Oxford Open Economics, 2024, **3**, i906–i929
https://doi.org/10.1093/ooec/odad082
Dimensions of Inequality: The IFS Deaton Review
The labour market

A task-based approach to inequality

Daron Acemoglu[1,*] and Pascual Restrepo[2]

[1]Department of Economics, Massachusetts Institute of Technology, Cambridge, MA 02142, USA
[2]Department of Economics, Boston University, Boston, MA 02215, USA
*Correspondence address. E-mail: daron@mit.edu

ABSTRACT

This article reviews recent work on how automation and task displacement have contributed to labour share declines and inequality in the US labour market. We summarize the basic building blocks of a task-based framework in which a set of tasks is allocated between capital, skilled labour and unskilled labour. Automation, which corresponds to the use of new technologies expanding the set of tasks that can be performed by capital, always reduces the labour share in value added and may depress overall wages and employment. The negative effects of automation on labour share and its potentially adverse consequences for labour demand can be counterbalanced by the creation of new labour-intensive tasks, which can reinstate labour into the production process. We also show that when automation displaces unskilled labour from the tasks in which they used to specialize (which has been its modal impact so far), it increases the demand for skills and inequality. New tasks may or may not limit the increase in the demand for skills depending on whether they are mostly targeted at skilled workers. We then provide a range of evidence supporting the basic predictions and implications of this framework. Most importantly, the decline in the share of labour in national income and the increase in the demand for skills appear to be related to an acceleration in the pace of automation and a deceleration in technological changes complementing humans during the last 30 years. We end with a discussion of the potential bias towards automation in the development and adoption of digital technologies, and how this will affect the nature of work in the face of recent advances in artificial intelligence.

Key words: automation; capital; demands for skills; displacement effect; labour demand; labour share; inequality; productivity; reinstatement effect; tasks; technology; wages

1. Introduction

Recent breakneck-paced advances in artificial intelligence (AI) have intensified concerns about the future of work and inequality. The optimistic scenario is that tools such as large language models can increase productivity and provide resources for workers to perform their jobs more successfully. More concerning are scenarios where AI technologies will automate a large number of jobs, boost inequality and even reduce wages for many worker groups. Any understanding of whether these concerns are well placed and whether there is any validity to the optimistic scenario must start with a clear conceptual framework in which new technologies can simultaneously substitute for and complement human work.

In this article, we review such a framework that we and other researchers have developed, how this framework accounts for changes in the wage structure in the industrialized world over the last several decades and draw out its implications for the adjustment of labour markets to existing and oncoming advances in AI and other digital technologies.

The background to this discussion is by now familiar. Labour market inequality has risen significantly in many industrialized economies, and especially in the USA and the UK, over the last 4 decades (Acemoglu and Autor 2011; Peterson Institute 2020). Simultaneously, many of these economies have also experienced a decline in the labour share in national income.[1] Despite a voluminous literature on both topics, these trends remain imperfectly understood. The leading explanations for both relate to the changing nature of technological progress. For example, new technologies, such as computers, are argued to be skill-biased and to have raised the productivity of skilled workers (e.g. those with college or post-graduate degrees) more than those of less skilled workers (see Krueger 1993; Autor et al. 1998).

In the most canonical approach to this problem, which dates back to the path-breaking work of Tinbergen (1974), output is assumed to be produced with an aggregate production function of the form: $G(A_H H, A_L L)$. Here, H and L are employment levels of skilled and unskilled (or high- and low-skill) workers, respectively, and A_H and A_L represent technologies augmenting these two types of workers. Skill-biased technical change (SBTC) in this framework corresponds to an increase A_H (relative to A_L).[2] Likewise, most economists think of capital–labour substitution and the effects of productivity on the shares of these two factors using a similar production function with factor-augmenting technologies, $F(A_K K, A_L L)$, where K is capital and L is now total labour.

These popular and influential frameworks face several shortcomings, however. First of all, they lack descriptive realism. The assumption that technologies take a factor-augmenting form as posited in these production functions does not have a clear empirical counterpart. Most technologies improve the productivity of

[1] See, for example, Elsby et al. (2013), Karabarbounis and Neiman (2013), Piketty (2014), Dao et al. (2019) and Autor et al. (2020). However, Gutierrez and Piton (2020) argue that the decline in the labour share is a US phenomenon.

[2] This requires that the elasticity of substitution between H and L is greater than one. This framework is developed in detail in pioneering work by Katz and Murphy (1992) and Goldin and Katz (2008).

Received: September 9, 2023. **Revised:** September 9, 2023. **Accepted:** October 29, 2023

a factor in some tasks (e.g. a better paintbrush makes a worker better at painting, but not necessarily in other tasks), improve the productivity of some industry, reallocate some tasks from one factor to another (as with the spinning and weaving technologies that started the British Industrial Revolution in the middle of the 18th century), or create new tasks, invent new goods or introduce completely new ways of combining existing tasks or intermediates. None of these easily fits into the factor-augmenting framework (see Acemoglu and Restrepo 2019).

Second, and perhaps more importantly, these frameworks make a range of counterfactual predictions. For example, assuming that there is no technological regress, new technologies should not reduce the productivity of any of the factors. For example, skill-biased technological change can benefit college graduates more than high school graduates, but should never reduce the real wages of high school graduates. But declining real wages of low-educated men has been a persistent trend in the US labour market over the last 4 decades (as we show further below; see also Acemoglu and Autor 2011; Autor 2019). Similarly, with just factor-augmenting technological changes, it is difficult to have instances in which new technologies reduce labour demand, employment and wages. Yet again, we have important examples of new technologies, especially automation technologies such as industrial robots, that have been associated with lower wages and employment (e.g. Acemoglu and Restrepo 2020*a*).

Third, and relatedly, a framework based on factor-augmenting technologies does not generate meaningful changes in the labour share (as we will argue below). For realistic values of the elasticity of substitution, generating the changes in the share of labour experienced in manufacturing in the USA via factor-augmenting technologies would necessitate huge changes in technology. These improvements would be associated with very large and counterfactual increases in total factor productivity (TFP). Likewise, a problem of the standard SBTC model, which is omitted, is that to match the observed changes in the skill premium, the model would need unrealistically large TFP changes (Acemoglu and Restrepo 2020*b*).

Fourth, new advances in AI also highlight the inadequacy of a framework that relies on factor-augmenting technologies. The possibility that different types of AI applications can automate work, or remedy for expertise inadequacies for some specialized workers (e.g. Noy and Zhang 2023; Brynjolfsson et al. 2023), or create complementarities for the highest-skill employees cannot be incorporated into this type of benchmark model, and instead calls for a framework in which new technologies can have very different effects on specific tasks, as well as introduce new tasks.

In a series of papers, we have argued that these problems can be overcome with a more micro-based approach in which we start modelling the production process at the task level and allow for technologies that automate certain tasks, displacing workers and replacing them with machines or algorithms (see Acemoglu and Restrepo 2018, 2019, 2020*a*, 2022).[3] While the SBTC framework and the standard capital–labour aggregate production function focus on how technology complements various factors, our task-based approach emphasises how several major types of technologies replace and displace labour from the production process. In addition to its stronger microfoundations, our framework makes a range of more realistic predictions—for example, linking the impact of automation technologies on labour to whether the negative displacement effect is outweighed by the positive productivity effect. This framework also allows us to think about new labour-intensive tasks, which reinstate labour into the production process, and explains how the evolution of labour demand depends on the balance between the displacement effect created by automation technologies and the reinstatement brought by new tasks. Large changes in the labour share (at the industry or the economy level) are telltale signs in this framework of an imbalance between displacement and reinstatement.

In this article, we present a tractable version of the framework that we have developed in past work. In Section 2, we start with a version of this framework with three factors: capital, skilled labour and unskilled labour. In Section 3, we specialize this framework to include only capital and one type of labour, and study how different types of technologies affect the demand for labour and the share of labour in value added. We also emphasise how a countervailing set of technological changes—which introduce new labour-intensive tasks—are important both conceptually and empirically, and we explore their implications. In Section 4, we return to the full framework from Section 2 and explore how automation and new tasks affect inequality. We emphasize in this section how these types of changes, which alter the allocation of tasks to factors and change the ‘task content of production’, can have major effects on inequality without generating much productivity or TFP changes. This contrasts with the implications of factor-augmenting technologies as we also document.

In Section 5, we consider a multi-sector extension of our framework in order to clarify how the extent of displacement and reinstatement can be measured. In Section 6, we estimate changes in the task content of production, and especially the extent of displacement and reinstatement, in the USA since World War II, based on the approach outlined in Section 5. The most important results in this section are as follows. First, we find that measures of displacement and reinstatement are meaningful and strongly related to proxies for automation and new tasks, respectively. Second, we uncover a striking change in the extent of displacement and reinstatement in the US economy. The extent of displacement and reinstatement used to be balanced in the 4 decades after World War II. Starting in the 1980s, a very different pattern emerged: rapid displacement and very little countervailing reinstatement. It is this imbalance that accounts for the decline in the share of labour in the aggregate economy and in manufacturing. Finally, we also substantiate the claim that without major changes in the task content of production, it would take gargantuan changes in TFP to account for the observed changes in the labour share.

In Section 7, we turn to the empirics of the changes in the demand for skills. Using the measures of displacement and reinstatement estimated in Section 6, we show that displacement has been far from neutral. Industries undergoing rapid displacement have been exactly the ones increasing their demand for skills. Interestingly, we find that reinstatement was counterbalancing the increase in the demand for skills coming from automation until the 1980s, but has itself become an additional force towards greater demand for skills and inequality since the 1980s. We provide a number of possible explanations for this pattern. We conclude in Section 8 by summarizing the main elements of our framework and drawing out its implications for future work and inequality in the face of AI advances.

[3] These papers build on Zeira (1998), Acemoglu and Zilibotti (2001), Hellwig and Irmen (2001), Autor et al. 2003 and Acemoglu and Autor (2011).

2. Conceptual framework

At the centre of our conceptual framework are tasks that need to be performed for production, and they can be performed using different types of labour or capital. For this article, we focus on two types of labour: unskilled and skilled (low- and high-skill) labour. Automation corresponds to an expansion of the set of tasks that can be produced by capital, which can come at the expense of either skilled or unskilled labour. We also allow the introduction of new tasks in which (skilled or unskilled) labour has a comparative advantage relative to capital. We start with a model for a single sector. We then embed this structure in a multi-sector set-up.

The unique final good is produced from a mass M of tasks $x \in \mathcal{T}$ combined via a CES aggregator,

$$Y = \left(\frac{1}{M}\int_{\mathcal{T}} (My(x))^{(\lambda-1)/\lambda} dx\right)^{\lambda/(\lambda-1)},$$

where $\lambda \geq 0$ is the elasticity of substitution between tasks. Tasks are performed by unskilled labour, $\ell(x)$, skilled labour $h(x)$ or capital $k(x)$,

$$y(x) = \psi_L(x)\ell(x) + \psi_H(x)h(x) + \psi_K(x)k(x),$$

where $\psi_j(x) \equiv A_j \cdot \gamma_j(x)$ for $j \in \{L, H, K\}$ denotes the productivity of factor j at task x. If a particular task cannot be performed by a factor, then we set the corresponding $\psi_j(x) = 0$.

We assume that skilled and unskilled labour are supplied inelastically, with market-clearing conditions $L = \int_{\mathcal{T}} \ell(x)dx$ and $H = \int_{\mathcal{T}} h(x)dx$, and also take the supply of capital, K, as given and impose the market-clearing condition $K = \int_{\mathcal{T}} k(x)dx$. We denote by $\mathcal{T}_L$, $\mathcal{T}_H$ and $\mathcal{T}_K$ the set of tasks performed by each factor.[4] A competitive equilibrium is represented by an allocation of tasks to factors and production of capital goods that maximizes output. It can be shown that equilibrium output can then be expressed as

$$Y = \left(\Gamma_K^{1/\lambda}(A_K K)^{(\lambda-1)/\lambda} + \Gamma_L^{1/\lambda}(A_L L)^{(\lambda-1)/\lambda} + \Gamma_H^{1/\lambda}(A_H H)^{(\lambda-1)/\lambda}\right)^{\lambda/(\lambda-1)},$$

where the share parameters, Γ_K, Γ_L and Γ_H, are endogenously determined and summarize the task content of production. More specifically, they represent the range of tasks performed by the two types of labour and are given as

$$\Gamma_j = \frac{1}{M}\int_{\mathcal{T}_j} \psi_j(x)^{\lambda-1} dx \text{ for } j \in \{K, L, H\}.$$

The effects of various technologies on the skill premium can be expressed as

$$d\ln\left(\frac{w_H}{w_L}\right) = -\frac{1}{\sigma} d\ln\left(\frac{H}{L}\right) + \frac{\sigma-1}{\sigma} d\ln\left(\frac{A_H}{A_L}\right) + \frac{1}{\lambda} d\ln\left(\frac{\Gamma_H}{\Gamma_L}\right)\Bigg|_{(A_H H)/(A_L L)}, \quad (1)$$

where the last term represents changes in the task content of production coming from technological changes. Indeed, the canonical model generates a relationship between technology and the skill premium given by the first line of equation (1).[5] Moreover, as opposed to the canonical model, the elasticity of substitution between skilled and unskilled labour, σ, is not an exogenous parameter, but reflects the competition between skilled and unskilled labour for tasks (and the corresponding changes in the allocation of these tasks between skilled and unskilled labour in response to changes in factor-augmenting technologies and relative supplies). Namely,

$$\sigma = \lambda \Big/ \left(1 - \frac{\partial \ln(\Gamma_H/\Gamma_L)}{\partial \ln(A_H H/A_L L)}\right) \geq \lambda.$$

This elasticity reflects two types of substitution: substitution between tasks, represented by λ (with more productive skilled labour, there is greater production of skill-intensive tasks); and substitution at the extensive margin whereby some tasks are reallocated from unskilled labour and capital to skilled labour. It is because of this second type of substitution that $\sigma \geq \lambda$.[6]

In addition to factor-augmenting changes—the A_L, A_H and A_K terms—that increase the productivity of a factor in all tasks, this framework enables us to analyse the impact of technologies that affect the productivity of a factor in some tasks. Particularly relevant is 'automation'—changes that enable capital to be used in tasks that were previously performed by labour (or equivalently increase the productivity of capital in such tasks). For example, robots can become more productive in welding, a task previously performed by human welders, who will now be displaced from the tasks in which they specialized. The effects of automation and other technological changes affecting the allocation of tasks to factors work through the last term in equation (1). Formally, automation will correspond to an increase in $\psi_K(x)$ for a set of tasks currently not in $\mathcal{T}_K$. This type of advance in automation technology will lead to an expansion in the set of tasks allocated to capital, $\mathcal{T}_K$, and a contraction in the set of tasks allocated to workers. The left panel of Fig. 1 illustrates the direct effect of automation on the allocation of tasks to workers and capital. The figure shows the effects of automating a subset of tasks $\mathcal{D}$ performed by low-skill labour, though one could model the automation of tasks performed by high-skill labour in the same way.

In addition to automation, we are interested in the effects of new labour-intensive tasks. The addition of new tasks that are performed by high-skill labour can be thought of as an expansion of the set $\mathcal{T}$ to $\mathcal{T}' = \mathcal{T} + \mathcal{N}$ such that $\mathcal{N} \subset \mathcal{T}'_H$, as shown in the right panel of Fig. 1. The addition of new tasks performed by low-skill labour can be modelled analogously.

Before we study the effects of automation on inequality, we clarify how automation affects labour as a whole. To do this, in the next section we specialize the model to include only one type of labour, and we discuss in some detail how automation affects wages and employment—and labour demand.

3. Effects of automation and new tasks on labour demand

Suppose that there is only one type of labour, denoted by L. The equilibrium in this case is simply a special instance of the one

[4] Formally, these sets are given from cost minimization as

$$\mathcal{T}_L = \left\{x : \frac{w_L}{\psi_L(x)} < \frac{w_H}{\psi_H(x)}, \frac{w_L}{\psi_L(x)} < \frac{1}{\psi_K(x)}\right\},$$

$$\mathcal{T}_H = \left\{x : \frac{w_H}{\psi_H(x)} \leq \frac{w_L}{\psi_L(x)}, \frac{w_H}{\psi_H(x)} < \frac{1}{\psi_K(x)}\right\},$$

$$\mathcal{T}_K = \left\{x : \frac{1}{\psi_K(x)} \leq \frac{w_L}{\psi_L(x)}, \frac{1}{\psi_K(x)} \leq \frac{w_H}{\psi_H(x)}\right\}.$$

[5] As derived in Acemoglu and Restrepo (2020b). A more general framework is analysed in Acemoglu and Restrepo (2022). For brevity, we refer the reader to these two papers and do not repeat the proofs of the claims here.

