Oxford Review of Economic Policy, Volume 34, Number 3, 2018, pp. 349–361

Technology and the labour market: the assessment

Abi Adams*

Abstract: This article reviews the evidence on how technology will affect labour markets in advanced democracies over the next 10–15 years. In addition to assessing evidence on the impact of technology on the aggregate number of jobs, we take a wider view to consider how the 'Digital Revolution' will affect the quality and composition of jobs, how and where we work, and whether labour is supplied as part of a larger firm or as a self-employed micro-entrepreneur. This assessment further aims to situate the papers covered in this issue of the *Oxford Review of Policy* within the broader literature.

Keywords: technological change, automation, labour markets, online labour markets, offshoring

JEL classification: O30, O33, J23, J01

> *Technological advances continue to revolutionalise the possibilities for humanity . . . we have to remember that the challenges we face do not outweigh the opportunities*. Theresa May, Davos Speech 2017

> *Technological progress is going to leave behind some people, perhaps even a lot of people, as it races ahead . . . , there's never been a worse time to be a worker with only 'ordinary' skills and abilities to offer, because computers, robots, and other digital technologies are acquiring these skills and abilities at an extraordinary rate*. Brynjolfsson and McAfee (2014, p. 11)

I. Introduction

The relationship between technological progress and the fate of workers has a fraught history. While technology is considered to be the engine of economic progress, its advancement is often portrayed as coming at the expense of jobs, generating risks of widespread technological unemployment and leaving some 'to vegetate in the backwaters of the stream of progress' (Landes, 2003). Popular anxiety about the consequences of the current 'digital revolution' is widespread and being expressed with increasing urgency.

*New College, Oxford, e-mail: abi.adams@economics.ox.ac.uk

The author would like to thank Christopher Adam, Alison Gomm, Cameron Hepburn, Ken Mayhew, Alex Teytelboym, and David Vines for very helpful comments and encouragement. Funding from the Economic and Social Research Council (grant ES/N017099/1) is gratefully acknowledged.

doi:10.1093/oxrep/gry010

Seventy-two per cent of surveyed Americans fear a future in which computers and robots can do more human jobs (Pew, 2017). A study predicting that 47 per cent of total US employment is at risk from digitalization 'went viral' in 2013 (Frey and Osborne, 2017).

In this issue of the *Oxford Review of Economic Policy*, we review the evidence on how technology has shaped labour market change since the 1980s and make some tentative predictions about how it might continue to do so in the short to medium term. Media portrayal of the topic has often focused on the impact of technology on the number of jobs. Yet technology has many effects on the labour market beyond its potential effect on the aggregate number of jobs in the economy. Technological change shapes the quality and composition of jobs, how we perform them and where we do them, and whether we work as part of a larger firm or as a self-employed micro-entrepreneur. The papers in this issue cover a broad range of topics in an attempt to provide a comprehensive assessment of our current understanding of how technological change impacts on the labour market.

The purpose of this Assessment is to draw out the key themes that are explored in detail in the papers in the issue. We begin with a fundamental question: what do economists mean by 'technology' before reviewing the key mechanisms at the heart of modern economic models of the labour market. We then move on to consider how the 'wiring of the labour market' is affecting how and where we work. We end by discussing the wider social and political factors influencing the speed at which technological progress deploys.

II. How do economists model technology?

Technology changes what we can produce from a given set of inputs and how we can substitute one input for another. It influences which factors get employed by which firms at what prices. At a high level, the papers in this issue consider the following ways in which technology might have affected the labour market in recent years:

(i) the relative productivity of different inputs to the production process;
(ii) the ease with which different inputs can be substituted for one another;
(iii) the nature and efficiency of the search and matching process;
(iv) organizational design.

The typical starting point for discussions of the interaction between labour and technology is the production function. The production function summarizes what level of output it is possible to achieve from a given amount of capital and labour (potentially with different skill levels). Technology alters the form of this function with consequences for employment and wages as organizations alter the mix of workers and capital they employ in order to maximize profits or minimize costs.

