ECONSTOR
Make Your Publications Visible.

A Service of
ZBW
Leibniz-Informationszentrum Wirtschaft
Leibniz Information Centre for Economics

Graetz, Georg

Working Paper

Labor Demand in the Past, Present, and Future

IZA Discussion Papers, No. 13142

Provided in Cooperation with:
IZA – Institute of Labor Economics

Suggested Citation: Graetz, Georg (2020) : Labor Demand in the Past, Present, and Future, IZA Discussion Papers, No. 13142, Institute of Labor Economics (IZA), Bonn

This Version is available at:
https://hdl.handle.net/10419/216454

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

WWW.ECONSTOR.EU

Mitglied der Leibniz-Gemeinschaft

I Z A Institute
of Labor Economics
Initiated by Deutsche Post Foundation

DISCUSSION PAPER SERIES

IZA DP No. 13142

Labor Demand in the Past, Present, and Future

Georg Graetz

APRIL 2020

I Z A Institute
of Labor Economics

Initiated by Deutsche Post Foundation

DISCUSSION PAPER SERIES

IZA DP No. 13142

Labor Demand in the Past, Present, and Future

Georg Graetz
Uppsala University and IZA

APRIL 2020

Any opinions expressed in this paper are those of the author(s) and not those of IZA. Research published in this series may include views on policy, but IZA takes no institutional policy positions. The IZA research network is committed to the IZA Guiding Principles of Research Integrity.

The IZA Institute of Labor Economics is an independent economic research institute that conducts research in labor economics and offers evidence-based policy advice on labor market issues. Supported by the Deutsche Post Foundation, IZA runs the world's largest network of economists, whose research aims to provide answers to the global labor market challenges of our time. Our key objective is to build bridges between academic research, policymakers and society.

IZA Discussion Papers often represent preliminary work and are circulated to encourage discussion. Citation of such a paper should account for its provisional character. A revised version may be available directly from the author.

ISSN: 2365-9793

IZA – Institute of Labor Economics

Schaumburg-Lippe-Straße 5–9
53113 Bonn, Germany

Phone: +49-228-3894-0
Email: publications@iza.org

www.iza.org

ABSTRACT

Labor Demand in the Past, Present, and Future

Since the beginning of the Industrial Revolution, technological change has led to the automation of existing tasks and the creation of new ones, as well as the reallocation of labor across occupations and industries. These processes have been costly to individual workers, but labor demand has remained strong, and real wages have steadily increased in line with productivity growth. I provide evidence suggesting, however, that in recent decades automation has outpaced the creation of new tasks and thus the demand for labor has declined. There is strong disagreement about the future of labor demand, and predictions about technological breakthroughs have a poor track record. Given the importance of overall labor demand for workers' standard of living as well as their ability to adjust to a changing labor market, obtaining accurate forecasts should be a priority for policy makers.

JEL Classification: J23, O33

Keywords: automation, labor demand, labor share, technology, wages

Corresponding author:
Georg Graetz
Economics Department
Uppsala University
P.O. Box 513
75120 Uppsala
Sweden

E-mail: georg.graetz@nek.uu.se

Labor Demand in the Past, Present, and Future*

Georg Graetz
Uppsala University

1 Introduction

Concerns that new technologies may lead to large-scale job destruction and mass unemployment are not new, but have recently resurfaced with renewed force in academic and policy debates as well as in the media (Autor, 2015; Shiller, 2019). There is a consensus that for most of the past 250 years—since the beginning of the Industrial Revolution—technology has been a blessing overall, driving the spectacular rise in incomes and standards of living over this period (Jones, 2016). However, recent advances in robotics and artificial intelligence lead some to suggest that this time is different—prospects for less-skilled workers in particular are deteriorating, as automation threatens to proceed at a much higher pace (Brynjolfsson and McAfee, 2014; Ford, 2015), and the creation of new tasks appears to slow down (Acemoglu and Restrepo, 2019b). Others suggest that the demand for middle-skill workers in particular—who have lost out from recent technological change—may well pick up again (Autor, 2015). Yet others question the ability of machine learning and robotics to deliver sustained productivity growth (Gordon, 2012, 2014).

Technology affects the labor market in at least two distinct ways. First, the extent of automation compared to the rate at which new types of jobs are created, determines overall labor demand relative to the demand for capital (Acemoglu and Restrepo, 2019a). If the two forces proceed at the same pace, overall labor demand will remain stable, and workers will share the gains from increased productivity. Second, technology leads to the reallocation of labor across industries and occupations, imposing costs on affected workers even when overall labor demand is unchanged. In this paper, I review existing evidence and present new findings on both of these points.

In Section 2, I investigate how technology has affected the evolution of labor demand over the past 30 years. As is well-known, the share of GDP accruing to workers has been declining across the world (Karabarbounis and Neiman, 2014), but there is no consensus yet about the driving forces of this decline. I use the decomposition method developed by Acemoglu and Restrepo (2019a) to isolate the component of changes in the labor share that is due to changes in task content, as opposed to changes in industry composition or changes in factor inputs.

I find that in the US and across five large European economies, the change in task content from 1987-2007 is negative, implying that automation has outpaced the creation of new tasks over this period.

*This paper was written for the European Commission Directorate-General for Economic and Financial Affairs as part of the research fellowship "The productivity challenge: jobs and incomes in the dawning era of intelligent robots," and was first circulated by the European Commission as European Economy Discussion Paper 114. I thank Adrian Adermon, Michael Böhm, Guy Michaels, and Oskar Nordström Skans for helpful comments. Any remaining errors are my own. Email: georg.graetz@nek.uu.se.

Moreover, the change in task content is of similar magnitude across countries, unlike raw changes in the labor share. Acemoglu and Restrepo (2019a) establish that their estimated task-content changes for the US correlate with more direct measures of automation such as robot adoption. In addition, I find that changes in task content in European industries are positively correlated with changes in the same industries in the US; and again, this is not true for raw industry-level changes in labor shares. I also discuss alternative explanations for the decline the in labor share, especially those related to a rise in firms' market power.

Even in an economy in which the rate of automation equals that of the creation of new tasks, technological change can be costly to individual workers. I discuss the reasons for this in Section 3, and review some of the existing evidence. I report on ongoing research on the individual costs of occupational decline (Edin, Evans, Graetz, Hernnäs, and Michaels, 2019). A theme common to most research on distributional aspects of technological change is that workers' ability to cope with a changing labor market depends critically on overall labor demand being strong.

While governments have for a long time made efforts to predict employment growth at the level of industries and particularly occupations, there is a lack of systematic efforts in forecasting the determinants of overall labor demand, such as the rate of automation and the rate of creation of new tasks. This is a formidable challenge as the relevant information—for instance, knowledge about imminent technological breakthroughs and their likely applications—is widely dispersed. I argue in Section 4 that prediction markets devoted to the forecasting of productivity growth, wage growth, and related variables may help remedy the situation.

I offer concluding remarks, including a brief discussion of recent trends in labor supply, in Section 5.

2 Has demand for labor declined? If so, is technology to blame?

This section is concerned with changes in the share of GDP accruing to labor over the past few decades. Movements in the labor share imply that growth in average wages diverges from productivity growth. Denote average wages by w, output by Y, the size of the labor force by N, and the labor share by s_N. By definition, $s_N \equiv wN/Y$. Therefore, wage growth is the sum of productivity growth and growth in the labor share,

$$\underbrace{\Delta \log w}_{\text{wage growth}} = \underbrace{\Delta \log(Y/N)}_{\text{productivity growth}} + \underbrace{\Delta \log s_N}_{\text{growth in the labor share}} . \tag{1}$$

A falling labor share implies that a higher rate of productivity growth is required to achieve a given rate of wage growth. When productivity growth is slow by historical standards, a decline in the labor share is thus especially bad news for workers' welfare.

A falling labor share is also a strong indicator that the demand for labor is declining relative to the demand for capital. The question then is where this decline in labor demand comes from.

I document a secular decline in the labor shares across five European countries and the US in Section 2.1. I provide evidence suggesting that this decline is at least partly driven by automation in Section 2.2, and discuss alternative explanations for the decline in Section 2.3.

2.1 *The declining labor share in Europe and the United States*

Figure 1 plots the labor share over the period 1970-2007 for France, Germany, Italy, the Netherlands, the United Kingdom, and the United States. The choice of countries, as well as the highlighting of the base year 1987 in the graphs, is dictated by data requirements for the analysis of Section 2.2, which will decompose changes in the labor share from 1987-2007. The source for all data used in this paper is the March 2011 release of EUKLEMS (Timmer, van Moergastel, Stuivenwold, Ypma, O'Mahony, and Kangasniemi, 2007; O'Mahony and Timmer, 2009).[1] I present two different measures of the labor share. One is based on the earnings of employed workers ('Employees'), while the other also includes the earnings of the self-employed ('All labor'). It is conceptually and practically difficult to divide the earnings of the self-employed into labor and capital income (Krueger, 1999; Elsby, Hobijn, and Sahin, 2013), so the employee-only measure may be preferable. For completeness, I will report results for both measures throughout.

Notes: The figures plot the shares of GDP received by employees ('Employees') and by all labor, including the self-employed ('All labor'), over time. Source: EUKLEMS.

Figure 1: The labor share over time

Figure 1 shows a secular decline in the labor share in all six countries, regardless of the measure considered. The decline is not monotone, and its timing differs across countries, nevertheless a downward

[1] An exception are the population data, which I obtain from the World Bank.

trend is visible everywhere.[2]

Notes: The figures plot the shares of GDP received by employees ('Employees') and by all labor, including the self-employed ('All labor'), over time. Each series is normalized to zero in 1987. Source: EUKLEMS.

Figure 2: Changes in the labor share, in percentage points relative to 1987

Due to differences in levels and volatility of the labor share series, the precise magnitudes of the decline are easier to discern when plotting percentage point (pp) changes relative to a base year. This is done in Figure 2. The labor share, including only employees' compensation, has declined by 2-3pp from 1987-2007 in all countries considered except Germany, where the decline was much larger at around 8pp. However, Germany's labor share was higher in 1987 than in 1970, whereas for the other countries there was also a decline from 1970-1987.[3] The labor share series that includes self-employed as well mostly shows similar patterns, except in the cases of France and Italy, where the broader measure shows a substantially larger decline.