[6] Put differently, if we kept Γ_H/Γ_L constant, then the first line of this equation would also have λ instead of σ.

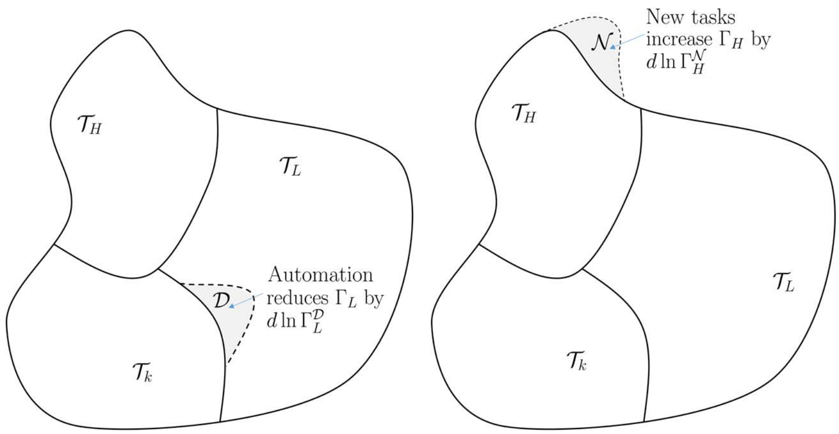

Figure 1: Effects of automation and new task creation on the share of tasks performed by workers and capital

characterized in the previous subsection. To further simplify our derivations, we assume that the allocation of tasks to factors is fully determined by technology and does not respond to changes in factor supplies (see Acemoglu and Restrepo (2022), for a complete treatment of the more general case).[7] Because the allocation of tasks remains unchanged following small changes in factor supplies, the elasticity of substitution between capital and labour is equal to λ.

To understand the effects of automation, consider an improvement in the productivity of capital $\psi_k(x)$ as a set of tasks $\mathcal{D}$ (which we assume is small) currently in $\mathcal{T}_L$, as shown in Fig. 1. Suppose also that, following these improvements, the tasks in $\mathcal{D}$ are reallocated from labour to capital. As a result of this displacement, the share of tasks performed by labour falls by

$$d\ln\Gamma_L^{\mathcal{D}} = -\frac{\int_{\mathcal{D}} \psi_L(x)^{\lambda-1}dx}{\int_{\mathcal{T}_L} \psi_L(x)^{\lambda-1}dx} \leq 0,$$

and the share of tasks performed by capital rises to

$$d\ln\Gamma_K^{\mathcal{D}} = \frac{\int_{\mathcal{D}} \psi_K(x)^{\lambda-1}dx}{\int_{\mathcal{T}_K} \psi_K(x)^{\lambda-1}dx} \geq 0.$$

The superscript $\mathcal{D}$ in both expressions indicates that these are tasks directly lost to automation.

Likewise, the creation of new tasks in $\mathcal{N}$, where labour has a comparative advantage, will expand the share of tasks performed by labour by

$$d\ln\Gamma_L^{\mathcal{N}} = \frac{\int_{\mathcal{N}} \psi_L(x)^{\lambda-1}dx}{\int_{\mathcal{T}_L} \psi_L(x)^{\lambda-1}dx} \geq 0.$$

The superscript $\mathcal{N}$ indicates that these are new tasks.

As we will see, the effects of technology on labour demand will be mediated by the labour share in value added. In this framework, the labour share can be written as

$$s^L(K,L) = \frac{\Gamma^{1/\lambda}(A_L L)^{(\lambda-1)/\lambda}}{\Gamma^{1/\lambda}(A_L L)^{(\lambda-1)/\lambda} + (A_K K)^{(\lambda-1)/\lambda}}, \qquad (2)$$

where

$$\Gamma = \frac{\Gamma_L}{\Gamma_K}$$

captures the task content of production – the relative importance of tasks allocated to labour in the production process. Γ, and thus the labour share, is decreasing in automation and increases with the creation of new labour-intensive tasks.

Labour demand in this economy can be written as

$$W^d(L,K) = \frac{Y(L,K)}{L} \times s^L(K,L).$$

Naturally, labour demand $W^d(L,K)$ is decreasing in L and increasing in K. We next analyse the effects of different types of technologies on labour demand.[8]

We start with automation:

$$\begin{aligned} d\ln W^d(L,K) = \; & d\ln Y && \text{(productivity effect)} \\ & + \frac{1}{\lambda}\left(1-s^L\right) d\ln\Gamma_d^{\mathcal{D}} && \text{(displacement effect)}. \end{aligned}$$

This formula shows that automation has two distinct effects on labour demand. First, there is a 'productivity effect', as automation increases productivity—i.e. $d\ln y > 0$—and raises the demand for labour in non-automated tasks. If nothing else happened, this increase in productivity would directly, and by the same amount, increase labour demand. However, the 'displacement effect', which automation creates, reduces labour demand. This reflects the fact that automation displaces labour from certain tasks and squeezes it into fewer non-automated tasks. Automation raises labour demand when the productivity effect dominates displacement, but reduces it otherwise.

We can also compute the productivity effect as

$$d\ln y = s^L\left(-d\ln\Gamma_L^{\mathcal{D}}\right)\pi > 0,$$

[7] Formally, this entails assuming that capital is sufficiently abundant as to perform all tasks where it has a positive productivity. In general, as emphasized in the previous section, $\mathcal{T}_K$ and $\mathcal{T}_L$ are endogenously determined given capital and labour productivities. Here, we are simplifying the notation and exposition by varying the realized allocation of tasks to factors.

[8] Once the effects of technology on labour demand are determined, how this translates into employment and wage changes is partly regulated by labour supply and partly by labour market imperfections, neither of which we model explicitly in this article (see Acemoglu and Restrepo 2018). It suffices to note that with an upward-sloping (quasi-) labour supply schedule, lower labour demand will translate into both lower employment and lower wages.

where $\pi > 0$ are the average cost-saving gains generated by automation in all tasks in $\mathcal{D}$:

$$\pi = \int_{\mathcal{D}} \frac{\psi_L(x)^{\lambda-1}}{\int_{\mathcal{D}} \psi_L(u)^{\lambda-1} du} \frac{1}{1-\lambda} \left[\left(\frac{W}{A_L \psi_L(x)} \right)^{1-\lambda} - \left(\frac{R}{A_K \psi_K(x)} \right)^{1-\lambda} \right] dx.$$

This expression is intuitive. The productivity gains from automation depend on the difference in the cost of producing the automated tasks with labour, $W/[A_L\psi_L(x)]$, and the cost of producing them with capital, $R/[A_K\psi_K(x)]$. The productivity effect will be stronger when automation significantly increases productivity (because capital is more productive than labour in these tasks).[9] This last point, though simple, is important. Not only does it show that when the effective wage is similar to the effective rental rate, the productivity effect will be small and thus automation will reduce labour demand, but it also implies that, contrary to a common presumption in popular debates, it is not 'brilliant' automation technologies but those that are 'so-so' and generate only small productivity improvements that will tend to worsen the prospects of labour (see Acemoglu and Restrepo 2019).[10]

If the world was one of just automation technologies, this analysis shows that the share of labour in value added would be continuously declining. This, of course, has not been the case for much of the last 200 years of growth in the industrialized world. Acemoglu and Restrepo (2018) argued that this is because of the introduction of new labour-intensive tasks, which have also been a major driver of productivity growth. For example, design tasks, most manufacturing engineering tasks, most back-office activities and all programming occupations are new relative to the first half of the 20th century, and have been major drivers of the growth of labour demand. We now study the labour demand implications of new tasks in which labour has a comparative advantage.

The effects of creation of new tasks can be determined similarly:

$$\begin{aligned} d\ln W^d(L,K) = {} & d\ln Y && \text{(productivity effect)} \\ & + \frac{1}{\lambda}\left(1-s^L\right) d\ln \Gamma_d^{\mathcal{N}} && \text{(reinstatement effect)}. \end{aligned}$$

The new item here is the 'reinstatement effect', which reinstates labour into additional tasks and, via this channel, increases labour demand and the labour share.

Finally, we can also determine the effects of factor-augmenting technologies:

$$\begin{aligned} d\ln W^d(L,K) &= s^L d\ln A_L && \text{(productivity effect)} \\ &\quad + \frac{\lambda-1}{\lambda}\left(1-s^L\right) d\ln A_L && \text{(quality substitution effect)}, \\ d\ln W^d(L,K) &= \left(1-s^L\right) d\ln A_K && \text{(productivity effect)} \\ &\quad + \frac{1-\lambda}{\lambda}\left(1-s^L\right) d\ln A_K && \text{(quality substitution effect)}. \end{aligned}$$

Critically, there is no displacement or reinstatement effect here because there is no reallocation of tasks to factors and hence no change in the task content of production. The new items are the 'quality substitution effects', which capture the change in the pattern of capital–labour substitution resulting from changes in technology. This is because factor-augmenting technologies affect the 'quality' (effective productivity) of the factors, inducing a substitution between capital-intensive and labour-intensive tasks and production when $\lambda \neq 1$—but crucially 'without' a change in the allocation of tasks to factors. Whether this substitution increases or reduces labour demand (and the labour share) depends on whether the elasticity of substitution λ is greater than or less than 1.

4. Effects of automation and new tasks on inequality

Let us now return to the inequality implications of automation and new tasks. Automation can, in principle, displace skilled or unskilled labour. In the context of industrial robotics technology, the evidence presented in Acemoglu and Restrepo (2020*b*, 2022) suggests that most of the automated tasks used to be performed by less skilled workers, and we start with this case.

To do this, suppose that there is an improvement in automation technologies such that the productivity of capital in a set of tasks in $\mathcal{D} \subset \mathcal{T}_L$ leads to the displacement of labour from these tasks, as in Fig. 1. Then we can show (see the appendix of Acemoglu and Restrepo 2020*b*) that inequality changes by

$$d\ln\left(\frac{w_H}{w_L}\right) = -\frac{1}{\sigma} d\ln \Gamma_L^{\mathcal{D}} > 0.$$

Moreover, following such an automation advance, w_L may increase or decrease (which again reflects the countervailing displacement and productivity effects, this time focusing on unskilled workers).[11]

Several points are worth noting. First, the effect of automation technologies on the skill premium is completely driven by the set of tasks (weighted by their effective productivity) unskilled labour loses relative to the entire set of tasks previously performed by these workers (this effect is not mediated by the elasticity of substitution, and σ does not need to be greater than one). This close connection between the set of task reallocations and factor price changes is the main conceptual insight of this class of models.

Second, advances in automation technologies increase TFP, but this impact, coming from cost savings due to automation, may be small (Acemoglu and Restrepo 2020*b*, 2022). In fact, we will see below that, in our estimates, this boost to TFP is often quite small, and the reason for this is explained in the next paragraph.

Third, the magnitude of the change in the skill premium is decoupled from productivity increases. Specifically, in the canonical model (with no changes in the task content of production), we have

$$\left.\frac{d\ln TFP}{d\ln(w_H/w_H)}\right|_{A_L} = s_H \cdot \sigma/(\sigma-1),$$

[9] In addition to the productivity effect, automation may generate additional countervailing forces, raising labour demand. First, automation is likely to induce additional usage of capital in the sector or additional capital accumulation, which can increase labour demand (Acemoglu and Restrepo 2018). Second, there could be 'deepening of automation', meaning increases in the productivity of capital and tasks already automated, which also increases labour demand (Acemoglu and Restrepo 2019). Even factoring in these changes, in this framework, automation always reduces the labour share.

[10] An implication of this analysis is that the productivity effects of different types of technologies can have potentially very different magnitudes, and thus, we cannot generally presume that one set of automation technologies will affect labour demand in exactly the same way as another set—this will depend on their respective productivity effects.

[11] Notice, again, that the relevant elasticity is σ, which reflects the fact that changes in technology will trigger a reallocation of tasks between skilled and unskilled labour. If factor supplies do not alter the allocation of tasks to factors of production, then $\sigma = \lambda$.

where s_H is the share of skilled labour in value added. Thus, to get the demand for skilled labour to increase by 1%, one needs a 0.83% increase in productivity. In contrast, in our task-based framework, the effect of automation on TFP is

$$\frac{d\ln TFP}{d\ln w_H/w_L} = \sigma \cdot s_L \cdot \pi,$$

where $\pi > 0$ is the average proportional cost reduction in automated tasks. This expression shows that when $\pi \rightarrow 0$, our model generates large swings in the skill premium from very small changes in TFP. Because of this difference, our framework generates sizeable changes in the skill premium for reasonable changes in TFP. For example, if automation reduces the cost of producing tasks by $\pi = 30\%$, as in the case of industrial robots (Acemoglu and Restrepo 2020*b*), then the increase in the college premium between 1963 and 1987 can be explained with as little as 0.6% per annum growth in TFP.

Fourth, the unskilled wage may decline, and this happens precisely when the increase in TFP is small, but the skilled wage always increases because tasks produced by other factors, which are q-complements to those produced by skill workers, are becoming cheaper.[12]

To analyse the effects of new tasks on inequality, suppose a small set of new tasks is introduced, as in the right panel of Fig. 1. If skilled workers have comparative advantage in these tasks—i.e. $w_H/\psi_H(x) < w_L/\psi_L(x)$ at current wages—then the skill premium increases by

$$d\ln\left(\frac{w_H}{w_L}\right) = \frac{1}{\sigma} d\ln\Gamma_H^{\mathcal{N}} > 0.$$

If, however, unskilled workers have comparative advantage in these tasks—i.e. $w_L/\psi_L(x) < w_H/\psi_H(x)$ at current wages—then the skill premium would decline.

To interpret these expressions, note that the effect on the skill premium is again a function of the set of tasks reallocated across factors. Analogously, these changes always increase TFP. Also notable is that new tasks may increase or reduce the skill premium, depending on whether they are allocated to skilled or unskilled labour.

Two other types of technological changes can be studied in this framework. The first is 'standardization', which involves the simplification of previously complex, skilled tasks so that they can now be more cheaply performed by unskilled workers (Acemoglu and Restrepo 2018). The second is 'skill upgrading', which involves the transformation of unskilled tasks so that they can be more productively performed by skilled workers (see the appendix of Acemoglu and Restrepo 2020*b*).

5. The multi-sector economy

In order to measure displacement and reinstatement—driven by automation and new tasks—we need to consider a multi-sector version of our economy, which we now develop. Once again, we simplify the exposition by focusing on the case with a single type of labour and where the task allocation does not respond to changes in factor supplies. Relative to the previous sections, we now assume that there are sectors, each one of which has a production function—with its own task-specific productivity for different factors and the possibility of automation, introduction of new tasks and factor-augmenting technologies.

We index sectors by subscript i and let Γ_i represent the task content in industry i. We denote the price of the goods produced by sector i by P_i, while its factor prices are denoted by W_i and R_i. Note that we can write the labour share in a sector as a function of technology and factor prices:

$$s_i^L = \frac{\Gamma_i\left(W_i/A_{i,L}\right)^{1-\lambda}}{\Gamma_i\left(W_i/A_{i,L}\right)^{1-\lambda} + \left(R_i/A_{i,K}\right)^{1-\lambda}}.$$

This shows that sectoral labour shares are decreasing in automation in that sector and increase with the introduction of new tasks (holding factor prices constant).

Total value added (GDP) in the economy is $Y = \sum_{i\in\mathcal{I}} P_i Y_i$, and we define $\chi_i = (P_i Y_i)/Y$ as the share of sector i in total value added.

