		0.109	(0.015)	
	$Var(\gamma_o\|a\to o)$	0.136	(0.012)		0.107	(0.018)	

Notes: Standard errors in parentheses.

The second and third rows show the variances observed in each task. The variances are higher in the UK than in Germany, although the relative size of the variance across tasks is similar in both countries. The final four rows show the variances for each of the four types of switchers, which we average in the estimation. Each of these moments is empirically very similar, providing support for our underlying approach. However, the variances for the switchers are higher than for those observed in each task, shown just above. This relationship contradicts the theoretical moments given in expressions 4 and 5 which imply that the cross-sectional variances must be higher. However, this relationship also implies that the conditional variance in equation 5 presumably tends to zero in the tails of the underlying distribution. Finally, notice that all the variances are estimated much more precisely than the employment gradient.

We turn next to the fit. First notice the results from the second and third rows. These show that the model matches the variances of productivities in each sector almost exactly. However, the fit to the other two sets of moments is less exact. In particular, the fitted model implies a higher employment growth than is seen in the data. Symmetrically, the fitted model implies a lower dispersion in wages for switchers than we see in the data, especially in Germany. One way to interpret these results is as follows: the high variance of wages for switchers implies a high correlation of returns across broad task groups. This high correlation then implies that we should observe more switching across sectors than was seen in reality. Despite this, it should be emphasized that the framework overall does a good job of matching the features of the data we examine. Moreover, the estimates shown in table 2 not only broadly fit the moments used here, they are also consistent with strong attenuation in observed wages over time, shown in section 4.

To understand our findings further it is helpful to re-interpret them using a transformed model. We can rewrite the wage equation 1 in section 2 in terms of orthogonal components:

$$w_{ijt} = \delta_{jt}X_{it} + \theta_{jt} + \zeta_i + \xi_{ij} + \nu_{it}$$

such that γ_{ij} is now captured by ζ_i, ξ_{ij} and ξ_{ik}, which are all orthogonal to each other for $k \neq j$. Here we can also normalize the mean of each of these components to zero without loss of generality. These orthogonal components are in some ways simpler to interpret: ζ_i picks up a 'universal' measure of unobserved productivity, while ξ_{ia}, for example, captures a pure abstract component. The variances of these components relate directly to the parameter estimates shown in table 2: a higher correlation coefficient ρ implies a higher variance, and greater importance, of the universal productivity component. The pure task-specific components are therefore residuals from their raw counterparts.

Table 4: Covariance Estimates of Orthogonal Components: 2 Occupational Groups

		UK		Germany	
Variance returns in abstract component	$\sigma^2_{\xi,a}$	0.09	(0.01)	0.11	(0.01)
Variance returns in non-abstract component	$\sigma^2_{\xi,o}$	0.05	(0.02)	0.04	(0.02)
Variance universal returns	σ^2_{ζ}	0.07	(0.01)	0.03	(0.02)

Notes: Standard errors in parentheses.

Estimates of the variances of these orthogonal components are given in table 4. They show that, in both countries, the pure abstract component is much more variable across the population than either of the other two components, and particularly the non-abstract component. These results indicate, therefore, that sorting into occupational groups is based much more on abstract productivity than on any other characteristic. They also indicate that a greater demand for abstract labour is generally associated with greater inequality, because these skills are much more variable across the population.

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 have 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 and 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.

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 also use the panel data structure to estimate further important features of the Roy model of choice across sectors. In particular, we estimate the covariance matrix of latent productivities/returns across sectors and individuals. Estimating these features is important when considering issues of welfare, and specifically considering the effects of on-going occupational change. This feature has been of interest since Roy (1951) itself. We find evidence in both countries that heterogeneity in pure abstract productivities is the largest component of persistent worker differences. This implies that productivity in abstract tasks is the most important factor in sorting across sectors, and that a greater demand for abstract labour is associated generally with higher inequality, holding fixed other factors.

This paper prompts several threads of future research. First, it would be instructive to extend the structural analysis to a finer level of occupational aggregation. In a related way, it would be useful to consider the multiplicity of tasks required in each job, a la Autor and Handel (2013). Finally, it is noteworthy that we find that growth in employment in the abstract sector was observed to be lower than the model suggests. This may be because of stickiness in occupational choice that persists over the long term. Taking this into account requires an explicitly dynamic model of occupational choice. We leave this future research.

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.

AUTOR, D. AND D. DORN (2009): “This Job is ’Getting Old’: Measuring Changes in Job Opportunities using Occupational Age Structure,” *American Economic Review*, 99, 45–51.

AUTOR, D. AND M. HANDEL (2013): “Putting tasks to the test: Human capital, job tasks, and wages,” *Journal of Labor Economics*, 31, S59–S96.

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 (2016): “The UK wage premium puzzle: how did a large increase in university graduates leave the education premium unchanged?” *IFS Working Paper*, W16/01.

BÖHM, M. (2015): “The price of polarization: Estimating task prices under routine-biased technical change,” *mimeo*.

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.

CARRILLO-TUDELA, C., B. HOBIJN, P.-W. SHE, AND L. VISSCHERS (2016): “The extent and cyclicality of career changes: evidence for the UK,” *European Economic Review*, 84, 18–41.

COMBES, P., G. DURANTON, AND L. GOBILLON (2008): “Spatial wage disparities: Sorting matters!” *Journal of Urban Economics*, 63, 723–742.

CORTES, G. AND G. GALLIPOLI (2017): “The costs of occupational mobility: An aggregate analysis,” Forthcoming *Journal of the European Economic Association*.

CORTES, G., N. JAIMOVICH, C. NEKARDA, AND H. SIU (2014): "The micro and macro of disappearing routine jobs: A flows approach," *NBER working paper*, wp20307.

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.

FRENCH, E. AND C. TABER (2011): "Identification of Models of the Labor Market," in *Handbook of Labor Economics*, vol. 4, 537–617.

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.

HECKMAN, J. AND B. HONORE (1990): "The empirical content of the Roy model," *Econometrica: Journal of the Econometric Society*, 58, 1121–1149.

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.

LONGHI, S. AND M. BRYNIN (2009): “Occupational Change in Britain and Germany,” Tech. rep., Social Science Research Network, Rochester, NY.

MICHEALS, G., A. NATRAJ, AND J. VAN REENEN (2014): “Has ICT Polarized Skill Demand? Evidence from Eleven Countries over Twenty-Five Years,” *Review of Economics and Statistics*, 96, 60–77.

MUTHEN, B. (1990): “Moments of the censored and truncated bivariate normal distribution,” *British Journal of Mathematical and Statistical*, 43, 131–143.

ROY, A. (1951): “Some thoughts on the distribution of earnings,” *Oxford economic papers*, 3, 135–146.

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 cyclicality of real wages: how important is composition bias,” *Quarterly Journal of Economics*, 109, 1–25.

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.

