Contents

Summary

As machines become increasingly capable of performing tasks once thought to be the sole preserve of people, some commentators have raised the spectre of mass unemployment and profound economic disruption. Yet despite the growing capability of robots and artificial intelligence (AI), we are not on the cusp of a 'post-human' economy. Automation will produce significant productivity gains that will reshape specific sectors and occupations. In aggregate, however, these gains are likely to be recirculated, with jobs reallocated rather than eliminated, economic output increased, and new sources of wealth created.

The critical challenge of automation is likely to be in distribution rather than production. If the benefits are fairly shared, automation can help build an economy where prosperity is underpinned by justice, with a more equitable distribution of wealth, income and working time. But there is no guarantee that this will occur. Managed poorly, automation could create a 'paradox of plenty': society would be far richer in aggregate, but, for many individuals and communities, technological change could reinforce inequalities of power and reward.

The pace, extent, and distributional effects of automation will be determined by our collective choices and institutional arrangements, and the broader distribution of economic power in society. The future will not be technologically determined; it will be what we choose to create. Public policy should therefore actively shape the direction and outcome of automation to ensure its benefits are fairly shared. Three broad steps can help achieve this.

First, realising the benefits of technological advances will require the managed acceleration of automation and the adoption of digital technologies throughout the economy. Government needs to provide greater support for firms in all sectors and parts of the country to integrate new technologies, improve management, achieve higher rates of investment, and enable a stronger voice for employees in shaping the use of technology at work. The skills system also needs to be reformed to ensure people can thrive in an era of greater human-machine collaboration.

Second, new public institutions are needed to inform and regulate how automating technologies are used and to ensure that society responds proactively to the profound ethical issues raised by robotics and artificial intelligence.

Third, new models of collective ownership are required to ensure that everyone has a claim on the dividends of technological change, to enable automation to work for the common good.

These conclusions are based on five key propositions.

1. **Work will be transformed by automation, not eliminated.** Automation is likely to lead to the steady redeployment of labour over a period of decades, rather than a sudden and rapid elimination of employment. The task contents of most jobs will evolve, changing the nature of work.
    - **Aggregate effects on employment are likely to offset negative sectoral impacts**. As automating technologies are increasingly introduced, rising labour productivity will cause a decline in the numbers of some kinds of jobs in some sectors. These are likely to be offset by an increase in the demand for labour in other sectors, and in other kinds of jobs. The overall

impact in any sector will depend on the elasticity of demand for that sector's output. Within the economy as a whole, employment is likely to be reallocated rather than eliminated.

- **Automation is likely to change the composition of the tasks that people perform.** An estimated 60 per cent of occupations have at least 30 per cent of activities which could be automated with already-proven technologies. As tasks are automated, work is likely to be redefined, focusing on areas of human comparative advantage over machines.
- **Polarisation between 'lovely' jobs and 'lousy' jobs is a serious risk**. Automation could increase the demand for work in creative, cognitive, planning, decision-making, managerial and caring roles, where humans still outperform machines. New jobs and ways of working, often in close partnership with machines, will emerge. However, some emerging technologies will risk reducing autonomy at work and intensifying exploitation. The quality of work should therefore be a key focus of policy.

2. **In the absence of policy intervention, the most likely outcome of automation is an increase in inequalities of wealth, income and power.** The economic dividends of automation are likely to flow to the owners of technologies and businesses, and the highly skilled, as income shifts from labour to capital and the labour market polarises between high- and low-skilled jobs.
   - **In the UK, the total level of wages associated with jobs with the technical potential to be automated is £290 billion per annum.** If automation leads to lower average wages or working hours, or loss of jobs in aggregate, a significant amount of national income could be transferred from labour to capital. Even if wages do not decline, if relative rewards to capital rise more quickly, the share of national income going to capital would increase.
   - **Jobs with the highest potential for automation typically have lower wages.** On average, low wage jobs have five times the technical potential to be automated of high paid jobs. Technological change is likely to increase the incomes of highly skilled labour in roles which augment machines.
   - **Different regions and sectors are variably susceptible to automation.** London has the highest proportion of jobs assessed as more resilient to automation trends; poorer regions have a larger number of jobs with greater technical potential for automation. Transportation (63 per cent), manufacturing (58 per cent), and wholesale and retail trade (65 per cent) sectors have the greatest number of such jobs.
   - **Automation risks increasing gender and race inequality.** Of all the jobs in the UK, a greater proportion of those held by women compared to men are likely to be technically automatable, and women make up a smaller proportion of people in high-skill occupations that are resilient to automation or complemented by technology. Some ethnic groups are also more likely to work in low-skill 'automatable' occupations. The impact on inequality will depend on the skill-level of new jobs created and individuals' ability to access opportunities.

3. **A managed acceleration of automation is needed to reap the full productivity benefits and enable higher wages.** Due to the UK's low investment rates, poor management practices, and long tail of low-wage, low-productivity firms, it is the relative absence of robots in the UK economy, not their imminent rise, that is the biggest challenge.

- **Automation can deliver a powerful boost to UK productivity.** An accelerated trajectory of automation could raise productivity growth by between 0.8 to 1.4 per cent annually, boosting GDP by 10 per cent by 2030.
- **Productivity growth among ordinary firms has stalled, partly due to slow rates of technological adoption, low investment, and weak management.** While the top 1 per cent of firms have seen average productivity growth of around 6 per cent per year since 2000, one-third of UK companies have seen no rise in productivity at all. This is partly explained by a very slow rate of adoption of new and digital technologies. Overall business spending on ICT, machinery and other equipment has hardly grown in real terms since 2000.
- **The more rapid adoption of digital technologies, including automation, should become one of the national 'missions' of the Government's industrial strategy.** The goal should be to make the UK the most digitally advanced economy in the world by 2040.
- **A new partnership body, Productivity UK, should be established with the goal of raising firm-level productivity, including the acceleration of investment in automation technologies.** It should focus on the wider adoption of digital and other technologies throughout the economy. Working with sector bodies and local economic partnerships, it would provide support to businesses to access diagnostic services, advice, management training and skills development.
- **The UK skills system needs to better equip people with skills to complement automating technologies and retrain where jobs are lost**. The apprenticeships levy should be turned into a 'productivity and skills levy' that firms can use for wider skills training and utilisation. A personal retraining allowance should be introduced for workers made redundant, or with below level 3 qualifications, in communities or sectors facing economic decline or transition.

4. **An Authority for the Ethical Use of Robotics and Artificial Intelligence should be established to regulate the use of automating technologies.** The ethical and regulatory framework governing the use of technologies should precede, rather than follow in the wake of, technological development.
    - **Increasingly powerful and widespread 'autonomous technologies' will raise profound ethical questions.** Without action, the ethical and social norms framing the use of new technologies will be determined by technology companies, which are not accountable to society as a whole.
    - **The proposed authority should make recommendations to Government and business on the governance and use of robots and AI.** It should be modelled on the Human Fertilisation and Embryology Authority (HFEA) that regulates embryonic technologies, ensuring that society determines the rules and ethical frameworks governing autonomous technologies.