Denoting the average wage by W and aggregate employment by L, total labour demand is

$$WL = \sum_{i\in\mathcal{I}} W_i L_i = \sum_{i\in\mathcal{I}} Y \times \chi_i \times s_i^L.$$

The effects of a change—of any type—in technology can then be summarized as follows:[13]

$$\begin{aligned} d\ln(WL) = {} & d\ln Y && \text{(productivity effect)} \\ & + \sum_{i\in\mathcal{I}} \ell_i (1-s_i^L)\left[d\ln\Gamma_i^{\mathcal{D}} + d\ln\Gamma_i^{\mathcal{N}}\right] && \text{(change in task content)} \\ & + \sum_{i\in\mathcal{I}} \ell_i (1-\lambda)\left(1-s_i^L\right)\left(d\ln W_i - d\ln R_i\right) && \text{(price substitution effect)} \\ & - \sum_{i\in\mathcal{I}} \ell_i (1-\lambda)\left(1-s_i^L\right)\left(d\ln A_{i,L} - d\ln A_{i,K}\right) && \text{(quality substitution effect).} \end{aligned} \tag{3}$$

Here, $\ell_i = (W_i L_i)/(WL)$ is the share of the wage bill generated in sector i. This decomposition is formally derived in the appendix of Acemoglu and Restrepo (2019) and showcases the several distinct impacts of technology on labour demand. First, there is the multi-sector equivalent of the 'productivity effect': technology raises productivity, which tends to increase aggregate value added, Y, raising the demand for labour.[14] Second, there is a 'composition effect' resulting from sectoral reallocation in response to changes in technology (and this reallocation in turn depends on consumer preferences, among other things). The composition effect increases labour demand when economic activity is reallocated towards labour-intensive sectors (those with $s_i^L > s^L$) and has the opposite effect when the reallocation is towards capital-intensive sectors (those with $s_i^L < s^L$). Third, we come to the main notable feature of our framework: the 'change in task content' resulting from changes in the allocation of tasks to factors. Mathematically, this change in task content is captured by changes in the share parameters represented by the Γ_j terms. Finally, there are changes

12 Some of the automated tasks in $\mathcal{D}$ may be previously performed by skilled workers: artificial intelligence may replace tasks currently employing skilled workers, and many of the iconic innovations of the Industrial Revolution automated the spinning, weaving and knitting tasks previously performed by skilled artisans. If so, automation may have the opposite effect on the skill premium.

13 Because, as in Section 3, the allocation of tasks between capital and labour is also independent of factor supplies, the relevant elasticity of substitution between capital and labour is the elasticity of substitution between tasks, λ.

14 More generally, $d\ln Y = d\ln TFP + s^L d\ln L + \sum_i\left[(R_i K_i)/Y\right] d\ln K_i$. In our set-up, technological improvements increase TFP but their overall impact on GDP depends on the adjustment of labour and capital as well.

resulting from variations in capital–labour substitution; these are themselves a consequence of the same 'quality substitution effect', already emphasized in Section 3, as well as a multi-sector generalization of this effect, the 'price substitution effect', which results from changes in wage to rental rate ratio at the sectoral level. The direction of the impact of these last two effects on labour demand depends on the elasticity of substitution across tasks, λ.

This general decomposition can be applied to study the impact of specific technologies on labour demand. For illustration purposes, consider the introduction of a new automation technology in sector j. This will first generate a displacement effect in the same sector, given by

$$\left(1-s_i^L\right) d\ln\Gamma_j^{\mathcal{D}} < 0,$$

which reduces labour demand. In addition, the substitution of (effectively) cheaper capital for labour increases TFP by

$$d\ln TFP = \chi_j s_j^L\left(-d\ln\Gamma_j^{\mathcal{D}}\right)\pi_j,$$

where π_j are the average cost-saving gains per task generated by automation in this sector. This change in TFP produces a productivity effect partially restoring labour demand, and typically also generates a series of sectoral reallocations in response to changes in sectoral prices.

The implications of different types of technologies can be analysed similarly. The creation of new tasks in a sector continues to generate the reinstatement effect, increasing the task content of production (i.e. $\left(1-s_j^L\right)d\ln\Gamma_j^{\mathcal{N}} > 0$) and thus labour demand; it also generates similar productivity and reallocation effects. Factor-augmenting technological changes generate productivity and relocation effects as well, but do not affect the task content of production.

In summary, the implications of 'any' technological change will work through, and can be decomposed into, a productivity effect, composition effects, price and quality substitution effects and changes in the task content of production. We next proceed to implement this decomposition.

6. Measuring displacement and reinstatement

In this section, we estimate the extent of displacement and reinstatement in the US economy during the last 70 years. Our methodology builds on the approach outlined in the previous section and uses a number of different sources of data, which we describe below.

6.1. Inferring changes in the allocation of tasks to factors

Our point of departure is equation (3). We use a discrete approximation to this equation using yearly changes; i.e. we approximate dX by $\Delta X_t = X_{t+1} - X_t$. On the basis of this, we construct:

$$\textbf{observed change in wage bill}_t = \Delta\ln\left(\frac{W_t L_t}{Pop_t}\right);$$

$$\textbf{productivity effect}_t = \Delta\ln\left(\frac{Y_t}{Pop_t}\right);$$

$$\textbf{composition effect}_t = \sum_{i\in\mathcal{I}}\left(\frac{s_{i,t}^L}{s_t^L}-1\right)\Delta\chi_{i,t};$$

$$\textbf{price substitution effect}_t = (1-\lambda)\sum_{i\in\mathcal{I}}\ell_{i,t}\left(1-s_{i,t}^L\right)\Delta\ln\left(\frac{W_{i,t}}{R_{i,t}}\right).$$

Here, Pop_t denotes the US population in year t, Y_t is GDP and $W_t L_t$ is the total wage bill, which is an inclusive measure of overall labour demand and thus our main object of interest. Relative to equation (3), we are normalizing the wage bill and GDP by population to account for population growth during our sample period. Note also that we are using sector-specific measures of wages and returns to capital from the US Bureau of Labor Statistics (Acemoglu and Restrepo 2019).

We take a baseline value for λ of 0.8 (which is in line with the estimates in Oberfield and Raval (2021)).[15] We discuss below how the overall qualitative and even quantitative implications of our approach are very similar for different values of λ (see also the appendix of Acemoglu and Restrepo 2019). Throughout, we impose 'no technological regress', meaning that no component of θ_i will become worse over time. Furthermore, in the text, we start with the assumption that $A_{i,L}/A_{i,K}$ in all sectors improves at the rate of GDP per worker (e.g. by 1.5% a year between 1987 and 2017) so that without any changes in the task content and capital-augmenting technologies, labour-augmenting technological change can account for the entire growth of productivity. We can then compute the quality substitution effect as

$$\begin{aligned}\textbf{quality substitution effect}_t &= (1-\lambda)\sum_{i\in\mathcal{I}}\ell_{i,t}\left(1-s_{i,t}^L\right)\Delta\ln\left(A_{i,L,t}/A_{i,K,t}\right)\\ &= 0.015\,(1-\lambda)\sum_{i\in\mathcal{I}}\ell_{i,t}\left(1-s_{i,t}^L\right).\end{aligned}$$

Under these assumptions, we can compute an estimate for the change in task content at the industry level as

$$\begin{aligned}\textbf{change in task content}_{i,t} &= \left(1-s_i^L\right)\left[\Delta\Gamma_i^{\mathcal{D}} + \Delta\ln\Gamma_i^{\mathcal{N}}\right]\\ &= \Delta\ln s_{i,t}^L - (1-\lambda)\left(1-s_{i,t}^L\right)\\ &\quad\times\left[\left(\Delta\ln\left(W_{i,t}/R_{i,t}\right)-\left(\Delta\ln\left(A_{i,L,t}/A_{i,K,t}\right)\right)\right].\end{aligned}$$

That is, changes in task content are obtained from the behaviour of the labour share, once we adjust for the influence of factor prices and factor-augmenting technologies.

The change in the task content of the entire economy is then given by

$$\textbf{change in task content}_t = \sum_{i\in\mathcal{I}}\ell_{i,t}\textbf{change in task content}_{i,t}.$$

Using this approach, we can decompose observed changes in labour demand (wage bill) during any sample period into a productivity effect, a composition effect, a change in task content, a price substitution effect and a quality substitution effect. We now proceed to apply this decomposition to various sample periods.

We further note that our estimates should be interpreted as upper bounds for the quality substitution effect (because, in general, growth in GDP per worker will be driven not just by labour-augmenting technological changes) and thus for changes in the task content of production (meaning that when our estimates are negative, the actual changes may be even larger). Nevertheless, reasonable variations on the magnitude of relative labour-augmenting technological change have very small effects on our decomposition results, as we discuss below.

[15] The relevant λ in our model is the elasticity of substitution between capital and labour at the industry level. This tends to be greater than the firm-level elasticity because of output substitution between firms.

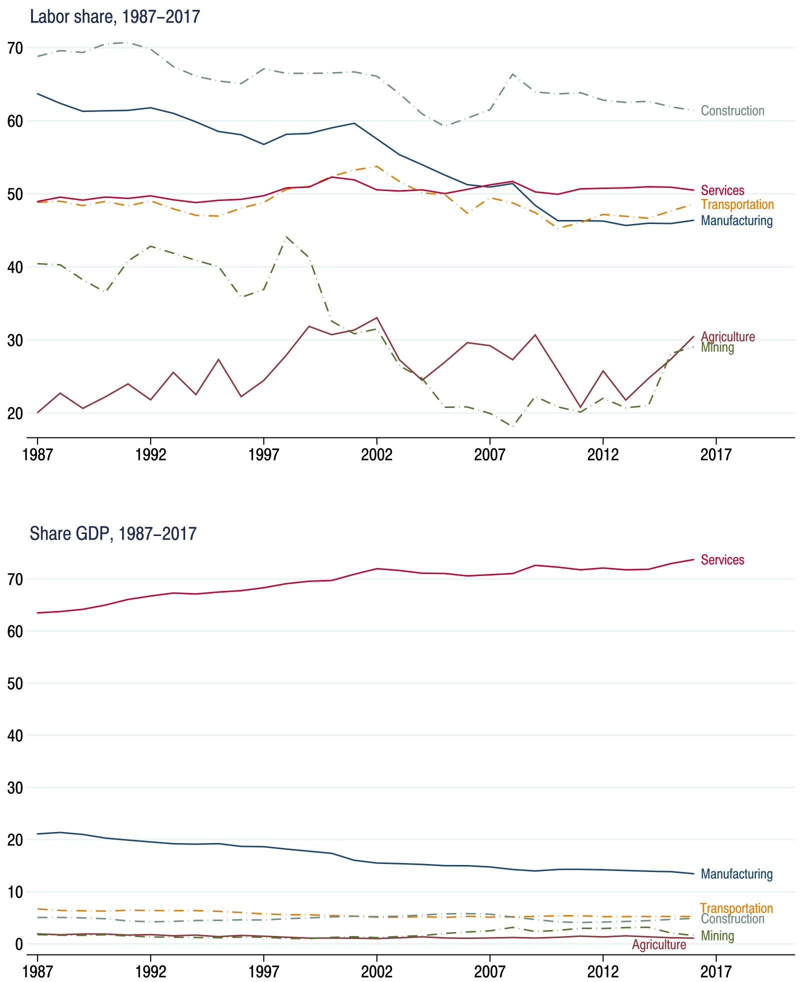

Figure 2: The labour share and sectoral evolutions, 1987–2017.
Note: The top panel shows the labour share in value added in services, manufacturing, construction, transportation, mining and agriculture between 1987 and 2017, while the bottom panel shows the share of value added in these sectors relative to GDP.
Source: Data from the BEA industry accounts.

6.2. Changes in the task content of production: 1987–2017

Let us start with the most recent 30-year period, 1987–2017, for which we have detailed data at the sectoral level (and where we will also be able to relate changes in the task content of production to measures of automation and creation of new tasks). For this period, we use data from the Bureau of Economic Analysis (BEA) for 61 NAICS industries. We start in the top panel of Fig. 2 by presenting the evolution of the labour share at the level of (roughly) one-digit sectors—for construction, services, transportation, manufacturing, agriculture and mining. We see a sharp decline in the labour share for manufacturing and mining, with much less change for the other industry groupings. The bottom panel of the figure shows the evolution of the share of value added of these sectors, highlighting the reallocation of economic activity away from manufacturing.

The top panel of Fig. 3 reports the implied decomposition for the entire economy. Several points are worth noting. First, comparing this figure with Fig. 10 for the period 1947–87, we see that overall labour demand grows much more slowly during the more recent 30 years—its annual growth rate is 1.33% compared with 2.44% between 1947 and 1987. Second, labour demand follows productivity fairly closely until the late 1990s, so the slow growth of labour demand in the first half of the sample is in large part because of the slow growth of productivity. Third, after the late 1990s, the gap between labour demand and productivity opens up sharply. Fourth, our estimates of composition and price substitution effects are quite small (and so are the quantity substitution effects implied by factor-augmenting technological changes, which are not shown in the figure). The small magnitude of the composition effect is particularly noteworthy because several popular mechanisms work entirely through sectoral reallocation captured by this composition effect.[16] Finally and most importantly, we see a sizeable decline in the task content of production, reflecting the fact that production is becoming less labour-intensive. The figure makes it clear that it is this change in task content that accounts for the decoupling of labour demand and productivity after 2000.

The large decline in labour share in manufacturing depicted in Fig. 2 suggests that changes in the task content of production in manufacturing may be playing a particularly important role. To investigate these changes, the bottom panel of the figure

[16] These include any effects from international trade in final goods, mechanisms emphasizing the Baumol effect (Aghion et al. 2017), and any non-homotheticities in preferences and structural transformation (Hubmer 2020).

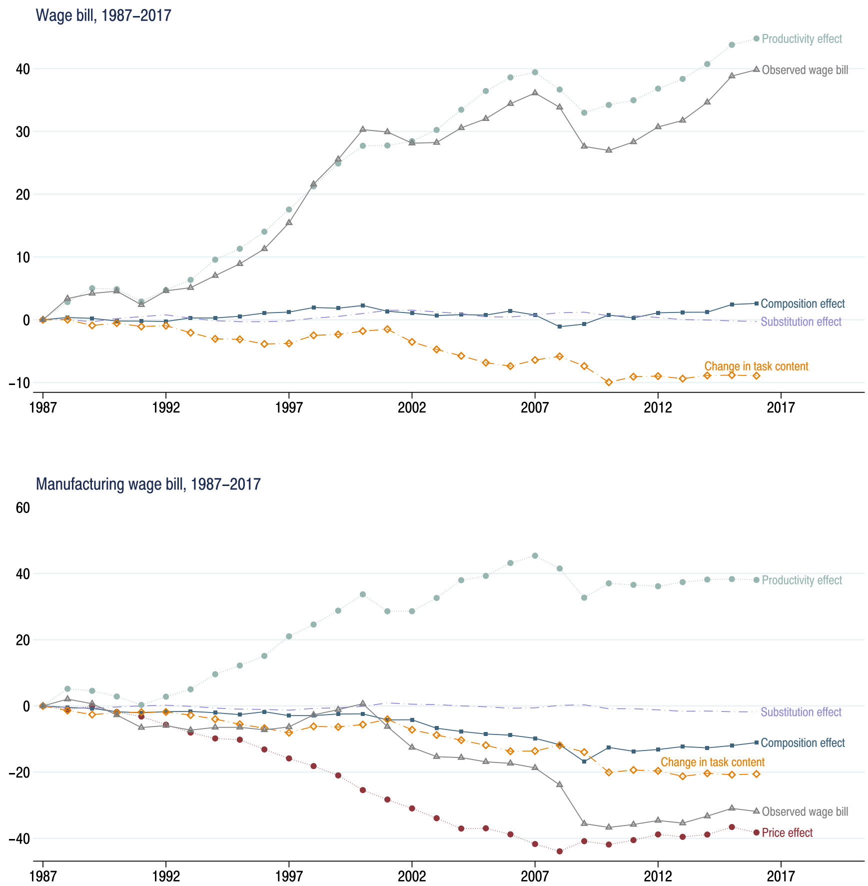

Figure 3: Sources of changes in labour demand, 1987–2017.
Note: This figure presents the decomposition of labour demand (wage bill) between 1987 and 2017 based on equation (3). The top panel is for the entire economy and the bottom panel is for the manufacturing sector. In both panels, we assume an elasticity of substitution between capital and labour equal to $\lambda = 0.8$ and relative labour-augmenting technological change at the rate of 1.5% a year.

applies the same decomposition to just manufacturing industries, with the only difference being that there is now an additional term corresponding to the price effect, which captures the movement in the relative price of manufacturing. The results from this decomposition for manufacturing are similar but even more pronounced. Notably, during this period, labour demand in manufacturing exhibits an absolute decline, which is in stark contrast to what we see in the previous 40 years in Fig. 10. Our decomposition shows that this is accounted for by sizeable negative changes in the task content of manufacturing production, again with a very limited role for composition, price substitution and quality substitution effects. In addition, during this period there is no contribution to labour demand from the productivity effect in manufacturing, reflecting the fact that manufacturing output has grown at roughly the same rate as the rest of the economy, but the relative prices of manufacturing goods have declined sharply, creating a sizeable negative price effect.[17]

Finally, Fig. 4 shows that the pattern of within-manufacturing changes is similar when we focus on 452 four-digit industries, for which we estimate our decomposition using data from the BEA input–output tables for 1977–2007.