A Appendix: Further Information and Results

A.1 Summary Statistics

Table A.1: Summary Statistics by task

BHPS			GSOEP		
	Mean	s.d.		Mean	s.d.
Abstract					
Age	40.77	8.75	Age	42.53	8.96
Married or cohabiting	0.84	0.37	Married or cohabiting	0.75	0.43
University degree	0.36	0.48	University degree	0.42	0.49
Monthly earnings (2010 GBP)	3,064.04	1,951.01	Monthly earnings (2010 GBP)	2,775.97	1,226.71
Nr of waves	12.74	4.12	Nr of waves	17.23	6.79
Routine Manual					
Age	40.72	9.37	Age	40.83	9.54
Married or cohabiting	0.85	0.36	Married or cohabiting	0.76	0.43
University degree	0.03	0.16	University degree	0.01	0.07
Monthly earnings (2010 GBP)	1,962.64	906.40	Monthly earnings (2010 GBP)	1,887.60	626.24
Nr of waves	12.46	4.08	Nr of waves	16.53	7.14
Routine Cognitive					
Age	39.60	9.80	Age	41.92	9.49
Married or cohabiting	0.75	0.43	Married or cohabiting	0.71	0.45
University degree	0.11	0.31	University degree	0.08	0.27
Monthly earnings (2010 GBP)	1,778.80	832.60	Monthly earnings (2010 GBP)	2,199.71	925.74
Nr of waves	12.29	4.08	Nr of waves	17.01	6.63
Non-routine Manual					
Age	40.49	9.38	Age	40.71	9.33
Married or cohabiting	0.80	0.40	Married or cohabiting	0.72	0.45
University degree	0.03	0.17	University degree	0.01	0.09
Monthly earnings (2010 GBP)	1,764.96	870.87	Monthly earnings (2010 GBP)	1,706.41	540.04
Nr of waves	12.00	4.08	Nr of waves	16.43	7.26
Tot observations	24,364		Tot observations	36,918	

A.2 Strict Exogeneity Test

Here we perform a specification test of the fixed effect framework described in section 2. Validity of the fixed effect model depends on a strict exogeneity condition. This condition requires that sectoral choice depends on task prices, observables and unobserved fixed skills only, and not the model residual v_{it}. We test this condition for the UK by examining the correlation between the estimated residuals and sectoral switches. The computed correlations are displayed in table A.2. It shows correlations between residual wages at time t and dummies for each of the twelve types of sectoral switch. The dummies take value 1 when a switch occurs between time t and time $t+1$, and zero otherwise. The table shows that for all types of switch, the observed correlation is both indistinguishable from zero, and small in absolute size. This corroborates the strict exogeneity assumption.

Notice that we have chosen to examine the residual wage in the exiting sector rather than in the entering sector.

We do this because wages in the first year in a new sector are likely to be suppressed owing to switching effects. Nevertheless, when we perform the same calculation as reported in table A.2 , but examining switches between time $t-1$ and time t, we find similar results. Correlations are insignificant for all switches.

Table A.2: Test of Strict Exogeneity: the UK

		Correlation of $\hat{v}_t$ with t-to-$t+1$ switch (%)	Standard error	95% C.I. on correlation
Switch up	RM → abstr.	-0.3%	1.2%	[-2.7,2.1]
	RC → abstr.	0.6%	1.6%	[-2.5,3.7]
	NRM → abstr.	-1.1%	1.2%	[-3.5,1.2]
	RC → RM	-2.9%	1.9%	[-6.5,0.8]
	NRM → RM	-3.4%	2.1%	[-7.4,0.7]
	NRM → RC	-1.5%	2.2%	[-5.7,2.7]
Switch down	abstr. → RM	0.0%	0.9%	[-1.7,1.7]
	abstr. → RC	-0.1%	0.7%	[-1.5,1.3]
	abstr. → NRM	0.1%	0.4%	[-0.6,0.8]
	RM → RC	-1.1%	1.0%	[-3.0,0.8]
	RM → NRM	-2.3%	1.2%	[-4.7,0.1]
	RC → NRM	-1.5%	1.3%	[-4.1,1.1]

Notes: RM stands for routine manual. RC stands for routine cognitive. NRM stands for non-routine manual

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 the coefficients of the specifications discussed in section 4 and presented in table 1.

Figures A.1 and B.6 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.

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). 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.

As the figures show, the results are consistent with those in the main analysis.

Table A.3: Changes in task price in UK: controls in the different 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)	-0.000	(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)
1991 Abstract ∗ Job Tenure							-0.000	(0.00)			0.002	(0.00)
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)
Non union									-0.073***	(0.01)		
1991 Abstract ∗ Non union									0.027	(0.02)		
1991 R Cog ∗ Non union									0.040	(0.03)		
1991 NR Man ∗ Non union									-0.001	(0.03)		
Non union ∗ Age									0.001	(0.00)		
Constant	2.169***	(0.03)	2.484***	(0.06)	2.491***	(0.06)	2.485***	(0.06)	2.518***	(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. The omitted category for education is mid secondary.

Table A.4: Changes in task price in Germany: controls in the different specifications

	OLS		Benchmark FE		Educ. Returns		Job Tenure		TU		35-60 y-o	
Secondary ed.	-0.044**	(0.02)	-0.012	(0.02)	-0.010	(0.03)	-0.011	(0.02)	0.026	(0.02)	-0.019	(0.02)
High secondary	0.092**	(0.04)	0.099***	(0.03)	0.049	(0.05)	0.101***	(0.03)	0.079**	(0.03)	0.009	(0.03)
High vocational	0.108***	(0.02)	0.073***	(0.02)	0.058**	(0.03)	0.072***	(0.02)	0.059**	(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)	0.143***	(0.04)	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)
1991 Abstract ∗ Job Tenure							-0.002	(0.00)			-0.003*	(0.00)
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)
Non union									-0.051***	(0.02)		
1991 Abstract ∗ Non union									0.061**	(0.03)		
1991 R Cog ∗ Non union									0.020	(0.04)		
1991 NR Man ∗ Non union									0.037	(0.03)		
Non union ∗ Age									0.002	(0.00)		
Constant	2.674***	(0.03)	2.767***	(0.07)	2.767***	(0.06)	2.764***	(0.06)	2.670***	(0.09)	2.790***	(0.08)
Observations	36918		36918		36918		36918		21791		27124	
Adjusted R^2	0.309		0.144		0.142		0.144		0.101		0.058	

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. The omitted category for education is mid secondary.