The traditional literature has typically abstracted from many features of working environments in order to examine the most salient aspects of worker–technology interactions. Discussions of skill-biased technological change, for example, usually consider changes in the relative productivity of 'skilled' and 'unskilled' labour, leaving details on how these skills are deployed in the market implicit (Bound and Johnson, 1992; Katz and Murphy, 1992). Abstracting from the use of capital, many contributions assume a 'constant elasticity of substitution' aggregate production function in which output, Y, is a function of skilled labour (L_s) and unskilled labour (L_u):[1]

$$Y = \left[\left(A_s L_s \right)^{\sigma} + \left(A_u L_u \right)^{\sigma} \right]^{1/\sigma}.$$

In this simple framework, technological change influences the labour market by changing relative factor specific productivities, A_s and A_u, as well as the elasticity of substitution between factors, σ. In particular, under the plausible assumption that $\sigma > 0$, technological change is skill biased if the productivity of highly skilled workers increases at a faster rate than that of less skilled workers. An increase in the relative productivity of skilled workers increases the relative demand for skill, causing the wages of high-skilled workers to increase at a faster rate than those of low-skill workers.

However, as discussed in the contribution to this issue by Hugh Lauder, Phillip Brown, and Sin-Yi Cheung (Lauder *et al.*, 2018), a downside of this approach is that the impact of technology on labour and production is somewhat of a black box. Using data on the relative supply of college graduates in the US and assuming an elasticity of substitution of 1.4 between college-educated and non-college-educated workers, Katz and Murphy (1992) estimate that relative productivities of college labour grew at approximately 10 per cent per year though the 1960s, 1970s, and 1980s. However, it is unclear what exactly this relative shift in productivities represents. Why and how has technological change manifested itself in this way? Furthermore, generalizing from simple intuitions about the elasticity of substitution in a two-factor world to cases with multiple inputs is non-trivial (Blackorby and Russell, 1989).

As developed in greater depth in Maarten Goos's contribution (Goos, 2018, this issue), modelling the interactions between workers and technologies today usually takes a more granular 'task-based' approach that makes the connection between technology, skills, and tasks more explicit. These task-based models make a distinction between a worker's skills and the tasks that they perform in a job (Autor *et al.*, 2003; Goos and Manning, 2007; Acemoglu and Autor, 2011). Workers have different endowments of different types of skill (e.g. creativity, reasoning, strength, etc.). These then combine to imply a certain capability for performing tasks. Relaxing the one-to-one mapping between skills and tasks allows for new insights into how workers with different skill endowments end up performing different sorts of jobs as technology, demand conditions, and factor supplies vary. It is with this framework in mind that we now turn to discuss how technology substitutes and complements different types of skill in the performance of particular tasks in more detail.

III. Technology–skill complementarities

The impact of a particular technology on employment and wages depends both on whether technology substitutes or complements for labour and also on equilibrium impacts that manifest themselves through changes in labour supply and product demand.

[1] See for example, Goldin and Katz (2009).

Many workplace technologies are designed to save labour—both of the brawn and brain variety. Historical examples often centre on the direct substitution of manual labour for machines. During the Industrial Revolution some prominent economists voiced concern at the scale and pace of the mechanization of production. As discussed by Mokyr *et al.* (2015), even Ricardo appeared to have a change of heart:

> in the past, Ricardo (1821) noted he had been convinced that an application of machinery to any branch of production was a general good, but he had more recently concluded that the 'substitution of machinery for human labour is often very injurious to the interests of the class of labourers. . . . [It] may render the population redundant and deteriorate the condition of the labourer.

Non-manual tasks have also not been immune in the past to direct automation. Human 'calculators' played a vital role in the Second World War effort before being displaced by 'electric calculation machines' that did not tire (Light, 1999).

To assess the likelihood that a particular task may be automated, economists have moved beyond a simple 'skilled' versus 'unskilled' or 'manual' versus 'non-manual' distinction to consider how 'routine' a task is (Autor *et al.*, 2003). As laid out in detail in Frank Levy's contribution to this issue (Levy, 2018), assessments of the routineness of a task, and thus its ability to be automated,

> begin with two observations:
>
> - All human work involves the processing of information. A financial analyst reading a report, a chef tasting a sauce, a farmer looking to the sky for signs of rain: these are each examples of processing information into an understanding of what to do next or to update a picture of the world.
> - A computer processes information by executing instructions.
>
> It follows that for [artificial intelligence] to automate a task, it must be possible to model the required information processing using a set of instructions.