5. **New models of capital ownership are needed to ensure automation broadens prosperity rather than concentrates wealth**. Spreading capital ownership will help ensure automation creates an economy where prosperity is underpinned by justice.
    - **Capital is highly unevenly owned in the UK.** The wealthiest 10 per cent of UK households own an estimated 77 per cent of all equity stocks. The richest 10 per cent own five times more wealth than the poorest half.
    - **As capital in the form of machines becomes more important in the economy, 'Who owns the robots?' becomes a vital determinant of the distribution of prosperity.** If the share of national income flowing to the

owners of capital increases, then unequal levels of capital ownership will accelerate inequality.

- **To make sure that the dividends of automation are broadly shared, we need new models of ownership that hold wealth in common and democratise capital at scale.** These could include a Citizens' Wealth Fund that owns a broad portfolio of assets on behalf of the public and pays out a universal capital dividend. It could also include the creation of employee ownership trusts to give workers a stronger stake in the firms for which they work, and an ownership claim on the value they help create.
- **Automation should enable us to work better, but less**. The productivity gains of automation should promote a debate on how best to organise working time, including how to redistribute the productivity gains in the form of a shorter working week.

Introduction

Automation is the substitution of labour by capital, reducing or eliminating the need for people to perform specific tasks in the production process. As well as replacing the need for human labour, it can augment the capabilities of, and demand for, human effort and ingenuity. From the Industrial Revolution onwards, this process has transformed how we live and work, produce and consume. In aggregate, though often unevenly, automation has immensely benefited society.

The growing capabilities of artificial intelligence and robotics have led to claims we are on the cusp of a new machine age that will dwarf previous waves of automation in terms of the scale, speed and scope of the disruption it causes (Brynjolfsson and McAfee 2014; Ford 2015; CitiBank 2016; Susskind and Susskind 2015; World Economic Forum 2015; Avent 2016; Srnicek and Williams 2015). Whereas the technologies that drove automation in the past required clear instructions in controlled environments to substitute for human endeavour, new technologies are now increasingly able to act and problem-solve independently, inferring the appropriate solution or actions on the basis of the external inputs, and 'learning' as they do so. As a result, machines (whether hardware or software) are increasingly able to perform both routine *and* non-routine tasks, physical *and* cognitive work. Tasks once thought to be the sole preserve of humans can now often be performed better, and increasingly more cheaply, by machines.

This combination – of rising performance and the falling cost of technology – has the potential to drive radical productivity gains, both eliminate and reallocate employment in service sectors and professional occupations as well as in manufacturing, and reshape the nature of work. There is likely to be an unequal distribution of the gains and losses of this process across sectors, regions and socioeconomic groups. The integration of autonomous technologies will also challenge foundational assumptions about the economy: the role of employment as the primary means of distributing economic reward, labour's position as the central factor in production, notions of scarcity, and how we organise working time, among others.

Of course, predictions about the transformative effect of automation have been made repeatedly in the past, with both dystopian and utopian promises made. What is clear is that technology is not destiny. How automation reshapes the economy, and to whose benefit, will depend on the choices society makes, the policies we adopt, and the institutions we create. Automation managed for the common good could help create a society where material plenty is sustainably generated and widely shared, and economic power is evenly distributed. Yet without policy intervention, automation risks reproducing and amplifying existing inequalities within the economy. There is no liberation in technology itself; only in the types of institutions, policies and politics we build to shape its use and outcomes.

This discussion paper argues that public policy should seek to accelerate automation to reap the productivity benefits, while building new institutions to ensure the dividends of technological change are broadly shared.

To that end, it sets out five propositions. The first two are primarily analytical, relating to the likely macroeconomic effects of automation and its potential to accelerate inequalities of wealth and income. The final three set out how we believe public policy should make sure that automation works for the common good. These propositions discuss how best to manage the acceleration of automation, the public institutions needed to manage the ethical and regulatory challenges that autonomous technologies will create, and the new models of common ownership needed to ensure the fruits of automation are fairly shared. Together, we believe these ideas can ensure a new machine age helps us all lead fuller human lives.

1. **Work will be transformed by automation, not eliminated.** Automation is likely to lead to the steady redeployment of labour over a period of decades, rather than a sudden and rapid elimination of employment. The task contents of most jobs will evolve, changing the nature of work.
2. **In the absence of policy intervention, the most likely outcome of automation is an increase in inequalities of wealth, income and power.** Without policy intervention, the economic dividends of automation are likely to flow to the owners of technologies and businesses and the highly skilled, as income shifts from labour to capital and the labour market polarises between high- and low-skilled jobs.
3. **A managed acceleration of automation is needed to reap the full productivity benefits and enable higher wages.** Due to our low investment rates, poor management practices and long tail of low-wage, low-productivity firms, it is the relative absence of robots in the UK economy, not their imminent rise, that is the biggest challenge.
4. **An Authority for the Ethical Use of Robotics and Artificial Intelligence should be established to regulate the use of automating technologies.** The ethical and regulatory framework governing the use of technologies should precede, rather than follow in the wake of, technological development.
5. **New models of capital ownership are needed to ensure automation broadens prosperity rather than concentrates wealth**. Spreading capital ownership will help ensure automation creates an economy where prosperity is underpinned by justice.

1. Work will be transformed by automation, not eliminated

THE THEORY OF TECHNOLOGICAL JOBLESSNESS, OR NOT

From the invention of the spinning jenny to today's multifunctional manufacturing robots and algorithmic management, machines have substituted for labour in the production process. At the same time, many technologies, such as computers, have complemented humans, augmenting their productive potential rather than replacing them. The rate at which machines can replace workers, or the 'elasticity of substitution', is determined by both the relative price and capabilities of each. Should machines be able to perform all of the tasks and roles of workers cheaper and better, labour would be forced to specialise in a shrinking set of tasks, progressively squeezed out of the economy by advanced capital with wages driven lower and lower, and 'in the limit, labour is fully immiserated' and 'technological unemployment' follows(Susskind 2017). But this is quite an unlikely scenario. After all, most jobs involve a whole range of tasks, some defined, some tacit, many of which machines still struggle to do.

Historically, many occupations have only experienced partial automation (of certain tasks), while the number of jobs in that occupation has actually increased. This is because, if a particular task is automated so that it can be performed more quickly or cheaply, the demand for workers to do the other non-automated tasks around it may increase. This experience has been repeated in many industries affected by automation (Autor 2015a). In roles where there is partial automation, jobs are typically redesigned with a focus on the tasks that are harder to automate and where humans retain comparative advantages over machines. But, even if jobs are re-defined and workers are able to adapt to high automation rates, it is still possible for overall working hours in that occupation to be reduced.

If the full automation of human work is unrealistic, the increasing use of machines can still lead to an overall reduction in jobs.