6.3. Estimating displacement and reinstatement effects

Under the assumption of no technological regress, negative changes in the task content of production of an industry indicate that there is faster automation than creation of new tasks, and likewise positive changes are evidence of faster creation of new tasks than automation. We can thus estimate the extent of displacement (automation) and reinstatement (new task) effects at the industry level under the additional assumption that when there is faster automation, there will be no creation of new tasks in that industry during that same time period, and vice versa. To reduce the influence of measurement error, here we compute estimates for displacement and reinstatement effects for 5-year time windows using the following equations:

$$\text{displacement}_t = \sum_{i \in \mathcal{I}} \ell_{i,t} \min\left\{0, \frac{1}{5}\sum_{\tau=t-2}^{t+2} \text{change in task content}_{i,\tau}\right\}; \tag{4}$$

$$\text{reinstatement}_t = \sum_{i \in \mathcal{I}} \ell_{i,t} \max\left\{0, \frac{1}{5}\sum_{\tau=t-2}^{t+2} \text{change in task content}_{i,\tau}\right\}.$$

The resulting estimates are depicted in Fig. 5.[18] In the top panel, we see that both the displacement and reinstatement effects are

[17] As pointed out in the previous section, in our framework, rapid automation can go hand-in-hand with slow productivity growth if new automation technologies are 'so-so', or take place in tasks where the effective wage of labour is not much higher than the effective cost of capital—hence again implying little cost-reduction gain from automation.

[18] These estimates should be interpreted as 'lower bounds' because, within a 5-year time window, there are likely to be both automation and new tasks created in some industries, and this procedure only considers the difference between these two. Indeed, when we analogously estimate displacement and

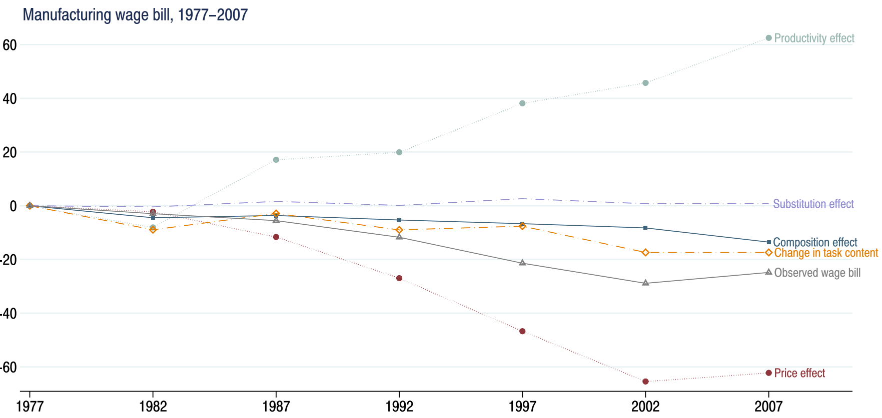

Figure 4: Sources of changes in manufacturing labour demand, 1977–2007.
Note: This figure presents the decomposition of labour demand (wage bill) between 1987 and 2007 based on equation (3). The figure uses detailed data for 452 manufacturing industries from the BEA input–output tables available every 5 years. We assume an elasticity of substitution between capital and labour equal to $\lambda = 0.8$ and relative labour-augmenting technological change at the rate of 1.5% a year.

sizeable, so both automation and new tasks appear to be ongoing at the industry level at all times. Nevertheless, the displacement effect is significantly larger, explaining the net negative change in the task content of production. In the bottom panel, we show the same decomposition for just manufacturing industries. Now the displacement effect is considerably larger, plausibly reflecting the greater extent of automation within manufacturing.[19]

6.4. Robustness and the role of factor-augmenting technologies

The patterns reported in the previous two subsections are robust and fairly insensitive to the assumptions on the elasticity of substitution and the rate of factor-augmenting technological change we have imposed. In the appendix of Acemoglu and Restrepo (2019), we verified that the results are very similar for different values of the elasticity of substitution (in particular, with $\lambda = 0.6$, $\lambda = 1$ and $\lambda = 1.2$). They are also very similar when we assume different rates of factor-augmenting technological changes.

Even more telling is a complementary exercise on the importance of factor-augmenting technologies we report here; we compute the extent of factor-augmenting technological change at the industry level that would be necessary to account for the changes in the labour share we observe without any change in the task content of technology. These results, depicted in Fig. 6, show that explaining industry-level changes in labour shares with constant task content of production would require huge changes in technology, with TFP growth several-folds larger than what is observed over the same time period. This again underscores the need for major changes in the task content of production to account for the evolution of labour demand during recent decades.

There is a simple reason why the exact rate of factor-augmenting technological change does not affect the labour share by much (while at the same time the magnitude of changes necessary to account for observed task contents is huge). The formula for the quality substitution effect in equation (3) implies that a 1% increase in labour-augmenting technologies reduces the labour share by $(1 - \lambda)\left(1 - s_i^L\right)$%. This implies a very small elasticity (between −0.08 and 0.08), given plausible values for the elasticity of substitution between capital and labour (between 0.8 and 1.2) and the observed labour share in most industries (~60%). Thus, only very large changes in technology—accompanied by sizeable changes in TFP—can move industry–labour shares meaningfully.

6.5. What does the change in task content capture?

As we are computing the change in task content as a residual, a natural concern is that it corresponds to something completely different from the displacement and reinstatement effects. In this subsection, we provide suggestive evidence to support our interpretation. We show that our measure of change in task content at the industry level is correlated negatively with several measures of the introduction of automation technologies, and positively with some proxies of new tasks.

The results are presented in Figs 7 and 8, and in Table 1.[20] Figure 7 provides the bivariate cross-industry associations between change in task content 1987–2017 and proxies for industry-level automation technologies. The first one is the adjusted penetration of robots measure from Acemoglu and Restrepo (2020*a*) for our 61 industries (matched to 19 industries as classified by the International Federation of Robotics). A strong negative correlation is visible in the top-left panel, and Table 1 verifies this relationship. The coefficient estimate is −1.40 (SE = 0.38) and this variable accounts for 18% of cross-industry variation in change in task content. The second column of the table confirms that this relationship is not driven by the contrast of manufacturing to non-manufacturing sectors; the coefficient estimate is similar, −0.99 (SE = 0.37), when we control for a manufacturing dummy. The third column further controls for import competition from China (Autor et al. 2013) and for the extent of offshoring (Feenstra and Hanson 1999), with very similar results.[21] Because industrial robots are a clear and important exemplar of automation technologies, this negative association is reassuring for our interpretation.

The top-right panel uses a broader measure of the potential for automation technologies. Acemoglu and Autor (2011) measure the share of routine jobs in our 61 industries using their distribu-

reinstatement effects at the yearly frequency, these are larger than the 5-year averaged estimates presented in Fig. 5.

19 Recent papers have applied this methodology to European countries with comparable results. See Graetz (2019) and Nardis and Parente (2021).

20 Further details on all of the variables discussed in this subsection are provided in the appendix of Acemoglu and Restrepo (2019).

21 This reflects the fact that import competition from China does not predict changes in the task content of production, which is noteworthy in and of itself. Instead, imports from China affect aggregate labour demand via the composition and productivity effects (Acemoglu and Restrepo 2019).

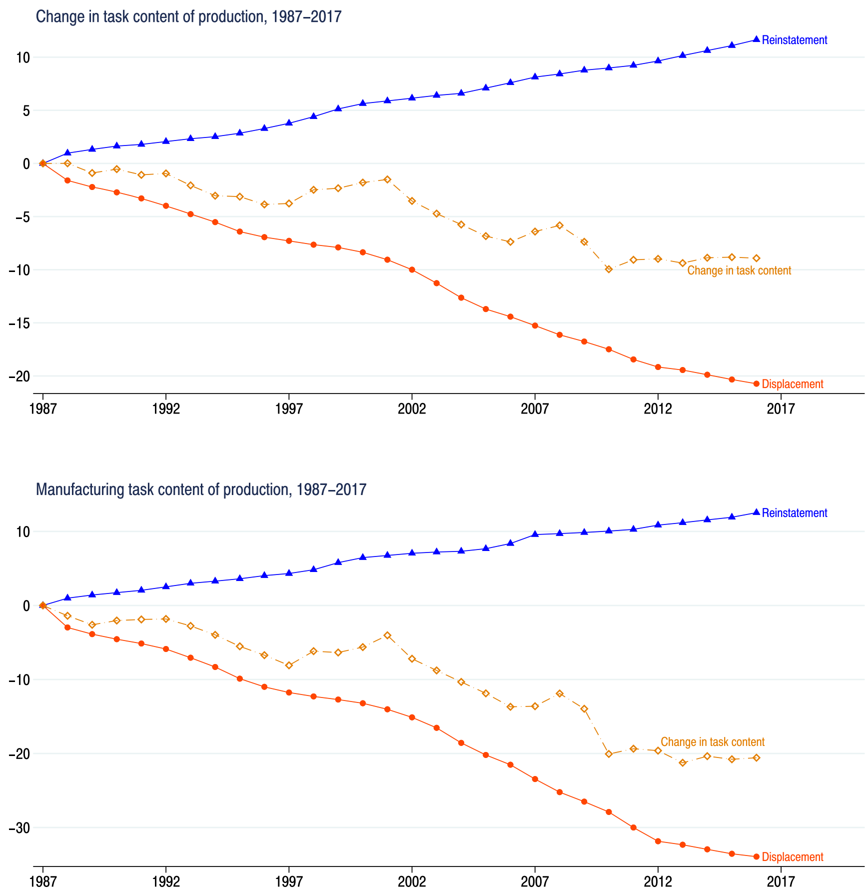

Figure 5: Estimates of the displacement and reinstatement effects, 1987–2017.
Note: This figure presents our baseline estimates of the displacement and reinstatement effects based on equation (4). The top panel is for the entire economy and the bottom panel is for the manufacturing sector. In both panels, we assume an elasticity of substitution between capital and labour equal to $\lambda = 0.8$ and relative labour-augmenting technological change at the rate of 1.5% a year.

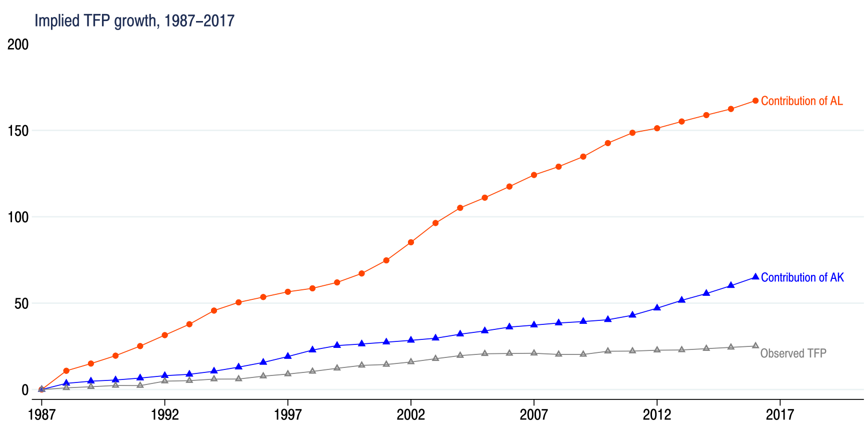

Figure 6: Counterfactual TFP changes.
Note: This figure presents the counterfactual TFP changes that would be implied if our estimates of the displacement and reinstatement effect in 1987–2017 were accounted for by industry-level changes in labour-augmenting and capital-augmenting technological changes alone, respectively. For comparison, the figure also reports the observed increase in TFP for both periods. We assume an elasticity of substitution between capital and labour equal to $\lambda = 0.8$.

tion of employment across occupations in 1990. There is a similar negative relationship, even conditioning on covariates.

The bottom-left panel uses measures of other automation technologies from the Survey of Manufacturing Technologies (SMT) for 1988 and 1993 (specifically the share of firms using automation technologies). These technologies include automatic guided vehicles, automatic storage and retrieval systems, sensors on machinery, computer-controlled machinery, programmable controllers and industrial robots (see Doms et al. 1997). These technology measures are available only for 148 'technology-intensive' manufacturing industries (which are all part of the two-digit manufacturing industries: fabricated metal products, industrial machinery, electronics, transportation equipment and controlling instruments). This panel therefore uses estimates of changes in task content for 1987–2007 for these 148 more detailed, four-digit SIC industries. There is once again a strong negative

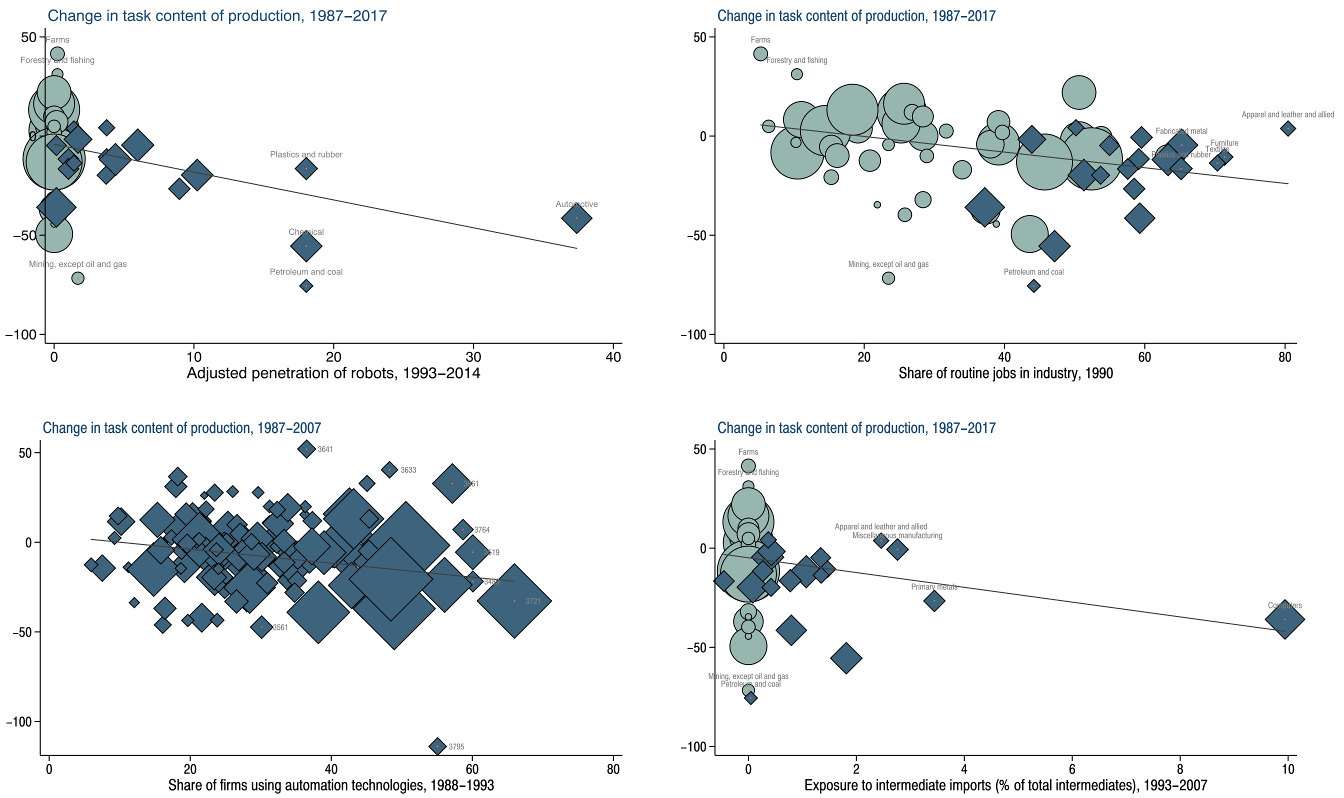

Figure 7: Automation technologies, offshoring and changes in the task content of production.
Note: Each panel presents the bivariate relationship at the industry level between change in task content and the indicated proxy for automation technologies or offshoring. Diamond markers designate manufacturing industries and circles denote non-manufacturing industries. The proxies are: adjusted penetration of robots, 1993–2014 (from Acemoglu and Restrepo (2018)), share of employment in routine occupations in 1990 (Acemoglu and Autor 2011), share of firms (weighted by employment) using automation technologies, from the 1988 and 1993 SMT and exposure to imports of intermediate goods (from Feenstra and Hanson (1999)). See text for details.

association, which is confirmed in Table 1. Finally, the bottom-right panel shows a similar negative relationship for offshoring, which could create the same type of displacement as automation.