'Routine' tasks were originally considered to be those that could be *fully* specified as a sequence of instructions that together fully enumerate what actions a machine will perform in what order in every state of the world. Non-routine tasks on the other hand have until very recently been thought of as the manifestation of Polanyi's paradox (Autor, 2014). In the words of philosopher Michael Polanyi (1966):

> We can know more than we can tell. . . . The skill of a driver cannot be replaced by a thorough schooling in the theory of the motorcar; the knowledge I have of my own body differs altogether from the knowledge of its physiology.

The key idea being that much of our understanding of the world is tacit and our performance of many tasks evades explicit enumeration into a set of deductive instructions. This conceptualization of (non-)routineness led Autor *et al.* (2003) to predict that truck-driving, for example, would not be automated in the near future.

Yet, recent technological developments in machine learning and natural language processing increasingly enable us to automate tasks without necessarily being able to articulate the underlying information-processing structure. Thus, by allowing us to discover information-processing rules that we cannot articulate, machine learning enables us to unravel at least a part of Polanyi's paradox. This is leading to a shift in the scope and nature of tasks that might be considered ‘routine’ and thus what tasks we believe technology can directly substitute for. Indeed, the original prediction that trucking would be immune to automation in the near term has already been proven false. Multiple companies are currently piloting driverless trucks and the International Transport Forum now believes that driverless trucks will be a regular presence on many roads within the next decade.[2]

Discussions of this type seem to suggest that we are, indeed, doomed. And perhaps doomed sooner than we think. As we collect more and more data upon which to ‘train’ our machine-learning algorithms, the range of processes that can be automated will increase as our tacit understanding of the world gets solidified in the predictions of increasingly accurate algorithms. However, it is important to bear in mind two important facts in such discussions. First, machine-learning models have difficulties in predicting outcomes for cases that lie beyond the data on which they were estimated and in scenarios where the structure connecting the features of a particular decision-making situation to outcomes is unstable.

Second, and importantly, as David Autor (2015) writes, ‘tasks that cannot be substituted by automation are generally complemented by it’. That is, automating some set of tasks in a job can increase human productivity on the remaining elements. Giving the example of lawyers, Levy notes that many of the tasks lawyers perform involve unstructured cognitive work that is aided by automation of quick document discovery and classification that is now possible through greater use of machine-learning techniques in the field. The rise of automatic teller machines (ATMs) in retail banks is another classic example. The roll out of ATMs in the 1970s actually coincided with a rise in the number of human bank tellers employed in retail banks (Bessen, 2015). With the essential elimination of the routine, cash-dispensing tasks from the job, bank tellers were able to spend more time on a wider range of advice and sales roles to customers. Further, the fall in the cost of operating a bank branch facilitated a rise in the overall number of branches, again mitigating any negative impacts on the overall number of people employed. It is, however, unclear how useful this analogy is to predictions of the future of employment in retail banking as we move to a phase where digitization and online banking is starting to reduce branch networks.

Thus, to understand the direction of the effect of technological change, it is important to consider the balance of such tasks within a particular job that are substitutable for or complementary with digital capital. Indeed, as explored by Susskind (2017), it is also important to bear in mind that labour might not be best placed to perform the tasks that are complementary to automation; might not ‘advanced capital’ erode the comparative advantage of workers in performing these complementary tasks? A more nuanced approach is explored in Levy’s contribution as he works through a set of case studies to demonstrate the likely heterogeneity of employment effects. Given what is likely to be feasible to automate in the coming decade, software is likely to substitute for a maximum of 13 per cent of a corporate lawyer’s time. By contrast, automated coding of a radiologist’s report can substitute for 100 per cent of the radiology coder’s time.

[2] See https://www.ft.com/content/7686ea3e-e0dd-11e7-a0d4-0944c5f49e46

IV. Equilibrium impacts

Technological progress does not only affect relative labour demand. Its direct effect on employment can be mitigated or amplified by the wider effects that technological change might have on the economy, either through changes in product demand or labour supply. As explained by Goos (2018, in this issue), as technological progress lowers costs and product prices in some sectors, one would expect a shift in product demand to reflect this. These effects are found to be an important attenuating factor in a set of studies (Goos *et al.*, 2014; Gregory *et al.*, 2016; Acemoglu and Restrepo, 2017), although the effect has not been strong enough to fully compensate for the direct effect of technological change on employment as product demand is relatively inelastic in many of the affected sectors.