Figure A.1: Task prices for the UK

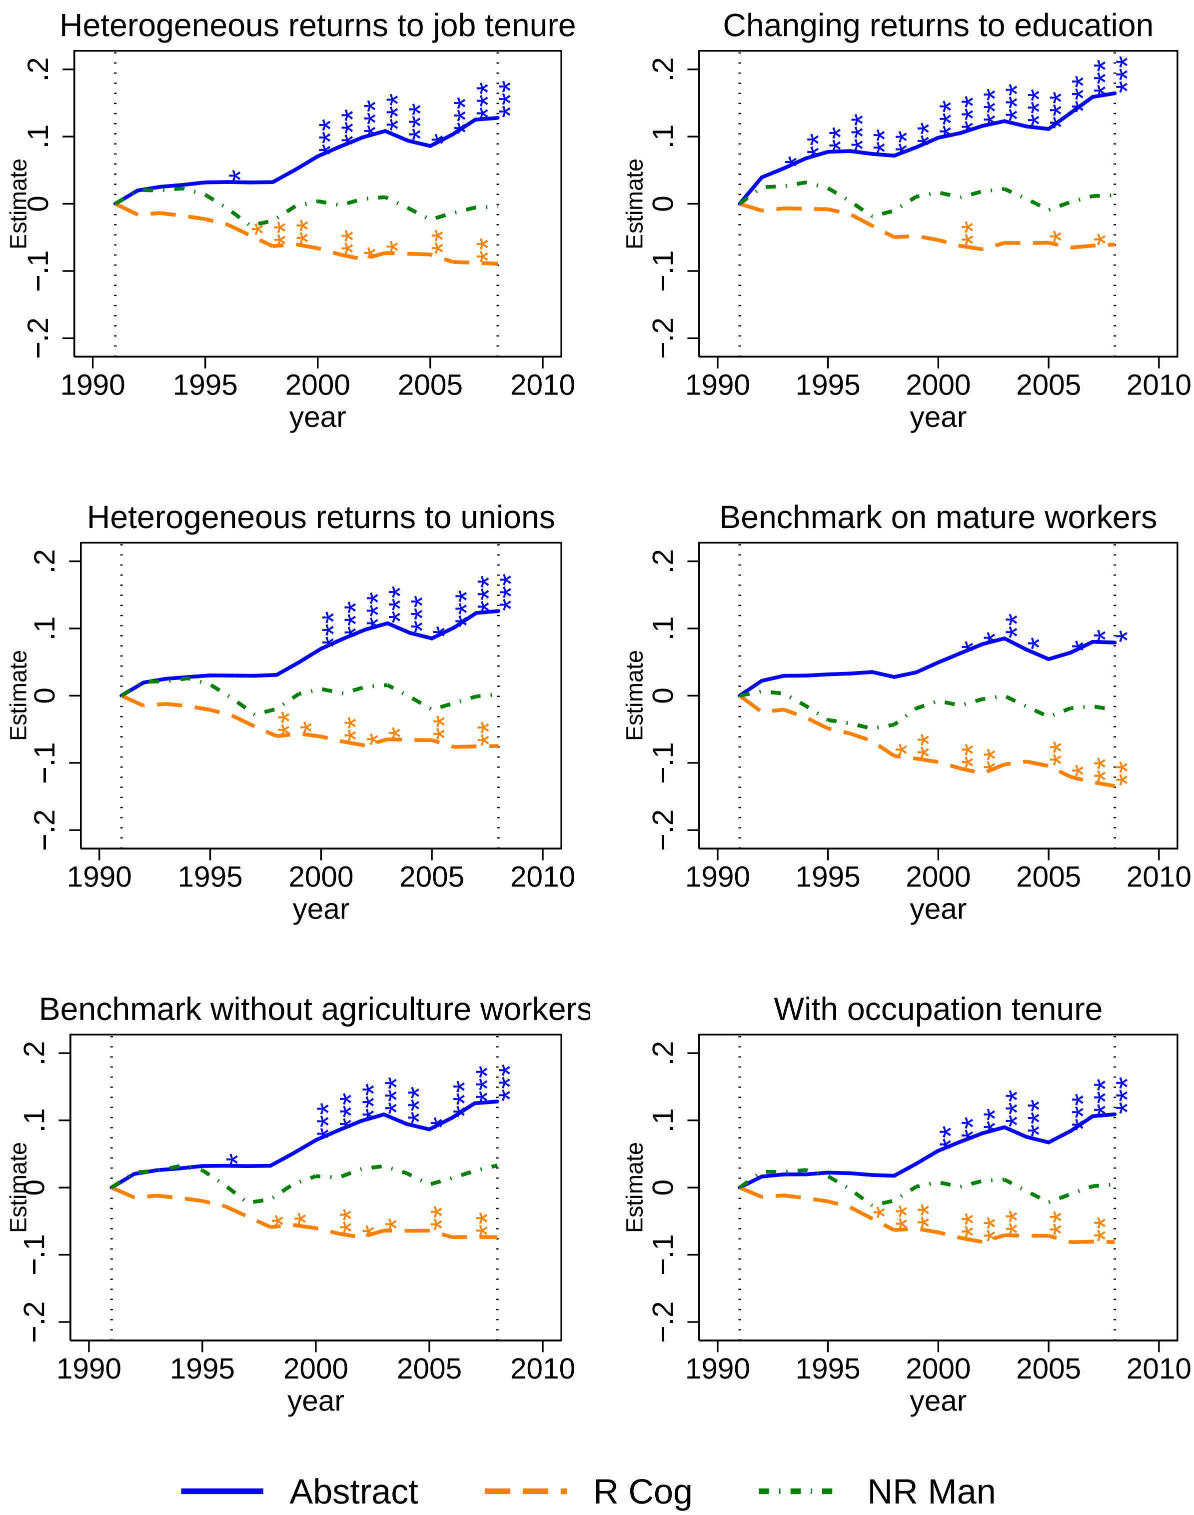

Picture is based on a sample of 25 to 60 y-o men.
'Abstract' stands for abstract task. 'NR Man' indicates non-routine manual, and 'R Cog' indicates routine cognitive.
The vertical dashed line indicates the common period in the two data sets, from 1991 to 2008.

