KATHOLIEKE UNIVERSITEIT
LEUVEN

Faculty of Economics and Applied Economics

Department of Economics

Explaining job polarization: the roles of technology, offshoring and institutions

by

Maarten Goos, Alan Manning & Anna Salomons

Monetary and information Economics

Center for Economic Studies
Discussions Paper Series (DPS) 11.34
http://www.econ.kuleuven.be/ces/discussionpapers/default.htm

December 2011

DISCUSSION PAPER

Explaining Job Polarization: The Roles of Technology, Offshoring and Institutions

By MAARTEN GOOS, ALAN MANNING AND ANNA SALOMONS*

This paper develops a simple and empirically tractable model of labor demand to explain recent changes in the occupational structure of employment as a result of technology, offshoring and institutions. This framework takes account not just of direct effects but indirect effects through induced shifts in demand for different products. Using data from 16 European countries, we find that the routinization hypothesis of Autor, Levy and Murnane (2003) is the most important factor behind the observed shifts in employment but that offshoring does play a role. We also find that shifts in product demand are acting to attenuate the impacts of recent technological progress and offshoring and that changes in wage-setting institutions play little role in explaining job polarization in Europe.

JEL: J21, J23, J24
Keywords: Labor Demand, Technology, Globalization, Polarization

* Goos: Department of Economics, University of Leuven, Naamsestraat 69, 3000 Leuven, Belgium, maarten.goos@econ.kuleuven.be. Manning: Centre for Economic Performance and Department of Economics, London School of Economics, Houghton Street, London WC2A 2AE, United Kingdom, a.manning@lse.ac.uk. Salomons: Utrecht University School of Economics, Kriekepitplein 21-22, 3583 EC Utrecht, The Netherlands, a.m.salomons@uu.nl.
Acknowledgements: We thank David Autor, Larry Katz, Thomas Lemieux, Stephen Machin, Guy Michaels, John Van Reenen and numerous seminar participants for excellent suggestions. [Date: 23 November 2011]

I. Introduction

Economists and non-economists alike have long been fascinated by the ever-changing structure of employment. Economists have developed a number of hypotheses about the driving forces behind these changes. The most popular emphasize the importance of technological change, globalization (partly driven by technology, but perhaps partly also an independent force from declining man-made barriers to trade), and labor market institutions (e.g. that alter the relative wages of different types of labor). Until recently, the most common way to summarize the changes in the structure of employment was that demand was shifting in favor of skilled relative to unskilled workers and that it was technological change rather than globalization that was the most important force behind these shifts, leading to the hypothesis of Skill-Biased Technological Change (SBTC) – for an excellent survey of this see Autor and Katz (1999). But, recently, the need for a more nuanced view has become apparent, primarily because of the observation of job polarization (see Autor, Katz and Kearney 2006, for the US; Goos and Manning 2007, for the UK) that while there has been a rapid rise in the share of employment in the highest-paid occupations, this has primarily been at the expense of occupations in the middle of the pay distribution – the employment share of the lowest-paid occupations has been constant or even rising. Although the phenomenon of job polarization has been documented for the US, the UK and Germany (Spitz-Oener 2006; Dustmann, Ludsteck and Schönberg 2009) we still do not know how pervasive it is (though Goos, Manning and Salomons 2009; Michaels, Natraj and Van Reenen 2010, suggest it is happening in almost all European countries). One of the contributions of this paper is to document that it really is widespread across Europe.

Another contribution of the paper is to shed light on the reasons for job polarization. The most popular explanation is the 'routinization' hypothesis first proposed by Autor, Levy and Murnane (2003), that technological change displaces human labor in tasks that can be described as routine (essentially so that a computer program can be written to mimic what a human would do). But, some authors have also pointed to 'offshoring' as a potential cause of polarization – see, for example, Grosmann and Rossi-Hansberg (2008) and Blinder and Krueger (2009) who estimate that approximately 25% of US occupations might become offshorable within the next 20 years. One would expect both of these factors to be at work in all developed economies so there is value in our consideration of every European economy – our empirical evidence suggests routinization is more important than offshoring. The paper also shows how one can separately identify whether technological change is labor- or capital-augmenting – our preferred estimates suggest, plausibly, that routinization is best modeled as capital-augmenting technical change.

The final, and perhaps most important contribution of this paper, is to provide and estimate a conceptual framework capable of a more complete explanation of the impact of routinization and offshoring on the demand for different types of labor. Our model is one in which there are different industries and industry output is produced from a common set of 'tasks', used in different proportions by different industries. Routinization and offshoring impact the task-level production functions. It is a model with roots in the way in which SBTC was modeled using an aggregate production function whose inputs were skilled and unskilled workers, possibly with the addition of age as a separate factor – see, for example, Katz and Murphy (1992) or Card and Lemieux (2001) – this is what Acemoglu and Autor (2011) call the 'canonical' model. With its two basic factors – skilled and unskilled labor – the canonical model obviously cannot explain a phenomenon like job polarization and models that aim to do so (Autor, Levy and Murnane 2003; Autor, Katz and Kearney 2006; Autor and Dorn 2011; Acemoglu and Autor 2011) are necessarily more complicated.

But our model is also more complex in another, more important, way. As the paper shows, one cannot obtain an adequate understanding of the changing structure of employment if one ignores the channels by which a change affecting the demand for one type of labor is likely to spill over to every other type of labor. A simple example of such spillover effects, inspired by one of the popular works of Krugman (1999) – though none the worse for that – will illustrate. A hamburger requires one bun and one patty. Suppose there is an improvement in the technology of patty-making so that one now only needs half the number of workers to produce one patty. This change obviously only directly affects patty-making so a simple-minded approach would choose an empirical specification in which the technological change variable only appears in the equation for the number of patty-makers. But, if the empirical specification assumes that technological change does not affect the employment of the bun-makers and the number of buns and patties must remain in the same proportion, the only possible conclusion is that the innovation reduces the employment of patty-makers and employment overall. Krugman's point is that this is a serious mistake.

Specifically, the reduction in the cost of producing patties also implies a reduction in the cost of producing hamburgers. This leads to a reduction in the price of hamburgers causing an increase in the demand for them (assuming they are non-Giffen). The employment of the bun-makers then rises and the employment of the patty-makers is higher than one would predict if one assumed the production of hamburgers remained constant but not necessarily so large as to prevent an overall fall in employment. Employment in bun-making is affected by innovation in patty-making and we have a clear idea of the channel – through changes in product demand induced by changes in costs and prices. And if preferences are non-homothetic, induced changes in the level and distribution of income will also induce changes in the structure of employment. Although none of these ideas are new – they date back to at least the work of Baumol (1967) – existing models are largely silent on the question of how important these derived demand effects are[1].

The paper is organized as follows. Section II describes the data and shows how the employment structure in 16 European countries is polarizing. Section III then presents our simple theoretical framework of the demand for occupations within industries to organize our thoughts about how the hypotheses outlined above affect the demand for labor. Section IV describes the variables we use to capture these hypotheses. In section V we estimate the model of within industry changes in occupational labor demand across

[1]One recent notable exception is Autor and Dorn (2011) who test the labor market implications of unbalanced productivity growth in a general equilibrium framework using US data. Arguably, their model is richer than ours – because it also endogenizes labor supply to react to technological progress – but otherwise both models are qualitatively similar and there is a tighter link between our theory and empirical specification.

countries. Section VI then seeks to move beyond within industry equations to consider the importance of changes in relative product demand through the introduction and estimation of product demand curves. Finally, Section VII evaluates to what extent economy wide job polarization can be explained by our model.

Our main conclusion is that the routinization hypothesis of Autor, Levy and Murnane (2003) has the most explanatory power for understanding job polarization but offshoring does play a role. We also find that induced changes in relative output prices and the induced changes in demand for different products acts to attenuate job polarization, that the evidence tends to support a capital-augmenting rather than labor-augmenting view of the impact of technological change and that income and institutional effects are relatively unimportant.

II. A Picture of Changes in the European Job Structure

A. Employment Data[2]

In this paper we model employment by industry and occupation. Our main data source for employment is the harmonized individual level European Union Labour Force Survey (ELFS) for the period 1993-2006. The ELFS contains data on employment status, weekly hours worked, 2-digit International Standard Occupational Classification (ISCO) codes and 1-digit industry codes from the Classification of Economic Activities in the European Community (NACE revision 1). Throughout this paper, we use weekly hours worked as a measure for employment, although our results are not affected by using persons employed instead. Out of the 28 countries available in the ELFS, we exclude 11 new EU member countries[3] and Iceland because of limited data availability. We also discarded Germany from the ELFS because of its too small sample size and limited time span and replaced it with data from the German Federal Employment Agency's IABS data set.[4] Data for the remaining 15 European countries (Austria, Belgium, Denmark, Finland, France, Greece, Ireland, Italy, Luxembourg, the Netherlands, Norway, Portugal, Spain, Sweden and the United Kingdom) is used in the analysis.

We drop some occupations and industries from the sample: those related to agriculture and fishing because they do not consistently appear in the data and because OECD STAN industry output data, that we will use later in this paper, is not suited for comparison across countries for these sectors; and those related to the public sector (public administration and education) because German civil servants are not liable to social se-

[2]See Appendix A for details.

[3]Cyprus, the Czech Republic, Estonia, Hungary, Lithuania, Latvia, Poland, Slovenia, Slovakia, Bulgaria, Romania.

[4]The IABS dataset is a 2% random sample of German social security records for the period 1993-2002. For each individual it contains data on occupation and industry, as well as several demographic characteristics (among others, region of work, full-time or part-time work). We drop workers who are not legally obliged to make social security contributions (some 9% of all observations) because for them the IABS is not a random sample. Lacking a measure of hours worked, we use time-varying information on average weekly hours worked for full-time and part-time workers in both East- and West-Germany, obtained from the European Foundation for the Improvement of Living and Working Conditions to proxy for total hours worked in IABS occupation-industry-year cells (though our results are robust to restricting the sample to full-time workers). We then manually convert the German occupation and industry codings to match ISCO and NACE in the ELFS.

curity and therefore not included in the IABS, and because OECD STAN net operating surplus data is not reliable for these two sectors. Our results are never driven by the exclusion of these occupations and industries.