Whether a machine performs all, or some, of the tasks previously performed by workers, it will likely increase labour productivity (measured as output over total hours worked). With the aid of machines, workers can then produce the same amount of outputs as before but in less time. Whether higher productivity leads to fewer or more labour hours then depends on the level of demand for the product. If demand keeps pace with increased output, employment levels could stay the same, or even grow. For example, David Autor contrasts how technological change led to a dramatic reduction in employment in agriculture in the 20th century, compared to healthcare, where employment has risen, due to differing levels of demand for the outputs (ibid).

Whether jobs are lost or not depends on the elasticity of demand (the extent to which customers increase their demand following a price change)[1] for the output. For a particular good and the industry that produces it, the impact of automation on employment will depend on the relative strength of the elasticity

1 In a competitive environment, firms that benefit from increased productivity will be encouraged to reduce their prices.

of substitution between workers and machines compared to the elasticity of demand for the industry's output. Where robots and software can readily do what workers used to do, and consumers and business have limited appetite for more output, industry employment will fall.

But even in the case where employment in a specific industry or occupation falls, it does not necessarily follow that economy-wide employment will also decline. Even if demand for one industry's products declines, aggregate demand should rise as productivity improvements lead to rising wages, which are then spent in the wider economy. In addition, all being equal, consumers will be richer if the relative price of the automated-industry goods fall. So consumption of goods and services in other sectors may also increase, leading to higher employment in their production. Most notably, employment has risen hugely in the industries supplying information and digital services and automating technologies, which have expanded significantly over recent decades as technical capabilities have grown.

To understand the impact of localised automation in a particular industry on the entire economy, it is therefore necessary to look at the elasticity of substitution between the outputs of the automating industry and the outputs of other industries. If consumers redirect their spending from a machine-intensive to a labour-intensive industry, overall employment may actually go up, even as the number of workers in the machine-intensive industry falls. As the gains from automation are recirculated, this will have knock-on effects on industries entirely untouched by the new technologies. This is the point the economist William Baumol first made in 1967: technologically leading sectors may paradoxically see their employment fall while lagging sectors will see their employment rise because of rapid technological change in the leading sector (Baumol 1967).

Another possible response to both automation and demand 'saturation' is simply for people to work less but more productively. In the simplest, although still useful, model of labour supply, individuals face a trade-off between leisure and earning a wage, which allows them to consume. If automation raises worker productivity, then it should also raise their wages (in theory at least). With higher wages and/or lower prices, as well as choice over the number of hours worked, individuals can choose to reduce their hours to increase their leisure while maintaining their consumption.[2] If they choose to reduce their hours, the supply of labour would fall; partially, fully, or even over-compensating for the fall in the demand for workers in the automating industry. Collectively, society could choose to use the productivity gains of automation to reduce overall working hours.

In other words, there are multiple ways in which accelerated automation may not result in aggregate technological unemployment. Clearly, it is still theoretically *possible* for automation to lead to aggregate job losses, but many factors must align and push in the same direction. If automation were to spread throughout the entire economy, aggregate demand to saturate, and those in work to choose *not* to reduce their hours, then mass joblessness would be the most likely outcome. So far, this is not what has happened.

THE STORY SO FAR

Over the course of the 20th century, despite huge technological change, the number of people in work (the employment to population ratio) actually increased (Autor 2015a). Over the last 40 years (and probably before, though comparable data is not available), the primary predictor of aggregate

2 They could also choose to work more to further boost their consumption.

employment growth in developed countries has been population growth and female workforce participation, not technological change. In the US, for example, productivity and employment have both risen in more than two-thirds of the years since 1929 (Autor and Salomons 2017). In other words, rapid technological change, including automation, has been accompanied by a *positive* effect on employment growth (McKinsey 2011).

Though employment levels have not fallen as advanced economies have automated large swathes of production, there has been a relatively steady decline in hours per worker since the late 19th century, with average hours falling among OECD countries since the 1950s (see figure 1.1).[3] Though causality is hard to establish precisely, the fall in hours worked is both negatively correlated with productivity growth (people choose to spend productivity gains on working less) and positively correlated with consumption growth (people choose to spend some of the gain on more consumption). For example, average hours worked per adult are significantly higher in low-income countries than in high-income countries with higher productivity rates (Bick et al 2015). Boppart and Krusell interpret this as evidence that people's labour supply is somewhat elastic. As productivity and wages increase, labour supply decreases, but only marginally (2016). On average, when people are made richer, they choose both more consumption and leisure, but relatively more of the latter. Over the long term, then, automation has not had a significant effect on the total number of jobs but has likely contributed to reducing hours worked per worker.

FIGURE 1.1

Average hours worked have fallen over the last 65 years
Hours worked in selected OECD countries

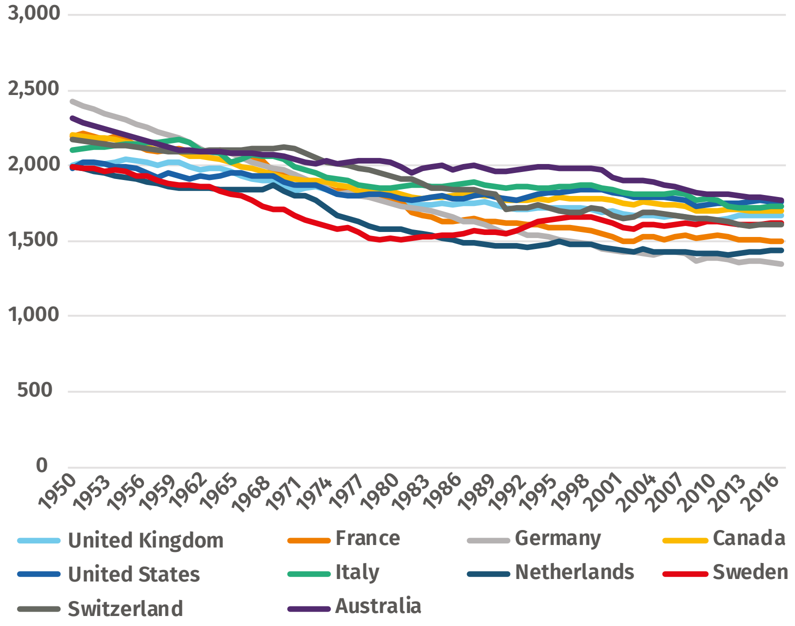

Source: IPPR analysis, 'Output, Labor, and Labor Productivity, 1950-2017' series, *Total Economy Database* **2017**

3 In most countries, total hours have also decreased steadily – though not in the US and Canada, where average annual hours have stayed roughly constant. Once you account for increased participation rates, however, the US and Canada also witnessed decreasing hours per worker (see Boppart and Krusell 2016).