Figure A.2: Task prices for Germany

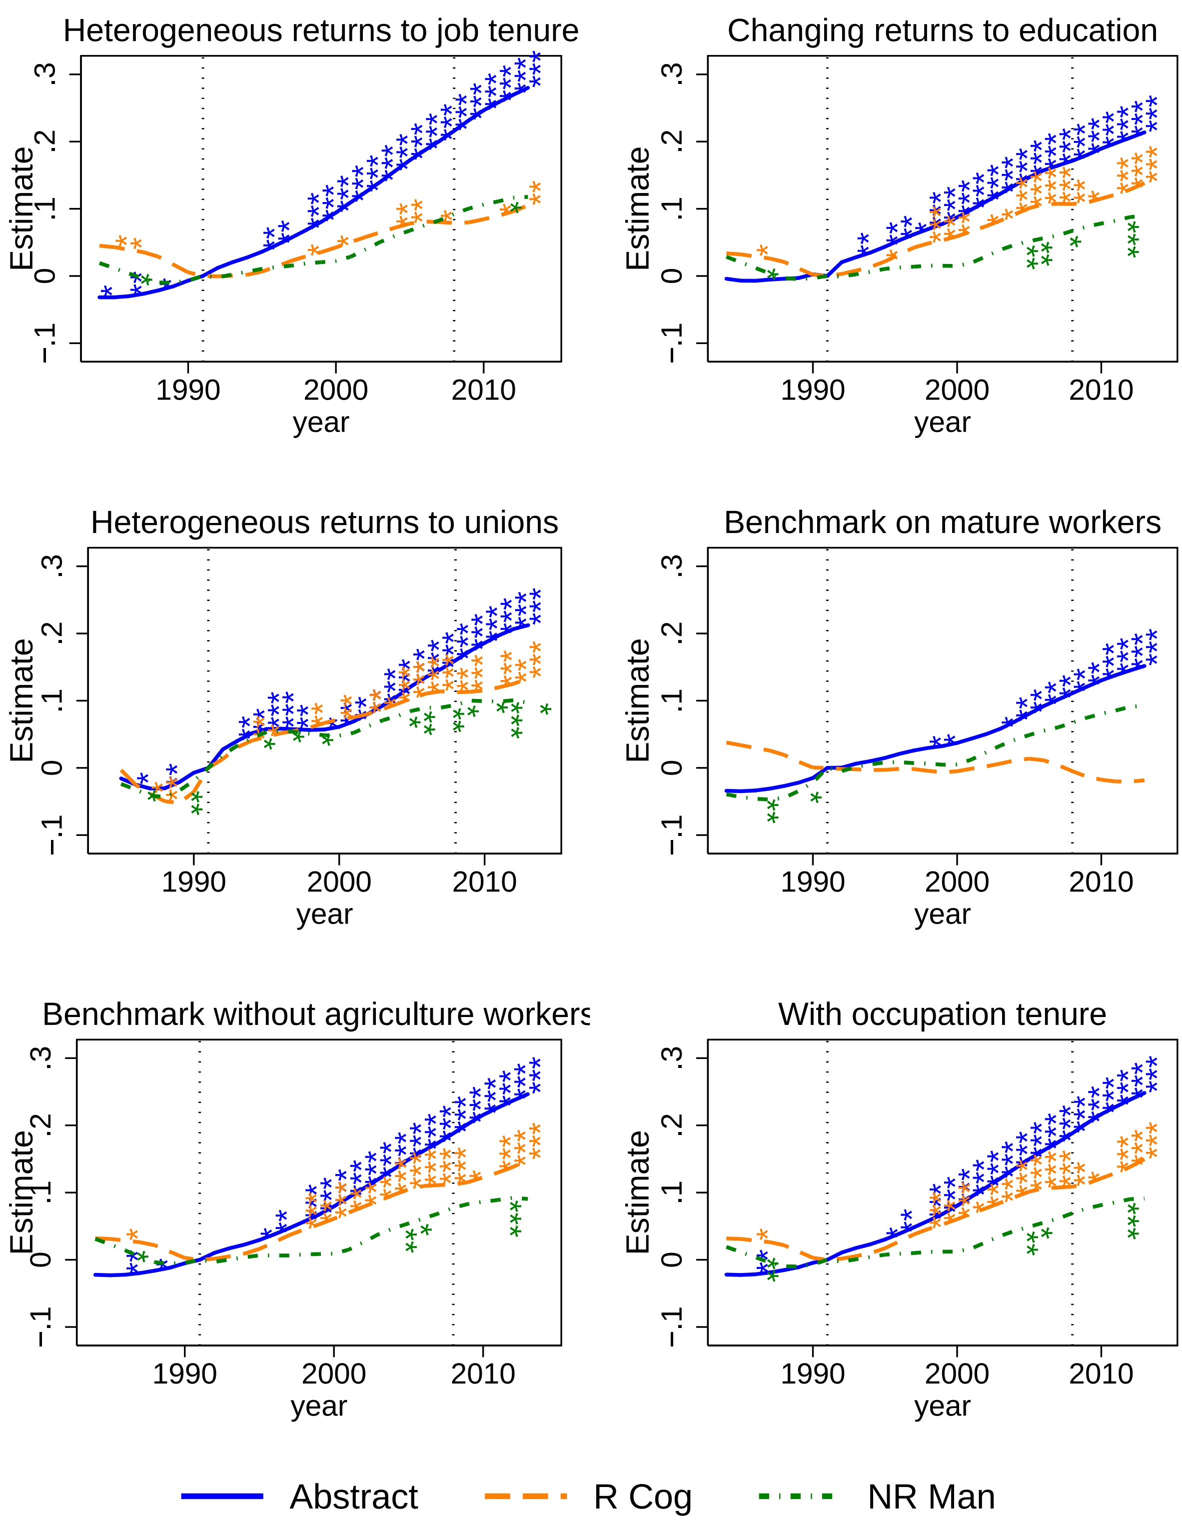

Picture is based on a sample of 25 to 60 y-o men.
'Abstract' stands for abstract task. 'NR Man' indicates non-routine manual, and 'R Cog' indicates routine cognitive.
The vertical dashed line indicates the common period in the two data sets, from 1991 to 2008.

B Appendix to Section 4: Robustness Exercises

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, which may differ systematically with the different type of interviewing technique. In Section 3 we explain how the interviews are conducted for BHPS and GSOEP and how that affects the measurement error. In this section we describe the method we use to address this issue.

With BHPS, we mainly base our correction on the variable indicating the year when a given job has started. 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."* We follow a series of steps. First, we correct for inconsistencies in this variable. For example, if the year is 1991 for the waves of 1995 and 1996, and the year is 1989 for the wave of 1997, we recode the latter as 1991. Second, we use the year when the respondent has changed job to correct the ISCO-88 code. For example, we check whether the respondent reported a change in his employment position with respect to the previous year and we compare the recorded occupation. We change the occupational code to that of the previous year if there are differences and the respondent has reported no changes. Finally, we correct for 1-year switches that were not affected by the previous steps. By 1-year switches we mean, for example, a worker with a given occupational code for at least a couple of years before and after 2001, and a different code in 2001. In that case, we recode the occupational code of the year 2001.

With GSOEP, we use a similar procedure. We combine a variable indicating the year since when the respondent is at the current employer with an indicator if anything in the respondent's job situation has changed since the previous interview. First, we check whether there are inconsistencies in the variable indicating employer's tenure. Then, we correct the ISCO-88 occupational codes according to the combination of these two variables, just as for the BHPS.

As mentioned in section 3, 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 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 independent interviews were introduced. In the GSOEP, 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: BHPS: Mobility rate computed on the original task measure and on the corrected measure

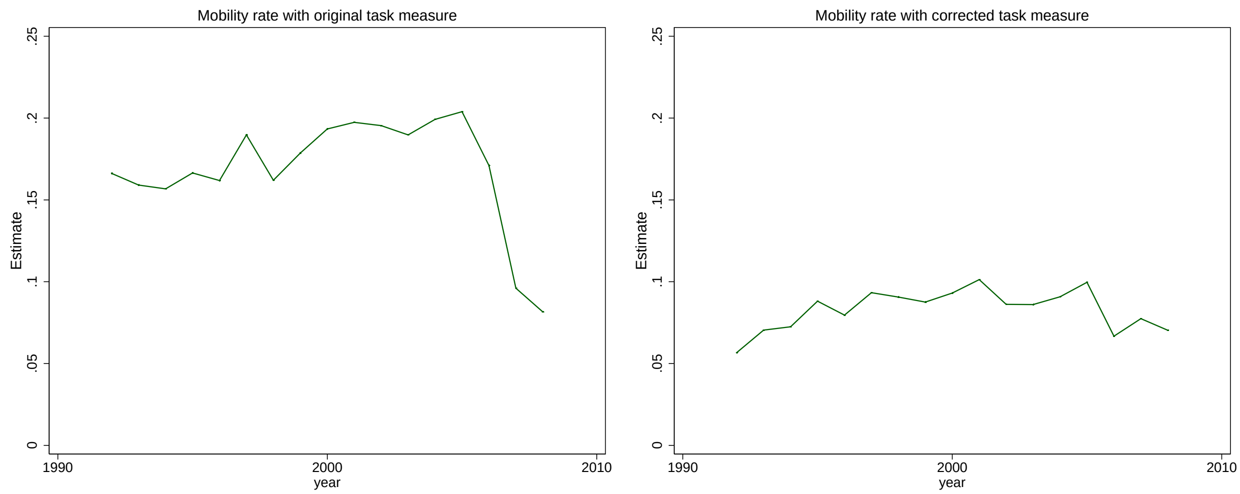

Notes: Mobility rate computed as the year-on-year switches between our four task groupings. Picture is based on a sample of 16 to 64 y-o men

Figure B.2: GSOEP: Mobility rate computed on the original task measure and on the corrected measure