B. Data Summary

To provide a snapshot of changes in the European job structure, Table 1 shows the employment shares of occupations and their percentage point changes between 1993 and 2006 after pooling employment for each occupation across our 16 European countries.[5] This table shows that the high-paying managerial (ISCO 12, 13), professional (ISCO 21 to 24) and associate professional (ISCO 31 to 34) occupations experienced the fastest increases in their employment shares. On the other hand, the employment shares of office clerks (ISCO 41), craft and related trades workers (ISCO 71 to 74) and plant and machine operators and assemblers (ISCO 81 to 83), which pay around the median occupational wage, have declined. Similar to patterns found for the US and UK, there has been an increase in the employment shares for some low-paid service workers (ISCO 51, of which the main task consists of providing services related to travel, catering and personal care; but not ISCO 52, of which the main task consists of selling goods in shops or at markets) and low-paid elementary occupations (ISCO 91, which are service elementary workers including cleaners, domestic helpers, doorkeepers, porters, security personnel and garbage collectors; and ISCO 93, which mainly includes low-educated laborers in manufacturing performing simple tasks that require the use of hand-held tools and often some physical effort). This is an indication that the existing evidence to date that there is job polarization in the US, UK or Germany is not an exception but rather the rule.[6]

There may, of course, be heterogeneity across countries in the extent of polarization. Table 2 groups the occupations listed in Table 1 into three groups: the four lowest-paid occupations, nine middling occupations and the eight highest-paid occupations. We then compute the percentage point change in employment share for each of these groups in each country between 1993 and 2006. Table 2 confirms that employment polarization is pervasive across European countries – the share of high-paying occupations increases relative to the middling occupations in all countries but Portugal, and the share of low-paying occupations increases relative to the middling occupations in all countries.[7]

As outlined in the introduction, there are a number of possible hypotheses – technological progress, globalization and induced effects operating through various channels such as product demand – that can explain these changes and the next section starts to

[5] Since all countries do not have data for the entire time-span of 1993-2006, we calculate average annual changes for each country and use these to impute the employment shares in 1993 and/or 2006 where they are not available.

[6] One might be concerned that 1993 is a recession and 2006 a boom so that the changes in Table 1 are cyclical not trends. However, it can be shown that time series of occupational employment growth are primarily trends and that the polarization found in Table 1 is not sensitive to endpoints. See Appendix A for details.

[7] This result is upheld when we add customer service clerks, a middle-paid service occupation, to the four lowest-paid occupations: indeed, in this case, we observe an increase in the share of high- and low-paying occupations relative to the middling occupations in all countries, i.e. including Portugal. In Acemoglu and Autor (2011) it is argued that the numbers in Table 2 show that job polarization is at least as pronounced in our sample of European countries as in the US.

outline a simple theoretical framework to help us to understand and estimate the relative importance of these factors.

III. A Simple Model of Within Industry Demands for Occupations

Our ultimate aim is to explain the changes in the aggregate occupational structure of employment documented in the previous section. In order to do this we develop a model of the demand for different occupations in different industries conditional on industry output and then later we model the demand for industry output.

A. The Production of Goods and the Demand for Tasks

Industry level production function

Assume that output in all industries is produced by combining certain common building blocks that we will call tasks. Some industries are more intensive users of some tasks than others. In particular, assume the following CES production function for industry i using tasks $T_1, T_2, ..., T_J$ as inputs:

$$Y_i(T_{i1}, T_{i2}, ..., T_{iJ}) = \left[\sum_{j=1}^{J} \left(\beta_{ij} T_{ij}\right)^{\eta}\right]^{\frac{1}{\eta}} \text{ with } \eta < 1. \tag{1}$$

The cost-minimizing demand for task j conditional on Y_i is:

$$\begin{aligned} T_{ij}(c_1^T, c_2^T, ..., c_J^T | Y_i) &= Y_i \frac{1}{\beta_{ij}} \left(\frac{c_j^T}{\beta_{ij}}\right)^{-\frac{1}{1-\eta}} \left[\sum_{j=1}^{J} \left(\frac{c_j^T}{\beta_{ij}}\right)^{-\frac{\eta}{1-\eta}}\right]^{-\frac{1}{\eta}} \\ &= Y_i \left(\frac{c_i^I}{c_j^T}\right)^{\frac{1}{1-\eta}} \beta_{ij}^{\frac{\eta}{1-\eta}} \end{aligned} \tag{2}$$

where c_j^T is the unit cost of using task j (derived below) and c_i^I industry marginal cost that is given by:

$$c_i^I = \left[\sum_{j=1}^{J} \left(\frac{c_j^T}{\beta_{ij}}\right)^{-\frac{\eta}{1-\eta}}\right]^{-\frac{1-\eta}{\eta}} \approx \prod_{j=1}^{J} \left(\frac{c_j^T}{\beta_{ij}}\right)^{\kappa_{j|i}} \tag{3}$$

with $\kappa_{j|i}$ the share of task j in the cost of producing one unit of good i. The final expression in equation (3) is an approximation if the task level production function is not Cobb-Douglas but is useful for producing log-linear estimating equations. Note that in the empirical work that follows we also have country and time subscripts but we ignore these for the moment to avoid excessive notation.

We assume that output of task j is produced using labor of one occupation and some other input which we will very generally refer to as capital.[8] The other input could be machinery to capture task-biased technological progress or offshored overseas employment to capture offshoring. We proceed by assuming the other input is one-dimensional but that is just for simplicity and explicitly accounting for multiple inputs only adds algebraic complication. In particular, assume that in industry i tasks are produced using domestic labor of occupation j, N_{ij}, and capital, K_{ij}, according to the following CES technology:

$$T_{ij}(N_{ij}, K_{ij}) = \left[\left(\alpha_{Nj} N_{ij}\right)^{\rho} + \left(\alpha_{Kj} K_{ij}\right)^{\rho}\right]^{\frac{1}{\rho}} \text{ with } \rho < 1 \tag{4}$$

where it is assumed that the technology to produce task j is common across industries so the subscript i only appears on the input factors.[9]

The cost-minimizing demand for labor of type j conditional on T_{ij} is given by:

$$\begin{aligned} N_{ij}(w_j, r_j | T_{ij}) &= T_{ij} \frac{1}{\alpha_{Nj}} \left(\frac{w_j}{\alpha_{Nj}}\right)^{-\frac{1}{1-\rho}} \left[\left(\frac{w_j}{\alpha_{Nj}}\right)^{-\frac{\rho}{1-\rho}} + \left(\frac{r_j}{\alpha_{Kj}}\right)^{-\frac{\rho}{1-\rho}}\right]^{-\frac{1}{\rho}} \\ &= T_{ij} \frac{1}{\alpha_{Nj}} \left(\frac{w_j}{\alpha_{Nj} c_j^T}\right)^{-\frac{1}{1-\rho}} \end{aligned} \tag{5}$$

where w_j is the wage in occupation j and r_j the price of the other input used in the production of task j and where the cost of producing one unit of task j is given by:

$$c_j^T = \left[\left(\frac{w_j}{\alpha_{Nj}}\right)^{-\frac{\rho}{1-\rho}} + \left(\frac{r_j}{\alpha_{Kj}}\right)^{-\frac{\rho}{1-\rho}}\right]^{-\frac{1-\rho}{\rho}} \approx \left(\frac{w_j}{\alpha_{Nj}}\right)^{\kappa} \left(\frac{r_j}{\alpha_{Kj}}\right)^{(1-\kappa)} \tag{6}$$

Note that the final expression in equation (6) is again an approximation if the task level production function is not Cobb-Douglas. In the final expression, κ is the share of do-

[8] This type of two-stage set-up for modeling the production process is increasingly standard in the task literature, although different papers do it differently. Autor, Levy and Murnane (2003) assume industry output uses routine and non-routine tasks that can be done by a continuum of differently skilled workers where capital can substitute for labor doing routine tasks. Autor, Katz and Kearney (2006) assume three tasks are used in aggregate output: abstract, manual and routine. Abstract tasks are done by college workers whereas manual and routine tasks are done by a continuum of differently skilled high-school workers that are competing with capital for doing routine tasks. Grossman and Rossi-Hansberg (2008) assume two goods are produced domestically using a continuum of offshorable tasks done by domestic or foreign low-skilled labor and a continuum of non-offhorable tasks done by domestic high-skilled labor. Autor and Dorn (2011) have two industries (goods and services), three tasks (abstract, manual and routine) and assume the goods industry production function uses abstract and routine labor as well as capital and the services sector uses only manual labor. In Acemoglu and Autor (2011) a single final good is produced from a continuum of tasks that can be produced by three different types of labor (high-, medium- and low-skilled) and capital. It is easy to show that all these models fit within our set-up.

[9] This is a strong assumption (though it has been used in other models e.g. Grossman and Rossi-Hansberg 2008) and Appendix B provides some evidence on its validity.

mestic labor in the unit cost of producing task j that we approximate as being constant across occupations.

B. *The Conditional Demand for Labor*

Combining equations (2), (5) and (6) and taking logs and adding country and time subscripts leads to the following expression for the demand for labor conditional on industry output and industry marginal costs:

$$
\begin{aligned}
\log N_{ijct} = & -\left[\frac{1-\kappa}{1-\rho}+\frac{\kappa}{1-\eta}\right]\log w_{jct} \\
& +\left[\frac{1-\kappa}{1-\rho}+\frac{\kappa}{1-\eta}-1\right]\log \alpha_{Njt} \\
& +\left[\frac{1}{1-\rho}-\frac{1}{1-\eta}\right](1-\kappa)\log\left(\frac{r_{jt}}{\alpha_{Kjt}}\right) \\
& +\frac{\eta}{1-\eta}\log \beta_{ijc}+\frac{1}{1-\eta}\log c^{I}_{ict}+\log Y_{ict}
\end{aligned}
\tag{7}
$$

The first term in this labor demand curve shows, unsurprisingly, that wages should be negatively related to labor demand with the elasticity being determined by the elasticity of substitution both within and across tasks, and the share of domestic labor in task costs. The second term is the effect of domestic labor-augmenting productivity – the sign of the effect of this on labor demand, depends, as is well-known, on whether the wage elasticity is greater than or less than unity. The third term is the effect of capital cost and productivity – these are never separated out and will be referred to henceforth as capital-augmenting productivity. The effect of this term on labor demand depends on the difference in the elasticity of substitution within and across tasks – a reduction reduces labor demand because labor becomes relatively expensive within tasks but increases because the overall cost of doing the task is reduced and the net effect depends on the relative size of these two effects. The final line shows that labor demand is increasing in industry output and also 'increasing' in the marginal industry cost. This last effect should be interpreted with some caution as a change in the wage and the task production function will have some effect on industry marginal costs.

This equation forms the starting point for the empirical investigation and the next section describes how we construct the relevant variables.