[2]Karabarbounis and Neiman (2014) perform statistical tests for a large set of countries, including the ones I focus on here, and find that the decline in the labor share in most cases is unlikely due to random variation around a constant long-run value.

[3]The lack of US data for much of the 1970s is due to efforts by the authors of EUKLEMS to make data comparable. I speculate that this also explains why the labor share decline for the US shown here is somewhat smaller than reported in the literature (Karabarbounis and Neiman, 2014, for instance). Reassuringly, my decomposition results for the US in Section 2.2 are however very similar to those of Acemoglu and Restrepo (2019a).

2.2 *Is technology to blame for the falling labor share?*

I will use the constant-elasticity-of-substitution (CES) production function, a familiar tool in macroeconomics, to illustrate the forces that could potentially cause changes in the labor share. The CES production function combines capital K and labor N to produce output Y as follows:

$$Y = \left[\alpha^{\frac{1}{\sigma}} (A_K K)^{\frac{\sigma-1}{\sigma}} + (1-\alpha)^{\frac{1}{\sigma}} (A_N N)^{\frac{\sigma-1}{\sigma}}\right]^{\frac{\sigma}{\sigma-1}}. \tag{2}$$

A_K and A_N denote factor-augmenting technologies, and σ is the elasticity of substitution. The interpretation of $\alpha \in (0,1)$ is to be thoroughly discussed below, but for now I treat it as a fixed parameter, as is standard in the literature. Suppose that there are many perfectly competitive firms hiring capital and labor at rates r and w—the rental price of capital and the wage, respectively—and combining these factors to make a homogenous good according to (2). Profit maximization then implies that the labor share equals

$$s_N = \frac{1}{1 + \frac{\alpha}{1-\alpha}\left(\frac{r/A_K}{w/A_N}\right)^{1-\sigma}}. \tag{3}$$

Equation (3) suggests that changes in the labor share are due to changes in the ratio of effective factor prices $\frac{r/A_K}{w/A_N}$. A change in this ratio induces firms to adjust their capital-labor ratios—as capital becomes relatively cheaper (more expensive) firms use more (less) capital relative to labor.[4] What this does to the labor share depends on the elasticity of substitution σ. If $\sigma > 1$, factors are easily substitutable, so that a fall in the (effective) rental rate is more than offset by the increased use of capital, and the labor share declines. But if $\sigma < 1$, firms only modestly increase the use of capital in response to a lower rental rate, and the labor share actually increases.

If $\sigma = 0$, a change in relative factor prices will have no effect on the labor share. In this case the production function (2) collapses to the Cobb-Douglas form $Y = (A_K K)^{\alpha}(A_N N)^{1-\alpha}$, and the labor share is equal to $1-\alpha$. This is why α is often referred to as a 'share parameter' in the literature. Indeed, the labor share is a fixed quantity in many macroeconomic models as $\sigma = 1$ is a common assumption. But this contradicts the evidence shown in Section 2.1. Karabarbounis and Neiman (2014) propose to allow for $\sigma \neq 1$, while keeping α fixed, and show that changes in factor prices—in particular, a falling rental rate—together with $\sigma > 1$ are capable of explaining the falling labor share. However, much independent econometric evidence suggests that $\sigma < 1$. The alternative, then, is to allow α to change. This is the

[4] Profit maximization implies the first-order condition (FOC)

$$\left(\frac{\alpha}{1-\alpha}\right)^{\frac{1}{\sigma}} \left(\frac{A_K}{A_N}\right)^{\frac{\sigma-1}{\sigma}} \left(\frac{K}{N}\right)^{\frac{1}{\sigma}} = \frac{r}{w},$$

and perfect competition ensures that all output accrues to capital and labor (there are no pure profits), so that $s_N = 1/(1 + rK/(wN))$. This together with the FOC then leads to (3). Alternatively, the labor share can be expressed in terms of the ratio of effective factor uses $A_K K/(A_N N)$,

$$s_N = \frac{1}{1 + \left(\frac{\alpha}{1-\alpha}\right)^{\frac{1}{\sigma}} \left(\frac{A_K K}{A_N N}\right)^{\frac{\sigma-1}{\sigma}}}.$$

The way in which effective factor uses affect the labor share is analogous to that of effective factor prices as discussed in the main text, which is due to the one-to-one relationship between relative factor prices and uses given by the FOC.

approach taken by Acemoglu and Restrepo (2019a), and the one that I will pursue here.

Acemoglu and Restrepo (2019a, henceforth AR) do not start with (2). Instead, their point of departure is the observation that production requires the completion of tasks (AR, p6):

> The production of a shirt, for example, starts with a design, then requires the completion of a variety of production tasks, such as the extraction of fibers, spinning them to produce yarn, weaving, knitting, dyeing, and processing, as well as additional nonproduction tasks, including accounting, marketing, transportation, and sales. Each one of these tasks can be performed by human labor or by capital (including both machines and software). The allocation of tasks to factors determines the task content of production.
>
> Automation enables some of the tasks previously performed by labor to be produced by capital.

At the heart of the task framework is the equation

$$Y = \left(\int_Z Y(z)^{\frac{\sigma-1}{\sigma}} dz \right)^{\frac{\sigma}{\sigma-1}}. \tag{4}$$

The different tasks are indexed by $z \in Z$, and task outputs $Y(z)$ are combined to make the final good Y via a CES production function. The set of tasks Z may change over time, as I discuss below.

How exactly are tasks completed? AR assume that all tasks can be performed by labor, but only a subset can be performed by capital (machines). The state of technology determines the size of this subset, that is, the fraction of tasks that can be automated. AR further assume that the effective rental price of capital is sufficiently below the effective wage, so that firms will always automate tasks whose automation is feasible. Once factors have been assigned to tasks, output can be expressed as a function of capital and labor exactly as in (2), up to a re-scaling. Crucially, AR show that α can now be interpreted as the share of tasks performed by machines, and hence $1-\alpha$ as the task share performed by labor.[5] If further automation becomes feasible due to technological progress, α increases. And so, according to (3), automation causes the labor share to decline.

There are several reasons why the task framework may be preferable to a model that directly links output to factor inputs as in (2). First, the task framework allows for a straightforward as well as realistic way of modeling automation, namely, by increasing the share of tasks that can be performed by capital. Second, this framework allows for technology to increase productivity of capital or labor differentially across tasks. In contrast, in the simple CES model of (2) an increase in A_K or A_N implies improved productivity of these factors regardless of their use. Third, as shown by Acemoglu and Restrepo (2018), wages could actually decrease as a result of technological progress—increased automation—in the task framework, whereas improved technologies will always lead to higher wages in the simple CES model.[6]

Fourth, the task framework allows for an additional force, namely the introduction of new tasks (an expansion of the set Z) due to technological progress. This is of great empirical relevance—think of tasks such as web design, piloting aircrafts, analysing large datasets, or performing transplant surgery. The introduction of new tasks will cause the fraction of tasks performed by labor, α, to increase assuming that humans have a comparative advantage in such new tasks, as seems empirically plausible. AR call this the *reinstatement effect*, which counteracts the *displacement effect* that is due to automation.[7]

[5] More precisely, α is the weighted share of tasks performed by machines, where the weights are a function of capital's and labor's productivity in each task.

[6] For related results on technological change and wages, see Caselli and Manning (forthcoming).

[7] The task framework also gives rise to a different interpretation of the substitution elasticity σ in the aggregate production

I now take the task framework to the data in order to investigate the drivers of the decline in the labor share in the six countries. I implement the method suggested by AR, which requires industry-level panel data on value added, labor and capital inputs, as well as factor prices and allows for a decomposition of wage bill growth into several distinct components. Suppose that the production function in industry i, country c, and year t is

$$Y_{ict} = \left[\alpha_{ict}^{\frac{1}{\sigma}} \left(A_{K,ct} K_{ict}\right)^{\frac{\sigma-1}{\sigma}} + (1-\alpha_{ict})^{\frac{1}{\sigma}} \left(A_{N,ct} N_{ict}\right)^{\frac{\sigma-1}{\sigma}} \right]^{\frac{\sigma}{\sigma-1}}. \tag{5}$$

Of course, I maintain the interpretation that this production function ultimately depends on an underlying task production function $Y_{ict} = \left(\int_{Z_{ict}} Y_{ict}(z)^{\frac{\sigma-1}{\sigma}} dz\right)^{\frac{\sigma}{\sigma-1}}$. The set of tasks Z may vary across industries, countries, and over time, as may be the case for the subset of tasks that can be automated. Hence, α_{ict} may vary accordingly. I restrict factor-augmenting technologies to be constant across industries in a given country and year, and the substitution elasticity to be constant across all countries, industries, and years.

Recall that according to (1), the growth (change in the log) of the economy-wide average wage can be written as the sum of productivity growth and growth in the labor share. Productivity growth may be affected by automation, factor-augmenting technological change, and the introduction of new tasks. Since my interest here is in explaining deviations of wage growth from productivity growth, I do not investigate the precise sources of productivity growth. Instead, I focus on changes in the labor share. These changes in the labor share, as shown by AR, can be decomposed into a composition effect, a substitution effect, and the change in task content. The composition effect is due to reallocation of economic activity across industries over time, and it will be important if the labor share differs across industries. The substitution effect comes from firms' responses to changes in relative effective factor prices; substitution effects are allowed to differ across industries, and the decomposition uses an employment-weighted sum of these industry-level effects. Finally, the change in task content is due to advances in automation as well as the introduction of new tasks. It is also allowed to vary across industries, and an employment-weighted sum is used in the decomposition. In sum, the decomposition of wage growth I will carry out can be written as[8]

$$\begin{aligned} \text{Change in log average wage} \quad = \quad & \text{Productivity effect} \\ & + \text{Composition effect} + \text{Substitution effect} \\ & + \text{Change in task content.} \end{aligned} \tag{6}$$

I leave the formal statement and derivation of this decomposition to the appendix, and provide here an informal description. For the left-hand side, the average wage is calculated as the labor share times GDP (value added) divided by population, and its growth rate (log change) is computed. On the right-hand side, the productivity effect is simply the growth in GDP per capita;[9] the composition effect is

function. If (2) is the starting point of the model, then σ captures the substitutability of capital versus labor. However, if (2) instead results from (4), then σ actually measures substitutability across different tasks. As the above quote as well as further introspection may suggest, task substitutability is likely to be low in reality. However, there currently exist no econometric estimates of σ when interpreted in this way.