In the last few centuries, agricultural employment has fallen dramatically across advanced countries. For example, farm employment fell from 40 per cent in 1900 to 2 per cent in 2000 in the United States, and by similar levels in the UK over a longer period (McKinsey 2017). At the same time, the share of service sector employment has increased steadily for most of the 20th century, and now forms the majority of employment in all advanced economies. If these sectoral shifts have been well documented, more recent work has improved our understanding of how these were affected by technological change. In particular, Autor and Salomons (2017) have investigated the impact of rapid productivity gains across five sub-sectors of the economy: mining, utilities and construction, manufacturing, education and health, low-tech services and high-tech[4] services. Using multiple estimation techniques to measure employment and productivity, they find robust and consistent negative effects of productivity growth on within-sector employment.

This is largely because the demand response to productivity gains (and lower prices) is not sufficient to fully counter-balance the increase in labour productivity, leading to an overall reduction in the number of workers needed for production. In their analysis, Autor and Salomons find that a 1 per cent increase in productivity only increases consumption by 0.5 per cent across these sectors. Using different data, covering 200 years and a more tightly defined set of industries (allowing for better identification),[5] Bessen (2017) finds that the demand response to productivity-increasing automation tends to follow an inverted U-shape; one that mimics and explains the inverted U-shape of employment in those industries. Initially, productivity growth increases employment, as demand is elastic (above unity), but this then leads to decreasing employment as it becomes inelastic. In other words, automation will boost demand for industry outputs where large untapped demand exists, but reduce employment when and if that demand eventually is saturated. At least in goods-producing industries like steel, textile and automotive, demand has long since been saturated in advanced economies.

Autor and Salomons also find that the spillover effects of productivity improvements are consistently positive across all sectors. Spillovers are the impacts that events in certain sectors or economies have on other sectors or economies, often apparently unrelated to them. These impacts have been sufficiently large to fully offset the negative within-sector effects, resulting in an aggregate positive, albeit small, effect of productivity on employment. Although the direction of both direct and spillover effects is consistent across sectors, their magnitude varies widely (see figure 1.2) – something that can help inform predictions of future automation scenarios.

4 Measured by capital intensity.

5 By focusing on a particular subset of industries (cotton, steel and automotive), Bessen is better able to identify the relationships between productivity, wages, prices, income, trade and demand than Autor and Salomons.

FIGURE 1.2

Low-tech services has the largest positive spillovers

Sizes of direct and spillover employment effects differ by sector

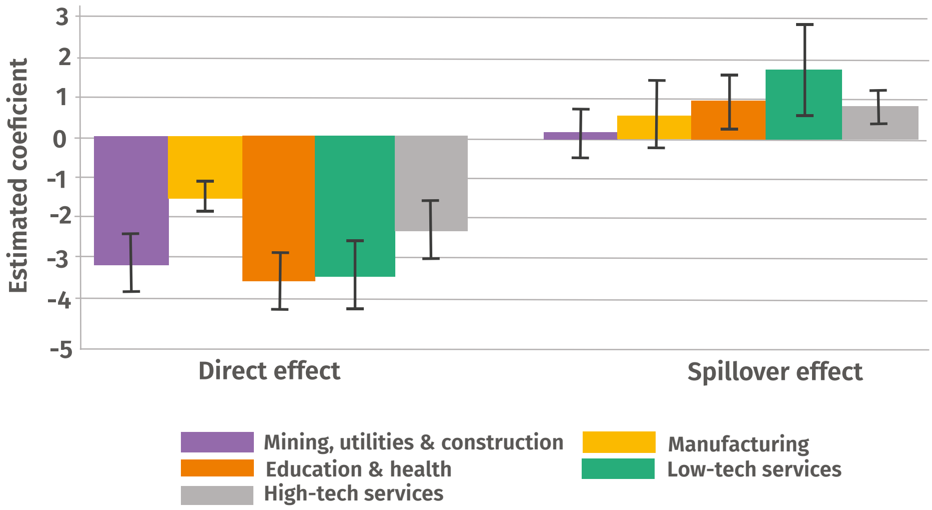

Source: Autor and Salomons 2017
Note: Model is estimated by OLS; includes country, industry and year FE; and controls for population growth. Productivity is gross output per worker.

If the analysis of direct and indirect effects helps explain why aggregate employment hasn't fallen, it is not sufficient to understand its impact on the structure of the economy. These sectoral shifts are, in large part, explained by the differential rates of productivity growth across sectors. The other significant factor relates to how spillovers have themselves been reallocated, with the evidence suggesting that individuals have spent most of their additional income on the purchase of services.

The manufacturing sector has exhibited the fastest productivity growth over the last 40 years, and its share of the labour force has decreased proportionally (see figure 1.3). At the other end of the spectrum, productivity growth has been minimal in the health and education sectors, even as their shares have increased. Meanwhile, low-tech services and mining, utilities and construction have experienced middling growth and their share has stayed roughly constant. These four sectors appear to confirm the Baumol hypothesis: that laggard sectors see the biggest gains in employment, as leading ones shrink (1967). However, high-tech services complicate the picture somewhat, as the sector is second only to manufacturing in terms of productivity while also increasing its share of labour the most.

While the direct effect on employment of automation in high-tech services is negative, the spillovers from other sectors have contributed toward increasing demand for these services. In fact, the demand for services appears to have increased across the board. The evidence suggests that consumers have used their higher incomes to purchase relatively more services – the demand for services is income elastic (Mazzorali and Ragusa 2013). Specifically, low-skill service occupations (typically also low-tech) such as food service workers, cleaners, or recreational workers, have grown rapidly in recent decades, especially in places where the automation of low-skill manufacturing jobs has been most pronounced (Autor and Price 2013). This can be observed in the UK labour market over the last two decades, where, at a time of rapid technological change, there has been a significant expansion in the share of employee hours in occupations with low skills, particularly in caring and personal services (though this is also partly due to an ageing population with a higher disease burden).

FIGURE 1.3

Manufacturing has experienced the strongest cumulative productivity growth and the sharpest decline in employment

Cumulative productivity growth and changes in employment by sector across 19 developed economies,6 1970–2007