IV. Data

In this section we describe the data sources of our measures of technological change, offshoring and wages at the level of occupations as well as our measures of industry marginal costs and industry output.

A. *Technological Progress*[10]

To measure the type of work done in occupations, we use data from the December 2006 version of the Occupational Information Network (ONET) database. ONET is a primary source of occupational information, providing comprehensive data on key attributes and characteristics of workers in US occupations. It is a replacement for the Dictionary of Occupational Titles (DOT) which has been used extensively in earlier research, (Autor, Levy and Murnane 2003; Goos and Manning 2007; Autor, Katz and Kearney 2008; and Autor and Dorn 2011). ONET data comes from job incumbents, occupational analysts and occupational experts and is collected for 812 occupations which are based on the 2000 Standard Occupational Code (SOC). We manually converted the 2000 Standard Occupational Code (SOC) used in the ONET data to ISCO codes used in the ELFS.

One part of ONET consists of some 100 variables related to worker characteristics, worker requirements and general work activities. We select 96 of these task content measures and, following Autor, Katz and Kearney (2008) and Autor and Dorn (2011), manually categorize each of the 96 ONET variables into one of three groups: Abstract, Routine or Service.[11] Routine tasks are those which computers can perform with relative ease, such as jobs that require the input of repetitive physical strength or motion, as well as jobs requiring repetitive and non-complex cognitive skills. The non-routine dimension is split up into Abstract and Service to capture the different skill content of these non-routine tasks: examples of Abstract tasks are "complex problem solving" (e.g. needed by engineers and medical doctors) and Service tasks are "caring for others" (e.g. needed by hairdressers and medical doctors).

Each of these three task content measures we collapse to the ISCO classification which we have in our European data by calculating an average across SOC occupations, weighted by US employment in each SOC cell taken from ONET. Columns (1) through (3) of Table 3 show the values, normalized to have mean zero and unit standard deviation, for 2-digit ISCO occupations ranked by their mean 1993 wage across the 16 European countries. Column (4) of Table 3 summarizes the information in columns (1)-(3) by constructing a one dimensional Routine Task Intensity (RTI) index, defined as Routine task importance divided by the sum of Abstract and Service task importances and again standardized to have mean zero and unit standard deviation.

Columns (1) to (3) of Table 3 allow us to categorize occupations into three broad groups. Firstly, some occupations are highly routine and relatively low in abstract and service task importance (craft and related trade workers (ISCO 71-74); plant and machine operators and assemblers (ISCO 81-83)). Secondly, some occupations are low in routine task importance and high in both abstract and service task importance (managers (ISCO 12,13); professionals (ISCO 21-24); associate professionals occupations (ISCO 31-34)). Thirdly, some occupations are low in routine and abstract but high in service task importance (customer service clerks (42); office clerks (41) although much less ser-

[10]See Appendix C for details.

[11]Our results are robust to constructing task dimensions by means of principal component analysis rather than by manual assignment. See Appendix C for details.

vice orientated compared to customer service clerks; low-paid service workers (ISCO 51,52); low-paid elementary occupations (ISCO 91,93) although these occupations are relatively low in service importance compared to low-paid service workers).

Our measures of task content only vary across occupations – to capture technological change we interact these measures with a linear time trend.[12] The assumption here is that occupations that are more routine are also subject to greater technological progress.

To see how the routinization hypothesis can yield different predictions from the SBTC hypothesis, column (6) of Table 3 also presents mean educational attainment by occupation. This variable derives from the education variable available in the ELFS (categorized with the International Standard Classification of Education, ISCED), which we average by occupation across countries.[13] It is clear that the technology variables in columns (1) to (4) do not map simply into the one-dimensional level of education in column (6).

B. Offshoring[14]

There are a number of existing approaches to measuring the impact of offshoring on the labor market. Typically, use is made of measures of foreign direct investment, imports in total GDP or the share of intermediate imports in total imports. At best, this type of data is available at the country-industry-time level but never at the occupation level.

We obtain a measure of offshorability from the European Restructuring Monitor (ERM) of the European Monitoring Centre on Change (EMCC), which is a part of Eurofound. ERM is available online and provides summaries of news reports (so-called 'fact sheets') since 2002 about companies located in Europe that announce offshoring plans. These fact sheets contain information on the company that is offshoring part(s) of its production process, such as the country and the industry in which it operates, how many workers are employed nationwide or in that particular location, how many jobs are being offshored and to which country, and, most importantly for our purposes, which occupations are being offshored.

We process 415 fact sheets (covering May 31st, 2002 up to June 30th, 2008), or cases of offshoring, to construct an index of how offshorable the different occupations are. We sum the number of cases for each ISCO occupation, and generate a rank by rescaling the number of cases across occupations to a distribution with mean zero and unit standard deviation.

Column (5) of Table 3 shows this occupation level measure of offshorability. It can be seen that some occupations that are high in routine task importance are offshored most often (metal, machinery and related trade workers (ISCO 72); plant and machine

[12]The lack of time variation in ONET would be problematic for the analysis below if the task composition within occupations is changing over time. However, using similar DOT measures across US occupations and over time, Goos and Manning (2007) find that most of the overall changes in mean task measures happened between and not within occupations. Also note that ONET does not contain any variation in job task measures within occupations. However, Autor and Handel (2009) use the individual level Princeton Data Improvement Initiative to show that occupation is the dominant predictor for the variation in the task measures that are also used in this paper.

[13]Occupational education levels are very highly correlated among countries, the average correlation coefficient being 0.93 with a standard deviation of 0.03.

[14]See Appendix D for details.

operators and assemblers (ISCO 81,82)). However, Table 3 also shows that other routine occupations are much less offshorable (construction workers (ISCO 71); precision, handicraft, craft printing and related trade workers (ISCO 73,74); drivers (ISCO 83)). Similarly, some occupations low in routine but high in abstract or service task importance (engineering (associate) professionals (ISCO 21,31); other associate professionals which includes call-centre workers (ISCO 34); office clerks (ISCO 41); low-paid elementary workers in goods production (ISCO 93)) are still much more offshorable than others (managers (ISCO 12,13), life science and health (associate) professionals (ISCO 22,32); customer service clerks (ISCO 42); low-paid service (elementary) workers (ISCO 51,52,91)). This all seems sensible.It is also reassuring that our measure corresponds closely to the preferred measure in Blinder and Krueger (2009) that is constructed by professional coders based on a worker's occupational classification.[15]

Our measure of offshorability also only varies across occupations – to capture changes over time we interact our occupation specific offshoring variable with a linear time trend.

C. Occupational Wages[16]

Since the ELFS does not contain any earnings information, we obtain time-varying country-specific occupational wages from the European Community Household Panel (ECHP) and European Union Statistics on Income and Living Conditions (EU-SILC). The ECHP contains gross monthly wages for the period 1994-2001, whereas the EU-SILC reports gross monthly wages for the period 2004-2006. For the UK, we use the gross weekly wage from the UK Labour Force Survey because it contains many more observations and is available for 1993-2006. All wages have been converted into 2000 Euros using harmonized price indices and real exchange rates.

To match our employment dataset, we construct an occupational wage measure weighted by hours worked. Because sample sizes in the ECHP and EU-SILC are relatively small, we smooth wages by pooling together all years for each occupation and estimating a model in which the dummy on occupation varies smoothly with a quadratic time trend. We also use this procedure to impute wages for years that are missing. The sample sizes are also too small to allow for occupation and time-specific variation in industry wage differentials so industry wage effects are subsumed into industry dummies.

Given the less than perfect nature of the wage data, it is reassuring that the wage rank of occupations is intuitive, and highly stable over time. Just as in Tables 1 and 3, Table 4 ranks occupations by their mean European wage in 1993 shown in column (1). The ranking is as expected, with managers and professionals being the most highly-paid, service workers and workers in elementary occupations the lowest-paid, and manufacturing and office workers somewhere in between. Column (2) provides mean wages in 2006 to show that this ranking is very stable over our sampling period. The final column of Table 4 reports growth rates – that is, $[(2) - (1)]/(1) \times 100$ – to suggest that wage changes in

[15] Using the preferred measure in Blinder and Krueger (2009) in the analysis below does not change any of our results. See Appendix D for details.

[16] See Appendix E for details.

Europe are not correlated to technological progress and globalization in any straightforward way. For example, coefficients of a regression of these numbers onto the RTI and offshorability measures in columns (4) and (5) of Table 3 are positive and only jointly statistically significant at the 32% level, whereas using the employment share changes from the last column of Table 1 as the dependent variable instead yields negative point estimates that are jointly statistically significant at the 2% level.

D. Industry Marginal Cost and Industry Output[17]

Measures of industry output and industry marginal costs are taken from the OECD STAN Database for Industrial Analysis. Each of our 16 countries except Ireland is included in STAN. This data covers the period 1993-2006 for all 15 of these countries. STAN uses an industry list for all countries based on the International Standard Industrial Classification of all Economic Activities, Revision 3 (ISIC Rev.3) which covers all activities including services and is compatible with NACE revision 1 used in the ELFS.[18]

The measure of output used in the analysis below is production, defined as the value of goods or services produced in a year, whether sold or stocked.[19] To obtain variation in output, production has been deflated using industry-country-year specific price indices also available from STAN. Finally, we approximate industry marginal costs by the difference between production and net operating surplus, divided by production. This gives an estimate of the average cost of using labor, capital and intermediate inputs per euro of output. This measure can be seen as a proxy for the variation in industry average cost, which in our model is identical to industry marginal cost.

V. Estimating Conditional Labor Demand

A. Basic Estimates

Our empirical estimation of equation (7) is based on the following models for labor-augmenting productivity:

$$\log \alpha_{Njt} = \gamma_j^N + \left[\gamma_R^N R_j + \gamma_O^N O_j\right] \times linear\ time\ trend \tag{8}$$

where γ_j^N is an occupation-specific constant; R_j is one of the occupation-specific measures of task content discussed above or a combination of those; O_j is the occupation-specific measure of offshorability discussed above; γ_R^N is the trend in labor productivity for a particular task content or a vector of those; and γ_O^N the trend in labor productivity for

[17] See Appendix F for details.

[18] Due to limited data on net operating surplus for the NACE industry "Private households with employed persons", we have one less industry when using STAN data in our regressions. Although this industry mainly employs low-paid service elementary workers and its employment share has increased from 0.82% in 1993 to 0.90% in 2006, it is too small to be important.