[8] I use the terms 'productivity effect' and 'productivity growth' interchangeably.

[9] To obtain wages and productivity I divide by population instead of the number of people employed, following AR. This

calculated using each industry's labor share and each industry's value added share in overall GDP; and computing the substitution effect requires industry-level labor shares, wages, rental prices, and quality-adjusted labor and capital inputs, as well as assumptions on the elasticity of substitution and the growth rate of factor-augmenting technology—following AR, I choose $\sigma = 0.8$ and set the growth rate of A_N/A_K equal to each country's average productivity growth.[10] Finally, and most importantly, the change in task content in each industry is estimated as the change in the industry's labor share minus the substitution effect. The industry-level changes are then added up (weighting by employment shares) to obtain the aggregate change in task content.[11]

The validity of the decomposition rests on several assumptions. First, product and factor markets must be perfectly competitive, although perfect factor mobility across industries is not required: factor prices are allowed to vary across industries. Second, the production function is of the CES form (5). Third, firms' factor use is consistent with profit maximization. Although trade is not explicitly modeled, the economy is not required to be closed, as long as changes in openness to trade are reflected in changes to goods prices and factor prices.[12] Finally, the framework abstracts from worker heterogeneity, but AR argue that this is second-order as long as the measured labor input is correctly adjusted for quality growth.

I implement the decomposition using the EUKLEMS data. In each of the six countries, there are 27 industries, covering the entire economy. Due to data availability constraints, I choose 1987 as the base year for my decomposition, and 2007 as the final year. Figure 3 presents the results, focusing on labor income as measured by compensation of employees (using all labor earnings, including those of the self-employed, yields very similar results, as shown in Figure A1). All series plotted represent differences in logs relative to the base year 1987. It is clear that wages in all six countries grew less than productivity, which of course reflects the falling labor share. Furthermore, the composition effect tends to be negative, implying that employment has moved into industries with lower labor shares. This is especially pronounced in Germany and Italy, and partly accounts for the larger labor share decline in these countries. There is a small positive substitution effect in all countries. Rental prices relative to wages have declined, and with an elasticity of substitution below one this contributes positively to wage growth. Finally, and most importantly, in all countries the change in task content is estimated to be negative. In terms of the theoretical framework, this means that automation has outpaced the creation of new tasks, in other words, that the displacement effect has dominated the reinstatement effect. Moreover, the magnitude of the task content change is similar in all countries at around minus ten percent, except in France, where the decline is less than five percent.[13]

Taking the results of the decomposition at face value, it appears that technology has indeed contributed to the decline in the labor share, and to a similar extent across countries. However, it is possible

is because declining labor shares and slow productivity growth may induce lower labor force participation, in which case per-worker measures understate the welfare implications of such secular trends. The calculations of composition effect, substitution effect, and change in task content are unaffected by this choice.

[10] See AR (p13) for a discussion of these choices, including estimates of σ from the literature.

[11] AR also show that under further and arguably strong assumptions, the change in the task content can be decomposed into an automation component (the displacement effect) and a new-tasks component (the reinstatement effect). Here I focus on the net effect only.

[12] An exception to this is offshoring, which may directly affect the task content of production. See the discussion below.

[13] The labor share in France appears to have declined more strongly when taking into account self-employed earnings, and using this measure leads to a larger (in magnitude) estimated change in the task content, see Figure A1.

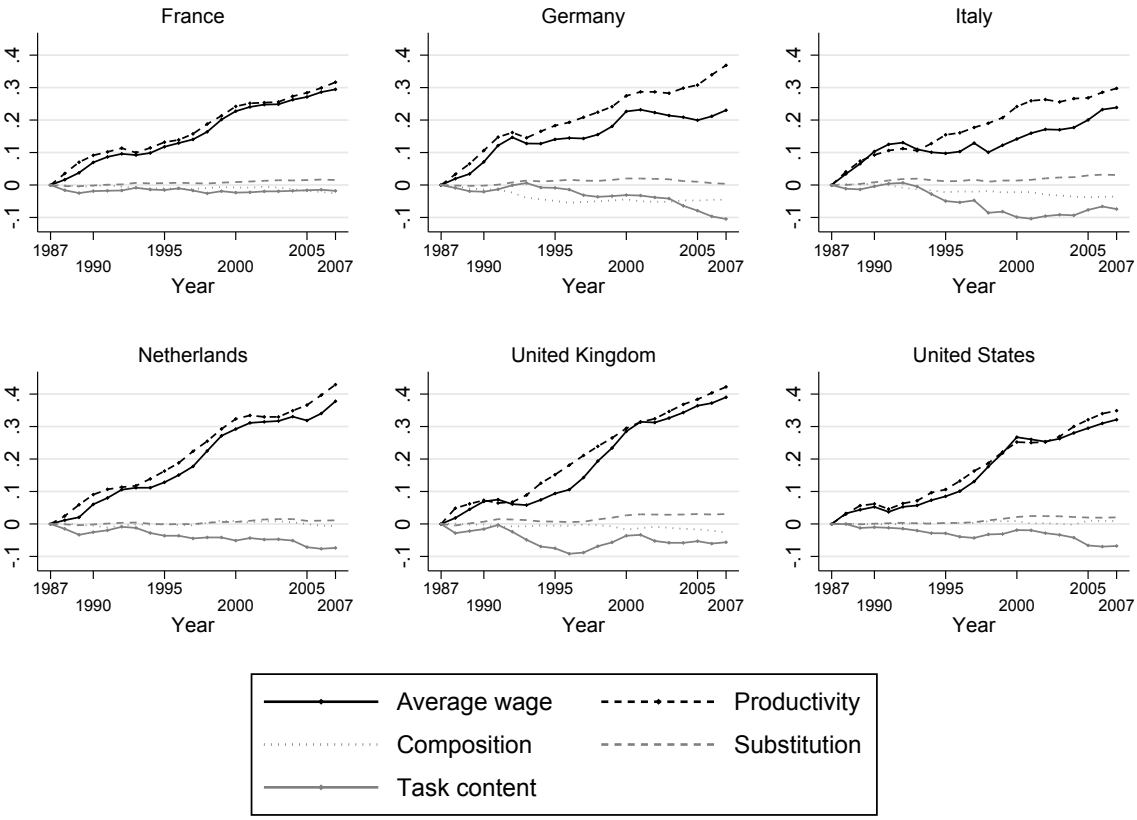

Notes: The figures plot the results from the decomposition (6) which is due to Acemoglu and Restrepo (2019a). The labor share used in the calculations is based on earnings of employees, excluding earnings of the self-employed.

Figure 3: Decomposition of wage growth (employees)

that the estimated change in task content does not in fact capture technological change but something else, especially given that it is essentially a residual (albeit with a clear theoretical interpretation). To address this concern, AR document that in the US, industries with larger declines (negative changes) in task content have also adopted industrial robots at a higher rate, and had a higher initial routine employment share, both widely accepted measures of exposure to automation (Graetz and Michaels, 2018; Acemoglu and Restrepo, 2017; Autor, Levy, and Murnane, 2003). They also find that offshoring is correlated with changes in task content, but controlling for offshoring only marginally affects the relationship between changes in task content and more direct measures of automation. The finding that since 1987, automation appears to have outpaced the creation of new tasks, marks a break with previous trends. AR find that from 1948-1987, the two forces cancelled each other out in the US.

Having implemented the decomposition for the same industries across different countries, I am able to carry out an important additional check. As technology adoption follows broadly similar patterns across countries, it is to be expected that the estimated changes in task content correlate positively within the same industry across countries. To test this, I use the US as a benchmark and pool the five European countries. I then regress for each of the 27 industries in the European countries the change in task content on the change in the corresponding industry in the US. The results are shown in panel A of Table 1. In

Table 1: Industry-level changes in labor share and task content, Europe versus the United States

	(1)	(2)	(3)	(4)
A. Changes in task content				
Change in US	0.50	0.32	0.23	0.21
	(0.20)	(0.090)	(0.087)	(0.081)
B. Changes in the labor share				
Change in US	0.14	0.030	-0.033	-0.040
	(0.18)	(0.093)	(0.059)	(0.058)
Outlier removed		✓		✓
Weighted by initial employment share			✓	✓
Observations	135	134	135	134

Notes: Results are shown from regressions of changes in task content and labor shares 1987-2007 in 27 industries across five European countries against changes in task content in the same industries in the US. Robust standard errors, clustered by industry, in parentheses.

the full sample, a 1-percent task content change in a US industry is associated with a 0.5-percent change in the same industry in the five European countries (column (1)). However, this association is likely sensitive to an extreme outlier (Figure A3). Removing this outlier indeed reduces the coefficient (column (2)). Alternatively, weighting the regression by each industry's initial within-country employment share also leads to a lower association. Conservatively, I conclude that a 1-percent task content change in a US industry is associated with a 0.2-percent change in the same industry in Europe (columns (3)-(4)). Remarkably, the raw industry-level changes in the labor share do not appear to co-move across countries in the same way (panel B). Thus, the changes in task content estimated by the AR method do seem to isolate a driving force that is related to technology and common across countries, more so than the raw changes in labor shares.[14]

The results presented here are also consistent with independent evidence on regional differences in robot adoption within Germany. Dauth, Findeisen, Suedekum, and Woessner (2018) find that in regions that adopted industrial robots at a higher rate from 1994-2014, the labor share in total regional income declined more.

Of the assumptions which are required for the decomposition to be valid, that of perfectly competitive output and factor markets is most likely to be violated in practice. In the following section I discuss whether changes in market power might explain the fall in the labor share instead.