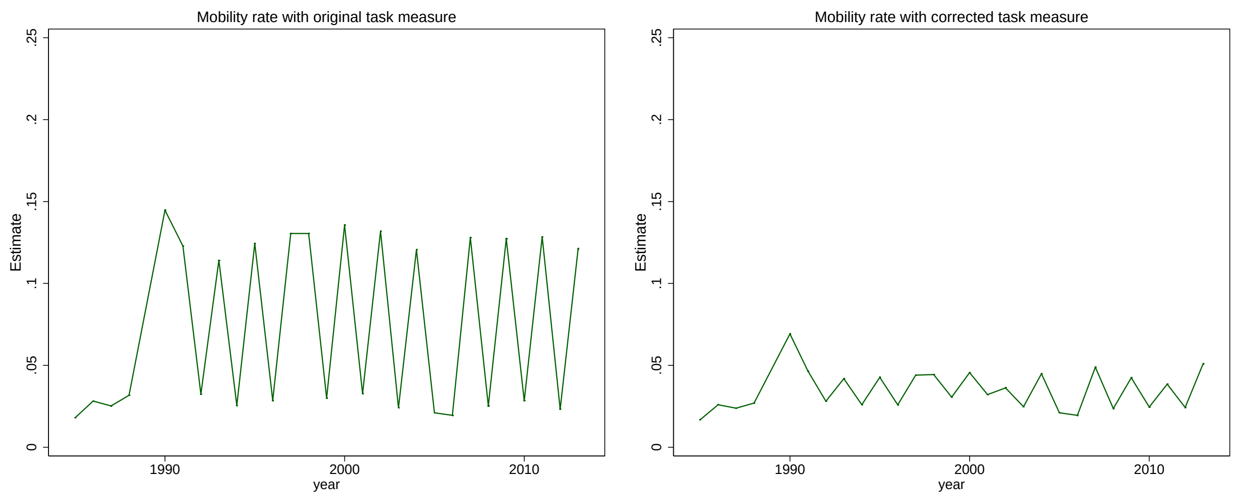

Notes: Mobility rate computed as the year-on-year switches between our four task groupings. Picture is based on a sample of 16 to 64 y-o men

Figure B.3: Task prices based on alternative catagorizations of tasks

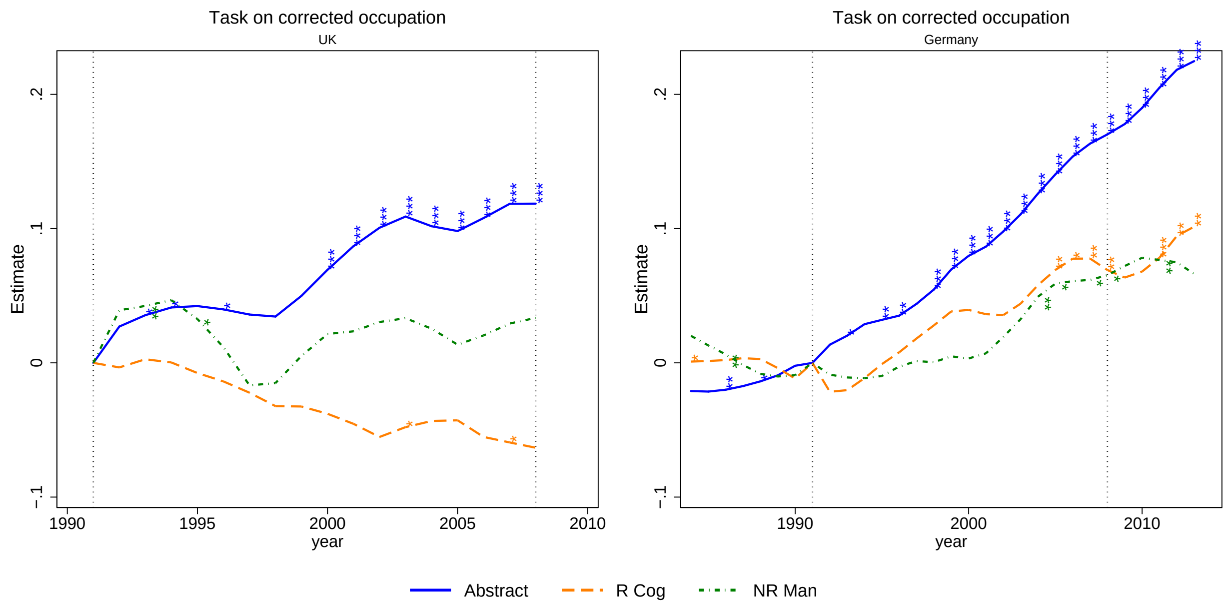

Picture is based on a sample of 25 to 60 y-o men.
'Abstract' stands for abstract task. 'NR Man' indicates non-routine manual, and 'R Cog' indicates routine cognitive.
The vertical dashed line indicates the common period in the two data sets, from 1991 to 2008.

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 GSOEP and BHPS dataset also contain classifications based on domestically-used classifications. The two classifications are explained below.

The classification in GSOEP was created by the German Federal Statistical Office in 1992 (KldB 92). The original structure was elaborated at the end of the 1960s. Since then, the Federal Statistical Office has updated it, in order to account for the technical and social development. 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 the sector. These are, in order : "Occupations in agriculture,veterinary, forestry and horticulture"; "Miners, mineral workers"; "Manufacturing occupations"; "Technical Jobs"; "Service workers"; "Other occupations". Service workers include sales and clerical occupations, and educational professionals. Teachers, university lecturers, psychologists are classified under the minor group "Other service occupations". All these groups list managerial occupations. This structure makes it difficult to derive a task categorization as 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.[33]

[33]For 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 GSOEP 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.

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 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.[34] 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.[35] We construct four broader groups by merging the nine occupational categories of the 1990 SOC, like we did with ISCO-88.[36] 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 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 figure 5. 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.

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, 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.

[34] This 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.