[19] We use production as it includes any intermediate inputs from offshore locations. Our results are robust to using value added instead.

a given level of offshorability. Similarly, we have for capital-augmenting productivity:

$$\log\left(\frac{r_{jt}}{\alpha_{Kjt}}\right) = \gamma_j^K + \left[\gamma_R^K R_j + \gamma_O^K O_j\right] \times \textit{linear time trend} \tag{9}$$

Note that when equations (8) and (9) are substituted into equation (7) one cannot separately identify the effects of task content and offshorability on labor- and capital-augmenting productivity – we return to this issue below.

Table 5 reports estimates of the conditional labor demand curve where industry output and marginal costs are modeled by industry-country-time dummies. Columns (1) to (3) report point estimates for each task content measure separately whereas column (4) adds them simultaneously. The point estimates suggest that, all else equal, employment increases by 1.33% and 1.28% faster annually for occupations one standard deviation more intense in abstract and service tasks, respectively, whereas employment in occupations one standard deviation more intense in routine tasks increases 1.33% slower annually. Although the point estimates in column (4) are generally smaller in absolute value, they have the expected sign and remain highly statistically significant for abstract and routine tasks. Using the RTI index, column (5) suggests that employment in occupations that are one standard deviation more routine has grown 1.52% slower each year, ceteris paribus. Columns (6) to (10) repeat the analysis in columns (1) to (5) but now adds our measure of offshorability. The estimated coefficients on the task content variables are very similar to those reported in columns (1) to (5) whereas the impact of offshorability is smaller in absolute value and significantly decreases in magnitude when task content measures – especially service task importance or the RTI index – are controlled for.[20]

To take our estimates further, we now replace the industry-country-year dummies with our industry output and industry marginal cost variables. To account for measurement error and transitory changes in industry output over time, columns (1) and (2) of Table 6 show 2SLS estimates using the logarithm of industry net capital stock wherever available in STAN as an instrument for log industry output.[21]

The point estimates on the task content measures in Table 6 are similar in magnitude and significance to those reported in columns (9) and (10) of Table 5 whereas the point estimates on offshorability are somewhat larger in absolute value and significant at the 5% level. The coefficient on log industry output (0.90 with a standard error of 0.13) is not significantly different from unity, consistent with the assumption of constant returns to scale. The coefficient on log industry marginal costs is positive and significant suggesting an elasticity of substitution between tasks in goods production, $1/(1-\eta)$, of 0.80 with a standard error of 0.23. To check for the robustness of this result in the larger sample of 16 countries, columns (3) and (4) show estimates using OLS while imposing constant returns to scale by constraining the coefficient on log industry output to be unity.

[20] Because we have over-identification, we can also test whether these estimates show heterogeneity at the country and industry level. It can be shown that this is not the case. See Appendix G for details.

[21] This restricts the number of observations because net capital stock data is not available for a number of countries (leaving Austria, Belgium, Denmark, Finland, France, Germany, Italy and Norway). The OLS estimate is 0.39 with a standard error of 0.04.

This gives an estimate of $1/(1-\eta)$ of 0.53 with a standard error of 0.13 and does not qualitatively affect our conclusion that equation (7) is a reasonable approximation of the employment variation observed in our data.

B. Endogenous Wages

The analysis so far has assumed that wages can be treated as exogenous in the estimation of the conditional labor demand curve. If this is not the case then the estimate of the wage elasticity will be biased. A prominent case in which there would be bias is if labor supply to particular occupations is not perfectly elastic and wages clear markets. However, it is not entirely clear what is the appropriate view to take about the elasticity in the supply of labor to different occupations.

Studies that segment the labor market by education typically assume that the supply of different skills is inelastic in the short-run because individuals primarily fix their education level at the beginning of their working life. But, occupational mobility is higher than educational mobility – especially so for those occupations without much specific human capital – so we would not expect the elasticity in the supply of labor to many occupations to be totally inelastic even in the short-run. In the long-run, studies like Goldin and Katz (2008) suggest that the supply of labor of different education levels is very elastic and the same is probably true of supply to different occupations.

Given these conceptual issues, it makes sense to look at evidence on whether technological progress and offshoring seem to affect occupational wages. Autor, Katz and Kearney (2008), Lemieux (2008) and Autor and Dorn (2011) find a positive correlation between employment polarization and wage growth across US occupations. In line with this finding, Firpo, Fortin and Lemieux (2010) and Acemoglu and Autor (2011) argue that routinization has had non-trivial effects on US wage inequality. However, the evidence is much less clear for European countries. Dustmann, Ludsteck and Schönberg (2009) find that occupational employment and wage growth during the 1990s in Germany are only weakly positively correlated across all occupations and even negatively correlated across occupations paying below the median. Goos and Manning (2007) report a similar finding for the UK for the period 1975-1999. So we investigate the potential endogeneity of wages in our data in a number of ways.

First, we estimate the specification in columns (3) and (4) of Table 6 while instrumenting the wage using demographic changes in labor supply as an exclusion restriction (Dustmann, Ludsteck and Schönberg 2009 suggest such supply changes might explain the wage changes they observe in Germany). Specifically, we use as an instrument counterfactual occupational employment only accounting for economy wide changes in employment by gender-migration while keeping the occupational composition within each gender-migration combination constant over time. This provides consistent estimates as long as economy wide changes in female labor force participation and immigration are not correlated with occupation specific changes in labor demand. The first-stage coefficients reported in columns (1) and (2) of Table 7 are negative and significant. Estimates of the impact of technological progress and offshoring, industry marginal costs (suggesting $1/(1-\eta) = 0.66$ with a standard error of 0.09) and wages are also as expected, except that the wage elasticity is estimated to be much larger, which perhaps is an indication of the poor quality of European wage data.

Secondly, we can simply estimate models excluding the wage and possibly including counterfactual labor supply which can be thought of – perhaps somewhat loosely – as a reduced form specification. For example, columns (3) and (4) of Table 7 exclude wages as well as industry marginal costs and industry output – as these are influenced by wages – showing it does not qualitatively affect the point estimates on technological progress and offshoring.

In sum, our data suggest that relative occupational wage movements in Europe are not strongly correlated with our technology and offshoring variables. This result differs from evidence for the US but is not necessarily inconsistent with it since many European countries have institutions (e.g. minimum wages and collective bargaining) that mute or stop a wage response, especially across middling and lower-paying occupations. Hence, we proceed by assuming that relative wages are exogenous.

C. Estimates of Structural Parameters

As estimates of the model's key structural parameters we use the estimates from column (4) of Table 6 or column (2) of Table 7 to give some idea of robustness. These two estimates primarily differ in terms of the estimated wage elasticity being below one from Table 6 and above one in Table 7 – as we shall see this has important consequences. Referring back to equation (7) the estimate of $1/(1-\eta)$, the elasticity of substitution between tasks in goods production, comes from the estimate of the coefficient on industry marginal costs. The elasticity of substitution between inputs in task production, $1/(1-\rho)$, can then be deduced using the estimated wage elasticity, the estimated $1/(1-\eta)$ and an estimate of the share of labor in task costs, κ. We set κ equal to 0.54 which is the share of labor costs in value added using STAN data.

Using point estimates from the last column of Table 6 or from the second column of Table 7, we get that the estimate of $1/(1-\eta)$ is 0.53 (with a standard error of 0.13) or 0.66 (with a standard error of 0.09). The respective imputed estimates of $1/(1-\rho)$ are 1.20 and 9.09. Note that the value of $1/(1-\rho)$ is less precisely estimated due to different estimates of the wage elasticity in Tables 6 and 7, as we explained above. In the analysis that follows, we will therefore allow for both $1/(1-\eta) = 0.53$ and $1/(1-\rho) = 1.20$ or $1/(1-\eta) = 0.66$ and $1/(1-\rho) = 9.09$.

Are these estimates reasonable? The problem is that there exist no other studies with estimates of these elasticities. At best, we can look at different but distinctly related numbers for guidance. For example, Katz and Murphy (1992) find an elasticity of substitution between high school and college equivalent men and women of 1.4. In line with this, Card and Lemieux (2001) find an estimate between 1.1 and 1.6 for men and women and between 2 and 2.5 for men only. Card and Lemieux (2001) also report an elasticity of substitution around 5 between more disaggregate five-year age groups with given levels of education. In this light, our estimates do not seem too unrealistic.

D. Are Technological Progress and Offshoring Labor- or Capital-Augmenting?

As discussed earlier, the estimates from the conditional labor demand equation do not allow us to separately identify the effects of technological progress and offshoring on labor- and capital-augmenting productivity changes. However, it would be helpful to be able to do so because the appropriate way to model technological progress and offshoring has been a subject of some discussion and, as we show later, one needs to be able to do this if one wants to take account of industry demand shifts in accounting for changes in the occupational structure of employment. This section discusses how we might be able to make progress on this question. The ideal way to do this would be to estimate demand curves for the other inputs (capital or foreign employment) apart from domestic labor. But it is hard to obtain the data necessary to do that as, for example, one would need the capital used in specific tasks. This section shows how one can identify the labor- and capital- augmenting change using industry marginal costs.

From equation (7), and the equations for labor- and capital-augmenting productivity (8) and (9) one can readily see that from the conditional labor demand equation one obtains an estimate for the effect of task content of:

$$\left[\frac{1-\kappa}{1-\rho}+\frac{\kappa}{1-\eta}-1\right]\gamma_R^N+\left[\frac{1}{1-\rho}-\frac{1}{1-\eta}\right](1-\kappa)\gamma_R^K \tag{10}$$

so that γ_R^N and γ_R^K are not identified even if, from previous arguments, we do have estimates of all other parameters. Similarly, from the conditional labor demand equation one can only get an estimate for the effect of offshorability of:

$$\left[\frac{1-\kappa}{1-\rho}+\frac{\kappa}{1-\eta}-1\right]\gamma_O^N+\left[\frac{1}{1-\rho}-\frac{1}{1-\eta}\right](1-\kappa)\gamma_O^K \tag{11}$$

which does not identify γ_O^N and γ_O^K separately.