2.3 Alternative explanations for the falling labor share

There are several alternative explanations for the decline in the labor share besides changes in task content. Recall that Karabarbounis and Neiman (2014) put forward an explanation based on the decline in the relative effective price of capital. Their explanation also puts technology center stage, as the declin-

[14] All results related to the decomposition exercise are very similar when using the labor share measure that includes self-employed earnings. See Figures A1, Table A1, and Figure A3.

ing price of equipment capital in particular is plausibly linked to technological advances such as personal computers and industrial robots (Nordhaus, 2007; Graetz and Michaels, 2018). However, recent research has suggested explanations that are less closely related to technological change, and these explanations have focussed on several aspects of changes in market power.

So far, my conceptual framework has assumed perfectly competitive firms, so that all income either accrues to labor or capital. However, in reality firms do earn pure profits.[15] According to Barkai (2017), the share of profits in US GDP has increased from 5 to 15 percent from 1985-2015, and this has not only come at the expense of labor, but even more so, at the expense of capital. Suppose that profits are due to product market power, and further suppose that such market power does not affect firms' relative factor demand curve. The expression for the labor share (3) then generalizes to

$$s_N = \frac{1 - s_\Pi}{1 + \frac{\alpha}{1-\alpha}\left(\frac{r/A_K}{w/A_N}\right)^{1-\sigma}}, \tag{7}$$

where s_Π is the share of profits in GDP. Everything else equal, the presence of pure profits implies a lower labor share. Furthermore, under commonly made assumptions the aggregate profit share is directly related to product market power, $s_\Pi = 1 - 1/\mu$, where μ is the markup that firms charge over marginal costs.[16] This suggests that changes in the profit share may be driven by changes in markups, which could be, for instance, due to increased market concentration.

Indeed, De Loecker and Eeckhout (2017) measure markups using firm-level data and find that markups in the US increased from 20 percent over marginal cost in 1980 to 65 percent today. De Loecker and Eeckhout (2018) document similar trends for most regions of the world. De Loecker and Eeckhout (2017) explore the implications of rising markups for the labor share, as do Eggertsson, Robbins, and Wold (2018). However, Traina (2018) and Karabarbounis and Neiman (2018) discuss measurement issues related to markup estimation using firm-level data, and argue that markups appear to be stable under alternative, not necessarily less defensible assumptions. Moreover, Karabarbounis and Neiman (2018) highlight that the estimation of profit shares requires accurate measures of both capital rental prices and capital services. They argue that rental prices have likely been mismeasured, and that the rise in the profit share is thus overstated.

While there appears to be no consensus yet about recent trend in markups, researchers largely agree that there has been a rise in concentration—both in terms of employment and sales—in most industries at the national level in the US since the early 1980s (Rossi-Hansberg, Sarte, and Trachter, 2018). Increased concentration seems to be driven by the largest firms in an industry becoming more dominant. Since these large firms tend to have lower payroll-to-sales ratios, increased concentration accounts for part of the fall in the labor share (Autor, Dorn, Katz, Patterson, and Van Reenen, 2017).

A related issue which has recently received attention is local concentration of employment and the

[15] Separating pure profits from capital income is however very challenging in practice. The EUKLEMS data, which I have used in this paper, actually treat all non-labor income as capital income. It is therefore not possible to study trends in profit shares using these data.

[16] For instance, suppose that consumers have CES utility $U = \left(\int C(j)^{\frac{\varepsilon-1}{\varepsilon}} dj\right)^{\frac{\varepsilon}{\varepsilon-1}}$ and that each variety j is produced by a monopolistically competitive firm facing the CES production function (2). Then equations (7) and $s_\Pi = 1 - 1/\mu$ hold, and the markup is $\mu = (\varepsilon - 1)/\varepsilon$, with $\varepsilon > 1$.

wage setting power of firms. While higher local concentration is associated with lower wages (Azar, Marinescu, Steinbaum, and Taska, 2018), local concentration has actually decreased in the US (Lipsius, 2018; Rinz, 2018). This appears to be driven by the same forces as the increase in concentration at the national level, as entry of large firms in a local labor market in fact leads to lower concentration (Rossi-Hansberg, Sarte, and Trachter, 2018) locally. Trends in local concentration and increased monopsony power thus cannot explain the falling labor share or rising inequality. Instead, what matters is that large firms, where the labor share tends to be lower, have grown even larger in relative terms (Autor, Dorn, Katz, Patterson, and Van Reenen, 2017).

2.4 Summary

There is evidence suggesting that technological change—in particular, automation outpacing the creation of new tasks—has driven the decline in the labor share. In this sense, the overall demand for labor relative to capital has indeed declined, and technology appears to be the culprit. While my estimates of changes in task content require several strong assumptions to be interpretable as resulting from technological change, I do not find it plausible to view them as driven by changes in market power instead, since they correlate with independent measures of automation and co-move across countries.

There also is evidence suggesting that the increasing dominance of large firms accounts for the decline in the labor share, but this is not necessarily in contradiction to technological change playing an important role. The increased market share of the largest firms may itself be driven by technological change, and reallocation of workers across firms within an industry may be a relevant mechanism for how changes in task content materialize in practice. The relationship between technological changes and increased market concentration is a matter for future research.

3 Technological change, reallocation of labor, and the distribution of labor earnings

The previous section focussed on changes in aggregate labor demand, which manifest themselves in changes in the economy-wide labor share. In this section I focus instead on how technology affects the composition of labor demand—what type of labor, in which industries or occupations, benefits or is harmed by new technologies.

Shifts in the demand for labor services between different industries, occupations, and firms have been a feature of economic development at least since the start of the industrial revolution. One of the most salient examples is the shift of employment out of agriculture into manufacturing, and from there into services, commonly referred to as structural change (Herrendorf, Rogerson, and Valentinyi, 2014). Technology has played a direct role in this process, for instance by automating many aspects of farming and industrial production; as well as an indirect one, as rising productivity caused rising incomes, which in turn induced consumers to shift their expenditure from food and goods to services. In any case, structural change implies a vast reallocation of workers across industries, and, because the nature of work differs between sectors, also across occupations. Reallocation of workers across firms is in fact a general feature of economic growth, even in the absence of sectoral shifts, since productivity growth comes about through a process of creative destruction, whereby some firms are forced out of business by more innovative competitors (Aghion, Akcigit, and Howitt, 2014).

Because labor markets in reality feature search frictions and asymmetric information, the process of reallocation implies that there are losers as well as winners from technological progress. I review some of the evidence on the costs of sudden job displacement, due to plant closures, in Section 3.1. In contrast to such mass layoffs, occupational decline represents a more gradual deterioration of the demand for a worker's skills. In ongoing research, my co-authors and I find that it can nevertheless be costly. I report on this research in Section 3.2.

Aside from reallocation of jobs, technological change affects inequality when it causes changes in the demand for tasks and skills. I discuss how this process has played out over the past 40 years in Section 3.3.

Throughout this part of the paper, my aim is to characterize not only the distributional aspects of technological change, but also the ways in which workers adjust to adverse shocks, and how workers' adjustment opportunities are affected by aggregate labor demand.

3.1 The costs of sudden job loss

A large literature has investigated workers' earnings losses following displacement due to plant closure. These losses are typically both large and persistent. Earnings losses following displacement are especially severe for workers with long tenure (Jacobson, LaLonde, and Sullivan, 1993) and older workers (Gathmann, Helm, and Schönberg, 2018). In contrast, displaced workers who manage to stay in the same firm or occupation suffer smaller losses (Huttunen, Moen, and Salvanes, 2011; Kambourov and Manovskii, 2009).

Importantly, the aggregate level of labor demand matters for the earnings losses of displaced workers. Davis and Von Wachter (2011) find that in the US, losses are twice as high in recessions than in expansions: when the national unemployment rate is less than 6 percent, men lose about 1.5 years of pre-displacement earnings (in present value terms) while they lose almost 3 years when the unemployment rate is above 8 percent. In a similar vein, Bana (2019) documents longer joblessness and larger earnings losses for displaced workers who face a shrinking demand for their occupation. These findings suggest that if technology were to further depress labor demand, the costs incurred by displaced workers would be substantially larger (at the same time, displacement would likely become more frequent).

3.2 The individual costs of occupational decline

In ongoing research (Edin, Evans, Graetz, Hernnäs, and Michaels, 2019) we explore the consequences of occupational decline for workers' careers. We begin by identifying occupations that have declined sharply during the last 30 years and determine whether their decline was due to technological replacement using the Occupational Outlook Handbook (Bureau of Labor Statistics, 1986, 2017, OOH). We classify occupations as having declined if their employment in the US contracted by more than 25 percent. We then map this information to Swedish occupations in order to study how occupational decline affects individual workers, using data on the entire Swedish population at annual frequency 1985-2013.[17] We

[17] It is much more challenging to track employment changes for hundreds of different occupations in Sweden than in the US, because Swedish occupational classifications have changed substantially. However, at the level of detail at which we are able to consistently measure occupational employment in Sweden, we do see a strong association between our US-based indicator of decline and actual Swedish employment changes. Furthermore, studying the effects of occupational decline in the US is challenging given the lack of large longitudinal data sets. Nonetheless, we replicate our analysis using data from the

are also able to assess to what extent occupational decline was anticipated, using forecasts contained in the OOH, as well as the size and past growth of Swedish occupations (which strongly predict growth 1985-2013).

Although occupational decline represents a more gradual fall in demand compared to say a mass layoff, we do find substantial costs for workers who in 1985 worked in a subsequently declining occupation. Over a period of 28 years, these workers have 2-5 percent lower cumulative earnings than comparable workers in non-declining jobs. And for workers at the bottom of the within-occupation earnings distribution, the losses are even larger at 8-11 percent.[18] Furthermore, workers exposed to occupational decline are less likely to still be working in their initial occupation in 2013. This is noteworthy because over a nearly 30-year period, occupations could decline dramatically simply by taking in fewer younger workers and via regular retirements. We also find that occupational decline is associated with increased unemployment and publicly sponsored retraining. Our baseline results focus on all occupations that have declined, but we find very similar results when focusing on occupations whose decline was directly linked to technological change.

Occupational mobility is in principle a mechanism that may help workers mitigate their earnings losses from occupational decline. However, workers in declining occupations may also be more exposed to displacement, and given labor market frictions, may find themselves making occupational moves that are associated with higher earnings losses than incurred by those who manage to stay. We do not find that movers out of declining occupations do better than stayers in those same occupations. However, it is likely that a high rate of occupational mobility helps to reduce earnings losses because of general equilibrium effects, as it implies an upward-sloping occupational supply curve.