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

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

Figure B.4: Employment patterns by SOC90 classification

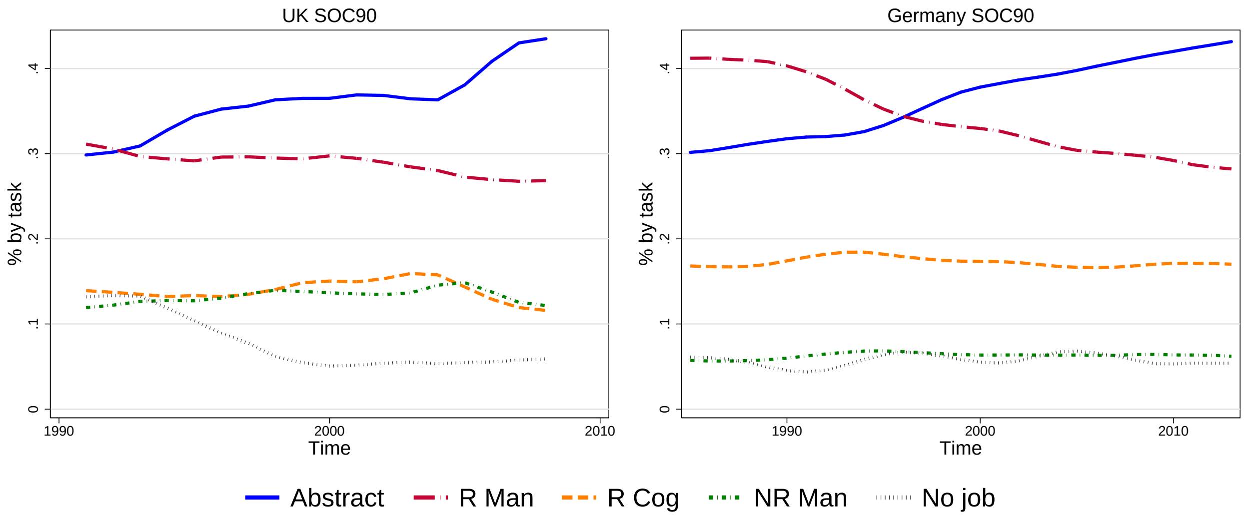

Picture is based on a sample of 16 to 64 y-o men

Figure B.5: Task prices based on alternative catagorizations of tasks

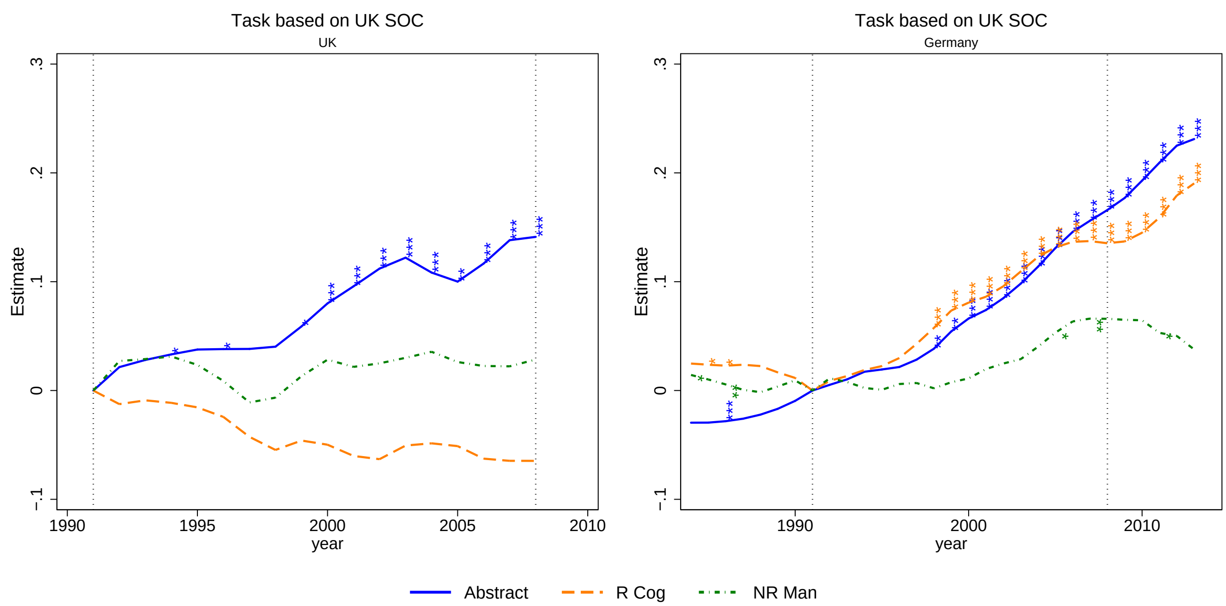

Picture is based on a sample of 25 to 60 y-o men.
'Abstract' stands for abstract task. 'NR Man' indicates non-routine manual, and 'R Cog' indicates routine cognitive.
The vertical dashed line indicates the common period in the two data sets, from 1991 to 2008.

Notice that for this exercise we only use the subsample of workers with standard ISCO-88 classifications.[37]

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). This work is not original. Our contribution, however, consists in unifying the process for different scales (DOT and O*Net) and different occupational classification (US Census Occupation Classification and ISCO).

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 Classifcation Occupations into an harmonised version of the 1990 Census Occupation. This harmonised classification, developed in Dorn (2009), is very useful for analyses of occupational patterns, as it harmonizes the occupation classifications of the 1950, 1970, 1980, 1990, and 2000 Censuses and of the 2005 American Community Survey. It can be particularly helpful, for example, on datasets that cover a long period, such as PSID, where the occupational classification changes over time (e.g. PSID is based on the 1970 and the 2000 Census Occupation Classifications).

Although we do not use it in this analysis, for completeness, we report the code to match the harmonized 1990 classification with the 1970 and 2000 standard COC, based on the crosswalk used for Autor and Dorn (2013).

For the European datasets we exploit the fact that most contain the ISCO classification. For both BHPS and GSOEP, we use ISCO88, which is available for all the waves. We match the ISCO classification with the harmonised 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. Specifically, for the DOT measures we rely on the 1990 Census classification, as modified in Dorn (2009). This classification is very useful for analyses of occupational patterns, as it harmonizes the occupation classifications of the 1950, 1970, 1980, 1990, and 2000 Censuses and the 2005 American Community Survey. 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

[37] Notice that both in BHPS and in GSOEP 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.

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. For 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.[38] 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) and define the content of abstract tasks with the average score from the variables GED-MATH (about quantitative reasoning requirements) and DCP (direction, control, and planning of activities). We identify routine tasks with the score on finger dexterity (FINGDEX) and an indicator of the adaptability to work that requires set limits or standards (STS). Finally, we use a variable that detects the importance of eye, hand and foot coordination (EYEHAND) for manual tasks.

We select the same O*Net variables of Acemoglu and Autor (2011) to reproduce a classification based on three main groups. Specifically, seven indicators of analytical and interpersonal skills are used to define abstract activities. We use six indicators of routine activities, among which variables that measure the importance of repetition, the role of structured versus unstructured work and the role of controlling machines or processes. Finally, 5 variables measure non routine manual activities.[39]

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. This is particularly apparent when looking at the graphs based on the US DOT and O*Net. The omitted category here is the group of occupations including routine manual and routine cognitive activities. Figure 5 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.

[38]Notice that excluding these categories does not affect the results as the number of occupational codes with missing values is low.

[39]In details they are the following. Abstract: Analytical 4.A.2.a.4 Analyzing data/information 4.A.2.b.2 Thinking creatively 4.A.4.a.1 Interpreting information for others 4.A.4.a.4 Establishing and maintaining personal relationships 4.A.4.b.4 Guiding, directing and motivating subordinates 4.A.4.b.5 Coaching/developing others; Routine 4.C.3.b.7 Importance of repeating the same tasks 4.C.3.b.4 Importance of being exact or accurate 4.C.3.b.8 Structured v. Unstructured work (reverse) 4.C.3.d.3 Pace determined by speed of equipment 4.A.3.a.3 Controlling machines and processes 4.C.2.d.1.i Spend time making repetitive motions. 4.A.3.a.4 Operating vehicles, mechanized devices, or equipment Manual 4.A.3.a.4 Operating vehicles, mechanized devices, or equipment 4.C.2.d.1.g Spend time using hands to handle, control or feel objects, tools or controls 1.A.2.a.2 Manual dexterity 1.A.1.f.1 Spatial orientation 4.A.4.a.4 Establishing and maintaining personal relationships.