But, now consider a log linearization of equations (3) and (6) which gives the following approximation for industry marginal costs:

$$\begin{aligned}\log c_{ict}^I \approx\ & \kappa\log w_{ict}+\left[(1-\kappa)\gamma_R^K-\kappa\gamma_R^N\right]\sum_{j=1}^{J}\kappa_{j|i}R_j\times \textit{linear time trend}\\ & +\left[(1-\kappa)\gamma_O^K-\kappa\gamma_O^N\right]\sum_{j=1}^{J}\kappa_{j|i}O_j\times \textit{linear time trend}\\ & -\kappa\sum_{j=1}^{J}\gamma_j^N+(1-\kappa)\sum_{j=1}^{J}\gamma_j^K-\log\beta_{ic}\end{aligned} \tag{12}$$

where $\kappa_{j|i}$ is modeled as the observed share for employment of occupation j in industry i and R_j is the RTI index. What this equation says is that if different industries use a different mix of routine tasks, this translates to differences in the average level of routinization across industries and this is something we can exploit in estimation as routine-intensive industries should, ceteris paribus, have lower increases in costs. Specifically, the coefficient on the industry-level routinization measure in equation (12) is an estimate of:

$$(1-\kappa)\gamma_R^K - \kappa\gamma_R^N \tag{13}$$

and the coefficient on the industry-level measure of offshorability is an estimate of:

$$(1-\kappa)\gamma_O^K - \kappa\gamma_O^N \tag{14}$$

Estimates of (10) and (13) together then allow to separate out the importance of labor- (γ_R^N) and capital-augmenting (γ_R^K) productivity changes due to technological progress. A similar argument applies for the impact of offshorability by combining estimates of (11) and (14) to see whether offshoring is predominantly labor-(γ_O^N) or capital-augmenting (γ_O^K).

Estimation of equation (12) leads to the following point estimates for (13) and (14):

$$\log c_{ict}^I = [\underset{(0.04)}{-0.12} \sum_{j=1}^{J} \kappa_{j|i} R_j \underset{(0.04)}{-0.01} \sum_{j=1}^{J} \kappa_{j|i} O_j] \times linear\ time\ trend + D_{ic,t} \tag{15}$$

with $D_{ic,t}$ a vector of industry-country dummies and time dummies. The estimates suggest that more routine and offshorable industries have had lower growth in marginal costs, although the impact of offshoring is small and not statistically significant suggesting that there is insufficient variation in offshorability across industries to precisely estimate its effect on industry marginal costs. We can now combine these estimates with the estimates of η and ρ described earlier to estimate the labor-augmenting and capital-augmenting changes.

The results are presented in Table 8 where column (1) is based on estimates from column (4) of Table 6 and column (2) is based on estimates from column (2) of Table 7. In column (1) – when the wage elasticity is less than one – the estimates imply that γ_R^N is negative i.e. that labor in more routinizable occupations is becoming relatively less productive over time. However, if the wage elasticity is less than one, the estimated impact of task content on conditional labor demand in equation (7) can only be explained by a relative increase and not decrease in the productivity of routine labor. A negative estimate of γ_R^N is also less plausible given (13) and the negative point estimate of routine task content in equation (15). The estimate of γ_R^K in column (1) of Table 8 is large and negative implying that capital used in routine tasks is becoming relatively cheaper over time. Given the elasticities of substitution, this is consistent with the estimated impacts of routine task content in equations (7) and (15). There are similarly signed effects for the offshoring variables though these estimates are not precise so not much weight should be put on them.

Column (2) of Table 8 does the same exercise but using a wage elasticity larger than one associated with the estimates of the conditional labor demand curve in column (2) of Table 7. The estimated impact of γ_R^N now is very close to zero. The estimated impact of γ_R^K is much larger and implies that, just as in column (1), capital used in routine tasks is becoming relatively cheaper over time and this is in support of the evidence found in equations (7) and (15). Note that the magnitude of the effects in column (2) is generally smaller than in column (1) but this is simply the result of the fact that when the wage elasticity is large, smaller changes in factor-augmenting technological change are needed to explain the same change in employment. The estimates for the effects of offshoring are very small though, again, we do not have much precision here.

Overall the estimates in Table 8 suggest that it is best to model routinization as coming from capital- rather than labor-augmenting technological change. This is a believable and important conclusion that is consistent with, for example, the treatment in Autor, Levy and Murnane (2003). Therefore, the estimates in column (2) of Table 8 are more plausible than those in column (1). But we continue to report results based on the estimates in both columns of Table 8 to give some idea of robustness of conclusions.

This section has shown that routinization has effects on relative marginal costs. From this we would expect effects on relative prices and from that effects on the demands for different sorts of labor. We now turn to quantifying these effects.

VI. Accounting for Product Demand Effects

All our estimates so far have been of the demand for labor conditional on industry output. While that is an interesting exercise well-suited to isolating the effects of technological progress and globalization on production functions, it does have limitations when it comes to our ultimate aim – explaining the phenomenon of job polarization in the economy as a whole – because it cannot take account of the derived demand effects induced from shifts in the demand for the outputs of different industries that we would expect to follow from routinization and offshoring. For example, in case of technological progress, equation (15) suggests that routinization will result in larger falls in prices in industries that historically used a lot of routine labor, and this will tend to benefit all labor that is used in these industries.[22]

Another possibility is that relative product demand shifts because preferences are non-homothetic since different income elasticities for different goods or services yield structural change even if productivity growth is balanced across tasks or sectors. Non-homothetic preferences also imply that changes in the distribution of income will lead to changes in relative product demand. A related but distinct hypothesis, proposed by Manning (2004) and Mazzolari and Ragusa (2008), is that rising wage inequality will cause high-wage workers to demand more low-skill service work so as to free up more of their time for market work. However, Autor and Dorn (2011) fail to find much evidence in support of this hypothesis.

[22]In the hamburger example from the introduction, the improvement in patty-making reduced the price of hamburgers inducing a rise in the demand for patty-makers as well as bun-makers. That industry shifts can be quantitatively important also follows from decomposing aggregate occupational employment share changes into within- and between-industry components. See Appendix H for details.

A. Modeling Product Demand

To analyze these effects, we need to go further and not condition on industry output in equation (7). To do this, we need an industry demand curve. We will start by deriving this from a demand curve at the individual level. Assume that individual k has income Z_k and that the demand for the output of industry i by individual k is given by:

$$Y_{ik} = Z_k^{\theta_i} P_i^{-\frac{1}{1-\gamma}} \tag{16}$$

with θ_i the income elasticity of demand for good i which will be equal to unity if preferences are homothetic but not otherwise; P_i the relative price of good i; and $1/(1-\gamma)$ the elasticity of substitution between goods in consumption with $\gamma < 1$. Where $\theta_i \neq 1$ for all industries one should acknowledge that this formulation of the demand curve does not satisfy the budget constraint in which case equation (16) is best thought of as a local approximation to the demand curve that will not be too bad if, as we find below, departures from homotheticity are fairly small.

If there are L individuals in the economy and we assume that income has a log-normal distribution with variance σ^2, we can add up the individual demand curves over all individuals in the economy to arrive at the following log aggregate demand equation (also adding time (t) and country (c) subscripts and rearranging terms):

$$\begin{aligned} \log Y_{ict} &= -\frac{1}{1-\gamma}\log P_{ict} + \theta_i \log Y_{ct} + (1-\theta_i)\log L_{ct} \\ &\quad + \frac{1}{2}\theta_i(\theta_i - 1)\log\sigma^2_{ct} \end{aligned} \tag{17}$$

with Y_{ict} aggregate demand for good i and Y_{ct} aggregate income in country c at time t.

Equation (17) shows that there are a number of ways to test for homotheticity – that is, $\theta_i = 1$ – in the data. Firstly, homotheticity implies that the elasticity of industry demand with respect to aggregate income must be unity because expenditure shares on goods are constant. Secondly, homotheticity implies that, conditional on aggregate income, population does not affect product demand. The intuition for this is simple: if preferences would be non-homothetic instead and we compare two economies with the same aggregate GDP but with different populations, the economy with the lower population will have a higher demand for luxury goods as GDP per capita is higher there. Thirdly, only if preferences are homothetic, we would expect income inequality not to affect product demand. Again the intuition for this is straightforward: if preferences would be non-homothetic instead and we compare two economies with the same average GDP per capita but different income inequality, the economy with more inequality will have a higher demand for luxury goods. We test these three predictions below.

Given that equation (17) enters equation (7) additively, we can estimate it separately to derive the impact of changes in relative product demand on employment. Specifically, equation (17) shows that if technological progress and offshoring decrease relative marginal costs and therefore relative prices in routine and offshorable intensive industries, their relative product demand and therefore labor demand increases. Moreover, this relative increase is increasing in the elasticity of substitution between goods in consumption, $1/(1-\gamma)$, another key structural parameter of our model.

B. Estimating Product Demand Curves

Table 9 shows results of estimating the following product demand equation:

$$\begin{aligned} \log Y_{ict} &= \delta_0 + \delta_1 \log P_{ict} + \delta_{2i} \log Y_{ct}/L_{ct} + \delta_3 \log L_{ct} \\ &\quad + \delta_{4i} \log \sigma^2_{ct} + \xi_{ict} \end{aligned} \tag{18}$$

where δ_1 is an estimate of $-1/(1-\gamma)$; under homothetic preferences it must hold that $\delta_3 = 1$ and $\delta_{2i} = 1$, $\delta_{4i} = 0$ for all i; and ξ_{ict} is an error term.[23]

Column (1) of Table 9 shows estimates if we assume preferences are homothetic – that is, we set $\delta_3 = 1$; $\delta_{2i} = 1$ and $\delta_{4i} = 0$ for all i. As predicted by our model, the point estimate on the industry output price is negative and significant suggesting that product demand at our level of aggregation is inelastic with an elasticity of substitution between goods in consumption, $1/(1-\gamma)$, of 0.31 with a standard error of 0.13.

Column (2) shows estimates for the unconstrained specification, which includes interactions of log income per capita and a measure of income inequality with industry dummies. This allows us to test for homotheticity in the data rather than assuming it.

Firstly, we can test whether population has a neutral impact on product demand, which would be the case if $\delta_3 = 1$. The assumption of homotheticity cannot be rejected as the point estimate on population is 1.00 with a standard error of 0.05.

Secondly, column (2a) reports estimates of the income elasticities for the different industries, given by the estimated coefficients on interactions of GDP per capita with industry dummies (δ_{2i}). If these elasticities are not statistically different from unitary, we cannot reject the assumption that preferences are homothetic. It can be seen that there are some deviations from unitary, but these are generally small in size and not statistically significant. In particular, the only estimated elasticity that is significantly higher (namely 1.12) is for Manufacturing. Even taking the estimated elasticities at face value, however, does not provide evidence that non-homotheticity is a good candidate for explaining economy wide job polarization. To see why, note that the industries are ranked from high- to low-paying by their mean UK wage in 1994: if anything, the lowest- and highest-paying industries exhibit lower income elasticities than the middling ones.