We also investigate the timing of earnings losses due to occupational decline. Workers who started out in subsequently declining occupations had lower earnings in all years, with the difference tending to grow larger over time. However, earnings losses were especially severe during the 1990s recession.

3.3 Task-biased technological change and its implications for inequality

Since the early 1980s, highly-educated workers in most developed countries have enjoyed a sustained increase in relative wages, even as educated labor has become more abundant. A large literature has provided evidence that the introduction of information and communication technology (ICT)—for instance, the personal computer—is the main reason behind this increased demand for skill (Acemoglu, 2002). A similarly sizeable literature (Autor, Levy, and Murnane, 2003; Acemoglu and Autor, 2011) has demonstrated that ICT substitutes for some tasks but complements others, and that highly-educated workers benefit from ICT because they traditionally perform the tasks which are complemented by it.

When viewed up close, the impact of new technologies on the demand for tasks and skills appears more complicated than a simple story of skill-biased technological change would suggest. First, the occupations which have seen the largest declines in employment shares since 1980 are in fact middle-skill, middle-wage occupations such as office clerks and machine operators. This hollowing-out of occupational employment, commonly referred to as job polarization, has been observed in most industrialized

National Longitudinal Survey of Youth. These results are much less precise than the Swedish ones, but lead to broadly similar conclusions.

[18] We give a range of estimates based on a several reasonable regression specifications.

countries (Goos and Manning, 2007; Autor, Katz, and Kearney, 2006; Goos, Manning, and Salomons, 2014), and the leading explanation is that ICT is particularly suited to automate the tasks that workers traditionally performed in middle-wage occupations. A complementary explanation emphasizes the economic incentives of firms to automate tasks. Firms would like to replace workers who are expensive, but replacing the most-skilled is not yet technologically feasible. Since the technological complexity of tasks performed by the least-skilled does not differ very much compared to middle-skilled workers,[19] firms may be induced by economic incentives to adopt automation technologies mainly in middle-wage occupations. In Feng and Graetz (forthcoming), we illustrate this mechanism using a formal economic model, and provide evidence of its relevance for job polarization in the US.[20]

Second, there is a multitude of skills that are valued in the labor market, and their returns have evolved differentially. For instance, in the US the returns to social and non-cognitive skills have increased relative to the returns to cognitive skills (Deming, 2017), and similar patterns are found for Sweden (Edin, Fredriksson, Nybom, and Öckert, 2017). Returns to manual skills have decreased (Taber and Roys, 2017). These changes in returns are closely linked to changes in task demand (Cortes, 2016; Böhm, 2018). That is, occupations requiring social and non-cognitive skills have grown in importance, whereas occupations intensive in manual skills have declined. The rise of cognitive-skill-intensive jobs seems to have stopped (Beaudry, Green, and Sand, 2016).

Changing task and skill returns are an important driver of changes in wage structures. Thus, technological change does not only account for the increased demand for highly-educated workers (Michaels, Natraj, and Van Reenen, 2014), but also for part of the narrowing gender pay gap, as women tend to have higher non-cognitive and social skills (Black and Spitz-Oener, 2010; Beaudry and Lewis, 2014).

Although the development and adoption of new technology often appears to be a gradual process, technology-induced changes in skill and task demands tend to speed up during cyclical downturns. For instance, much of the decline in middle-wage employment in the US over the past 40 years was concentrated in the last three recessions (Jaimovich and Siu, forthcoming). This may make it more difficult for workers to adjust to a declining demand for their skills.

3.4 Summary

Reallocation of labor across industries and occupations, directly or indirectly linked to technological change, can be costly for individual workers. As might be expected, sudden job loss due to plant closures implies particularly large earnings losses. But even relatively gradual processes such as occupational decline lead to substantial losses. Part of the reason may be that such otherwise gradual declines in demand speed up during cyclical downturns, when workers are in a bad position to cope with displacement. Technological change creates winners and losers not only because of reallocation processes, but also because it affects the demand for different tasks and skills.

[19] Low-skilled tasks such as cleaning offices or waiting tables are quite complex from an engineering perspective (Moravec, 1988).

[20] Barany and Siegel (2018), using a multi-sector growth model, show that job polarization can also be explained by structural change.

4 Predicting better

Given striking technological advances such as machine learning and robotics, and given a declining labor share and associated mediocre wage growth, policy makers and the public more generally are understandably concerned about the adverse consequences of automation. The findings reported in Section 2 unfortunately do not alleviate such concerns. However, as noted in the introduction, experts disagree on the expected labor market impact of new technologies over the next few decades. Among those who expect rapid technological progress, some worry about the implications for the average worker (Brynjolfsson and McAfee, 2014; Ford, 2015), while others appear more optimistic (Autor, 2015). And there are those who are not even confident in the ability of machine learning and robotics to deliver sustained productivity growth (Gordon, 2012, 2014).

What is one to make of such disparate views? Is the future of labor demand really as uncertain as these disagreements suggest? Again, it is useful to distinguish between overall labor demand on one hand and shifts in labor demand, especially between different occupations, on the other.

Forecasting occupational demand has a long tradition in economic and policy research. Policy makers have sought to provide detailed and up-to-date projections of employment growth, as well as guidance about what types of skills can be expected to remain in demand. Examples include the forecasts by the Bureau of Labor Statistics (BLS) discussed above, the O*NET project sponsored by the US Department of Labor, and the Skills Forecast conducted by the European Centre for the Development of Vocational Training. The BLS forecasts have proved quite accurate in forecasting occupational employment trends not only in the US (Veneri, 1997) but also in Sweden (Edin, Evans, Graetz, Hernnäs, and Michaels, 2019). This is testament both to the quality of the forecasts as well as to the fact that labor demand shifts follow very similar patterns across countries.

But what about projections for overall labor demand? As argued in Section 3, workers' ability to adjust to shifts in labor demand—for instance, between occupations—depends critically on overall labor demand remaining strong. Thus, obtaining accurate projections on this question, in addition to occupation-level forecasts, should be a high priority for policy makers.

The task framework suggests that forecasting overall labor demand essentially requires predicting how the pace of automation will compare to the rate at which new tasks are created. Recent research has produced forecasts of the former. Frey and Osborne (2017) estimate that about 50 percent of employment in the US is threatened by automation over the coming decades, while Arntz, Gregory, and Zierahn (2017) arrive at only about 10 percent using a similar methodology that takes into account within-occupation variation in task content. Predictions about the arrival of human-level artificial intelligence—which would imply that all existing tasks could feasibly be automated (though firms would likely not find it profitable to do so)—have proved unreliable, possibly because experts face perverse incentives in making these predictions. Armstrong and Sotala (2015) document that experts most commonly predict AI to arrive 15-25 years from the date the prediction is made. This range is convenient, because it helps an expert to justify that they are working on something important, while at the same time not running the risk to soon be proven wrong.

Predictions about the creation of new tasks are implicit in predictions about the net effect of technology on labor demand. In 2017, a panel of expert economists were asked to evaluate the statement "Holding labor market institutions and job training fixed, rising use of robots and artificial intelligence

is likely to increase substantially the number of workers in advanced countries who are unemployed for long periods". 38 percent agreed, 33 disagreed, and 29 percent were uncertain.[21]

Producing more reliable predictions about the evolution of overall labor demand requires incorporating research findings, conceptual insights, and widely dispersed private information about imminent innovations. Surveying experts' forecasts or polling them directly is helpful, but even better would be to leverage market forces.

Imagine a contract paying out €1 if real wages in the European Union grow less than productivity on average from 2020-2025.[22] If such a contract can be freely traded, its price will be informative about market participants' beliefs. In particular, under risk neutrality and if the efficient market hypothesis holds, the price of the contract reveals the markets' expectation of the probability that wage growth will lag productivity growth.

The use of prediction markets for economic forecasting is advocated by Snowberg, Wolfers, and Zitzewitz (2013, SWZ). They argue that prediction markets produce accurate forecasts for three reasons, and note their superiority over alternative methods (SWZ, p661):

> First, the market mechanism is essentially an algorithm for aggregating information. Second, as superior information will produce monetary rewards, there is a financial incentive for truthful revelation. Third, and finally, the existence of a market provides longer term incentives for specialization in discovering novel information and trading on it. While these facets are inherent in any market, other forecasting mechanisms, such as polling, or employing professional forecasters, lacks one or more of them. For example, polling lacks incentives for truthful revelation, and professional forecasters may have other motivations than simply forecast accuracy.

SWZ document prediction markets' accuracy, their ability to quickly incorporate new information, their lack of arbitrage opportunities, and robustness to manipulation, despite the presence of some behavioral biases. While most existing prediction markets are devoted to politics and sports, they have also been used to forecast economic variables such as retail sales and unemployment claims. In these cases, prediction markets weakly outperformed survey forecasts.[23]

SWZ also explain how alternative contracts can be used to elicit a variety of statistical moments, for instance, the expected rate of real wage growth, or its standard deviation. The latter, in turn, is an indicator of the uncertainty of market-based forecasts.

For the purpose of forecasting labor demand, several contracts could be designed tied to future wage growth, productivity growth, occupational employment growth, or the emergence of new occupations. The latter could, for instance, relate to revisions of official occupational classifications, with the share of employment under new occupational titles being the object to be forecast.

Despite its advantages, prediction markets are not yet as common as economists would like them to be. One reason may be legal barriers related to the regulation of gambling (Arrow et al., 2008). Another challenge is that for prediction markets to function, they must attract both experts and noise traders in sufficient numbers, so the subject must be interesting to many. In the case of forecasting medium and long term labor demand, the public interest indeed appears to be enormous. Furthermore, one may hope that people engaged in the development of new technologies would be enthusiastic about such a project,

[21] See `http://www.igmchicago.org/surveys/robots-and-artificial-intelligence`.

[22] The contract would specify the precise data sources and release dates.