Figure B.6: Task prices based on US occupation classifications

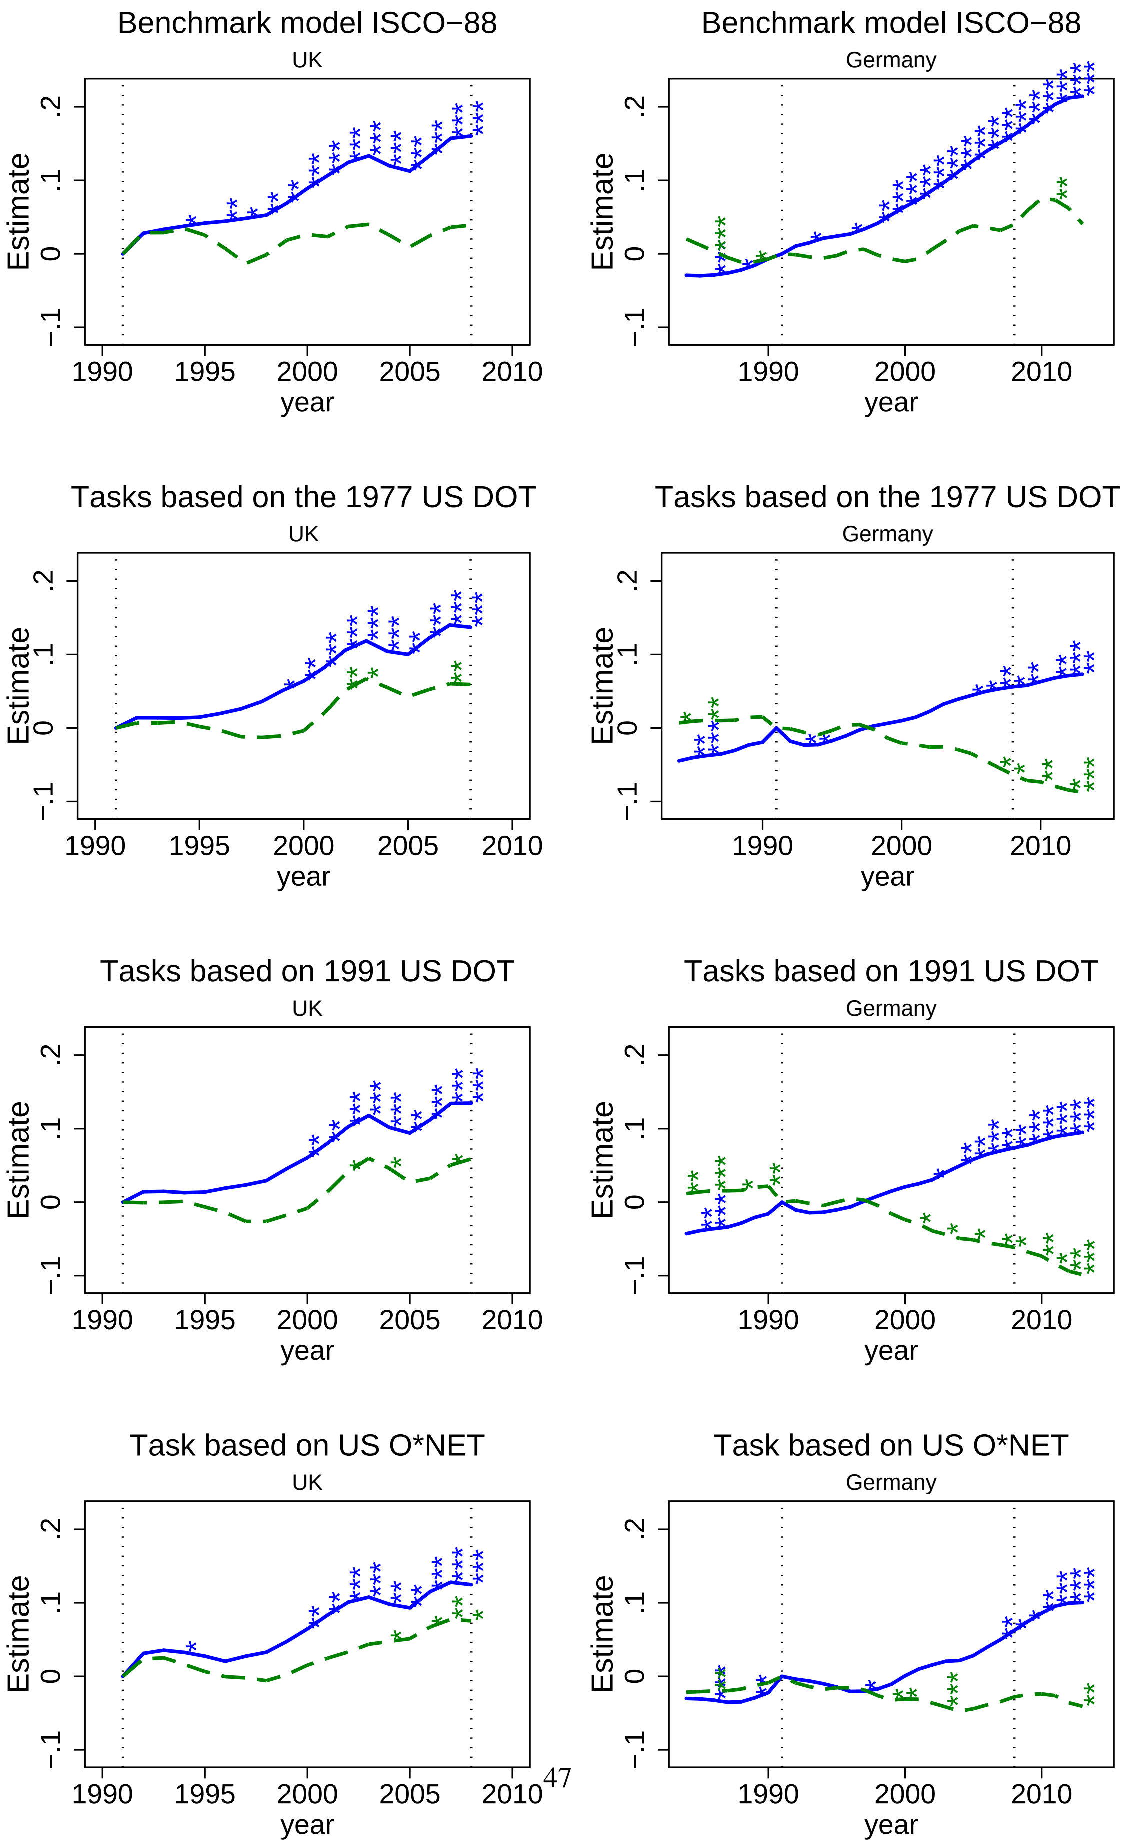

C Appendix to section 5: Allowing for shocks

In this appendix we examine the effect of allowing for shocks on the key identifying moments. We do this to allow for some (simple) mechanism other than price changes to generate occupational switches. The analysis is similar whether we allow for transitory shocks to (occupation-specific) wages or shocks to preferences. Here we model preference shocks. We also look exclusively at the variance of wages for switchers; this is a key moment in the identification of the covariance structure, and the moment that seems most susceptible to different modelling assumptions.

As in section 5, we base the analysis around 2 occupations. Suppose there are 2 periods $t = 1$ and $t = 2$. For simplicity we remove the effect of observable characteristics. These could be added as in section 5, but would change the analysis little.

We need only consider shocks ε_{it} which capture relative preferences for the abstract occupation at time t. We are interested in switchers both from non-abstract to abstract and vice versa. The switchers from non-abstract to abstract satisfy:

$$\gamma_{ia} - \gamma_{io} + \varepsilon_1 < \theta_{o1} - \theta_{a1}$$
$$\gamma_{ia} - \gamma_{io} + \varepsilon_2 > \theta_{o2} - \theta_{a2}$$

with a similar expression for the reverse switchers. We reduce the dimensionality of the problem by considering the following variables, $\eta_{i1} = \gamma_{ia} - \gamma_{io} + \varepsilon_1$ and $\eta_{i2} = \gamma_{ia} - \gamma_{io} + \varepsilon_2$, which capture the net preference for the abstract occupation in periods 1 and 2 respectively. Using the same notation for standardized random variables as in the main analysis we use the following orthogonalization:

$$\begin{pmatrix} \gamma_{ia}^* \\ \gamma_{io}^* \end{pmatrix} = B \begin{pmatrix} \eta_{i1}^* \\ \eta_{i2}^* \end{pmatrix} + \begin{pmatrix} \xi_{ia}^* \\ \xi_{io}^* \end{pmatrix}$$

where B contains functions of correlation coefficients, and $\begin{pmatrix} \xi_{ia}^* \\ \xi_{io}^* \end{pmatrix}$ is independent of $\begin{pmatrix} \eta_{i1}^* \\ \eta_{i2}^* \end{pmatrix}$. Intuitively, the initial and final periods give repeated measures of the underlying difference in skills.