Finally, Table 9 also tests the third implication of homotheticity: that income inequality has no impact on product demand. To that end, a measure of income inequality – that is, log(p90/p10) – interacted with industry dummies is included to capture the possibility that the demand for low-paid services partially reflects the higher demand for such

[23] The relative price of good i, P_{ict}, is the industry-country-year specific price index for production from STAN mentioned in Section IV relative to the country-year specific producer price index obtained from the OECD Main Economic Indicators data. Aggregate income or GDP, Y_{ct}, and population, L_{ct}, are also taken from the OECD Main Economic Indicators data. The measure used for inequality, $\log \sigma^2_{ct}$, is the country-year varying log(p90/p10) derived from the ECHP and EU-SILC data discussed in Section IV.

services from high-income earners when income inequality in a country increases over time.[24] Column (2b) reports the found elasticities by industry. It can be seen that six out of ten elasticities are not significantly different from zero, implying that product demand for those industries is not responsive to changes in income inequality. The elasticities that are significantly different from zero are for Financial intermediation; Transportation, storage and communication; Construction; and Health and social work. However, it should be noted that the magnitudes of the estimates in Table 9 are for a one log point increase in (p90/p10) whereas its standard deviation is only 0.21 in our data. Moreover, this includes cross-country variation in wage inequality whereas technological progress and offshoring can only predict variation in wage inequality over time: the standard deviation of wage inequality within a country over time is only 0.06 over our 14-year period. For example, a one standard deviation increase in inequality between 1993 and 2006 is associated with a $0.29 \times 0.06 = 1.74\%$ faster increase in the demand for health and social work over this period. These effects are small compared to the estimated impacts of technological progress and offshoring shown earlier. For instance, technological progress was found to increase relative labor demand within industries for occupations that are one standard deviation more intense in non-routine tasks by more than 1% per annum. In sum, the evidence does not strongly support the idea that changes in aggregate income or income dispersion – possibly following technological progress and globalization – play an important part in explaining changes in relative employment.

VII. How Much of Job Polarization Can Be Explained?

In this section we examine to what extent our framework can explain job polarization in Europe. To do so, we compare the observed changes in the job structure documented in Section II with a variety of counterfactuals constructed from our model in which we turn off and on the different channels of influence that we have discussed. What we are doing can be thought of as a form of shift-share analysis, something that is normally thought of as a useful accounting device but we root it in our theoretical model of labor demand. In all counterfactuals shown below we assume, in the interests of keeping results to a digestible length and in line with our findings above, that relative occupational wages in Europe have been rigid so can be treated as exogenous and that preferences are homothetic.

A. The Effects of Technology and Globalization on Occupational Employment

In this section, our aim is to compare the predictions of our framework for the changes in the occupational structure of employment with the actual employment share changes presented in Table 1. To do this we work out the predictions for the level of employment by industry-occupation, N_{ijct}, and then aggregate to get employment by occupation, N_{jct}. From this we can work out employment shares of different occupations for

[24]Our results are robust to using other measures of inequality (log(p90/p50) or log(p50/p10), although log(p90/p10) is our preferred measure since it better reflects changes in the buying power of high-income earners over the services provided by low-income earners.

each country-year. Note from equation (7) that one can write employment by industry-occupation as the product of an occupation-specific component (denote this by G_{jct}) and an industry-occupation specific component (denote this by G_{ijct}). So, total employment by occupation-country-year can be written as:

$$N_{jct} = \sum_{i=1}^{I} N_{ijct} = G_{jct} \sum_{i=1}^{I} G_{ijct} \tag{19}$$

Taking logs and differentiating over time gives:

$$\begin{aligned} \frac{\partial \log N_{jct}}{\partial t} &= \frac{\partial \log G_{jct}}{\partial t} + \frac{1}{\sum_{i=1}^{I} G_{ijct}} \sum_{i=1}^{I} G_{ijct} \frac{\partial \log G_{ijct}}{\partial t} \\ &= \frac{\partial \log G_{jct}}{\partial t} + \sum_{i=1}^{I} \kappa_{i|j} \left[\frac{\partial \log G_{ict}}{\partial t} \right] \end{aligned} \tag{20}$$

where $\kappa_{i|j}$ is the share for industry i of employment in occupation j. The first term in equation (20) captures changes in the demand for occupation j within each industry, which in our model are due to technological progress and offshoring. The terms in square brackets capture changes in the demand for occupation j between industries, which in our model exist because of changes in industry marginal costs and industry output – also ultimately due to technological progress and offshoring. It is this decomposition that we will use to examine whether or not the different channels in our model can go a substantial distance towards explaining job polarization.

To see this more clearly, let us start by considering the first term on the right-hand side of equation (20), $\partial \log G_{jct}/\partial t$, in more detail. Given equations (7), (8) and (9), we can write this as:

$$\begin{aligned} \frac{\partial \log G_{jct}}{\partial t} &= \left[\frac{1-\kappa}{1-\rho} + \frac{\kappa}{1-\eta} - 1 \right] \left[\gamma_R^N R_j + \gamma_O^N O_j \right] \\ &\quad + \left[\frac{1}{1-\rho} - \frac{1}{1-\eta} \right] (1-\kappa) \left[\gamma_R^K R_j + \gamma_O^K O_j \right] \end{aligned} \tag{21}$$

where all terms on the right-hand side of equation (21) are known from the analysis above.

Using equations (7) and (17), the terms in square brackets in equation (20) can be written as:

$$\begin{aligned} \frac{\partial \log G_{ict}}{\partial t} &= \frac{1}{1-\eta} \frac{\partial \log c_{ict}^I}{\partial t} + \frac{\partial \log Y_{ict}}{\partial t} \\ &= \frac{1}{1-\eta} \frac{\partial \log c_{ict}^I}{\partial t} - \frac{1}{1-\gamma} \frac{\partial \log c_{ict}^I}{\partial t} \end{aligned} \tag{22}$$

where the second equality makes the additional assumption that output prices are a constant mark-up on marginal costs. Note that the right-hand side of equation (22) is known given that $\partial \log c_{ict}^{I}/\partial t$ can be obtained by differentiating equation (15) over time.

B. Explaining Job Polarization

In this section we use the framework described above to quantify our predictions about the change in the structure of employment over our sample period, 1993-2006. The results are reported in Table 10. There we report the predicted changes in the employment shares for the 4 lowest-paid occupations, the 9 middling occupations and the 8 highest-paying occupations – just as in Table 2 and this is simply a way to make the results digestible. We also report predictions for the two sets of estimates of the structural parameters reported in columns (1) and (2) of Table 8 respectively.

Column (1) of Table 10 shows the predictions of the model assuming output does not respond to any of the changes in industry marginal costs. As one can see the combination of routinization and offshoring predicts polarization – a rising employment share for the lowest- and highest-paid occupations and a fall for the middling occupations. This is what one would expect given that Table 4 shows that more routine and offshorable occupations are concentrated around the median wage. The predicted polarization is larger than that observed for the structural parameters with a low wage elasticity of labor demand (the first row in each panel) but smaller for our preferred estimates with a higher wage elasticity (the second row in each panel). The predictions in column (2) of Table 10 turn off the offshoring effects so that all that remains is routinization. As the predictions do not vary very much from those reported in column (1), this suggests that it is routinization that is driving most of the observed changes.

Columns (3) and (4) of Table 10 quantify, in addition, the changes in employment shares due to the induced shares of output. Column (4) is based on a low price-elasticity of demand, column (4) on a higher one. Comparing to column (1) one observes that the predicted polarization is attenuated, because the decrease in the relative output prices of more routine and offshorable intensive goods, increases relative product and therefore relative labor demand for all occupations in those industries. However this attenuation effect is smaller when the wage elasticity is high, our preferred estimate. The model can explain 0.62 of the of 1.58 percentage point increase in the employment share of the low-paid occupations, 5.7 percentage points of the 7.8 percentage point decrease in the employment share of the middling occupations and 5.0 percentage points of the 6.2 percentage point rise in the employment share of the high-paid occupations. That is, our model can explain the bulk of observed polarization though there remains scope for a richer model to explain all of it.

VIII. Conclusions

The employment structure in Western European countries has been polarizing between 1993 and 2006 with rising employment shares for high-paid professionals and managers as well as low-paid personal services workers and falling employment shares of manufacturing and routine office workers. Together with existing evidence for the US, this establishes job polarization as a pervasive feature of labor markets in advanced economies.

In this paper we develop a model of the labor market with different tasks and industries that can explain job polarization. We also show how it can be used to assess the importance of technological change, globalization and institutions on the demand for different occupations and to decompose the influence into effects within industries and due to shifts in the demand for the products of different industries themselves caused by changes in relative industry marginal costs induced by technological progress and offshoring.

We find that the most important factor behind the observed changes is the routinization hypothesis of Autor, Levy and Murnane (2003), that the impact of routinization is capital-augmenting. We also find that offshoring does predict job polarization, although its impact is much smaller than that of routinization. We find that relative wages in Europe have not changed in response to polarization. We find that shifts in demand across industries act to attenuate but not eliminate job polarization. Overall we find that our simple model can explain the bulk – though not all – of the observed polarization.

References

Acemoglu, Daron and David Autor. 2011. “Skills, Tasks and Technologies: Implications for Employment and Earnings”, *Handbook of Labor Economics*, Orley Ashenfelter and David E. Card (eds.), Amsterdam: Elsevier, Vol 4B: 1043-1171.

Autor, David, and David Dorn. 2011. “Inequality and Specialization: The Growth of Low-Skill Service Jobs in the United States”, mimeo MIT.

Autor, David H. and Michael Handel. 2009. “Putting Tasks to the Test: Human Capital, Job Tasks and Wages”, mimeo MIT.

Autor, David H. and Lawrence F. Katz. 1999. “Changes in the Wage Structure and Earnings Inequality”, *Handbook of Labor Economics*, Orley Ashenfelter and David E. Card (eds.), Amsterdam: Elsevier, Vol 3A: 1463-1555.

Autor, David H., Katz, Lawrence F. and Melissa S. Kearney. 2006. “The Polarization of the US Labor Market”, *American Economic Review Papers and Proceedings*, 96: 189-194.

Autor, David H., Katz, Lawrence F. and Melissa S. Kearney. 2008. “Trends in U.S. Wage Inequality: Revising the Revisionists”, *Review of Economics and Statistics*, 90: 300-23.

Autor, David H., Levy, Frank, and Richard J. Murnane. 2003. “The Skill-Content of Recent Technological Change: An Empirical Investigation”, *Quarterly Journal of Economics*, 118: 1279-1333.