We next turn to proxies for new tasks. Even though these are almost certainly less well measured than our proxies for automation, Fig. 8 shows strong and very suggestive correlations between our measures of new tasks and the task content of reduction (especially reinstatement). The top-left panel uses the share of new job titles from the 1991 Dictionary of Occupational Titles as compiled by Lin (2011), which we then project to our 61 industries, again using their employment distribution across occupations in 1990. As expected, there is a positive correlation between this measure of new tasks and change in task content, and the relevant coefficient estimate is 1.60 (SE = 0.52).

Table 1 shows that this relationship is essentially unchanged when we control for manufacturing, imports from China and offshoring. The top-right panel uses a related proxy based on 'emerging tasks' as classified by the Occupational Information Network (O∗NET) projected to industries. The results are similar and equally strong. The two bottom panels use two measures of increased occupational diversity in an industry with very similar results. The first is the share of employment growth in an industry accounted for by 'new occupations' defined as four-digit occupations appearing for the first time in that industry in 2016, while the second is the percentage increase in the number of occupations in an industry between 1990 and 2016.

These patterns suggest that our measure contains valuable information about changes in task content of production and also support the interpretation that the rapid displacement effect of the last 3 decades is related to the introduction of modern automation technologies such as industrial robots and computer numerical control.

6.6. Changes in the task content of production: 1947–87 and 1850–1910

We next turn to the 4 decades following World War II, 1947–87. For this period, we have data for 60 SIC industries. Figure 9 shows changes in the labour share and value-added distribution for the same six sectors as in Fig. 2. Particularly noteworthy is that there are no significant changes in the labour share for any of these industries. Figure 10 depicts the observed changes in labour demand together with our decomposition, separately for the whole economy and for the manufacturing sector.[22] During this period, labour demand grew more rapidly than in the last 30 years (notice that the vertical scale here is different than in Fig. 3). Our decomposition shows that there is a more robust productivity effect and a tighter relationship between labour demand and productivity during this time period. This more pronounced productivity effect underscores our conceptual conclusion that rapid productivity growth is an important contributor to growth in labour demand, even if it comes from automation technologies.

[22] We now assume that $A^L_{i,t}/A^K_{i,t}$ grows at 2% a year to match the growth of GDP per worker during the sample period. The results are similar if we continue to assume an annual growth of 1.5%.

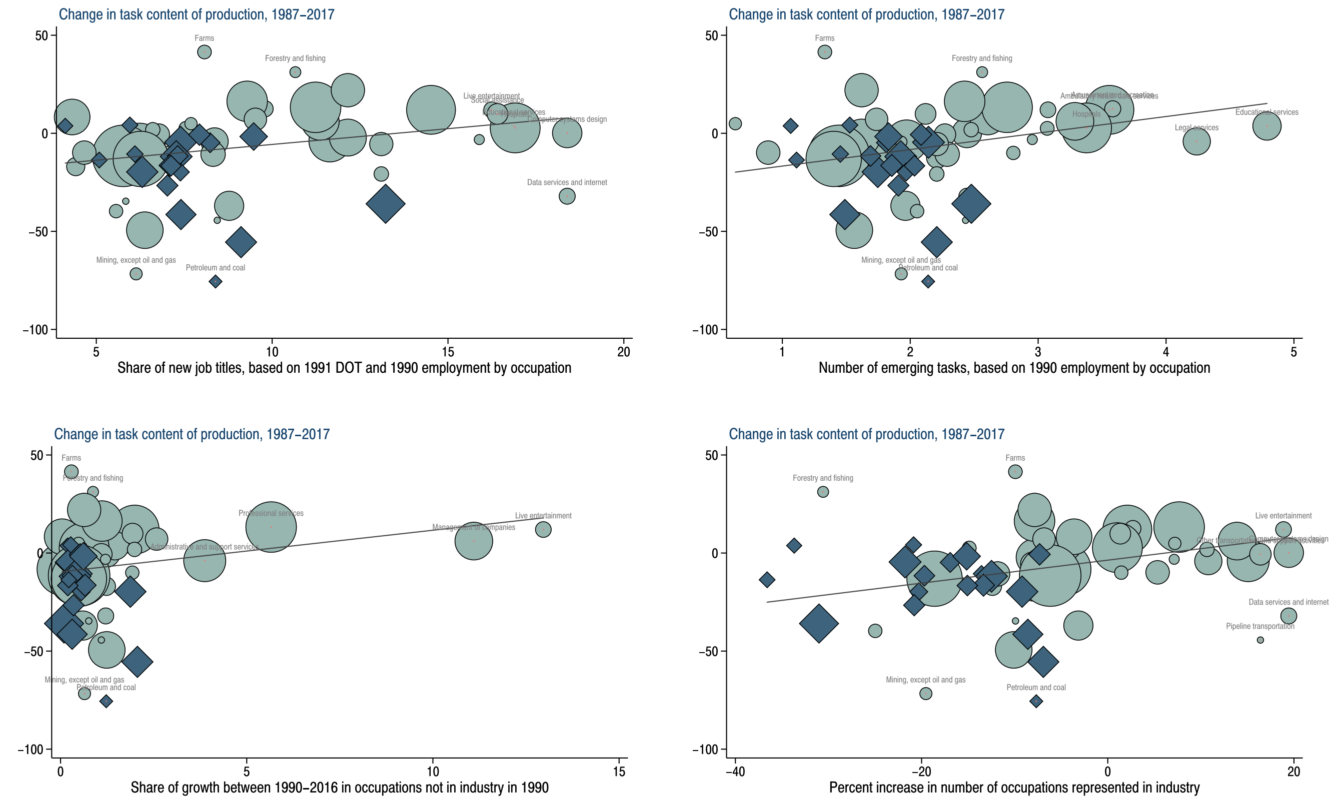

Figure 8: New tasks and change in task content of production.
Note: Each panel presents the bivariate relationship at the industry level between change in task content and the indicated proxy for new tasks. Diamond markers designate manufacturing industries and circles denote non-manufacturing industries. The proxies are: share of new job titles (from Lin 2011), number of emerging tasks (from O∗NET), share of employment growth between 1990 and 2016 in 'new occupations' (i.e. those that were not present in the industry in 1990) and the percentage increase in the number of occupations present in the industry between 1990 and 2016. See text for details.

Also noteworthy is the steady growth of labour demand in manufacturing, at least until the 1980s, which contrasts with its sharp contraction after the late 1990s in Fig. 3. Furthermore, consistent with the stable patterns of the labour share during this period, the change in the task content is small both for the entire economy and for manufacturing. Figure 11 confirms that this is because the displacement effect is more limited (0.5% per year compared with 0.66% per year in the later period, as shown in Fig. 5), and the reinstatement effect is more sizeable during this period than in the last 30 years (0.5% per year compared with 0.4% in the later period).[23]

Finally, we turn to the period 1850–1910, which witnessed rapid automation of a range of manual tasks in the context of the mechanization of agriculture. Figure 12 reports results from an analogous exercise during this period, but using only variation between agriculture and industry from data reported in Budd (1960). Because we do not have information on factor prices at the industry level for this period, in this figure we are forced to impose $\lambda = 1$, thus setting the quality and price substitution effects equal to zero. During this critical period of mechanization of agriculture, we see a decline in the labour share of agriculture—a telltale sign of automation in that sector—but a corresponding large increase in the labour share in industry. As a result, the change in the task content of production of the overall economy, though negative, is not very large. Our decomposition suggests that this in turn reflects the fact that the displacement effect in agriculture is being counterbalanced by a powerful reinstatement effect in manufacturing. In addition, in this case we estimate a composition effect that is somewhat larger, and this plausibly captures the sizeable reallocation of labour away from agriculture towards the more labour-intensive (manufacturing) industry.

The patterns reported in this subsection thus contrast with those of the last 3 decades and highlight the fact that the major difference setting the recent period apart from other epochs is not just the more anemic productivity effect but a sizeable displacement effect driven by automation and the absence of a powerful, countervailing reinstatement effect.

6.7. Summary of direct evidence

We have so far discussed how changes in displacement (due to automation) and reinstatement (due to the introduction of new tasks) can be inferred from data. We then depicted the evolution of these quantities in the USA over the last 70 years and provided evidence that our measures capture the relevant underlying economic concepts. In addition, the framework we have developed can also be used as the basis of empirical work directly estimating the effects of automation technologies (and new tasks, though there is much less work on estimating the impact of new tasks).

Here, we briefly discuss two lines of work. The first investigates the effects of automation technologies at the firm level. We pursued this strategy using detailed data on robot adoption in French manufacturing in Acemoglu et al. (2020*b*). The main prediction of our conceptual framework is that the adoption of automation technologies—in this instance, industrial robots—should be associated with a significant decline in the labour share

[23] These results are similar for different values of the elasticity of substitution and different assumed rates of factor-augmenting technological changes (Acemoglu and Restrepo 2019).

Table 1: Relationship between change in task content of production and proxies for automation and new tasks.

	Bivariate relationship (1)	Controlling for manufacturing (2)	Controlling for Chinese import and offshoring (3)
Proxies for automation technologies:			
Adjusted penetration of robots, 1993–2014	−1.404	−0.985	−1.129
	(0.377)	(0.369)	(0.362)
Observations	61	61	61
R^2	0.18	0.21	0.27
Share of routine jobs in industry, 1990	−0.394	−0.241	−0.321
	(0.122)	(0.159)	(0.164)
Observations	61	61	61
R^2	0.14	0.19	0.27
Detailed manufacturing industries (SMT):			
Share of firms using automation	−0.390		−0.397
technologies, 1988–93	(0.165)		(0.166)
Observations	148		148
R^2	0.08		0.09
Proxies for new tasks:			
Share of new job titles, based on 1991 DOT	1.609	1.336	1.602
and 1990 employment by occupation	(0.523)	(0.530)	(0.541)
Observations	61	61	61
R^2	0.12	0.23	0.32
Number of emerging tasks,	8.423	7.108	7.728
based on 1990 employment by occupation	(2.261)	(2.366)	(2.418)
Observations	61	61	61
R^2	0.14	0.25	0.33
Share of growth between 1990 and 2016	2.121	1.638	1.646
in occupations not in industry in 1990	(0.723)	(0.669)	(0.679)
Observations	61	61	61
R^2	0.08	0.20	0.26
Percent increase in number of occupations	0.585	0.368	0.351
represented in industry	(0.156)	(0.207)	(0.215)
Observations	61	61	61
R^2	0.14	0.19	0.25

Note: The table reports estimates of the relationship between the change in task content of production between 1987 and 2017, and proxies for automation technologies and new tasks. Column 1 reports estimates of the bivariate relationship between change in task content of production and the indicated proxy at the industry level. Column 2 includes a dummy for manufacturing industries as a control. In addition, Column 3 controls for the increase in Chinese imports (defined as the increase in imports relative to US consumption between 1991 and 2011, as in Acemoglu et al. 2016) and the increase in offshoring (defined as the increase in the share of imported intermediates between 1993 and 2007, as in Feenstra and Hanson (1999)). Except for the panels using the SMT, all regressions are for the 61 industries used in or analysis of the 1987–2017 period. When using the SMT, the regressions are for 148 detailed manufacturing industries. Standard errors robust against heteroscedasticity are given in parentheses. When using the measure of robot penetration, we cluster standard errors at the 19 industries for which this measure is available.

of value added. This is what Acemoglu et al. (2020*a*) find. In French manufacturing, growth adoption is associated with about a four percentage point decline in the labour share.

What about effects on employment? The impacts of robot adoption on employment (and wages) at the firm level can be very different from those characterized here. At the economy or industry level, as we have seen, labour demand responds to automation via a displacement effect that is negative and a productivity effect that is positive. At the firm level, there is an additional channel at work: a firm that adopts robots reduces its marginal cost and thus can expand at the expense of its competitors. The appendix of Acemoglu et al. (2020*a*) develops an extended model that incorporates this 'competition effect'. Empirically, they report significant increases in sales, value added and employment among firms adopting robots. However, this is accompanied by even larger declines in firms competing with these robot-adopters. Combining these two opposing effects, they find that the total impact on industry employment is negative and similar to industry-level estimates in the USA.[24]

The second strategy is developed and pursued in Acemoglu and Restrepo (2020*a*), and aims at estimating the local equilibrium effects of automation. In particular, this local equilibrium effect incorporates the displacement of workers due to automation as well as the creation of jobs in other tasks due to productivity benefits (cost reductions) brought about by automation. Acemoglu and Restrepo (2020*a*) consider a multi-industry version of the framework developed here and prove that local employment and wage effects of adoption of automation technologies (again, in

[24] There are several other papers using this firm-level strategy with data from other countries. Most notably, Dinlersoz and Wolf (2018), Bessen et al. (2019), Bonfiglioli et al. (2020), Humlum (2020) and Koch et al. (2021) find patterns that are broadly consistent with those discussed here. Graetz and Michaels (2018) use a related empirical strategy, focusing on cross-industry and cross-country variation. They find lower labour shares associated with robot adoption as well, and negative employment effects for unskilled workers.

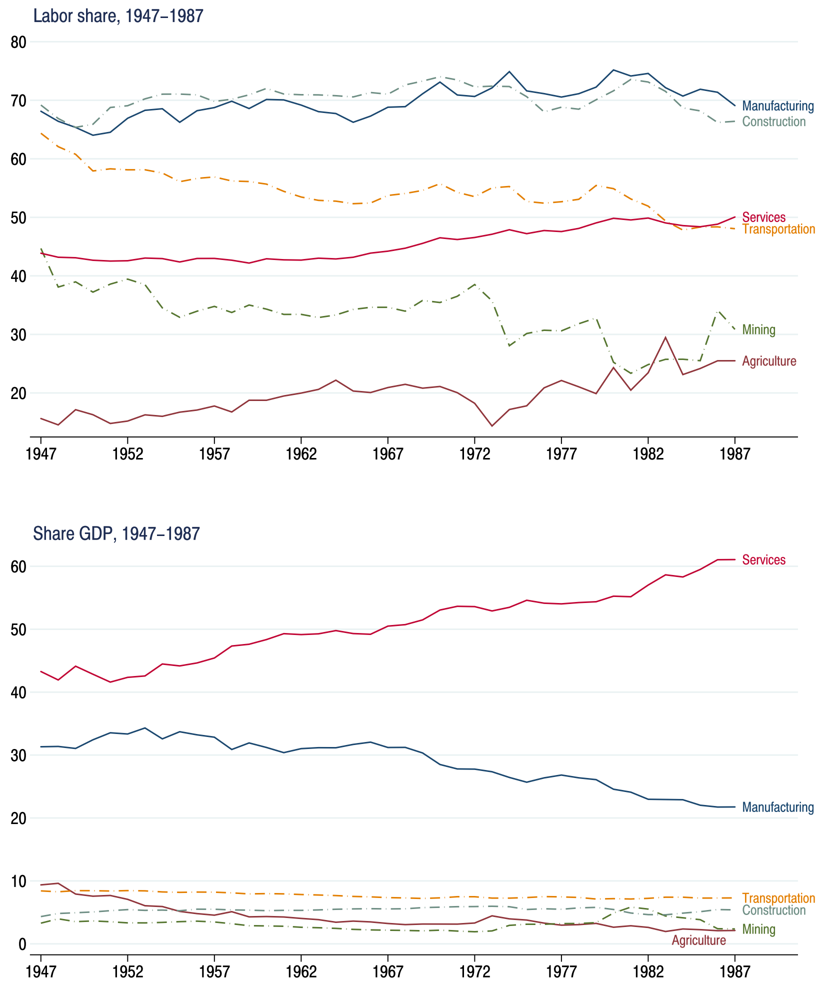

Figure 9: The labour share and sectoral evolutions, 1947–87.
Note: The top panel shows the labour share in value added in services, manufacturing, construction, transportation, mining and agriculture between 1947 and 1987, while the bottom panel shows the share of value added in the sectors relative to GDP.

this instance, industrial robots) can be estimated by linking local employment and wage outcomes to a measure of 'exposure to robots'. This measure of exposure to robots, for a local labour market/commuting zone c, is computed as

$$\text{exposure to robots}_c = \sum_{i \in \mathcal{I}} \ell_{ci} \cdot APR_i,$$

where ℓ_{ci} is the 'baseline share' of industry i in commuting zone c, and APR_i is the 'adjusted penetration of robots' in industry i.[25] Thus, this is a simple Bartik measure combining industry-level variation in the usage of robots and baseline employment shares. The adjusted penetration of robots, APR_i, measures how much robot adoption has been going on in that particular industry, but adjusted to account for the overall expansion in the output of an industry (which would normally increase the use of inputs). Because variation in US industries is likely to be endogenous to a variety of other factors that might affect labour demand, Acemoglu and Restrepo (2020*a*) compute this measure from robot adoption in several European economies that are somewhat ahead of the USA in terms of use of robotics in manufacturing.