[23] Occurrences of events that prediction markets considered to be unlikely, such as the outcome of the 2016 US presidential election, are not a refutation of prediction markets' accuracy. Accuracy here means that out of 100 events that market participants had judged to occur with a probability of five percent, five actually did happen.

given its innovative nature. On the other hand, long maturity of contracts may pose a challenge to maintaining liquidity. If this challenge can be overcome, however, prediction markets promise to deliver more accurate forecasts about the future evolution of labor demand than what is currently available.

5 Conclusion

In this paper I have reviewed existing evidence and presented new findings on the recent evolution of labor demand. For most of the past 250 years, wages have grown at the pace of productivity, and thus increased prosperity has been widely shared. However, recent decades have seen a declining labor share, and the evidence suggests that this is at least partly due to automation outpacing the creation of new tasks. I have also highlighted that reallocation of labor across industries and occupations imposes substantial costs on individual workers, and that these costs are particularly large when overall labor demand is weak. There appears to be no consensus among experts about the future evolution of labor demand. I have argued that prediction markets dedicated to medium-term wage and productivity growth could potentially provide accurate forecasts of labor demand.

I have focused exclusively on labor demand, but the wage bill, the tax base, and thus the viability of the welfare state, ultimately depend on labor supply as well. Hours worked typically decrease with GDP per capita, across countries as well as within countries over time (Bick, Fuchs-Schündeln, and Lagakos, 2018). Boppart and Krusell (forthcoming), who have inspired the title of this paper, explain these facts as resulting from stable preferences under steady productivity growth. That is, the income effect of wage growth is (slightly) larger than its substitution effect. Lower wage growth then implies a slower growth in leisure. On the other hand, labor force participation has trended downward recently in the US, a trend that is only partially explained by changes in demographics. Abraham and Kearney (2018) attribute much of this decline to lower wage growth due to changing labor demand. This of course implies that the substitution effect dominates the income effect, contrary to Boppart and Krusell (forthcoming), though more in line with Bick, Fuchs-Schündeln, and Lagakos (2018), who find that in rich countries, the relationship between wages and hours is flat or even increasing. Alternative explanations for lower labor force participation include an increased attractiveness of leisure (Aguiar, Bils, Charles, and Hurst, 2017) and the increased availability of opioid drugs (Krueger, 2017). More research is needed to understand how workers will adjust their hours in response to technology-driven changes in wages, and to assess the importance of independent drivers of labor supply.

References

ABRAHAM, K. G., AND M. S. KEARNEY (2018): "Explaining the Decline in the U.S. Employment-to-Population Ratio: A Review of the Evidence," Working Paper 24333, National Bureau of Economic Research.

ACEMOGLU, D. (2002): "Technical Change, Inequality, and the Labor Market," *Journal of Economic Literature*, 40(1), 7–72.

ACEMOGLU, D., AND D. AUTOR (2011): "Skills, Tasks and Technologies: Implications for Employment and Earnings," in *Handbook of Labor Economics*, ed. by O. Ashenfelter, and D. Card, vol. 4, chap. 12, pp. 1043–1171. Elsevier.

ACEMOGLU, D., AND P. RESTREPO (2017): "Robots and Jobs: Evidence from US Labor Markets," Working Paper 23285, National Bureau of Economic Research.

——— (2018): "Modeling Automation," *AEA Papers and Proceedings*, 108, 48–53.

——— (2019a): "Automation and New Tasks: How Technology Displaces and Reinstates Labor," *Journal of Economic Perspectives*, 33(2), 3–30.

——— (2019b): "The Wrong Kind of AI? Artificial Intelligence and the Future of Labor Demand," Working Paper 25682, National Bureau of Economic Research.

AGHION, P., U. AKCIGIT, AND P. HOWITT (2014): "Chapter 1 - What Do We Learn From Schumpeterian Growth Theory?," in *Handbook of Economic Growth*, ed. by P. Aghion, and S. N. Durlauf, vol. 2 of *Handbook of Economic Growth*, pp. 515 – 563. Elsevier.

AGUIAR, M., M. BILS, K. K. CHARLES, AND E. HURST (2017): "Leisure Luxuries and the Labor Supply of Young Men," Working Paper 23552, National Bureau of Economic Research.

ARMSTRONG, S., AND K. SOTALA (2015): "How we're predicting AI–or failing to," in *Beyond artificial intelligence*, pp. 11–29. Springer.

ARNTZ, M., T. GREGORY, AND U. ZIERAHN (2017): "Revisiting the risk of automation," *Economics Letters*, 159, 157 – 160.

ARROW, K. J., ET AL. (2008): "The Promise of Prediction Markets," *Science*, 320(5878), 877–878.

AUTOR, D., D. DORN, L. F. KATZ, C. PATTERSON, AND J. VAN REENEN (2017): "The Fall of the Labor Share and the Rise of Superstar Firms," Working Paper 23396, National Bureau of Economic Research.

AUTOR, D. H. (2015): "Why Are There Still So Many Jobs? The History and Future of Workplace Automation," *Journal of Economic Perspectives*, 29(3), 3–30.

AUTOR, D. H., L. F. KATZ, AND M. S. KEARNEY (2006): "The Polarization of the U.S. Labor Market," *American Economic Review*, 96(2), 189–194.

AUTOR, D. H., F. LEVY, AND R. J. MURNANE (2003): "The Skill Content of Recent Technological Change: An Empirical Exploration*," *The Quarterly Journal of Economics*, 118(4), 1279–1333.

AZAR, J. A., I. MARINESCU, M. I. STEINBAUM, AND B. TASKA (2018): "Concentration in US Labor Markets: Evidence From Online Vacancy Data," Working Paper 24395, National Bureau of Economic Research.

BANA, S. H. (2019): "Identifying Vulnerable Displaced Workers:The Role of State-Level Occupation Conditions," working paper.

BARANY, Z. L., AND C. SIEGEL (2018): "Job Polarization and Structural Change," *American Economic Journal: Macroeconomics*, 10(1), 57–89.

BARKAI, S. (2017): "Declining Labor and Capital Shares," mimeo.

BEAUDRY, P., D. A. GREEN, AND B. M. SAND (2016): "The Great Reversal in the Demand for Skill and Cognitive Tasks," *Journal of Labor Economics*, 34(S1), 199–247.

BEAUDRY, P., AND E. LEWIS (2014): "Do Male-Female Wage Differentials Reflect Differences in the Return to Skill? Cross-City Evidence from 1980-2000," *American Economic Journal: Applied Economics*, 6(2), 178–194.

BICK, A., N. FUCHS-SCHÜNDELN, AND D. LAGAKOS (2018): "How Do Hours Worked Vary with Income? Cross-Country Evidence and Implications," *American Economic Review*, 108(1), 170–99.

BLACK, S. E., AND A. SPITZ-OENER (2010): "Explaining Women's Success: Technological Change and the Skill Content of Women's Work," *The Review of Economics and Statistics*, 92(1), 187–194.

BÖHM, M. (2018): "The Price of Polarization: Estimating Task Prices under Routine-Biased Technical Change," working paper.

BOPPART, T., AND P. KRUSELL (forthcoming): "Labor Supply in the Past, Present, and Future: A Balanced-Growth Perspective," *Journal of Political Economy*.

BRYNJOLFSSON, E., AND A. MCAFEE (2014): *The second machine age: Work, progress, and prosperity in a time of brilliant technologies*. WW Norton & Company.

BUREAU OF LABOR STATISTICS (1986): *Occupational Outlook Handbook, 1986-87 Edition : Bulletin of the United States Bureau of Labor Statistics, No. 2250*United States Department of Labor.

——— (2017): *Occupational Outlook Handbook*United States Department of Labor.

CASELLI, F., AND A. MANNING (forthcoming): "Robot Arithmetic: New Technology and Wages," *American Economic Review: Insights*.

CORTES, G. M. (2016): "Where Have the Middle-Wage Workers Gone? A Study of Polarization Using Panel Data," *Journal of Labor Economics*, 34(1), 63–105.

DAUTH, W., S. FINDEISEN, J. SUEDEKUM, AND N. WOESSNER (2018): "Adjusting to Robots: Worker-Level Evidence," Working Papers 13, Federal Reserve Bank of Minneapolis, Opportunity and Inclusive Growth Institute.

DAVIS, S., AND T. VON WACHTER (2011): "Recessions and the Costs of Job Loss," *Brookings Papers on Economic Activity*, 42(2 (Fall)), 1–72.

DE LOECKER, J., AND J. EECKHOUT (2017): "The Rise of Market Power and the Macroeconomic Implications," Working Paper 23687, National Bureau of Economic Research.

——— (2018): "Global Market Power," Working Paper 24768, National Bureau of Economic Research.

DEMING, D. J. (2017): "The Growing Importance of Social Skills in the Labor Market*," *The Quarterly Journal of Economics*, 132(4), 1593–1640.

EDIN, P.-A., T. EVANS, G. GRAETZ, S. HERNNÄS, AND G. MICHAELS (2019): "Individual Consequences of Occupational Decline," working paper.

EDIN, P.-A., P. FREDRIKSSON, M. NYBOM, AND B. ÖCKERT (2017): "The Rising Return to Non-Cognitive Skill," IZA Discussion Papers 10914, Institute for the Study of Labor (IZA).

EGGERTSSON, G. B., J. A. ROBBINS, AND E. G. WOLD (2018): "Kaldor and Pikettys Facts: The Rise of Monopoly Power in the United States," Working Paper 24287, National Bureau of Economic Research.

ELSBY, M., B. HOBIJN, AND A. SAHIN (2013): "The Decline of the U.S. Labor Share," *Brookings Papers on Economic Activity*, 44(2 (Fall)), 1–63.

FENG, A., AND G. GRAETZ (forthcoming): "Training Requirements, Automation, and Job Polarization," *Economic Journal*.

FORD, M. (2015): *Rise of the Robots: Technology and the Threat of a Jobless Future*. Basic Books.

FREY, C. B., AND M. A. OSBORNE (2017): "The future of employment: How susceptible are jobs to computerisation?," *Technological Forecasting and Social Change*, 114, 254 – 280.

GATHMANN, C., I. HELM, AND U. SCHÖNBERG (2018): "Spillover Effects of Mass Layoffs," *Journal of the European Economic Association*.