Baumol, William J. 1967. “Macroeconomics of Unbalanced Growth: Anatomy of an Urban Crisis”, *American Economic Review*, 57: 415-426.

Blinder, Alan and Alan Krueger. 2009. “Alternative Measures of Offshorability: A Survey Approach”, NBER Working Paper No. 15287.

Card, David and Thomas Lemieux. 2001. "Can Falling Supply Explain The Rising Return To College For Younger Men? A Cohort-Based Analysis," *Quarterly Journal of*

Economics, 116: 705-746.

Dustmann, Christian, Ludsteck, Johannes and Uta Schönberg. 2009. “Revisiting the German Wage Structure”, *Quarterly Journal of Economics*, 124: 809-842.

Firpo, Sergio, Nicole Fortin and Thomas Lemieux. 2010. “Occupational Tasks and Changes in the Wage Structure”, University of British Columbia Working Paper.

Goldin, Claudia and Lawrence F. Katz. 2008. The Race Between Education and Technology, Cambridge, Mass.: Harvard University Press.

Goos, Maarten and Alan Manning. 2007. “Lousy and Lovely Jobs: The Rising Polarization of Work in Britain”, *Review of Economics and Statistics*, 89: 118-133.

Goos, Maarten, Alan Manning and Anna Salomons. 2009. “The Polarization of the European Labor Market”, *American Economic Review Papers and Proceedings*, 99:58-63.

Grossman, Gene M. and Esteban Rossi-Hansberg. 2008. “Trading Tasks: A Simple Theory of Offshoring”, *American Economic Review*, 98: 1978-1997.

Katz, Lawrence F. and Kevin M. Murphy. 1992. "Changes in Relative Wages, 1963-1987: Supply and Demand Factors," *Quarterly Journal of Economics*, 107: 35-78.

Krugman, Paul. 1999. The Accidental Theorist And Other Dispatches from the Dismal Science, W.W. Norton and Company, New York.

Lemieux, Thomas. 2008. “The Changing Nature of Wage Inequality”, *Journal of Population Economics*, 21: 21-48.

Manning, Alan. 2004. “We Can Work it Out: the Impact of Technological Change on the Demand for Low-Skilled Workers”, *Scottish Journal of Political Economy*, 51: 581-608.

Mazzolari, Francesca and Giuseppe Ragusa. 2008. "Spillovers from high-skill consumption to low-skill labor markets", University of California at Irvine Working Paper.

Michaels, Guy and Natraj, Ashwini and John Van Reenen. 2010. “Has ICT polarized skill demand?: evidence from eleven countries over 25 Years”. Centre for Economic Performance Discussion Paper, No. 987.

Spitz-Oener, Alexandra. 2006. “Technical Change, Job Tasks and Rising Educational Demand: Looking Outside the Wage Structure”, *Journal of Labor Economics*, 24: 235-270.

TABLE 1—LEVELS AND CHANGES IN THE SHARES OF HOURS WORKED FOR OCCUPATIONS RANKED BY THEIR MEAN 1993 EUROPEAN WAGE

	ISCO code	Average employment share in 1993 (%)	Percentage point change over 1993-2006
Corporate managers	12	4.47	1.23
Engineering professionals	21	2.94	1.02
Life science and health professionals	22	1.96	-0.12
Other professionals	24	2.80	0.65
Managers of small enterprises	13	3.53	1.25
Engineering associate professionals	31	3.96	0.87
Other associate professionals	34	6.85	2.15
Life science and health associate professionals	32	3.05	0.69
Drivers and mobile plant operators	83	5.37	-0.19
Stationary plant and related operators	81	1.71	-0.38
Metal, machinery and related trade work	72	8.15	-2.29
Precision, handicraft, craft printing and related trade	73	1.29	-0.40
Office clerks	41	11.96	-1.94
Customer service clerks	42	1.97	0.18
Extraction and building trade workers	71	7.98	-0.51
Machine operators and assemblers	82	6.55	-1.96
Other craft and related trade workers	74	3.13	-1.35
Personal and protective service workers	51	7.10	1.11
Laborers in mining, construction, manufacturing and	93	4.03	0.45
Models, salespersons and demonstrators	52	6.56	-1.38
Sales and service elementary occupations	91	4.65	0.89

Notes: Results are for employment pooled across all 16 European countries by occupation. The percentage point change over 1993-2006 is the long difference where employment shares in 1993 and 2006 are imputed on the basis of average annual growth rates for countries with shorter data spans. Occupations are ordered by their mean wage rank in 1993 across the 16 European countries.

TABLE 2—INITIAL SHARES OF HOURS WORKED AND PERCENTAGE CHANGES OVER TIME FOR GROUPS OF HIGH-PAYING, MIDDLING AND LOW-PAYING OCCUPATIONS

	4 Lowest paying occupations		*9 Middling occupations*		*8 Highest paying occupations*	
	Employment share in 1993 (%)	Percentage point change 1993-2006	Employment share in 1993 (%)	Percentage point change 1993-2006	Employment share in 1993 (%)	Percentage point change 1993-2006
Austria	22.72	-0.59	52.54	-14.58	24.74	15.17
Belgium	17.48	1.48	48.52	-9.50	34.00	8.03
Denmark	24.14	-0.96	39.70	-7.16	36.16	8.13
Finland	17.98	6.66	38.69	-6.54	43.33	-0.12
France	22.10	-0.74	47.91	-12.07	29.98	12.81
Germany	22.35	3.04	55.88	-8.72	21.77	5.67
Greece	21.65	1.75	47.80	-6.08	30.55	4.34
Ireland	19.32	6.19	46.11	-5.47	34.57	-0.72
Italy	26.95	-8.20	51.13	-9.08	21.92	17.28
Luxembourg	21.69	-1.66	50.07	-8.45	28.23	10.10
Netherlands	16.78	2.27	37.88	-4.68	45.33	2.41
Norway	22.64	4.96	38.87	-6.52	38.49	1.57
Portugal	25.75	2.39	47.49	-1.13	26.77	-1.26
Spain	28.02	0.96	48.68	-7.04	23.30	6.07
Sweden	21.50	1.91	41.94	-6.96	36.55	5.04
UK	16.86	5.77	43.68	-10.32	39.46	4.55
EU Average	*21.75*	*1.58*	*46.06*	*-7.77*	*32.20*	*6.19*

Notes: Occupational employment is pooled by occupation group within each country according to the mean 1993 European occupational wage rank.

TABLE 3—MEASURES OF TECHNOLOGY AND OFFSHORING FOR OCCUPATIONS RANKED BY THEIR MEAN 1993 EUROPEAN WAGE

	ISCO code	Abstract task importance (1)	Routine task importance (2)	Service task importance (3)	Routine Task Intensity (RTI) (4)	Offshorability (5)	Mean education level (6)
Corporate managers	12	1.80	-1.18	1.15	-1.29	-0.59	2.05
Engineering professionals	21	1.50	-0.86	-0.35	-0.80	-0.37	2.83
Life science and health professionals	22	1.47	-0.16	1.73	-0.81	-0.64	2.92
Other professionals	24	1.29	-1.63	1.14	-1.49	-0.51	2.69
Managers of small enterprises	13	1.80	-1.18	1.15	-1.29	-0.59	2.05
Engineering associate professionals	31	0.89	0.20	-0.44	-0.02	-0.27	2.22
Other associate professionals	34	0.75	-1.37	0.93	-1.25	-0.12	2.14
Life science and health associate professionals	32	0.36	0.21	0.86	-0.26	-0.64	2.40
Drivers and mobile plant operators	83	-0.59	1.33	0.01	0.90	-0.63	1.46
Stationary plant and related operators	81	-0.49	1.33	-1.21	1.38	1.63	1.56
Metal, machinery and related trade work	72	0.43	1.16	-0.29	0.65	0.29	1.68
Precision, handicraft, craft printing and related	73	-1.30	0.81	-1.79	1.51	-0.62	1.69
Office clerks	41	-0.42	-1.29	0.04	-0.89	1.21	1.91
Customer service clerks	42	-0.36	-0.82	0.74	-0.75	-0.27	1.89
Extraction and building trade workers	71	-0.23	0.98	-0.64	0.82	-0.59	1.55
Machine operators and assemblers	82	-0.46	1.31	-1.33	1.41	3.18	1.48
Other craft and related trade workers	74	-1.36	0.67	-1.30	1.18	-0.27	1.57
Personal and protective service workers	51	-0.37	-0.16	0.82	-0.35	-0.64	1.67
Laborers in mining, construction, manufacturing	93	-1.00	0.52	-0.53	0.64	0.87	1.41
Models, salespersons and demonstrators	52	-0.53	-0.94	1.00	-0.86	-0.64	1.66
Sales and service elementary occupations	91	-1.38	-0.11	-0.55	0.28	-0.37	1.40

Notes: Columns (1)-(4) show task content measures rescaled to mean 0 and standard deviation 1, a higher value means a task is more important. Values for ISCO 12 and 13 are identical because ONET SOC codes do not allow distinction. Routine Task Intensity (RTI) is defined as Routine task importance divided by the sum of Abstract and Service task importance, subsequently standardized. Column (5) shows our offshorability measure rescaled to mean 0 and standard deviation 1, a higher value means an occupation is more offshorable. Values for ISCO 12 and 13 have been made the same by taking the mean weighted by hours worked. Column (6) shows mean education levels weighted by hours worked and averaged across countries using the first year in which education data was available (typically 1999). Level 1 is up to and including lower secondary education, level 2 is upper secondary and post-secondary (non-tertiary) education, and level 3 is tertiary or post-graduate education. Values for ISCO 12 and 13 have been made the same by taking the mean weighted by hours worked. Occupations are ordered by their mean wage rank in 1993 across the 16 European countries.