Yet, even in such circumstances where the elasticity of demand for a sector is below unity, aggregate demand does not necessarily fall with improved productivity. As income rises through productivity improvements in technologically leading sectors, employment will be boosted in sectors that are income elastic. Beauty services, fitness, personal improvement services, and other 'technologically lagging' sectors might see a boost in employment if demand for such goods is strongly income-elastic. More than $16 billion were spent on yoga in the United States in 2016, for example (Yoga Journal, 2016).

Labour supply responses might also offset wage changes and unemployment effects. Keynes (1930) famously believed that technological unemployment would in the long run manifest itself as a technologically facilitated reduction in labour supply and an increase in leisure: 'all [technological unemployment] means in the long run is that mankind is solving its economic problem'. However, in this phase of 'temporary maladjustment', task-based models predict that wage inequality is likely to widen with the automation of middle-skilled jobs. It is often assumed that the comparative advantage of low-skilled to medium-skilled workers in the performance of non-routine low-paid tasks is weak, whereas the comparative advantage of high-skilled to medium-skilled workers in the performance of non-routine high-paid tasks is strong. Thus, labour supply to lower-wage occupations is relatively more elastic than to higher-wage occupations, attenuating any wage increases in non-routine low-paid tasks and increasing the relative wage of those in high-paid non-routine jobs.

V. Distributional consequences

This discussion highlights that even if technology is not necessarily labour decreasing, it will have big distributional consequences, some of which can already be seen in modern-day labour markets. This is the focus of the active literature on job polarization. As shown in the contribution by Erik Buyst, Maarten Goos, and Anna Salomons in this issue (Buyst *et al*, 2018), the period since the 1980s has been characterized by increased employment in high- and low-paid jobs at the expense of 'rapidly declining' medium-paid jobs in administration and manufacturing. This feature of recent labour market change has also been noted in many studies on other advanced democracies (Goos and Manning, 2007; Acemoglu and Autor, 2011; Goos *et al.*, 2014).

The presumption is that many 'middling' jobs are relatively intensive in the routine-tasks that are most at risk of automation or off-shoring. This interacts with technological change to produce a U-shaped relationship between employment share changes over time and jobs ranked by their wage level (as product demand effects have not fully offset the direct effect of technology on middle-income jobs). The relatively higher elasticity of labour supply into low-wage occupations also implies an increase in wage inequality along with this hollowing-out process. As discussed above, the scope of tasks in which computers might feasibly be substituted for workers looks set to be wider than previously envisaged as the tacit knowledge required in many upper-middle class occupations, such as solicitors, becomes codified.

Based on these insights, Goos (2018) identifies several policy areas that require greater development. While investment in STEM education (science, technology, engineering, and mathematics) is important to raise the productivity of high-skilled workers who interact directly with new forms of technology, there should also be a concerted effort to build skills that will remain difficult to automate in the near future. There should be an effort better to measure non-routine social, motivational, and interactional skills that are at the core of what it means to be human to move forward thinking on how or if these skills can be taught. Given the process of job polarization, Goos also discusses the importance of different labour market and income redistribution policies to ensure that the benefits of the Digital Revolution are broadly shared. Furthermore, to help those displaced from middle-income jobs, better mature education and retraining programmes should be developed.

VI. Technology and imperfect information

Further routes through which technology can influence the labour market appear once one relaxes the assumption of perfect information in markets. As is discussed in more detail below, and developed fully in Adeline Pelletier and Catherine Thomas's article in this issue (Pelletier and Thomas, 2018), labour markets are characterized by imperfect information about the location of workers and of vacancies, and by asymmetric information on the ability of workers and the quality of their work. Here technology can change both the search and hiring process and also facilitate new monitoring and performance management schemes once an employment relation is initiated

(i) Hiring

Labour market structure and efficiency are shaped by technological developments in search and hiring processes. The *Oxford Review of Economic Policy* has an established history of predicting the shape of these developments. As noted by Pelletier and Thomas in this issue, in 2002 Richard Freeman discussed in this journal how the internet was redefining job search and recruitment by allowing information about jobs and workers to be submitted at low cost (Freeman, 2002).