GOOS, M., AND A. MANNING (2007): "Lousy and Lovely Jobs: The Rising Polarization of Work in Britain," *The Review of Economics and Statistics*, 89(1), 118–133.

GOOS, M., A. MANNING, AND A. SALOMONS (2014): "Explaining Job Polarization: Routine-Biased Technological Change and Offshoring," *American Economic Review*, 104(8), 2509–26.

GORDON, R. J. (2012): "Is U.S. Economic Growth Over? Faltering Innovation Confronts the Six Headwinds," Working Paper 18315, National Bureau of Economic Research.

——— (2014): "The Demise of U.S. Economic Growth: Restatement, Rebuttal, and Reflections," Working Paper 19895, National Bureau of Economic Research.

GRAETZ, G., AND G. MICHAELS (2018): "Robots at Work," *The Review of Economics and Statistics*, 100(5), 753–768.

HERRENDORF, B., R. ROGERSON, AND A. VALENTINYI (2014): "Chapter 6 - Growth and Structural Transformation," in *Handbook of Economic Growth*, ed. by P. Aghion, and S. N. Durlauf, vol. 2 of *Handbook of Economic Growth*, pp. 855 – 941. Elsevier.

HUTTUNEN, K., J. MOEN, AND K. G. SALVANES (2011): "How Destructive is Creative Destruction? Effects of Job Loss on Job Mobility, Withdrawal, and Income," *Journal of the European Economic Association*, 9(5), 840–870.

JACOBSON, L. S., R. J. LALONDE, AND D. G. SULLIVAN (1993): "Earnings Losses of Displaced Workers," *The American Economic Review*, 83(4), 685–709.

JAIMOVICH, N., AND H. E. SIU (forthcoming): "Job Polarization and Jobless Recoveries," *Review of Economics and Statistics*.

JONES, C. (2016): "Chapter 1 - The Facts of Economic Growth," vol. 2 of *Handbook of Macroeconomics*, pp. 3 – 69. Elsevier.

KAMBOUROV, G., AND I. MANOVSKII (2009): "Occupational specificity of human capital," *International Economic Review*, 50(1), 63–115.

KARABARBOUNIS, L., AND B. NEIMAN (2014): "The Global Decline of the Labor Share," *The Quarterly Journal of Economics*, 129(1), 61–103.

——— (2018): "Accounting for Factorless Income," NBER Working Papers 24404, National Bureau of Economic Research, Inc.

KRUEGER, A. B. (1999): "Measuring Labor's Share," *American Economic Review*, 89(2), 45–51.

——— (2017): "Where Have All the Workers Gone? An Inquiry into the Decline of the U.S. Labor Force Participation Rate," *Brookings Papers on Economic Activity*, 48(2 (Fall)), 1–87.

LIPSIUS, B. (2018): "Labor Market Concentration does not Explain the Falling Labor Share," 2018 Papers pli1202, Job Market Papers.

MICHAELS, G., A. NATRAJ, AND J. VAN REENEN (2014): "Has ICT Polarized Skill Demand? Evidence from Eleven Countries over Twenty-Five Years," *The Review of Economics and Statistics*, 96(1), 60–77.

MORAVEC, H. (1988): *Mind Children*. Harvard University Press.

NORDHAUS, W. D. (2007): "Two Centuries of Productivity Growth in Computing," *The Journal of Economic History*, 67(01), 128–159.

O'MAHONY, M., AND M. P. TIMMER (2009): "Output, Input and Productivity Measures at the Industry Level: The EU KLEMS Database," *Economic Journal*, 119(538), F374–F403.

RINZ, K. (2018): "Labor Market Concentration, Earnings Inequality, and Earnings Mobility," CARRA Working Papers 2018-10, Center for Economic Studies, U.S. Census Bureau.

ROSSI-HANSBERG, E., P.-D. SARTE, AND N. TRACHTER (2018): "Diverging Trends in National and Local Concentration," Working Paper 25066, National Bureau of Economic Research.

SHILLER, R. J. (2019): "Narratives about technology-induced job degradation then and now," *Journal of Policy Modeling*.

SNOWBERG, E., J. WOLFERS, AND E. ZITZEWITZ (2013): "Chapter 11 - Prediction Markets for Economic Forecasting," in *Handbook of Economic Forecasting*, ed. by G. Elliott, and A. Timmermann, vol. 2 of *Handbook of Economic Forecasting*, pp. 657 – 687. Elsevier.

TABER, C., AND N. ROYS (2017): "Skill Prices, Occupations and Changes in the Wage Structure," Discussion paper.

TIMMER, M., T. VAN MOERGASTEL, E. STUIVENWOLD, G. YPMA, M. O'MAHONY, AND M. KANGASNIEMI (2007): "EU KLEMS Growth and Productivity Accounts Version 1.0," mimeo, University of Groningen.

TRAINA, J. (2018): "Is Aggregate Market Power Increasing? Production Trends Using Financial Statements," mimeo.

VENERI, C. M. (1997): "Evaluating the 1995 occupational employment projections," *Monthly Labor Review*, 120.

Theory appendix

Here I formally state and derive the decomposition (6) developed by AR. In the multi-sector economy, total value added (GDP) equals $Y = \sum_i P_i Y_i$, where P_i and Y_i are the price level and value added in industry i. Sectoral value added shares are denoted by $\chi_i \equiv P_i Y_i / Y$. Sectoral factor levels and prices are N_i, K_i and w_i, r_i, respectively. Sectoral wage bill shares are denoted by $\ell_i \equiv w_i N_i / (wN)$ and sectoral labor shares are $s_{N,i} \equiv w_i N_i / (P_i Y_i)$. The production function for sectoral output is given by (5). Since I carry out the decomposition separately for each country, I omit the subscript c throughout this appendix. To keep the notation light, I also omit the time subscript t.

Output markets are perfectly competitive, as are factor markets. However, no assumptions on factor mobility are made, so that factor prices are allowed to vary across sectors. Profit maximization implies the first-order condition (FOC)

$$\left(\frac{\alpha_i}{1-\alpha_i}\right)^{\frac{1}{\sigma}} \left(\frac{A_K}{A_N}\right)^{\frac{\sigma-1}{\sigma}} \left(\frac{K_i}{N_i}\right)^{\frac{1}{\sigma}} = \frac{r_i}{w_i},$$

and the FOC can be used to derive an expression for sectoral labor shares analogous to (3),

$$s_{N,i} = \frac{1}{1 + \frac{\alpha_i}{1-\alpha_i}\left(\frac{r_i/A_K}{w_i/A_N}\right)^{1-\sigma}}. \tag{A1}$$

Formally, the decomposition (6) is written as

$$\begin{aligned} d\log w \;=\; & d\log(Y/N) && \text{Productivity effect} \\ & + \textstyle\sum_i \frac{s_{N,i}}{s_N} d\chi_i && \text{Composition effect} \\ & + \textstyle\sum_i \ell_i (1-\sigma)(1-s_{N,i})(d\log(w_i/A_N) - d\log(r_i/A_K)) && \text{Substitution effect} \\ & + \textstyle\sum_i \ell_i \frac{1-s_{N,i}}{\alpha_i} d\log(1-\alpha_i). && \text{Change in task content} \end{aligned} \tag{A2}$$

The derivation of (A2) starts by noting that the economy-wide labor share $s_N \equiv wN/Y$ equals the sum over the products of sectoral value added and labor shares, and so $wN = \sum_i Y \chi_i s_{N,i}$. Totally differentiating this equation yields

$$dw \times N + w \times dN = \frac{dY}{Y} \times wN + \sum_i Y \times d\chi_i \times s_{N,i} + \sum_i Y \times \chi_i \times ds_{N,i},$$

and after dividing by wN and rearranging we have

$$\frac{dw}{w} = \frac{dY}{Y} - \frac{dN}{N} + \sum_i \frac{s_{N,i}}{s_N} d\chi_i + \sum_i Y \times \frac{\chi_i}{wN} \times ds_{N,i}.$$

Recalling that $dx/x = d\log x$, it is clear that the first two terms after the equality sign make up the productivity effect. The third term is the composition effect. It remains to develop the last term. By

definition, $\sum_i Y \times \frac{\chi_i}{wN} \times ds_{N,i} = \sum_i \ell_i \times \frac{ds_{N,i}}{s_{N,i}}$. Let $x_i \equiv \dfrac{r_i/A_K}{w_i/A_N}$. Totally differentiating (A1) yields

$$ds_{N,i} = -s_{N,i}^2 \times \left[\frac{d\alpha_i}{(1-\alpha_i)^2} x_i^{1-\sigma} + \frac{\alpha_i}{1-\alpha_i}(1-\sigma)x_i^{-\sigma} dx_i \right],$$

and rearranging further (noting that $s_{N,i}\frac{\alpha_i}{1-\alpha_i}x_i^{1-\sigma} = 1 - s_{N,i}$) leads to

$$\frac{ds_{N,i}}{s_{N,i}} = -(1-s_{N,i}) \times \left[-\frac{1}{\alpha_i}\frac{d(1-\alpha_i)}{(1-\alpha_i)} + (1-\sigma)\frac{dx_i}{x_i} \right].$$

Rearranging once more gives the decomposition of changes in the sectoral labor share,

$$\frac{ds_{N,i}}{s_{N,i}} = \underbrace{-(1-s_{N,i})(1-\sigma)\frac{dx_i}{x_i}}_{\text{Substitution effect}} + \underbrace{\frac{1}{\alpha_i}\frac{d(1-\alpha_i)}{(1-\alpha_i)}}_{\text{Change in task content}}.$$

The aggregate decomposition is obtained by computing the employment-weighted sum of the sectoral decompositions.

The mapping from this decomposition of infinitesimal changes to one involving discrete changes that can be implemented empirically involves a number of Taylor expansions, among other things. See the online appendix of AR for details.