TABLE 4—OCCUPATIONAL WAGES IN 1993 AND 2006 RANKED BY THEIR MEAN 1993 EUROPEAN WAGE

	ISCO code	*Real monthly wages (in 2000 Euros)* 1993 (1)	2006 (2)	Percentage growth over 1993-2006 [(2)-(1)]/(1)
Corporate managers	12	3,472	3,724	7.26
Engineering professionals	21	3,038	3,170	4.34
Life science and health professionals	22	2,720	3,164	16.32
Other professionals	24	2,712	2,910	7.30
Managers of small enterprises	13	2,653	2,685	1.21
Engineering associate professionals	31	2,150	2,324	8.09
Other associate professionals	34	2,115	2,227	5.29
Life science and health associate professionals	32	1,915	2,018	5.38
Drivers and mobile plant operators	83	1,793	1,916	6.86
Stationary plant and related operators	81	1,789	1,954	9.22
Metal, machinery and related trade work	72	1,748	1,927	10.24
Precision, handicraft, craft printing and related	73	1,733	1,968	13.56
Office clerks	41	1,679	1,865	11.08
Customer service clerks	42	1,624	1,732	6.65
Extraction and building trade workers	71	1,613	1,750	8.49
Machine operators and assemblers	82	1,565	1,728	10.41
Other craft and related trade workers	74	1,504	1,598	6.25
Personal and protective service workers	51	1,424	1,538	8.00
Laborers in mining, construction, manufacturing	93	1,402	1,518	8.27
Models, salespersons and demonstrators	52	1,237	1,344	8.65
Sales and service elementary occupations	91	1,112	1,242	11.69

Notes: Mean occupational wages weighted by weekly hours worked in each country in 1993 and 2006 and averaged across countries.

TABLE 5—EFFECTS OF TECHNOLOGY AND OFFSHORING ON OCCUPATION SPECIFIC TREND CHANGES IN EMPLOYMENT
DEPENDENT VARIABLE: LOG(HOURS WORKED/1000)

Time trend interacted with:	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
Abstract	1.33*			0.80*		1.22*			0.81*	
	(0.12)			(0.15)		(0.12)			(0.15)	
Routine		-1.33*		-0.74*			-1.22*		-0.75*	
		(0.12)		(0.15)			(0.12)		(0.15)	
Service			1.28*	0.31				1.19*	0.18	
			(0.13)	(0.18)				(0.14)	(0.19)	
RTI					-1.52*					-1.44*
					(0.12)					(0.13)
Offshorability						-0.45*	-0.40*	-0.20	-0.25	-0.21
						(0.13)	(0.12)	(0.13)	(0.13)	(0.13)
Log wage	-0.19	-0.08	-0.14	-0.13	-0.08	-0.18	-0.08	-0.13	-0.13	-0.09
	(0.12)	(0.12)	(0.12)	(0.12)	(0.12)	(0.12)	(0.12)	(0.12)	(0.12)	(0.12)

Notes: Point estimates on the task content and offshorability measures have been multiplied by 100. All 16 countries are included in each regression (34,816 observations). All regressions have an R^2 of 0.96. Each regressions contains dummies for industry-country-year and industry-occupation-country. Standard errors are clustered by industry-occupation-country. *Significant at the 5% level or better.

TABLE 6—ESTIMATES CONDITIONAL LABOR DEMAND
DEPENDENT VARIABLE: LOG(HOURS WORKED/1000)

	2SLS		*Constrained OLS*	
Linear time-trend interacted with:	(1)	(2)	(3)	(4)
Abstract task importance	0.76*		0.85*	
	(0.23)		(0.16)	
Routine task importance	-0.45*		-0.53*	
	(0.22)		(0.17)	
Service task importance	0.40		0.32	
	(0.29)		(0.22)	
Routine Task Index (RTI)		-1.28*		-1.34*
		(0.19)		(0.14)
Offshorability	-0.66*	-0.68*	-0.23	-0.24
	(0.17)	(0.16)	(0.15)	(0.14)
Log industry marginal costs	0.80*	0.81*	0.53*	0.53*
	(0.23)	(0.23)	(0.13)	(0.13)
Log industry output	0.90*	0.90*	1	1
	(0.13)	(0.13)		
Log wage	-1.05*	-1.04*	-0.87*	-0.84*
	(0.23)	(0.23)	(0.09)	(0.10)
R^2	0.95	0.95		
	First-stage (dependent variable: log output)			
Log net industry capital stock	0.79*	0.79*		
	(0.03)	(0.03)		
R^2	0.99	0.99		

Notes: Point estimates on the task content and offshorability measures have been multiplied by 100. Data for Austria, Belgium, Denmark, Finland, France, Germany, Italy and Norway are used in columns (1)-(2) (16,004 observations) and for all countries in columns (3)-(4) (33,702 observations). In columns (3)-(4), the coefficient on log industry output is constrained to 1. Each regression includes dummies for industry-occupation-country. Standard errors are clustered by industry-occupation-country. *Significant at the 5% level or better.

TABLE 7— TWO-STAGE AND REDUCED FORM ESTIMATES OF CONDITIONAL LABOR DEMAND
DEPENDENT VARIABLE: LOG(HOURS WORKED/1000)

	Constrained 2SLS		*Reduced form*	
Linear time trend interacted with:	(1)	(2)	(3)	(4)
Abstract task importance	1.20*		0.81*	
	(0.10)		(0.17)	
Routine task importance	-0.09		-0.69*	
	(0.11)		(0.17)	
Service task importance	0.38*		0.19	
	(0.13)		(0.21)	
Routine Task Index (RTI)		-1.12*		-1.38*
		(0.09)		(0.14)
Offshorability	-0.23*	-0.30*	-0.25	-0.22
	(0.09)	(0.09)	(0.15)	(0.14)
Log industry marginal costs	0.66*	0.66*		
	(0.09)	(0.09)		
Log industry output	1	1		
Log wage	-4.67*	-4.54*		
	(0.14)	(0.14)		
R^2			0.96	0.96
	First-stage (dependent variable: log wage)			
Log supply shift	-0.37*	-0.39*		
	(0.02)	(0.02)		
R^2	0.98	0.98		

Notes: Point estimates on the task content and offshorability measures have been multiplied by 100. Data for all countries except Germany and Italy are used in columns (1)-(2) (30,475 observations) and for all countries in columns (3)-(4) (34,816 observations). In columns (1)-(2) coefficients on log industry output are constrained to 1. Each regressions contains dummies for industry-occupation-country cells. Columns (3)-(4) additionally contain dummies for industry-country-year cells. Standard errors are clustered by industry-occupation-country in columns (3)-(4). *Significant at the 5% level or better.

TABLE 8—ESTIMATES OF THE MODEL'S STRUCTURAL PARAMETERS

	(1) Based on estimates from column (4) of Table 6	(2) Based on estimates from column (2) of Table 7
	A. Directly estimated parameters	
$\left[\frac{1-\kappa}{1-\rho}+\frac{\kappa}{1-\eta}\right]$	0.84* (0.10)	4.54* (0.14)
$1/(1-\eta)$	0.53* (0.13)	0.66* (0.09)
	B. Imputed parameter values	
$1/(1-\rho)$	1.20	9.09
γ_R^N	-6.30	-0.01
γ_R^K	-7.65	-0.28
γ_O^N	-1.18	-0.03
γ_O^K	-1.40	-0.05

Notes: Notes: All estimates have been multiplied by 100. Panel A gives point estimates for the wage elasticity and the coefficient on industry marginal costs reported in column (4) of Table 6 and column (2) of Table 7 respectively. Panel B shows imputed values for other structural parameters assuming a cost share of labor in task production of 0.54 and using estimates of equations (7) and (11). *Significant at the 5% level or better.

TABLE 9—ESTIMATES OF DERIVED LABOR DEMAND
DEPENDENT VARIABLE: LOG(INDUSTRY OUTPUT)

	(1)	(2)	
Log producer price index	-0.31*	-0.24	
	(0.13)	(0.13)	
Log GDP/capita	1		
Log population	1	1.00*	
		(0.05)	
	Measure:	Log income /capita	Log income inequality
Measure interacted with:		(2a)	(2b)
Electricity, gas and water supply		0.82*	0.30
		(0.24)	(0.18)
Financial intermediation		0.90*	0.41*
		(0.26)	(0.12)
Real estate, renting and business activity		0.96*	0.41
		(0.10)	(0.22)
Transport, storage and communication		1.03*	0.76*
		(0.13)	(0.18)
Construction		0.88*	0.26*
		(0.13)	(0.10)
Manufacturing		1.12*	0.10
		(0.01)	(0.20)
Wholesale and retail		1.10*	0.25
		(0.12)	(0.13)
Health and social work		0.84*	0.29*
		(0.12)	(0.13)
Other community, social and personal services		0.70*	0.17
		(0.19)	(0.12)
Hotels and restaurants		0.85*	0.41
		(0.10)	(0.22)

Notes: 1,400 observations are used in column (1) and 740 observations in column (2). The number of observations in column (2) is smaller due to incomplete time series in the inequality measure for several countries. In column (1), the coefficients on log income per capita and log population are constrained to 1. Column (2) has the constraint that the weighted (by industries' output shares in total output) sum of industry interactions with income per capita equals 1. Industries are ranked by their mean gross real hourly UK wage in 1994. Industry "Private household with employed persons" is included in "Other community, social…" for France, Portugal, Spain and the UK. Log income inequality is country- and time-varying log(p90/p10). Each regression includes dummies for industry-country cells. Standard errors are clustered by industry-country. *Significant at the 5% level or better.

TABLE 10—ACTUAL AND COUNTERFACTUAL CHANGES IN LOW-PAYING, MIDDLING AND HIGH-PAYING EMPLOYMENT SHARES BETWEEN 1993 AND 2006

		(1)	(2)	(3)	(4)
Routinization		Yes	Yes	Yes	Yes
Offshoring		Yes	No	Yes	Yes
Output Endogenous		No	No	Yes	Yes
				$1/(1-\gamma)=0.31$	$1/(1-\gamma)=0.50$
$1/(1-\eta)$	$1/(1-\rho)$	*A.*	*Percentage point change in employment share for 4 lowest-paid occupations (actual=1.58)*		
0.53	1.20	2.88	2.46	1.58	0.86
0.66	9.09	0.68	0.29	0.65	0.62
		B.	*Percentage point change in employment share for middling occupations (actual=-7.77)*		
0.53	1.20	-10.89	-9.85	-8.33	-6.61
0.66	9.09	-5.82	-4.73	-5.75	-5.71
		C.	*Percentage point change in employment share for high-paying occupations (actual=6.19)*		
0.53	1.20	7.92	7.29	6.67	5.83
0.66	9.09	5.09	4.38	5.06	5.04

Notes: Results are percentage point changes in employment shares between 1993 and 2006 for employment pooled across all 16 European countries by occupation groups, where counterfactual changes by occupation group are calculated from counterfactual changes by occupation. Conditional on industry output, column (1) allows for all channels in our model. Predictions in column (2) turn off the offshoring channel. Columns (3) and (4) allow for all channels – just as in column (1) – but also account for product demand effects using different values for the price elasticity of product demand (0.31 in column (3) and 0.5 in column (4)).