Job search has, indeed, been transformed since the early 2000s. Over 100m job advertisements were posted online in the US between 2007 and 2015. A Pew Research Center study in 2015 found that 79 per cent of Americans looking for work had used online resources and 34 per cent of respondents said that online resources were the most important resource available to them (Pew, 2015). In 1998, only 15 per cent of job seekers reported using the internet to find a job (Kuhn and Skuterud, 2000). However, both Hershbein and Kahn (2018) and Balgova (2018) note that jobs posted online are disproportionately higher skill. While the reason for this has not yet been comprehensively established, it is plausible that local labour markets for skilled labour are relatively thin and that the wage gains to moving are more likely to dominate relocation costs for these individuals, increasing the likelihood of someone geographically remote accepting a job offer, and thus increasing the benefit to firms of advertising widely for these roles.

These developments should have increased the efficiency with which workers and jobs are matched. Workers are now able to find out about a larger quantity of jobs openings and sort through the information on openings in a quicker fashion. Firms have more mechanisms with which to attract and screen workers. There is emerging evidence that online job search is associated with higher efficiency and match quality as predicted. Kuhn and Mansour (2014) and Faberman and Kudlyak (2016) find that online job search is associated with significantly shorter unemployment durations and longer tenures once in a job (Prakash, 2014).

(ii) Production

Developments in this area have, however, now gone far beyond the matching process itself. Recent years have seen the development of online labour markets in which the matching of firms and workers is done online, the work itself is performed online, and the output of this works is also delivered electronically. Platforms of this type include Amazon Mechanical Turk, Upwork, and peopleperhour.com. The growth rate of the sector is estimated to be in excess of 25 per cent per year (Kässi and Lehdonvirta, 2016), although the number of people working in this way in developed economies is still small relative to the overall size of the workforce.

These new platforms obviate the need for firms and workers to be collocated. Even though the majority of employers that make use of platform work are located in the US, over half of workers are located in developing economies, notably India (Pelletier and Thomas, 2018). Insights from trade suggest that the integration of geographically separate labour markets has the potential to bring about large benefits for employers (Autor, 2001). By making use of workers in areas where labour is relatively cheap, labour costs can fall and firms can also realize greater scale economies than might be possible if their workforce were limited to that available in the local labour market. As labour demand and supply to a particular region get more elastic, one would expect regional wage differentials to diminish in jobs that can be performed via the internet.

However, while the potential benefits and scale of labour market change that might arise from online labour markets is large, there are significant hurdles holding back the development of these markets. As reviewed in Pelletier and Thomas (2018), gathering and transmitting the information required to enable efficient production in this form seems prohibitively costly for many firms. New employers currently have little information about the operation of these markets and pervasive difficulties in transmitting information about production limit the type of work than can be effectively performed online, especially when production must be coproduced with teammates. Further, it is difficult for workers to transmit 'high bandwidth' information via the internet that employers need to assess the fit of applicants for certain types of jobs.

There are reasons to believe that this missing information may have larger consequences on online labour markets than their offline counterparts. As work is typically done on a 'gig-by-gig' basis, employers have little incentive to uncover missing information if the contracts are short term whereas they might if contracts were longer term. As Pelletier and Thomas conclude, online labour markets have not yet brought about the 'death of distance' (Cairncross, 1997) because of the information issues that still develop with institutional and geographic distance. New ways to improve the ease of information exchange about the organization of online production and hiring is required to unlock further growth in these markets.

(iii) Monitoring

While not reviewed in this issue, a final area in which technology is playing an increasing role is the monitoring of workers engaged in a job itself. Nearly a third of large employers surveyed in 2016 by Willis Towers Watson, a professional services group, had staff use wearable devices such as activity monitors or 'mood rings' that measure emotional intensity.[3] A recent survey by the Economist Intelligence Unit found that half of human resources departments surveyed reported an increased use of workplace analytics compared with 3 years ago. Both manual and non-manual jobs are affected by the trend. For example, Tesco uses electronic armbands to track the movements of stock pickers in its warehouses. The armbands direct staff around the warehouse, instruct them to select particular items, and automatically log inventory. The company Steelcase puts sensors in office furniture to monitor how employees interact and how long they stay at their desk, while Evolv takes screenshots to monitor how individuals are working in real time.