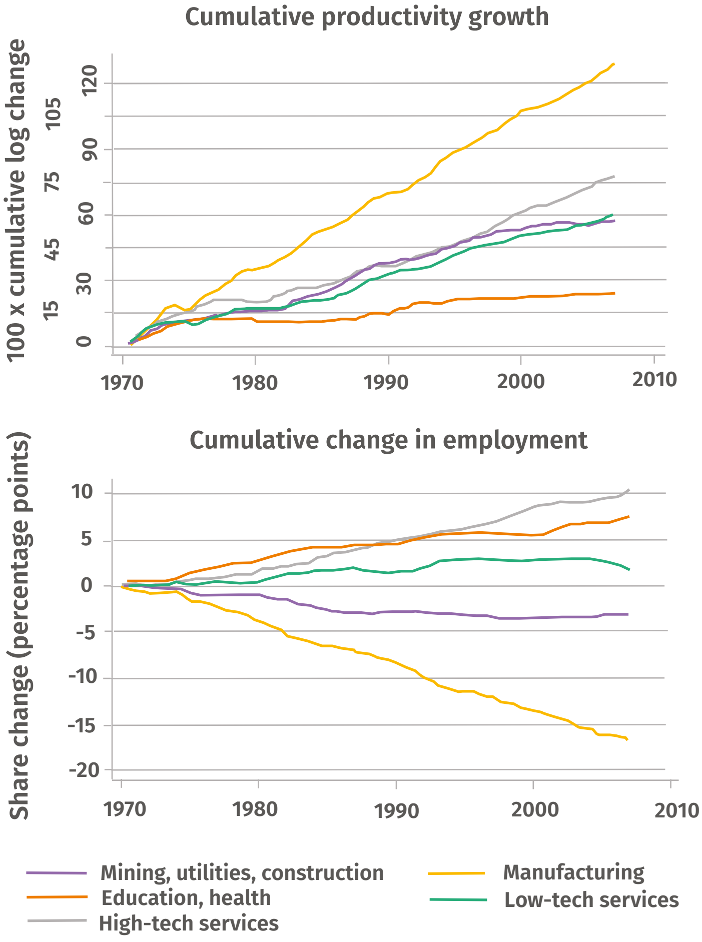

Source: Autor and Salomons 2017

6 The countries are Australia, Austria, Belgium, Denmark, Spain, Finland, France, Germany, Greece, Ireland, Italy, Japan, Korea, Luxemburg, Netherlands, Portugal, Sweden, UK, and the USA.

FIGURE 1.4

Hours worked in high- and low-skilled occupations have expanded as middle-skilled jobs have declined

Changes in share of employee hours by occupation in the UK: 1993–2014

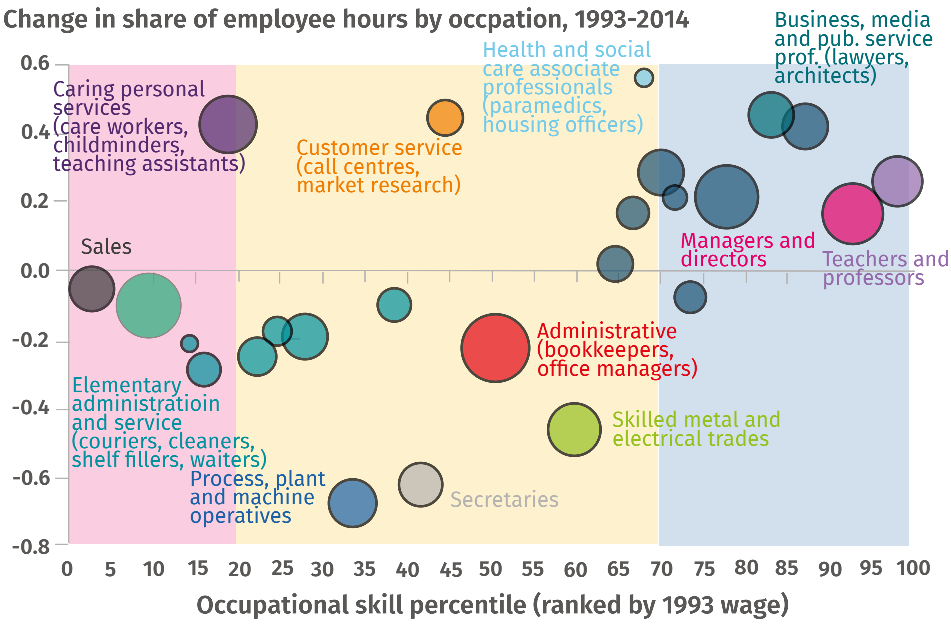

Source: Corlett and Gardiner 2015

Notes: The final quarter of 2014 is not included because data was not available at the time of the analysis. Bubble size reflects the average labour share between 1993 and 2014. See annex for other methodological details. Source: Resolution Foundation analysis of Labour Force Survey, ONS

FIGURE 1.5

Job polarisation has occurred across the OECD in recent decades

Percentage point change in share of total employment, 1995–2015

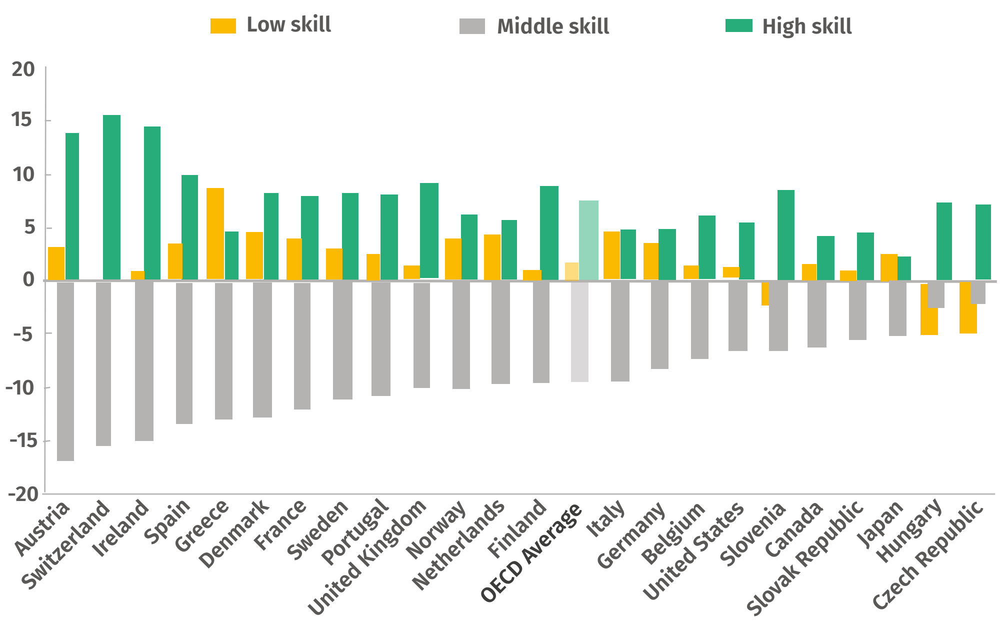

Source: OECD 2017

The polarising effect of technological change on the labour market in the UK more broadly can be seen in the fact that the share of employment in high-skilled occupations has increased in the last two decades, alongside the share of low-skilled jobs, while the share of middle-skilled employment has declined (Michaels et al 2014). This trend is widely observed: Spitz-Oener (2006) and Goos, Manning, and Salomons (2010) use wage and employment changes in occupations based on task content to show that routine occupations in the middle of the wage distribution in the UK are in decline. More broadly, employment polarisation can be observed in other advanced economies (see figures 1.4 and 1.5).

Summing up, the historical evidence suggests that in established industries, the demand response to increased productivity (and lower prices) is typically not sufficient to counteract the negative impact of these gains on the demand for labour. This will usually lead to decreasing employment in these mature industries. But the ultimate impact of the increased automation of a sector will also depend on the relative speed and distribution of productivity gains across sectors, with sectors exhibiting the largest production gains typically experiencing decreasing employment, consistent with the Baumol hypothesis. This unbalanced growth pattern can be counteracted (as it has in the case of high-tech services) or reinforced (in the case of manufacturing), depending on how consumers choose to allocate their relatively higher incomes and how public policymakers respond. In recent decades at least, these spillovers appear to have disproportionately flown toward service-industries.