Their estimates using data from 722 commuting zones in the USA show precisely estimated and sizeable (but not huge) negative employment and wage effects from the adoption of robots. For example, the adoption of one more industrial robot per 1000 workers in a commuting zone is associated with a 0.38 percentage point decline in employment to population ratio and a 0.71% decline in average wages relative to a commuting zone with no increase in robots (though the effects, once we take national adjustments into account, are somewhat less). In terms of the framework in their paper (and here), this is because the displacement effects are larger than the productivity effects. Acemoglu and Restrepo (2020*a*) document that these results are robust, they are not driven by other changes in technology, trade or offshoring, have no equivalent in the period before 1990 (before robots started being adopted in large numbers in US industry) and have been driven by declines in blue-collar occupations and especially in industries most exposed to robotics.[26]

[25] Namely, this measure is defined as

$$APR_i = \frac{d\text{Robots}_i}{L_i} - \frac{dY_i}{Y_i}\frac{\text{Robots}_i}{L_i},$$

where the last term is a correction factor for the increase in the number of robots per worker that would be implied by the expansion of the industry's value added.

[26] Dauth et al. (2021) adopt the same strategy in Germany. They estimate declines in manufacturing employment but not in overall employment across German regions.

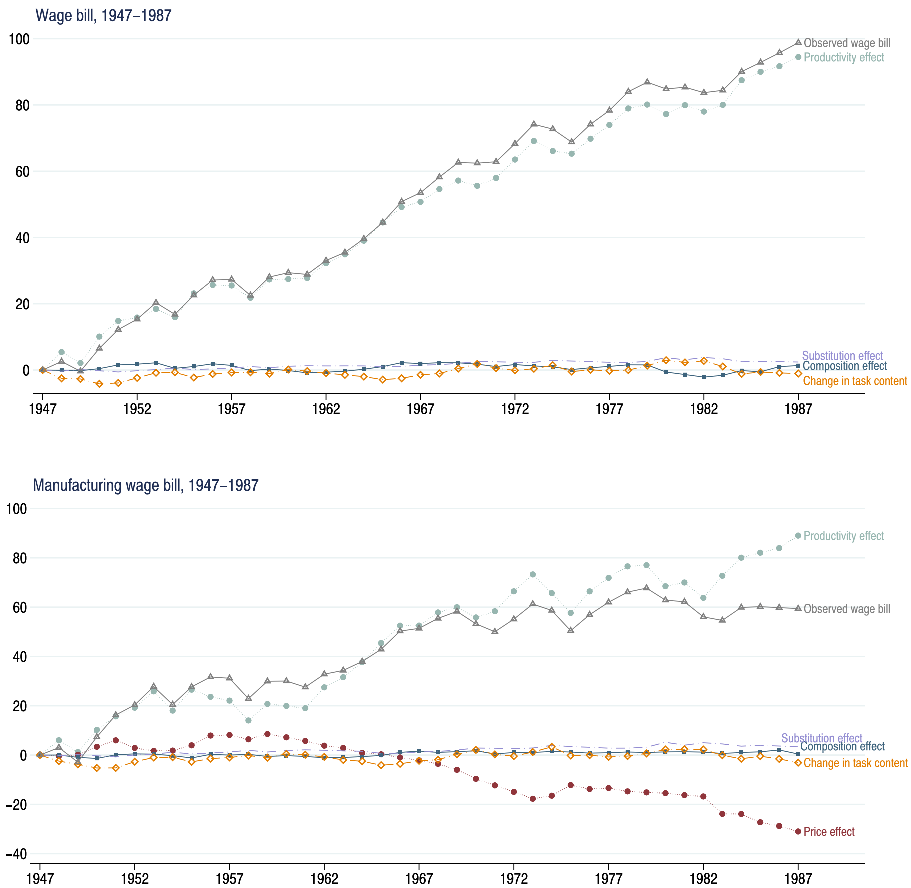

Figure 10: Sources of changes in labour demand, 1947–87.
Note: This figure presents the decomposition of labour demand (wage bill) between 1947 and 1987 based on equation (3). The top panel is for the entire economy and the bottom panel is for the manufacturing sector. In both panels, we assume an elasticity of substitution between capital and labour equal to $\lambda = 0.8$ and relative labour-augmenting technological change at the rate of 2% a year.

Overall, our conceptual framework lends itself to various different types of empirical strategies, and several of these avenues have already been used fruitfully. Most of this work indicates that the task-based approach provides a useful set of lenses to understand the effects of new technologies on the labour market. It also confirms that automation technologies create large displacement effects, which typically outweigh their productivity benefits, thus reducing labour demand—in stark contrast to the common presumption based on the standard framework in which new technologies typically tend to benefit labour.

7. Displacement, reinstatement and inequality

We now return to a discussion of inequality, linking the measures of displacement and reinstatement estimated in the previous section to the demand for skills. Our main objective is to uncover the statistical association between changes in task content of production—due to automation and new tasks—and the demand for skills.

7.1. Displacement, reinstatement and industry demand for skills

Using the industry-level measures of displacement and reinstatement obtained in the previous section, we estimate the following model separately for the same two subperiods:

$$\Delta\text{demand for skills}_i = \beta_d \text{displacement}_i + \beta_r \text{reinstatement}_i + \varepsilon_i. \quad (5)$$

Here, Δdemand for skills$_i$ (i.e. our measure of industry-level increase in the demand for skills) is the change in the log of the college wage bill relative to the high school wage bill in each industry during the relevant period. All regressions are weighted by the average share of the aggregate wage bill accounted by the industry during the period. Regression results are presented in Table 2, and we also depict them visually in Fig. 13.

Figure 13 summarizes the most important patterns. It shows a strong association between industry-level demand for skills and our measures of displacement (due to automation) and reinstatement (due to new tasks). During both subperiods, displacement is associated with increases in the demand for skills of the industry, though displacement changes are larger and the relationship becomes steeper in 1987–2016, shown in Panel (b). A 10% increase in displacement during 1987–2016 is associated with an increase of 8% in the relative demand for college workers (SE = 0.015). This estimate implies that displacement alone explains ∼30% of the variation in the demand for skills across industries during this period. We also note that the 0.55% increase in displacement per annum at the aggregate level during this period could account for an increase of as much as 0.44% in the demand for college skills (out of an estimated shift in the relative demand of 2.4 per annum; see Acemoglu and Autor 2011). As commented above, this

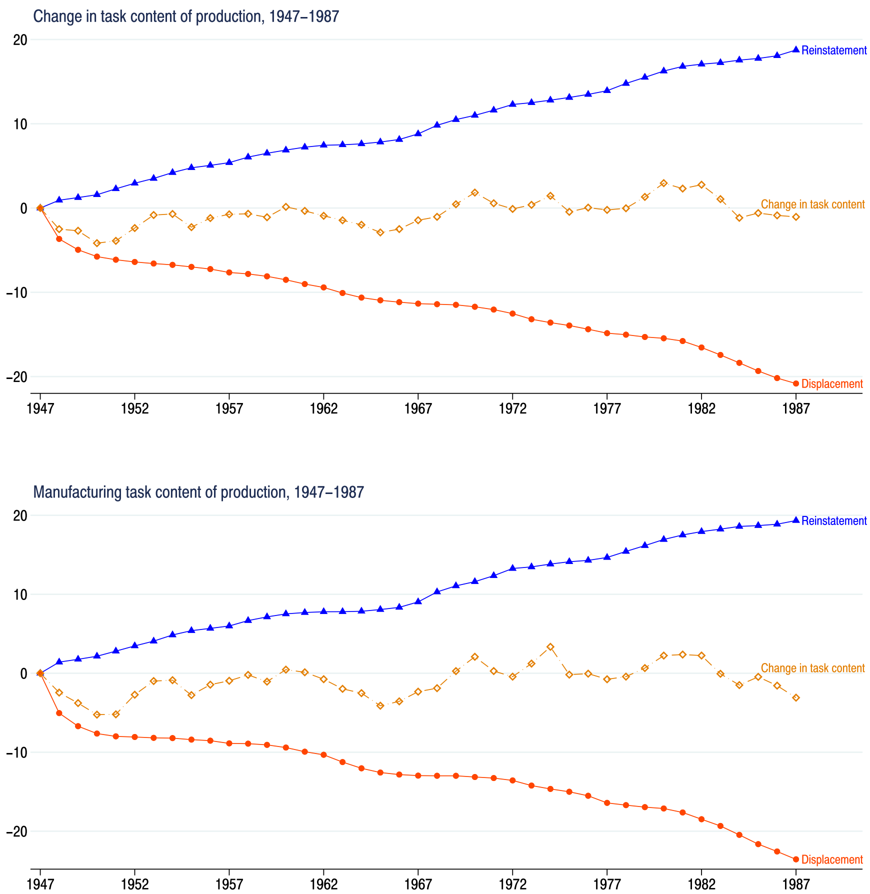

Figure 11: Estimates of the displacement and reinstatement effects, 1947–87.
Note: This figure presents our baseline estimates of the displacement and reinstatement effects based on equation (4). The top panel is for the entire economy and the bottom panel is for the manufacturing sector. In both panels, we assume an elasticity of substitution between capital and labour equal to $\lambda = 0.8$ and relative labour-augmenting technological change at the rate of 2% a year.

large effect on the demand for skills can happen with minimal changes in TFP. For example, when $\pi = 30\%$ (consistent with the evidence in Acemoglu and Restrepo (2020*a*)), this substantial increase in the relative demand for college skills is consistent with new technologies increasing TFP by as little as 0.16% per annum between 1987 and 2016.

Figure 13c and d depict the relationship between reinstatement and the demand for skills. Greater reinstatement is associated with lower demand for skills during 1947–87, presumably because unskilled labour had a comparative advantage in many of the new tasks introduced during this period. In contrast, reinstatement goes hand-in-hand with greater demand for skills in 1987–2016, which we interpret as new tasks being allocated more to skilled workers during the last 3 decades. Our estimates suggest that during this latter period, an increase of 10% in reinstatement is associated with a 7% increase in the relative demand for college workers (SE = 0.035).

Table 2 provides more details on relationships summarized in Fig. 13. Panels A–C provide estimates for 1947–87 and Panels D–F provide estimates for 1987–2016. In Panels A and D, we use the wage bill of college workers relative to high school workers as our measure for the demand for skills in an industry. In Panels B and E, we use the hours worked by college workers relative to high school workers as our measure for the demand for skills in an industry. In Panels C and F, we use the number of college workers relative to high school workers as our measure for the demand for skills in an industry. Columns 1–3 present estimates of equation (5) for all workers, and columns 4–7 present estimates separately for men, women and workers in different age groups.

Tables 3 and 4 provide estimates using alternative measures of changes in the task content of industries and the resulting measures of displacement and reinstatement. For this exercise, we use relative wage bill (columns 1–3) and relative hours (columns 4–6) as our measures of skill demand. Table 3 focuses on the 1947–87 period. Panel A provides results obtained by setting $\lambda = 1$ in our computation of the displacement and reinstatement effects, instead of $\lambda = 0.8$ as in the baseline estimates in Table 2. Panel B reverts to $\lambda = 0.8$, but we now use a 10-year moving average, rather than a 5-year moving average in our calculation of the displacement and reinstatement effects. Finally, in Panel C, we implement both changes simultaneously.

Table 4 focuses on the 1987–2016 period. Panel A provides results obtained by setting $\lambda = 1$ in our computation of the displacement and reinstatement effects. Panel B reverts to $\lambda = 0.8$, but we now use a 10-year moving average, rather than a 5-year moving average in our calculation of the displacement and reinstatement effects. In Panel C, we implement both changes simultaneously. In Panels D–F, we repeat these exercises, but now we use data from the BEA KLEMS accounts for 1987–2016. These data provide the labour share for each industry inclusive of self-employment.

Overall, the results in Tables 2, 3 and 4 confirm the patterns we see visually in Fig. 13. Automation is associated with significant declines in the demand for skills in both periods, regardless of the

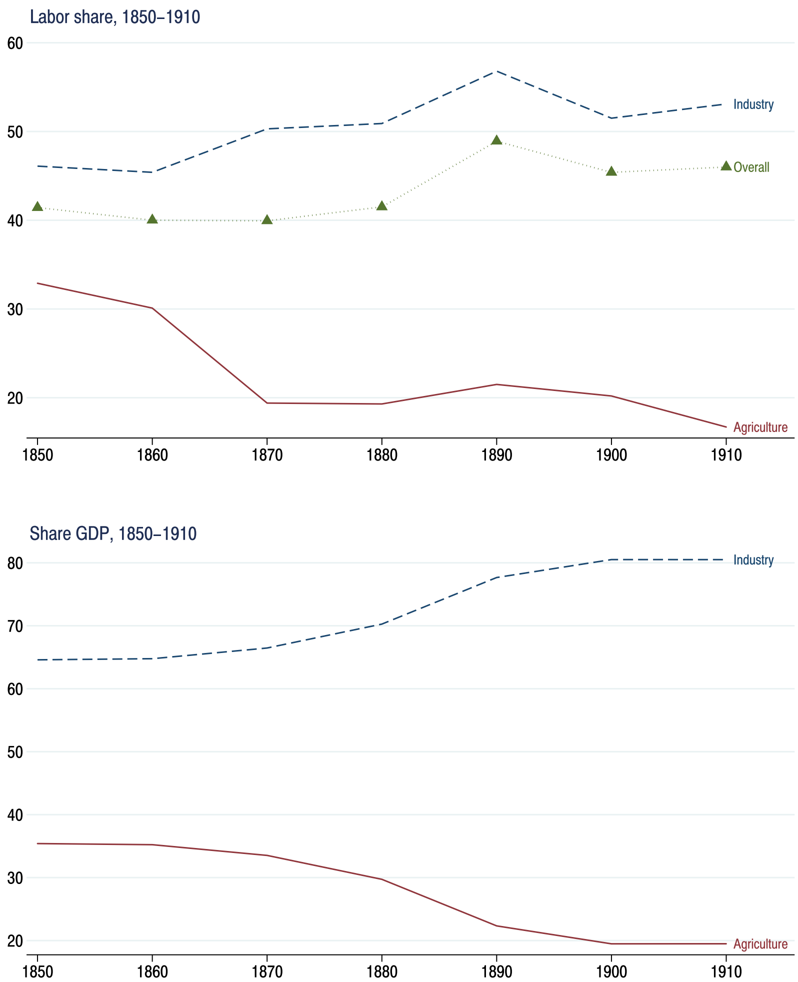

Figure 12: Labour share and sectoral evolutions during the mechanization of agriculture, 1850–1910.
Note: The top panel shows the labour share in value added in industry (services and manufacturing) and agriculture between 1850 and 1910, while the bottom panel shows the share of value added in these sectors relative to GDP.
Source: Data from Budd (1960).

specification or measure we use (and for different subgroups, such as men, women and younger workers). Reinstatement between 1947 and 1987 is associated with lower demand for skills, whereas between 1987 and 2016, it is associated with higher demand for skills. This pattern is robust as well. One additional finding is worth noting: even between 1947 and 1987, reinstatement does not appear to increase the demand for unskilled men by much, likely reflecting the fact that less-skilled women may have been the ones with comparative advantage in new tasks introduced during this period. We do not know why the reinstatement effect has become more favourable to skill workers during the last 30 years. One possibility is that many new tasks have been introduced in more skill-intensive sectors, which is consistent with the fact that there is very little reinstatement in manufacturing (as shown in the bottom panel of Fig. 5). Another possibility is that various institutional and other changes have made firms and researchers completely turn away from workers without college degrees, and thus firms have also been discouraged from introducing new tasks complementary to these workers.