The increased ease with which firms can monitor their employees should help to raise productivity by limiting moral hazard in the workplace. Evolv claims that its monitoring can help to improve productivity by at least 5 per cent in two-thirds of jobs, although this claim has not been independently verified. Insights from wearables can also allow firms to restructure jobs in a way that benefits its workers. Bank of America, for example, found that workers were more productive when they were allowed to take breaks together. Upon rolling out this policy universally, performance improved by 23 per cent and the 'amount of stress in workers' voices fell 19 per cent.'[4]

However, there are reasons to be cautious about developments in this area. If not all aspects of a job can be monitored and performance managed to the same degree, employers will need to be careful not to distort their employees' efforts across tasks that are more or less captured by the technology. Further, some studies have found that wearables and monitoring increase workplace stress with potentially detrimental effects on productivity and retention.

[3] See https://www.jiff.com/content/uploads/2016/11/Wearables-at-Work-Exec-Summary.pdf

[4] See https://www.ft.com/content/d56004b0-9581-11e3-9fd6-00144feab7de

VII. Organizational structure

The increased ability of firms and workers to meet and for production to be more closely monitored is shifting incentives for firms to adopt particular legal forms. Milgrom and Roberts (1990) argue that new technologies that have reduced the cost of data storage, communication, and monitoring activities within a firm have triggered a shift in organizational structures. Specifically, they argue that these developments have facilitated a flattening of the hierarchical structures, allowing highly skilled workers to perform a wide range of tasks in teams. Thus, in addition to direct substitution effects, technology might also induce skill-biased organizational change.

While there is some evidence in support of this hypothesis (Garicano and Rossi-Hansberg, 2004), a more salient feature of recent labour market change has been the growth of self-employment. Forty per cent of employment growth in the UK between 2008 and 2017 was attributable to increases in self-employment and companies with a single owner-manager. The rise of the so-called 'gig economy' is one salient aspect of this shift. This term is increasingly used to refer to individuals performing small tasks as independent contractors through digital platforms. Popular operators in the sector include Uber, Deliveroo, and TaskRabbit (Prassl, 2018).

Some have argued that technological change has reduced economies of scale and alleviated transactional failures, reducing costs to self-employment and making it easier for firms and workers to choose their legal form to better align with their preferences, to minimize their tax burden, and for employers to limit their exposure to employment law (Sundararajan, 2016). Online gig-economy applications ('apps') facilitate the efficient matching of engagers and labour suppliers. The provision of rating systems and the constant stream of data on individual performance discussed above are said to have created new mechanisms for control and provision of incentives beyond those found in traditional contracts of employment. As discussed in the article by Abi Adams, Judith Freedman, and Jeremias Prassl in this issue (Adams *et al.*, 2018), such developments have made it easier for firms and individuals to respond to long-standing pathologies in regulatory systems, making systematic and holistic reform of tax and employment law a pressing issue.

VIII. Concluding thoughts: what does this mean for policy and politics?

Technological developments will have a profound effect on labour markets. Even if new technologies do not necessarily destroy jobs, they will radically affect the quality and composition of jobs. They will affect how individuals find work and how they are monitored while performing it.

The impacts of these changes will not be limited to the labour market; there will be corresponding impacts upon society and democracy. The role of policy is not to prevent these impacts—this is neither possible nor desirable—but rather to cushion the transition. As Goos (2018) discusses in his piece, a wide-ranging policy response is required to prepare future generations and to help displaced workers find new high-quality work.

This might require new ways to build capacity in non-routine skills, as well as building STEM skills themselves.

There is also a role for policy now to shape the nature of future technological developments. Many authors in this issue caution against the 'science-push' view of the adoption and impact of technology on the labour market. As noted by Levy (2018), the ultimate impact of new advances will depend on the policy choices of governments, the production choices of firms, the labour supply and consumption behaviour of workers and consumers, and their voting behaviour as democratic citizens. Over time, some technologies will thus deploy faster than others and some occupations will be disrupted faster than others. The sequence and speed of developments and people's reactions to the developments will jointly determine how the economy evolves.