HOW IT MIGHT PLAY OUT THIS TIME

To attempt to predict the future path of automation, it is necessary to know which technologies are likely to drive the next wave, how they differ from previous ones, and so whether this time really might be different.

It is rapid advances in artificial intelligence (AI) that will radically reshape the future of work. AI describes a broad range of computing techniques that allow machines to infer the appropriate solution, or actions, on the basis of the external inputs it receives. Traditionally, machines had to be given a fully specified set of commands, describing exactly what actions needed to be performed at every step. With AI, computers are increasingly able to problem-solve, and 'learn', independently.

Not only can AI outperform humans in certain analytical tasks, it also has the potential to finally break Polanyi's Paradox, the fact that "we know more than we can tell" (Polanyi 1966). There are many tasks and actions that humans do intuitively, but for which we only understand the underlying skills tacitly. Tasks such as food preparation or cleaning often involve many of these skills and are, for this reason, difficult to automate. Aided by AI, modern robots may finally be able to perform some of these jobs. McKinsey's analysis of current AI and robotics technology suggests that these have the potential to perform a much higher proportion of these tasks than previous technologies (2017).

In both its ability to perform analytical decision-making tasks and overcome Polanyi's Paradox, AI technology differs critically from the previous IT-driven wave of automation (Autor 2015b). For the most part, IT has so far successfully automated routine, codifiable tasks. Traditional assembly-line manufacturing jobs, as well as clerical and simple administrative work, typically require routine tasks. It is these jobs – mainly performed by workers in the middle of the education, skills and wage distribution – that decreased the most in the 1990s and 2000s, increasing job polarisation (Acemoglu and Autor 2011). By contrast, the types of analytical tasks typically performed by high-skilled, high-wage workers were, for the most part, complemented rather than replaced by these advances in computation. Similarly, manual non-routine work was largely unaffected (ibid). In so far as AI could further substitute for human labour in both areas in the future, its impact could be quite different.

Should AI enable a wider range of tasks to be automated in the coming decades, its aggregate impact would, logically, be larger. Rather than being narrowly focused in certain industries like manufacturing, the productivity benefits have the potential to be reaped across a broader range of industries, professions and skills. These gains are potentially transformative. For example, McKinsey estimate that the widespread adoption of AI, machine learning and robotics could raise productivity growth by between 0.8 and 1.4 per cent between 2015 and 2065. By comparison, the introduction of the steam engine increased productivity by 0.3 per cent annually between 1850 and 1910, early robotics boosted productivity by 0.4 per cent per annum between 1993 and 2007, and ICT technologies created a 0.6 per cent annual acceleration between 1995 and 2005. In their central scenario, McKinsey estimate that automation could generate productivity growth in the largest economies in the world (G19 plus Nigeria) equivalent to an additional 1.1 billion–2.3 billion full-time workers by 2065 (McKinsey 2017).

Though it may affect a wider range of jobs, there is no compelling reason to think that the direct effect of AI-induced productivity gains in specific industries will, this time, be different in terms of employment effects. The direct effects should be at least as negative as they have been historically, since consumers in advanced economies are, overall, much richer today than they have ever been, and therefore demand for a particular industry's goods is more likely to be inelastic. This does not however mean that aggregate employment will fall; that depends on the spillover effects.

There are essentially two possible scenarios: one in which the spillovers are not large enough to counteract the negative direct effects, and the other in which they are, in line with historical precedent. There is some piecemeal evidence that spillovers today may not be as large as they were historically. When Autor and Salomons (2017) decomposed their analysis on a decade by decade basis, they found that, in the 2000s, the negative direct effect of productivity was higher and the positive spillover effect lower than in previous decades. As Bessen argues, it seems plausible that the demand for certain goods has saturated, as they have become more abundant (thanks to continuous productivity gains). This certainly appears to have been the case for both agricultural outputs and a significant proportion of manufacturing goods. Health, education and manual tasks have, so far, been stubbornly resistant to automation, but should that change and as quality/price improve, it is possible that demand for these services would also peak.

And yet, possible does not mean plausible. We may not have imagined 50 years ago that we would want or need smartphones; we are equally unaware of future products that could stimulate future demand. Unlike agricultural or industrial goods, services have human interactions at their core. And while there are limits to how much food individuals can consume, material limits to production given planetary boundaries, and diminishing marginal benefits from owning an additional phone, TV, or car, humans are naturally social creatures, who thrive on interactions with others. We are unlikely to satiate our demand for the type of capabilities like emotional intelligence and creativity at which humans excel.

Finally, most advanced economies are currently going through a demographic transition; one in which the ratio of retired to working population is set to increase (Lawrence 2016). More retired people means more demand for a whole variety of services, especially (but not only) in health and care. In other words, there is likely to be tremendous potential for the productivity dividends of technological change to be redirected to the consumption of social goods and infrastructure, and expanding employment in the provision of these services.

Of course, this may not happen. Habits change, norms evolve, and preferences shift. It may be that in the future we will not value human interaction in services where we currently do. It is possible that technological change will also negatively impact on employment in services. But it is unlikely.

AUTOMATION AND THE CHANGING NATURE OF WORK

If automation is unlikely to wholly displace human labour, it is likely to change the shape of the labour market, the occupations that individuals work in, and the type of work tasks humans perform. It is also likely to reshape the type of tasks we perform within those occupations. As some but not all of those tasks are automated, the nature of work will change. For example, a recent analysis of more than 2000 work activities across more than 800 occupations estimated that fewer than 5 per cent of all occupations could be automated entirely with existing technologies, but that 60 per cent of occupations have at least 30 per cent of constituent activities that could be automated today. The nature of work remaining would evolve significantly (McKinsey 2017) (see figure 1.6). A recent OECD study reached a similar conclusion: once job-level tasks are accounted for, the study estimated that only 9 per cent of jobs in the UK are susceptible to automation in the next decade, but that 35 per cent of jobs would change radically in the next two decades (Arntz et al 2016).

FIGURE 1.6

For 60 per cent of all occupations, at least 30 per cent of activities can be technically automated
Proportion of activities within an existing occupation that can be technically automated with existing technology, US

Source: McKinsey 2017

Partial automation of tasks, rather than occupations, has important implications for the type of work people are likely to perform. At present, nearly half of all time spent at work in advanced economies is on tasks involving the collection and processing of data, or predictable physical activities. These are the types of tasks most likely to be automated, leaving people to focus on those activities – managing and developing people, applying expertise, interacting and communicating, creative and caring work and unpredictable physical activity – which are most technically resistant to automation.