Appendix figures and tables

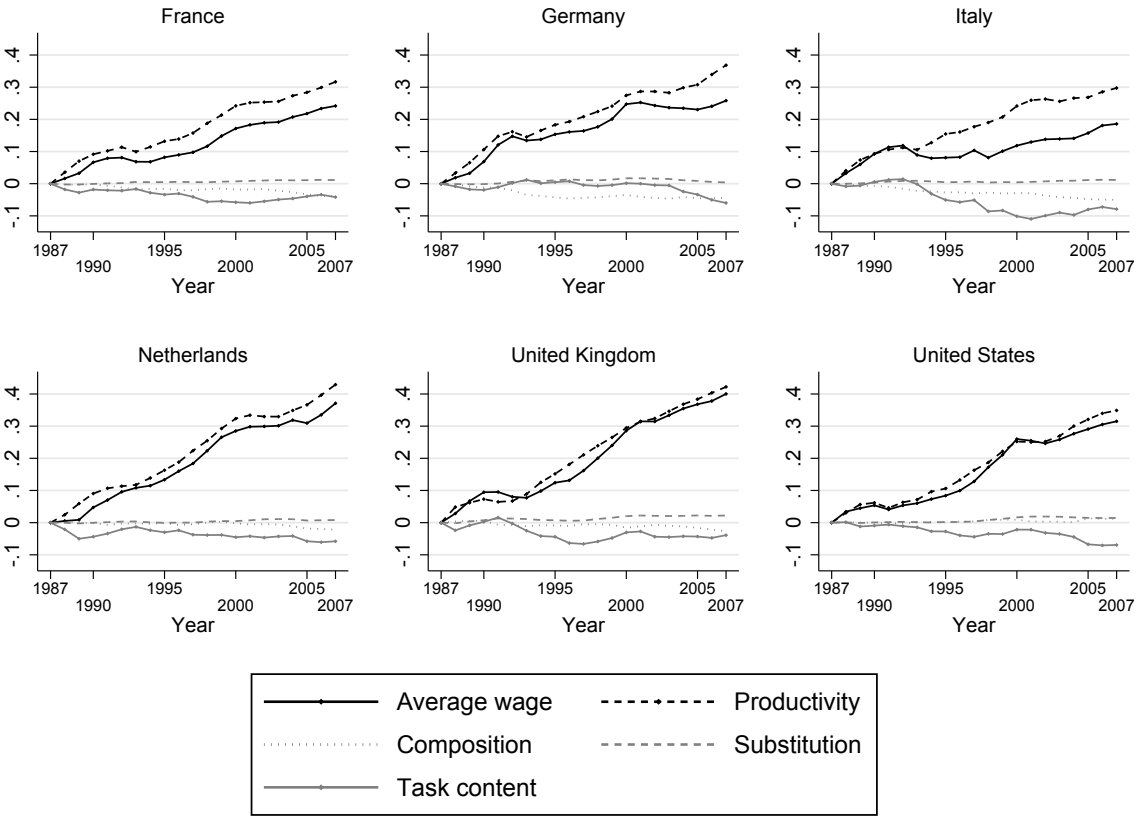

Notes: The figures plot the results from the decomposition (6) which is due to Acemoglu and Restrepo (2019a). The labor share used in the calculations is based on all labor earnings, including earnings of the self-employed.

Figure A1: Decomposition of wage growth (all labor)

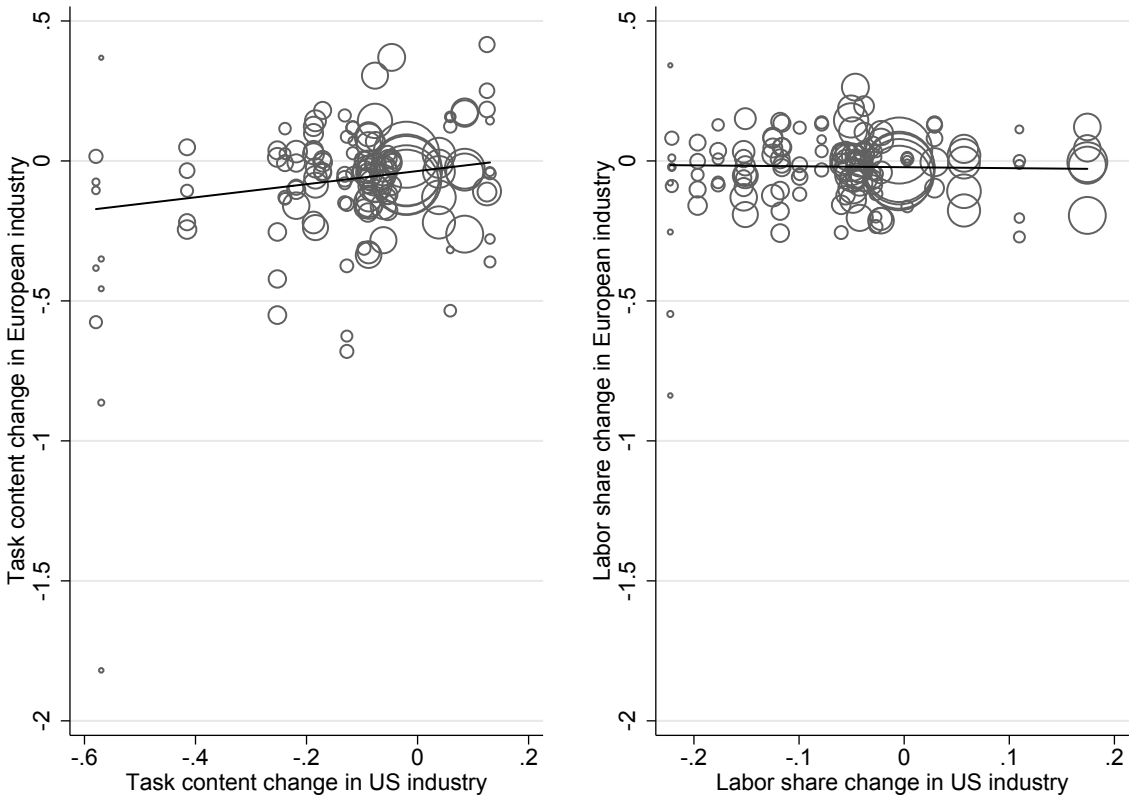

Notes: The figures plot changes in task content and labor shares 1987-2007 in 27 industries across five European countries against changes in task content in the same industries in the US. Bubbles are scaled according to initial within-country employment shares. The fitted lines correspond to the regressions in column (3), panels A and B of Table 1.

Figure A2: Industry-level changes in labor share and task content, Europe versus the United States (employees)

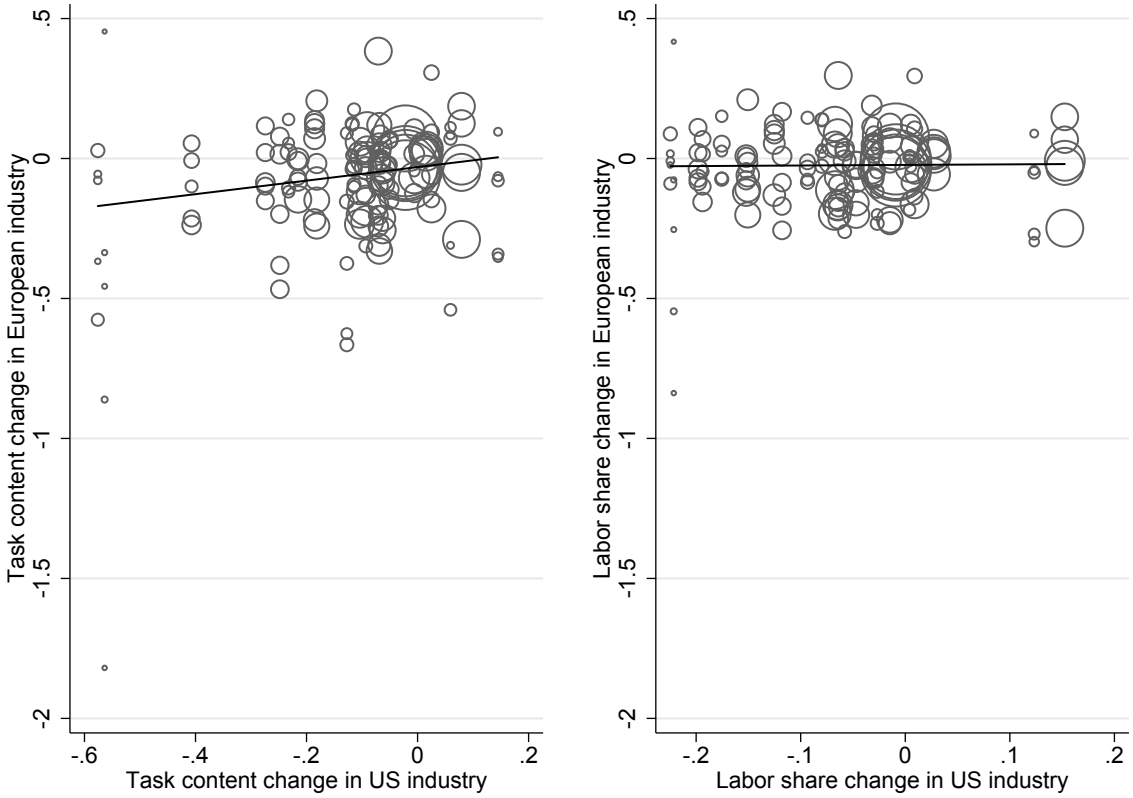

Notes: The figures plot changes in task content and labor shares 1987-2007 in 27 industries across five European countries against changes in task content in the same industries in the US. Bubbles are scaled according to initial within-country employment shares. The fitted lines correspond to the regressions in column (3), panels A and B of Table A1.

Figure A3: Industry-level changes in labor share and task content, Europe versus the United States (all labor)

Table A1: Industry-level changes in labor share and task content, Europe versus the United States (all labor)

	(1)	(2)	(3)	(4)
A. Changes in task content				
Change in US	0.41	0.23	0.24	0.22
	(0.21)	(0.093)	(0.091)	(0.085)
B. Changes in the labor share				
Change in US	0.080	-0.027	0.020	0.014
	(0.17)	(0.10)	(0.076)	(0.073)
Outlier removed		✓		✓
Weighted by initial employment share			✓	✓
Observations	135	134	135	134

Notes: Results are shown from regressions of changes in task content and labor shares (all labor) 1987-2007 in 27 industries across five European countries against changes in task content in the same industries in the US. Robust standard errors, clustered by industry, in parentheses.