Carl Benedikt Frey, Thor Berger, and Chinchih Chen consider whether technological disruption is already feeding through into voting behaviour (2018, in this issue). Their paper is inspired by the suggestion of Wassily Leontief (1953) that: 'If horses could have joined the Democratic party and voted, what happened on farms might have been different.' They find that electoral districts with a higher share of jobs exposed to automation were significantly more likely to support Trump in the 2016 US presidential election, concluding that to 'avoid further populist rebellion and a looming backlash against technology itself, governments must find ways of making the benefits from automation more widely shared'. Levy, too, urges a surge in political will to help local areas prepare for changes in occupational structure to prevent destabilization of the political and social structures of advanced democracies.

While it seems clear that large occupational shifts are on the horizon, one final point to note is that recent technological change is not yet showing up in improved productivity statistics in any meaningful way. As explored in Nicholas Crafts' contribution, a surprising and worrying feature of the economic environment in the advanced economies in recent years has been slow productivity growth. The slowdown started before the financial crisis and weak performance has continued through the last 10 years. Estimates of trend growth are now more pessimistic and projections of future economic growth have become less optimistic. A review of the best evidence in this area does not lead to clear conclusions concerning the degree to which this slowdown is real or reflects poor measurement of new digital outputs and services. This serves to underline the fundamental uncertainty that we face concerning the shape and pace of technological change and its impacts on the world of work.

To end then, it seems prudent to again return to the timeless advice of Keynes (1930). While we build our capacity for measurement of technological progress and its impacts on the labour market and beyond, 'there will be no harm in making mild preparations for our destiny, in encouraging, and experimenting in, the arts of life as well as the activities of purpose'.

References

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

Acemoglu, D., and Restrepo, P. (2017), 'Robots and Jobs: Evidence from US Labor Markets', NBER Working Paper No. 23285.

Adams, A., Freedman, J., and Prassl, J. (2018), 'Rethinking Legal Taxonomies for the Gig Economy', *Oxford Review of Economic Policy*, **34**(3), 475–94.

Autor, D. H. (2001), 'Wiring the Labor Market', *Journal of Economic Perspectives*, **15**(1), 25–40.

— (2014), 'Polanyi's Paradox and the Shape of Employment Growth', NBER Working Paper No. 20485, Cambridge, MA, National Bureau of Economic Research.

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

— Levy, F., and Murnane, R. J. (2003), 'The Skill Content of Recent Technological Change: An Empirical Exploration', *Quarterly Journal of Economics*, **118**(4), 1279–333.

Balgova, M. (2018), 'Essays in Search and Matching', DPhil thesis, University of Oxford.

Bessen, J. (2015), 'Toil and Technology', *Finance and Development*, **52**(1).

Blackorby, C., and Russell, R. (1989), 'Will the Real Elasticity of Substitution Please Stand Up?', *American Economic Review*, **79**(4), 882–8.

Bound, J., and Johnson, G. (1992), 'Changes in the Structure of Wages in the 1980s: An Evaluation of Alternative Explanations', *American Economic Review*, **82**(3), 371–92.

Brynjolfsson, E., and McAfee, A. (2014), *The Second Machine Age: Work, Progress, and Prosperity in a Time of Brilliant Technologies*, 1st edn, New York, W.W. Norton & Company.

Buyst, E., Goos, M., and Salomons, A. (2018), 'Job Polarization: An Historical Perspective', *Oxford Review of Economic Policy*, **34**(3), 461–74.

Cairncross, F. (1997), *The Death of Distance: How the Communications Revolution is Changing Our Lives,* Cambridge, MA, Harvard Business School Press.

Crafts, N. F. (2018), 'The Productivity Slowdown: Is it the "New Normal"?', *Oxford Review of Economic Policy*, **34**(3), 443–60.

Faberman, J., and Kudlyak, M. (2016), 'What Does Online Job Search Tell Us About The Labor Market?', *FRB Chicago Economic Perspectives*, **40**(1).

Freeman, R. B. (2002), 'The Labour Market in the New Information Economy', *Oxford Review of Economic Policy*, **18**(3), 288–305.

Frey, C. B., and Osborne, M. (2017), 'The Future of Employment: How Susceptible are Jobs to Computerisation?', *Technological Forecasting and Social Change*, **114**(C), 254–80.

— Berger, T., and Chen, C. (2018), 'Political Machinery: Did Robots Swing the 2016 US Presidential Election?', *Oxford Review of Economic Policy*, **34**(3), 418–42.