The evolving bundle of tasks performed by human labour has the potential to improve the experience of work for many. As David Autor has argued, the interplay between machine and human comparative advantage allows machines to substitute for labour in performing routine and programmable tasks, while amplifying the comparative of advantage of humans in supplying problem-solving skills, adaptability, emotional and caring labour, and creativity. For many, this will lead to improved quality of work as complex, difficult to automate tasks increase as a share of work, with labour specialising and increasing the value of the performance of tasks that people continue to perform (Autor 2017).

Currently, in the US it is estimated that only 2 per cent of work time is spent on creativity and generating novel patterns and categories, 9 per cent on social and emotional reasoning, 10 per cent on emotional and social output and coordinating between collective agents, and 13 per cent on problem-solving (McKinsey 2017). Spending more time performing and specialising in tasks in which we excel and retain comparative advantages over machines should be a positive outcome. At the same time, more repetitive tasks are more likely to be automated. Paradoxically, automation could make work more 'human'.

But if automation could increase the number of 'lovely' jobs, it also has the potential to create 'lousy' ones (Goos and Manning 2007). First, a low price of labour, or poor minimum employment conditions in tandem with an excess of labour, could lead firms to expand low-paid, poor quality work even if technically those jobs could be automated, because relatively labour is cheaper than capital. Second, humans working in tandem with machines could lead to greater monitoring and workplace stress. Sensors on Amazon stock pickers, for example, measure their productivity and location to a level of detail a human manager could not. New technologies could facilitate intensive "digital Taylorism" (O'Connor 2016) if we allow them. And third, our current labour market does not reward all of the jobs that involve human interaction and are therefore relatively resilient, such as social care work, with good wages or working conditions. Avoiding a scenario where automation concentrates employment in sectors where jobs may be 'human' but are also currently of poor quality and poorly paid, will require concerted effort. We may need these types of work to be revalued, and broader steps taken to ensure technological change improves work, rather than simply intensifying exploitation.

Automation and the UK economy

What proportion of jobs in the UK could technically be fully automated?

The impact of automation on the labour market is disputed and estimates of the likely impact vary significantly. Ground-breaking work by Frey and Osborne looked at the susceptibility of jobs to automation. They found that 47 per cent of jobs in the US had a high technical potential to be automated, and that these jobs "could be automated relatively soon, perhaps over the next decade or two" (Frey and Osborne 2013). The same authors have estimated that, in the UK, 35 per cent of jobs have a high potential to be automated over that period (Frey and Osborne 2014). Our analysis, using Frey and Osborne's data but with a different methodology, estimates that 44 per cent of jobs in the UK economy could feasibly be automated (see appendix). The Bank of England follow a similar method to estimate that 15 million jobs are at risk (Haldane 2015). By contrast, Arntz et al, using a different methodology, found that only 10 per cent of jobs in the UK are at high risk of automation (2016).

It is important to note that these analyses focus only on the potential for automation, and do not include other factors such as wage levels, investment and regulation that determine whether tasks or occupations are automated. Job losses are likely to be substantially lower than these headline figures; but there could be substantial adjustment costs for some to bear, and inequality could increase.

What proportion of wages in the UK could be automated?

Using Frey and Osborne's methodology, we estimate that £290 billion of wages annually are associated with jobs which have the technical potential for automation in the UK economy.[7] This represents 33 per cent of all wages and earnings. The actual amount of wages that might be automated is lower, as this estimate does not account for positive wage impacts and new employment (see chapter 2).

FIGURE 1.7

Elementary, administrative, and sales occupations are among the jobs with the highest technical potential to be automated

Automation potential among UK occupations

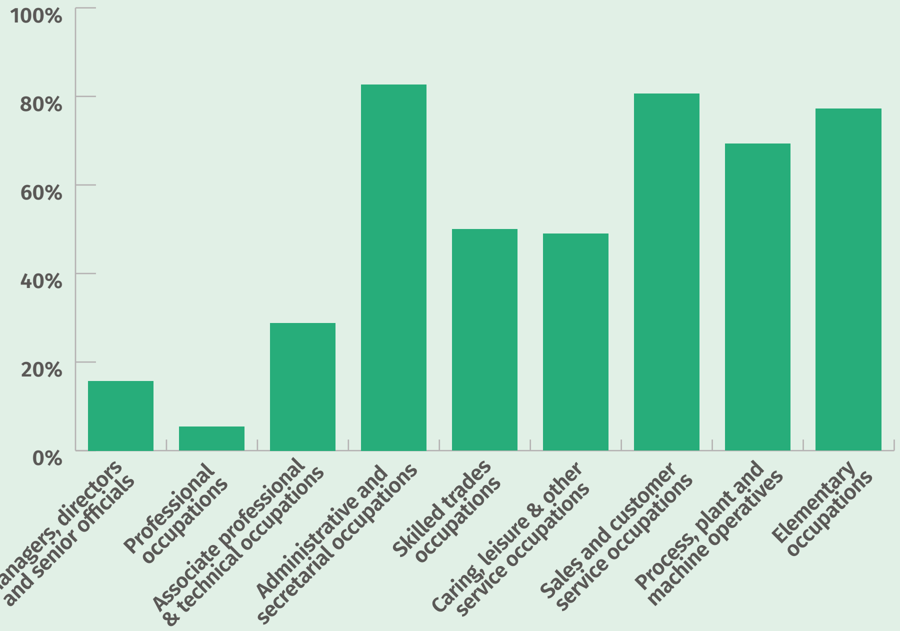

Source: IPPR analysis using Quarterly Labour Force Survey (ONS 2017d, 2015-16 data) and Frey and Osborne's probabilities of computerisation (2013). See appendix for methodology.

A broad range of occupations are likely to be affected by automation. High-skill and high-wage occupations are the least likely to be automated (Frey and Osborne 2013). However, though wage and skill levels are negatively correlated with technical automation potential (on average, occupations with lower wages and skill requirements have higher automation potential) there remains significant variation underlying the averages for broad occupation categories (see figure 1.8). For instance, while some skilled trades occupations such as roofers, roof tilers and slaters have high potential for automation, others such as electricians

7 IPPR analysis using ASHE 2016, Frey and Osborne and Quarterly Labour Force Survey 2015/16. See appendix for methodology.

and electrical fitters have low potential. And there is also variation within broad occupations regarding how much automation may be countered by increased demand for a services or goods; adult social care, for instance, is likely to grow despite having relatively high potential for automation.