Henceforth we drop i subscripts for ease of notation. It is useful to define some correlations. First define

$$\sigma_d^2 \equiv Var(\gamma_a - \gamma_o)$$

Further define:

$$\sigma_\eta^2 \equiv Var(\eta_t) = \sigma_d^2 + \sigma_\varepsilon^2$$

for $t = 1, 2$. Then define:

$$\rho_{12} \equiv Cov(\eta_1^*, \eta_2^*) = \frac{\sigma_d^2}{\sigma_\eta^2}$$

Finally define:

$$\rho_{a\eta} \equiv Cov(\gamma_a^*, \eta_1^*) = Cov(\gamma_a^*, \eta_2^*) = \frac{1}{\sigma_\eta}(\sigma_a - \rho\sigma_o)$$
$$\rho_{o\eta} \equiv Cov(\gamma_o^*, \eta_1^*) = Cov(\gamma_o^*, \eta_2^*) = \frac{1}{\sigma_\eta}(\rho\sigma_a - \sigma_o)$$

The first task is to compute the matrix B. In fact, it is easy to check that the columns of B are identical. Therefore let the rows of B be indicated by generic elements b_a and b_o. To compute B we use the covariance restrictions. For example:

$$\begin{aligned}\rho_{a\eta} \equiv Cov(\gamma_a^*, \eta_1^*) &= Cov(b_a\eta_1^* + b_a\eta_2^* + \xi_a^*, \eta_1^*)\\ &= b_a Var(\eta_1^*, \eta_1^*) + b_a Cov(\eta_2^*, \eta_1^*)\\ &\equiv b_a + b_a\rho_{12}\end{aligned}$$

Using similar restrictions we generate the following matrix equation:

$$\begin{pmatrix} \rho_{a\eta} & \rho_{a\eta} \\ \rho_{o\eta} & \rho_{o\eta} \end{pmatrix} = B\begin{pmatrix} 1 & \rho_{12} \\ \rho_{12} & 1 \end{pmatrix}$$

Inverting this equation gives:

$$B = \frac{1}{1+\rho_{12}}\begin{pmatrix} \rho_{a\eta} & \rho_{a\eta} \\ \rho_{o\eta} & \rho_{o\eta} \end{pmatrix}$$

We can then compute Σ_ξ, the covariance matrix of ξ. In fact, we only use the diagonal elements. To compute these we use, for example, the following identity:

$$\begin{aligned} Var(\gamma_a^*) &= Var(b_a(\eta_1^* + \eta_2^*) + \xi_a)\\ \Longrightarrow 1 &= b_a^2(2 + 2Cov(\eta_1^*, \eta_2^*)) + Var(\xi_a)\\ &= 2b_a^2(1+\rho_{12}) + Var(\xi_a)\\ &= \frac{2\rho_{a\eta}^2}{(1+\rho_{12})} + Var(\xi_a)\end{aligned}$$

We therefore have:

$$\sigma_{\xi_a}^2 = 1 - \frac{2\rho_{a\eta}^2}{(1+\rho_{12})}$$
$$\sigma_{\xi_o}^2 = 1 - \frac{2\rho_{o\eta}^2}{(1+\rho_{12})}$$

As in the main analysis we can then generate a generic description of the variance of productivities given some conditioning set. If Λ indicates the conditioning statement, and letting j indicate an occupation, and assuming ξ_j

is homoskedastic conditional on Λ, then

$$
\begin{aligned}
Var(\gamma_j|\Lambda) &= \sigma_j^2 Var(b_j(\eta_1^* + \eta_2^*) + \xi_j|\Lambda) \\
&= \sigma_j^2 \left(b_j^2 Var(\eta_1^*|\Lambda) + b_j^2 Var(\eta_2^*|\Lambda) + 2b_j^2 Cov(\eta_1^*, \eta_2^*|\Lambda) + Var(\xi_j) \right)
\end{aligned} \tag{6}
$$

As in the main analysis, we first quantify $Var(\xi_j) = 1 - \frac{2\rho_{j\eta}^2}{(1+\rho_{12})}$. Define $x = \frac{\sigma_\varepsilon^2}{\sigma_d^2}$, so that x indicates the size of preference shocks relative to differences in underlying productivities. For example, if there are no preference shocks then $x = 0$. Also let $\tilde{\rho}_{j\eta}$ be the value of $\rho_{j\eta}$ in the absence of preference shocks. In this case

$$
Var(\xi_j) = 1 - \frac{2\sqrt{1+x}}{2+x} (\tilde{\rho}_{j\eta})^2
$$

In a Taylor series expansion of $Var(\xi_j)$ around $1 - (\tilde{\rho}_{j\eta})^2$, the term $\frac{2\sqrt{1+x}}{2+x}$ disappears at order 1. Therefore

$$
Var(\xi_j) = 1 - (\tilde{\rho}_{j\eta})^2 + \mathscr{O}\left(\left(\sigma_\varepsilon^2\right)^2 \right)
$$

We conclude that for realistic values of σ_ε^2 (discussed below), then the presence of preference shocks has almost no effect on this moment.

We next examine the conditional variance moments, such as $Var(\eta_1^*|\Lambda)$. For those switching from non-abstract to abstract, the conditioning statement is

$$
\Lambda : \quad \eta_1^* < \beta \& \eta_2^* > \beta
$$

where β represents the standardized selection thresholds, which we now consider as constant over time. The conditional variances depend on the shape of the joint distribution $\begin{pmatrix} \eta_{i1}^* \\ \eta_{i2}^* \end{pmatrix}$. Here we use the example of the gaussian distribution. In this case we can characterize the conditional variances using the formulae from, for example, Muthen (1990) on moments of the truncated bivariate gaussian.

Let π represent the mass of individuals in set Λ. Furthermore, let $c = 1/\sqrt{\left(1 - \rho_{12}^2\right)}$. Then we have the following solution:

$$
\pi Var(\eta_1^*|\Lambda) = \pi - \beta\phi(\beta)[1 - \Phi[(\beta(1-\rho_{12}))c]] + \rho_{12}^2 \beta\phi(\beta)\Phi[(\beta(1-\rho_{12}))c] \tag{7}
$$

We can obtain similar expressions for $Var(\eta_2^*|\Lambda)$ and for $Cov(\eta_1^*, \eta_2^*|\Lambda)$.

We could similarly perform a Taylor-series expansion around $\sigma_\varepsilon^2 = 0$ in expression 7. It is simpler to use a calibration. To do this, first we need to quantify x, the ratio of preference shocks to permanent productivity differences. If $f()$ is the joint distribution of $\begin{pmatrix} \eta_{i1}^* \\ \eta_{i2}^* \end{pmatrix}$, then the mass of people switching as a function of the standardized threshold β is

$$
Proportion\ switch(\beta) = \left[\int_{-\infty}^{\beta} \int_{\beta}^{\infty} f(x,y)\,dxdy \right] + \left[\int_{\beta}^{\infty} \int_{-\infty}^{\beta} f(x,y)\,dxdy \right]
$$

In the data we observe at most 10% of individuals switching broad occupations in every year. Therefore this can be used as an upper bound on the quantity of switching, which is in fact generally smaller. To pin down the variance precisely we need to specify β. Here we calibrate at a standardized threshold of $\beta = 0$. This implies that $x = 0.05$, i.e. the variance of ε is only be 5% of the variance of $\gamma_a - \gamma_o$.

Plugging in these numbers gives a conditional variance, $Var(\eta_1^*|\Lambda)$, of around 2% of the variance of $\gamma_a - \gamma_o$, and a conditional covariance of around 0.01% of the variance of $\gamma_a - \gamma_o$. We conclude that that these variance terms are quantitatively small and that preference shocks do not contribute quantitatively to the variance of switchers.