What do these findings imply for the demand for skills, and hence for the increase in labour market inequality? To shed some light on this question, we now perform a very simple counterfactual exercise. We compute what the demand for skills in the USA would have been between 1987 and 2016 if: (1) displacement did not accelerate from the earlier period, 1947–87; (2) reinstatement did not decelerate again from 1947–87; and (3) the effect of reinstatement on the demand for skills remained the same as in 1947–87. Performing these exercises, out of the 35% increase in the skill premium between 1987 and 2016, we find that: (1) accounts for a 3.5% increase in the demand for skills, (2) accounts for a 3% increase in the demand for skills and (3) accounts for a 12% increase in the demand for skills.

7.2. Summary of direct evidence

More direct evidence on the importance of automation and changes in the task content of production for inequality is provided in Acemoglu and Restrepo (2022), who extend the framework presented here to include many types of labour and a richer menu of technologies.[27] The main conceptual breakthrough of that paper is to derive a simple equation linking the real-wage change of a group of workers (interpreted as a different factor of production that is imperfectly substitutable to other groups) to the task displacement it experiences. This task displacement is in turn predicted to be related to the baseline

[27] This paper in turn builds on Autor et al. (2003), who provided the first evidence that automation of routine jobs was an important factor in the changes in the demand for skills in the USA.

Table 2: Changes in task content and relative demand for skills, 1947–87 and 1987–2016.

	All employees			Men	Women	Aged 25–34	Aged 35–64
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	Panel A: college wage bill relative to high school wage bill, 1947–87						
Displacement	0.504		0.470	0.108	0.384	0.764	0.293
	(0.193)		(0.184)	(0.352)	(0.423)	(0.225)	(0.273)
Reinstatement		−0.585	−0.546	0.023	−0.639	−0.594	−0.544
		(0.306)	(0.278)	(0.482)	(0.501)	(0.430)	(0.261)
Observations	44	44	44	44	44	44	44
R^2	0.06	0.06	0.12	0.00	0.05	0.08	0.07
	Panel B: college hours relative to high school hours, 1947–87						
Displacement	0.686		0.644	0.315	0.458	0.738	0.608
	(0.219)		(0.165)	(0.301)	(0.401)	(0.252)	(0.194)
Reinstatement		−0.723	−0.670	−0.361	−0.630	−0.707	−0.633
		(0.343)	(0.304)	(0.431)	(0.434)	(0.463)	(0.234)
Observations	44	44	44	44	44	44	44
R^2	0.09	0.08	0.16	0.04	0.07	0.11	0.15
	Panel C: college employees relative to high school employees, 1947–87						
Displacement	0.873		0.834	0.587	0.536	0.941	0.769
	(0.204)		(0.158)	(0.323)	(0.337)	(0.224)	(0.206)
Reinstatement		−0.697	−0.629	−0.368	−0.575	−0.596	−0.644
		(0.352)	(0.292)	(0.363)	(0.415)	(0.422)	(0.256)
Observations	44	44	44	44	44	44	44
R^2	0.15	0.07	0.21	0.09	0.07	0.15	0.17
	Panel D: college wage bill relative to high school wage bill, 1987–2016						
Displacement	0.800		0.764	1.053	1.061	0.353	0.947
	(0.152)		(0.159)	(0.288)	(0.247)	(0.209)	(0.186)
Reinstatement		0.707	0.483	0.299	0.299	0.850	0.390
		(0.348)	(0.340)	(0.401)	(0.506)	(0.391)	(0.384)
Observations	44	44	44	44	44	44	44
R^2	0.31	0.06	0.34	0.34	0.40	0.16	0.37
	Panel E: college hours relative to high school hours, 1987–2016						
Displacement	0.558		0.520	0.754	0.778	0.185	0.697
	(0.137)		(0.141)	(0.220)	(0.227)	(0.179)	(0.169)
Reinstatement		0.658	0.506	0.196	0.404	0.768	0.431
		(0.310)	(0.317)	(0.329)	(0.431)	(0.349)	(0.371)
Observations	44	44	44	44	44	44	44
R^2	0.19	0.07	0.22	0.29	0.33	0.12	0.25
	Panel F: college employees relative to high school employees, 1987–2016						
Displacement	0.546		0.514	0.696	0.793	0.257	0.657
	(0.134)		(0.135)	(0.195)	(0.214)	(0.154)	(0.166)
Reinstatement		0.582	0.431	0.100	0.345	0.540	0.450
		(0.326)	(0.325)	(0.335)	(0.409)	(0.323)	(0.376)
Observations	44	44	44	44	44	44	44
R^2	0.19	0.05	0.22	0.29	0.34	0.11	0.24

Note: The table provides regression estimates of changes in the relative demand for skills across industries on measures of displacement and reinstatement. The appendix in Acemoglu and Restrepo (2020*b*) provides a description of the construction of these explanatory variables. Panels A–C provide estimates for 1947–87. Panels D–F provide estimates for 1987–2016. Each panel uses a different measure of changes in the relative demand for skills across industries. Panels A and D use the change in the log of the college wage bill relative to the high school wage bill in each industry as outcome. Panels B and E use the change in the log of college hours relative to high school hours in each industry as outcome. Panels C and F use the change in the log of the number of college employees relative to high school employees in each industry as outcome. In columns 1–3, the measures of changes in relative demand for skills are computed for all employed in an industry; in column 4, only for men; in column 5, only for women; in column 6, for employees aged 25–34 years; and in column 7, for employees aged 35–64 years. Standard errors robust against heteroscedasticity are given in parentheses.

distribution of employment of this demographic group across different industries and different types of jobs (in particular, routine vs non-routine occupations) and the overall amount of automation and task content change experienced by the industry in question.[28]

Acemoglu and Restrepo (2022) then estimate this new wage equation in the USA between 1980 and 2016. Their results indicate that between 50 and 70% of all changes in the US wage structure are directly related to automation-related changes in task displacement. These relationships are robust to changes in specification, to focusing on subperiods and to the inclusion of various

[28] More precisely, this relationship is $d\ln w_g$ = controls + $\beta\cdot$ task displacement$_g$, where

$$\text{task displacement}_g = \sum_{g\in\mathcal{G}} \omega_g^i \cdot \left(\omega_{gi}^R/\omega_i^R\right)\cdot\left(-d\ln s_i^L\right).$$

In this expression, ω_g^i is the share of wages earned by workers of group g in industry i, $d\ln s_i^L$ is the change in the labour share of industry i, which is related to automation (or this could be replaced by direct measures of industry-level automation) and ω_{gi}^R/ω_i^R measures the relative specialization of group g in industry i's routine jobs, where displacement takes place.

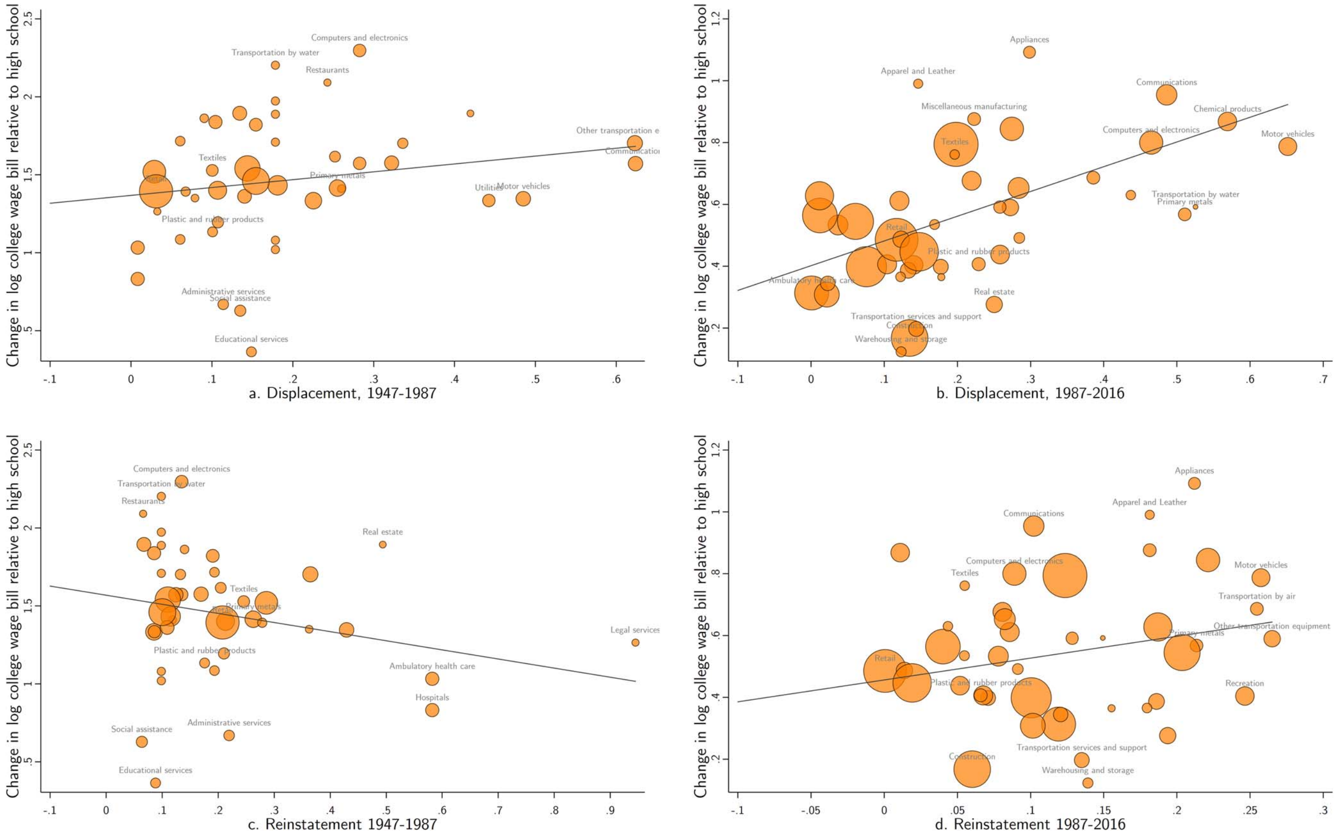

Figure 13: Change in relative demand for skills 1947–87 and 1987–2016 versus displacement and reinstatement.
Note: Relative demand for skills is measured as the log of the college wage bill relative to the high school wage bill. See the appendix of Acemoglu and Restrepo (2020*b*) for details and derivation of the estimates for displacement and reinstatement.

controls. Other factors, related to direct measures of skill-biased technological change, import competition from China, other types of technologies, changes in markups and economic concentration and de-unionization appear to be much less important. For example, according to their estimates, the reason why young men without a high school degree have experienced a 15% real-wage decline since 1980 is because they were concentrated in various routine occupations in manufacturing, mining, retail and wholesale industries. It is exactly these industries that experienced the largest declines in labour share or the most pronounced adoption of automation technologies, and these developments reduced the demand for labour for workers specializing in tasks that can be automated in these industries—proxied by routine occupations. Another implication of these findings is that previous work inferred a major role for SBTC (technologies increasing the relative productivity of more educated workers) precisely because it did not incorporate task displacement. Once we include an estimate of the effects of task displacement, we find no additional role for SBTC for college versus high school, though there is still some additional growth of post-graduate earnings, most likely reflecting their complementary skills to some of these technologies and perhaps winner-takes-all dynamics in some occupations.

Acemoglu and Restrepo (2022) also develop an extended framework in which general equilibrium interactions triggered by automation can be estimated. Briefly, these interactions incorporate both the productivity effects that are common across groups (and thus go into the constant term) and the ‘ripple effects’ that are created as one group is displaced from the tasks in which it has a comparative advantage and starts competing with other demographic groups for tasks in which they have secondary comparative advantage. We estimate that these ripple effects also contribute significantly to changes in the US wage structure. Productivity effects are important as well, but do not change the fact that automation leads to significant real-wage declines for groups who are experiencing task displacement.

8. Conclusion and implications for future work

This article has reviewed the task-based approach to production, capital–labour substitution and inequality. We have argued that this framework has better microfoundations and provides a more powerful and empirically accurate description of how technologies impact production in practice. Equally importantly, it leads to a range of new comparative statics (in particular, about when and how new technology will reduce labour demand, employment and wages) and to new empirical strategies. After reviewing some of these conceptual underpinnings, we have shown how this framework can be estimated and some of its basic economic objects can be inferred from data.

Most importantly, this framework and our past empirical work suggest that much of the decline in labour share in value added and slowdown of wage growth (and employment growth) in the USA is due to an acceleration in the adoption of automation technologies. It is not only that there has been a lot of automation. We also find that the introduction of new tasks and the associated reinstatement of labour has slowed down. We also estimate that automation technologies and the task displacement they generate account for the bulk—up to 70%—of all changes in US wage structure.

Table 3: Robustness to measures of task content, 1947–87.

	College wage bill relative to high school wage bill			College hours relative to high school hours		
	(1)	(2)	(3)	(4)	(5)	(6)
	Panel A: BEA data with λ = 1 and 5-year moving averages					
Automation	0.447		0.446	0.647		0.646
	(0.207)		(0.161)	(0.249)		(0.165)
Reinstatement		−0.484	−0.483		−0.580	−0.578
		(0.226)	(0.205)		(0.256)	(0.228)
Observations	44	44	44	44	44	44
R^2	0.06	0.07	0.13	0.10	0.08	0.18
	Panel B: BEA data with λ = 0.8 and 10-year moving averages					
Automation	0.536		0.410	0.774		0.624
	(0.224)		(0.219)	(0.220)		(0.183)
Reinstatement		−0.660	−0.595		−0.806	−0.708
		(0.265)	(0.262)		(0.303)	(0.294)
Observations	44	44	44	44	44	44
R^2	0.04	0.09	0.11	0.07	0.11	0.16
	Panel C: BEA data with λ = 1 and 10-year moving averages					
Automation	0.488		0.352	0.759		0.601
	(0.235)		(0.204)	(0.245)		(0.190)
Reinstatement		−0.577	−0.529		−0.698	−0.618
		(0.203)	(0.200)		(0.230)	(0.224)
Observations	44	44	44	44	44	44
R^2	0.04	0.10	0.12	0.08	0.12	0.17

Note: The table provides regression estimates of changes from 1947 to 1987 in the relative demand for skills across industries on measures of displacement and reinstatement. The appendix in Acemoglu and Restrepo (2020*b*) provides a description of the construction of these explanatory variables. Columns 1–3 use the change in the log of the college wage bill relative to the high school wage bill in each industry as outcome. Columns 4–6 use the change in the log of college hours relative to high school hours in each industry as outcome. Each panel presents results for a different construction of the displacement and reinstatement measures, as explained in the appendix in Acemoglu and Restrepo (2020*b*). Standard errors robust against heteroscedasticity are given in parentheses.

Empirical work on this approach is still in its infancy, and many directions of fruitful research are open. Most importantly, exercises similar to those reported here can be carried out with data from other economies, which would be very useful, both to validate the patterns found in US data and to start providing a comparative perspective on the effects of task displacement. There is also much more that can be done with microdata on robots and other technology adoption. Most importantly, both additional reduced-form and structural approaches to how automation technologies affect competition between different firms—and via this channel affect productivity, labour demand and inequality—are important new horizons for this research.

There are several new areas for theoretical research as well. Most importantly, our evidence suggests that there has been a marked shift in the nature of technological progress towards much greater automation and much less reinstatement over the last 3 decades. Why this has happened is still poorly understood. Acemoglu and Restrepo (2018) provide a framework in which the direction of technological change—whether it automates or introduces new tasks—is endogenous. This framework reveals several economic forces that can affect the direction of technology (factor supplies, new breakthroughs affecting the innovation possibilities frontier, but also labour market institutions, distortions, taxes and research fads). While they show that under some conditions the long-run equilibrium may feature balanced advances in automation and new tasks, the equilibrium is in general not efficient, and may feature excessive automation. For example, the presence of labour market distortions that increase equilibrium wages above the social opportunity cost of labour implies that the equilibrium direction of technology will be biased towards too much automation.