Garicano, L., and Rossi-Hansberg, E. (2004), 'Inequality and the Organization of Knowledge', *American Economic Review*, **94**(2), 197–202.

Goldin, C., and Katz, L. F. (2009), 'The Race Between Education and Technology: The Evolution of US Wage Differentials, 1890–2005', NBER Working Paper 12984.

Goos, M. (2018), 'The Impact of Technological Progress on Labour Markets: Policy Challenges', *Oxford Review of Economic Policy*, **34**(3), 362–75.

— Manning, A. (2007), 'Lousy and Lovely Jobs: The Rising Polarization of Work in Britain', *Review of Economics and Statistics*, **89**(1), 118–33.

— — Salomons, A. (2014), 'Explaining Job Polarization: Routine-biased Technological Change and Offshoring', *American Economic Review*, **104**(8), 2509–26.

Gregory, T., Salomons, A., and Zierahn, U. (2016), 'Racing With or Against the Machine? Evidence from Europe', ZEW Discussion Paper No. 16-053.

Hershbein, B., and Kahn, L. (2018), 'Do Recessions Accelerate Routine-biased Technological Change?', *American Economic Review*, forthcoming.

Kässi, O., and Lehdonvirta, V. (2016), 'Online Labour Index: Measuring the Online Gig Economy for Policy and Research', MPRA Paper 74943, University Library of Munich.

Katz, L. F., and Murphy, K. M. (1992), 'Changes in Relative Wages, 1963–1987: Supply and Demand Factors', *Quarterly Journal of Economics*, **107**(1), 35–78.

Keynes, J. M. (1930), 'Economic Possibilities for our Grandchildren', section **V**, ch. 2 in *Essays in Persuasion*, New York, Classic House Books.

Kuhn, P., and Mansour, H. (2014), 'Is Internet Job Search Still Ineffective?', *The Economic Journal*, **124**(581), 1213–33.

— Skuterud, M. (2000), 'Job Search Methods: Internet versus Traditional', *Monthly Labor Review*, **123**(3).

Landes, D. S. (2003), *The Unbound Prometheus: Technological Change and Industrial Development in Western Europe from 1750 to the Present,* Cambridge, Cambridge University Press.

Lauder, H., Brown, P., and Cheung, S.-Y. (2018), 'Fractures in the Education–Economy Relationship: The End of the Skill Bias Technological Change Research Programme?', *Oxford Review of Economic Policy*, **34**(3), 495–515.

Leontief, W. (1953), *Studies in the Structure of the American Economy,* Oxford, Oxford University Press

Levy, F. (2018), 'Computers and Populism: Artificial Intelligence, Jobs, And Politics In The Near Term', *Oxford Review of Economic Policy*, **34**(3), 393–417.

Light, J. S. (1999), 'When Computers were Women', *Technology and Culture*, **40**(3), 455–83.

Milgrom, P., and Roberts, J. (1990), 'The Economics of Modern Manufacturing: Technology, Strategy, and Organization', *American Economic Review*, **80**(3), 511–28.

Mokyr, J., Vickers, C., and Ziebarth, N. L. (2015), 'The History of Technological Anxiety and the Future of Economic Growth: Is This Time Different?', *Journal of Economic Perspectives*, **29**(3), 31–50.

Pelletier, A., and Thomas, C. (2018), 'Information in Online Labour Markets', *Oxford Review of Economic Policy*, **34**(3), 376–92.

Pew (2015), 'Searching for Work in the Digital Era', Washington, DC, Pew Research Center.

— (2017), 'Automation in Everyday Life', Washington, DC, Pew Research Center.

Polanyi, M. (1966), *The Tacit Dimension,* New York, Doubleday.

Prakash, A. (2014), 'Does Internet Job Search Result in Better Matches?', University of Arizona, mimeo.

Prassl, J. (2018), *Humans as a Service,* Oxford, Oxford University Press.

Sundararajan, A. (2016), *The Sharing Economy—The End of Employment and the Rise of Crowd-Based Capitalism,* Boston, MA, MIT Press.

Susskind, D. (2017), 'A Model of Technological Unemployment', Department of Economics Discussion Paper 819, University of Oxford.

Yoga Journal (2016), '2016 Yoga in America Study', conducted by *Yoga Journal* and Yoga Alliance.