Acemoglu et al. (2020*a*) argue, in addition, that US tax policy has created an inefficient bias towards automation—by taxing labour much more than capital, especially capital involved in automation, such as equipment and software. Whether the same is true for other industrialized economies is a question that requires further investigation. Even in the US context, much more can be done to estimate the implications of tax-induced changes in the direction of technology adoption and innovation.

Another interesting direction for future work is to explore the joint determination of rent-sharing and the task content of production. Bargaining or other sources of rent-sharing will be influenced by automation or even the threat of automation. In turn, high-rent tasks or workers may be more attractive targets for automation. The form of bargaining—and the power of unions—may also determine how easy it is for firms to automate certain tasks. There is little theoretical work and even less empirical work on these issues, and they may be important both for understanding changes in the direction of technological change in recent decades and the efficiency and distributional implications of these changes.

Equally important might be the direction of innovations when it comes to artificial intelligence (AI). AI is a broad technological platform that can be used for automation or for creating new labour-complementary tasks. Acemoglu and Restrepo (2020*c*) have argued that AI research has been inefficiently biased towards automation. In addition to the effects of tax policy and labour market imperfections, they have emphasized the business models of Big Tech companies and the 'values' and 'fads' in the AI industry. Acemoglu et al. (2022) provide some preliminary evidence that AI adoption has been targeted at replacing certain well-defined human tasks. This work, of course, does not deny

Table 4: Robustness to measures of task content, 1987–2016.

	College wage bill relative to high school wage bill			College hours relative to high school hours		
	(1)	(2)	(3)	(4)	(5)	(6)
	Panel A: BEA data with λ = 1 and 5-year moving averages					
Automation	0.620		0.535	0.412		0.335
	(0.138)		(0.154)	(0.120)		(0.136)
Reinstatement		0.931	0.606		0.755	0.551
		(0.333)	(0.350)		(0.301)	(0.329)
Observations	44	44	44	44	44	44
R^2	0.23	0.12	0.27	0.12	0.10	0.17
	Panel B: BEA data with λ = 0.8 and 10-year moving averages					
Automation	0.773		0.928	0.500		0.645
	(0.153)		(0.210)	(0.131)		(0.188)
Reinstatement		0.122	0.873		0.296	0.818
		(0.516)	(0.522)		(0.424)	(0.466)
Observations	44	44	44	44	44	44
R^2	0.22	0.00	0.27	0.11	0.01	0.17
	Panel C: BEA data with λ = 1 and 10-year moving averages					
Automation	0.630		0.807	0.385		0.537
	(0.149)		(0.195)	(0.130)		(0.169)
Reinstatement		0.593	1.195		0.627	1.028
		(0.557)	(0.563)		(0.486)	(0.512)
Observations	44	44	44	44	44	44
R^2	0.16	0.03	0.27	0.07	0.04	0.17
	Panel D: KLEMS data with λ = 0.8 and 5-year moving averages					
Automation	0.520		0.550	0.366		0.379
	(0.143)		(0.140)	(0.117)		(0.118)
Reinstatement		0.024	0.321		−0.072	0.132
		(0.368)	(0.333)		(0.355)	(0.344)
Observations	44	44	44	44	44	44
R^2	0.24	0.00	0.26	0.15	0.00	0.15
	Panel E: KLEMS data with λ = 1 and 5-year moving averages					
Automation	0.521		0.404	0.331		0.251
	(0.167)		(0.199)	(0.142)		(0.182)
Reinstatement		0.957	0.666		0.632	0.451
		(0.351)	(0.382)		(0.299)	(0.362)
Observations	44	44	44	44	44	44
R^2	0.14	0.11	0.19	0.07	0.06	0.10
	Panel F: KLEMS data with λ = 1 and 10-year moving averages					
Automation	0.444		0.558	0.243		0.322
	(0.200)		(0.199)	(0.170)		(0.165)
Reinstatement		1.196	1.535		0.865	1.060
		(0.716)	(0.719)		(0.670)	(0.673)
Observations	44	44	44	44	44	44
R^2	0.08	0.07	0.18	0.03	0.04	0.09

Note: The table provides regression estimates of changes from 1987 to 2016 in the relative demand for skills across industries on measures of displacement and reinstatement. The appendix in Acemoglu and Restrepo (2019) provides a description of the construction of these explanatory variables. Columns 1–3 use the change in the log of the college wage bill relative to the high school wage bill in each industry as outcome. Columns 4–6 use the change in the log of college hours relative to high school hours in each industry as outcome. Each panel presents results for a different construction of the displacement and reinstatement measures, as explained in the appendix in Acemoglu and Restrepo (2019). Standard errors robust against heteroscedasticity are given in parentheses.

that AI has also been used for technologies that increase productivity and labour demand as well as for reorganizing work. It certainly leaves open the possibility that its use can more strongly target more human complementary activities in the future.

The framework we have reviewed here and its implications become even more important in the context of rapid advances in AI technologies. AI is an example of a technological platform that will likely have pervasive effects throughout the economy and can be used for many different purposes. Our framework is well suited for incorporating these different types of uses of AI, for instance, in the form of automating certain tasks, improving worker decision-making in a range of knowledge tasks or creating new tasks for diverse skills.

Our framework emphasises that the consequences of these different uses of AI will be quite distinct, both for inequality and productivity. For one, greater automation powered by AI technologies could amplify the already high levels of inequality in industrialized labour markets. In contrast, however, AI has significant potential for creating new tasks for workers with diverse skills and remedying certain shortages of skills and expertise. For instance, AI-based technologies can generate new tasks for educators, nurses and even blue-collar workers, as detailed in Acemoglu and Restrepo (2020c) and Johnson and Acemoglu

(2023). Generative AI also has great potential for helping remedy skill shortages of knowledge workers, as already documented in small-scale settings by Noy and Zhang (2023) and Brynjolfsson et al. (2023). It may be particularly useful for building new platforms that bring together diverse skills and rich sets of tasks and demands together, as suggested by Acemoglu et al. (2021).

Our framework goes beyond clarifying the wage and inequality implications of these different uses of AI. We have also pointed out how certain types of automation may have disappointing productivity gains, especially if it is indeed the case that many digital technologies have been excessively focused on automation (and thus running into diminishing returns and not realizing productivity gains that exist via complementing humans). This may be driven by the business models of Big Tech companies or demands from automation. It may also be related to the powerful vision in the AI research community that places emphasis on 'autonomous machine intelligence' (Acemoglu and Restrepo 2020c; Johnson and Acemoglu 2023). If so, we may be concerned that the direction of development of new and promising AI technologies may go too much into the automation path as well. Acemoglu et al. (2022) provide some evidence consistent with this interpretation for (pre-generative) AI by documenting that AI activity has concentrated in establishments that have tasks that can be replaced by AI, and this activity has also come hand-in-hand with reduced hiring of non-AI workers in these establishments.

This evidence notwithstanding, our framework also stresses how new tasks created by AI could be quite transformative for the labour market and reduce, rather than increase, inequality. This may be especially relevant if generative AI has the capability to complement specialized workers with certain skill or expertise shortages (e.g. programmers who need help with some of the complex subroutines and programming steps, or electricians who can benefit from additional inputs for better diagnosis of problems). This type of use would both facilitate greater productivity and enable workers with diverse skills to start performing newer tasks, lowering inequality. By emphasizing the very different implications of AI depending on its path of development, and the endogeneity of this path, our framework raises several new areas for research and policy.

ACKNOWLEDGEMENTS

We gratefully acknowledge support from Accenture LLP, IBM Global Universities, Google, the National Science Foundation, the Sloan Foundation, Schmidt Futures and the Smith Richardson Foundation.

FUNDING

This article was written for the IFS Deaton Review of Inequalities, funded by the Nuffield Foundation (grant reference WEL/43603). The views expressed are those of the authors and not necessarily the Foundation.

REFERENCES

Acemoglu, D., and Autor, D. (2011) 'Skills, Tasks and Technologies: Implications for Employment and Earnings', in Card, D., and Ashenfelter, O. (eds) *Handbook of Labor Economics*, Vol. 4, Chapter 12. pp. 1043–171, Amsterdam: Elsevier.

Acemoglu, D., and Restrepo, P. (2018) 'The Race between Man and Machine: Implications of Technology for Growth, Factor Shares, and Employment', *American Economic Review*, **108**: 1488–542.

Acemoglu, D., and Restrepo, P. (2019) 'Automation and New Tasks: How Technology Displaces and Reinstates Labor', *Journal of Economic Perspectives*, **33**: 3–30.

Acemoglu, D., and Restrepo, P. (2020a) 'Robots and Jobs: Evidence from US Labor Markets', *Journal of Political Economy*, **128**: 2188–244.

Acemoglu, D., and Restrepo, P. (2020b) 'Unpacking Skill Bias: Automation and New Tasks', *AEA Papers and Proceedings*, **110**: 356–61.

Acemoglu, D., and Restrepo, P. (2020c) 'The Wrong Kind of AI? Artificial Intelligence and the Future of Labour Demand', *Cambridge Journal of Regions, Economy and Society*, **13**: 25–35.

Acemoglu, D., and Restrepo, P. (2022) 'Tasks, Automation, and the Rise in US Wage Inequality', *Econometrica*, **90**: 1973–2016.

Acemoglu, D., and Zilibotti, F. (2001) 'Productivity Differences', *Quarterly Journal of Economics*, **116**: 563–606.

Acemoglu, D. et al. (2016) 'Import Competition and the Great US Employment Sag of the 2000s', *Journal of Labor Economics*, **34**: S141–98.

Acemoglu, D., Lelarge, C., and Restrepo, P. (2020a) 'Competing with Robots: Firm-Level Evidence from France', *AEA Papers and Proceedings*, **110**: 383–8.

Acemoglu, D., Manera, A., and Restrepo, P. (2020b) 'Does the US Tax Code Favor Automation?' *Brookings Papers on Economic Activity*. Spring

Acemoglu, D., Jordan, M. I., and Weyl, G. (2021) 'The Turing Test Is Bad for Business', *Wired*. https://www.wired.com/story/artificial-intelligence-turing-test-economics-business/.

Acemoglu, D. et al. (2022) 'Artificial Intelligence and Jobs: Evidence from Online Vacancies', *Journal of Labor Economics*, **40**: S293–340.

Aghion, P., Jones, B. F., and Jones, C. I. (2017) 'Artificial Intelligence and Economic Growth', NBER Working Paper No. 23928.

Autor, D. H. (2019) 'Work of the Past, Work of the Future', *AEA Papers and Proceedings*, **109**: 1–32.

Autor, D. H., Katz, L. F., and Krueger, A. B. (1998) 'Computing Inequality: Have Computers Changed the Labor Market?' *Quarterly Journal of Economics*, **113**: 1169–213.

Autor, D. H., Levy, F., and Murnane, R. J. (2003) 'The Skill Content of Recent Technological Change: an Empirical Exploration', *Quarterly Journal of Economics*, **118**: 1279–333.

Autor, D. H., Dorn, D., and Hanson, G. H. (2013) 'The China Syndrome: Local Labor Market Effects of Import Competition in the United States', *American Economic Review*, **103**: 2121–68.

Autor, D. et al. (2020) 'The Fall of the Labor Share and the Rise of Superstar Firms', *Quarterly Journal of Economics*, **135**: 645–709.

Bessen, J. E. et al. (2019) 'Automatic Reaction – What Happens to Workers at Firms that Automate?' Boston University School of Law, Law and Economics Research Paper.

Bonfiglioli, A. et al. (2020) 'Robot Imports and Firm-Level Outcomes', CEPR Discussion Paper No. 14593.

Brynjolfsson, E., Li, D., and Raymond, L. R. (2023) 'Generative AI at Work', NBER Working Paper No. 31161.

Budd, E. C. (1960) 'Chapter 9: Factor Shares, 1850–1910', in *Trends in the American Economy in the Nineteenth Century*. pp. 365–406. Princeton, NJ: Princeton University Press.

Dao, M. C., Das, M., and Koczan, Z. (2019) 'Why Is Labour Receiving a Smaller Share of Global Income?', *Economic Policy*, **34**: 723–59.

Dauth, W. et al. (2021) 'The Adjustment of Labor Markets to Robots', *Journal of the European Economic Association*, **19**: 3104–53.

Dinlersoz, E. M., and Wolf, Z. (2018) 'Automation, Labor Share, and Productivity: Plant-Level Evidence from U.S. Manufacturing', Center for Economic Studies, Working Paper No. CES-WP-18-39.

Doms, M., Dunne, T., and Troske, K. R. (1997) ‘Workers, Wages, and Technology’, *Quarterly Journal of Economics*, **112**: 253–90.

Elsby, M. W., Hobijn, B., and Şahin, A. (2013) ‘The Decline of the US Labor Share’, *Brookings Papers on Economic Activity*, **2013**: 1–63.

Feenstra, R. C., and Hanson, G. H. (1999) ‘The Impact of Outsourcing and High-Technology Capital on Wages: Estimates for the United States, 1979–1990’, *Quarterly Journal of Economics*, **114**: 907–40.

Goldin, C. D., and Katz, L. F. (2008) *The Race Between Education and Technology*, Cambridge, MA: Harvard University Press.

Graetz, G. (2019) ‘Labor Demand in the Past, Present, and Future’, Technical report, Uppsala University.

Graetz, G., and Michaels, G. (2018) ‘Robots at Work’, *Review of Economics and Statistics*, **100**: 753–68.

Gutierrez, G., and Piton, S. (2020) ‘Revisiting the Global Decline of the (Non-Housing) Labor Share’, *American Economic Review: Insights*, **2**: 321–38.

Hellwig, M., and Irmen, A. (2001) ‘Endogenous Technical Change in a Competitive Economy’, *Journal of Economic Theory*, **101**: 1–39.

Hubmer, J. (2020) ‘The Race between Preferences and Technology’, Unpublished manuscript, University of Pennsylvania.

Humlum, A. (2020) ‘Robot Adoption and Labor Market Dynamics’, Working paper, University of Chicago.

Johnson, S., and Acemoglu, D. (2023) *Power and Progress: Our Thousand-Year Struggle Over Technology and Prosperity*, London: John Murray Press.

Karabarbounis, L., and Neiman, B. (2014) ‘The Global Decline of the Labor Share’, *Quarterly Journal of Economics*, **129**: 61–103.

Katz, L. F., and Murphy, K. M. (1992) ‘Changes in Relative Wages, 1963–1987: Supply and Demand Factors’, *Quarterly Journal of Economics*, **107**: 35–78.

Koch, M., Manuylov, I., and Smolka, M. (2021) ‘Robots and Firms’, *Economic Journal*, **131**: 2553–84.

Krueger, A. B. (1993) ‘How Computers Have Changed the Wage Structure: Evidence from Microdata, 1984–1989’, *Quarterly Journal of Economics*, **108**: 33–60.

Lin, J. (2011) ‘Technological Adaptation, Cities, and New Work’, *Review of Economics and Statistics*, **93**: 554–74.

Nardis, S. D., and Parente, F. (2021) ‘Technology, Task and Wage-Bill Changes in the EU Countries’, Luiss School of European Political Economy, Working Paper No. 6.

Noy, S., and Zhang, W. (2023) ‘Experimental Evidence on the Productivity Effects of Generative Artificial Intelligence’, Unpublished manuscript MIT.

Oberfield, E., and Raval, D. (2021) ‘Micro Data and Macro Technology’, *Econometrica*, **89**: 703–32.

Peterson Institute (2020) ‘How to Fix Economic Inequality? An Overview of Policies for the United States and Other High-Income Economies’, Technical report, Washington, DC: Peterson Institute for International Economics.

Piketty, T. (2014) *Capital in the Twenty-First Century*, Cambridge, MA: Harvard University Press.

Tinbergen, J. (1974) ‘Substitution of Graduate by Other Labour’, *Kyklos*, **27**: 217–26.

Zeira, J. (1998) ‘Workers, Machines, and Economic Growth’, *Quarterly Journal of Economics*, **113**: 1091–117.