FIGURE 1.8

Some sectors are almost three times as susceptible to automation as others

Proportion of jobs at with the highest technical potential for automation by industry in the UK (probability >0.7)

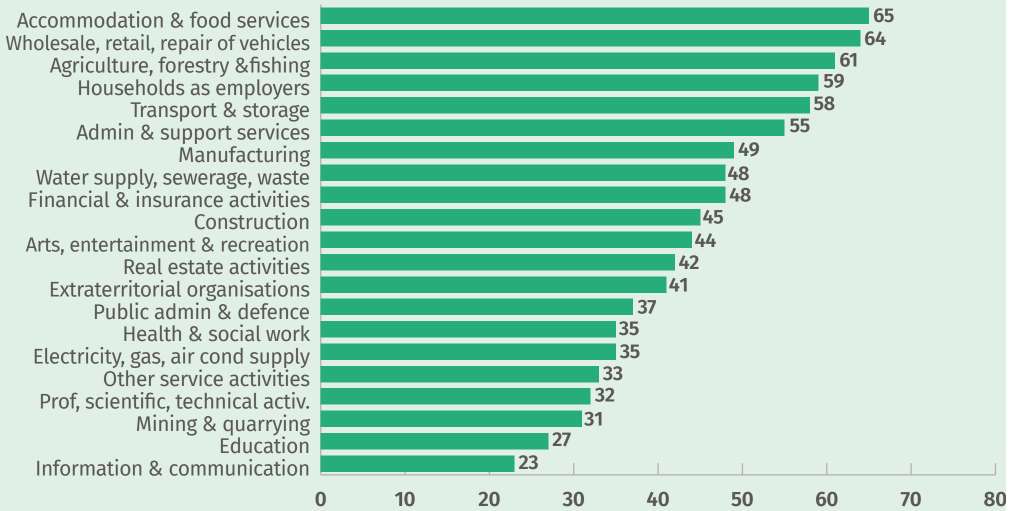

Source: IPPR analysis using Quarterly Labour Force Survey (ONS 2017d, 2015-2016 data) and Frey and Osborne's probabilities of computerisation (2013). See appendix for methodology.

FIGURE 1.9

Jobs in London and the South East are more resilient to automation

Proportion of jobs at with the highest and medium to lowest technical potential for automation by region (probability >0.7)

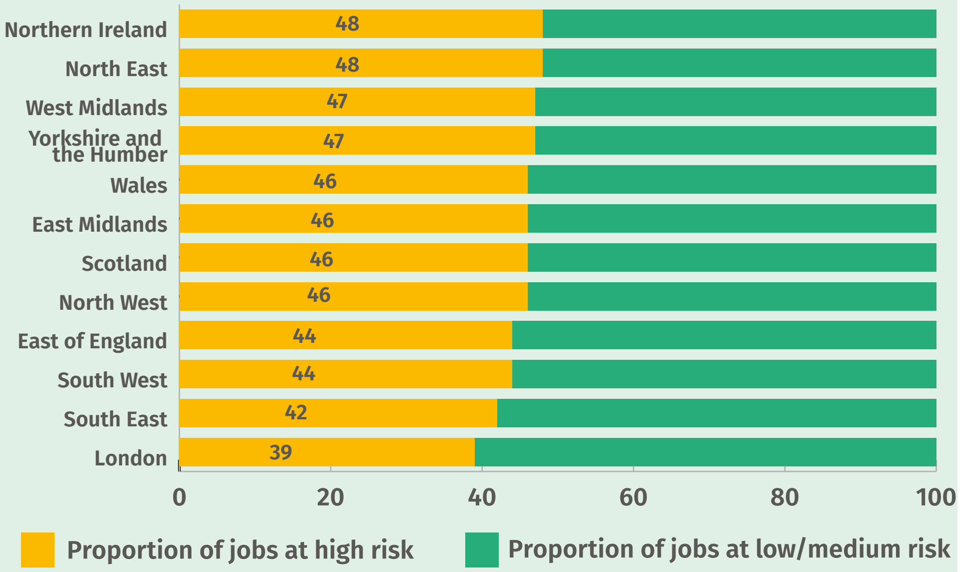

Source: IPPR analysis using Quarterly Labour Force Survey (ONS 2017d, 2015-2016 data) and Frey and Osborne's probabilities of computerisation (2013). See appendix for methodology.

Which regions will be most affected by automation?

There is evidence to suggest that the impact of automation will be geographically concentrated and may accentuate existing regional inequalities. Our analysis shows that on average jobs in London and the South East are more resilient to complete automation than those in the rest of the country (see figure 1.9). Whereas 39 per cent of jobs have a high potential to be automated in London, 47 per cent are at risk in Yorkshire and the Humber and the West Midlands, and 48 per cent are at high risk in the North East and Northern Ireland. These areas already have lower levels of employment than the national average, so higher levels of job losses as a result of automation may serve to undermine demand, increase unemployment, and exacerbate existing regional inequalities.

Within sectors, the impact is also likely to vary by area, as different sectors and industries have different occupational compositions. For example, in financial and insurance activities, 61 per cent of jobs in Wales have potential for automation compared to 36 per cent in London (IPPR analysis – see appendix).

Who will be most affected by automation?

Workers with lower levels of skills are more at risk of complete automation of their jobs according to the Frey and Osborne figures. PwC analysis of these suggests that educational levels is a key differentiating factor. For those with just GCSE-level education or lower, the estimated potential risk of automation is as high as 46 per cent in the UK, but this falls to only around 12 per cent for those with undergraduate degrees or higher (PwC 2017a).

As skill level is highly correlated with wage, automation is also likely to affect the low paid most. Combined with increasing wage inequality (see chapter 2), automation has the potential to substantially increase income inequality in the UK.

FIGURE 1.10

Lower paid occupations are more likely to have high automation potential
Index of probability of computerisation and employment proportion by mean annual earnings of occupation

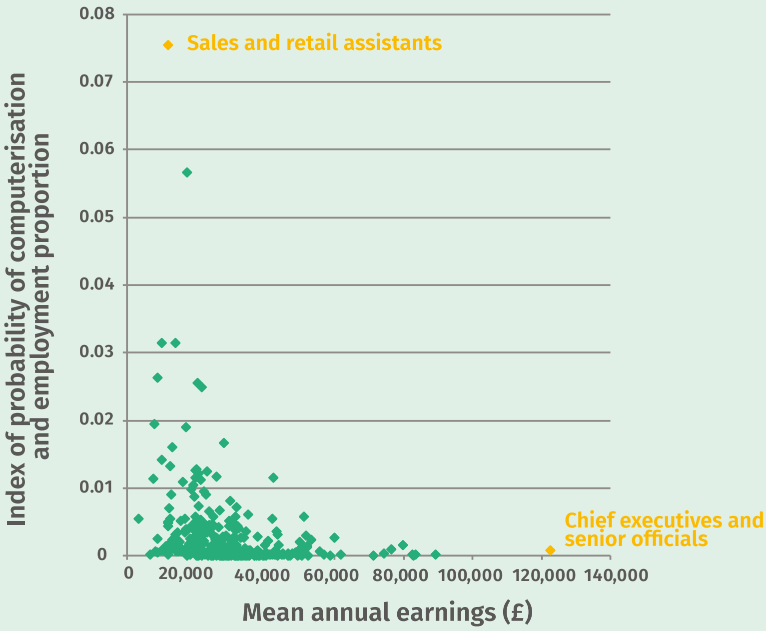

Index of probability of computerisation and employment proportion by mean annual earnings of occupation (excluding outliers)

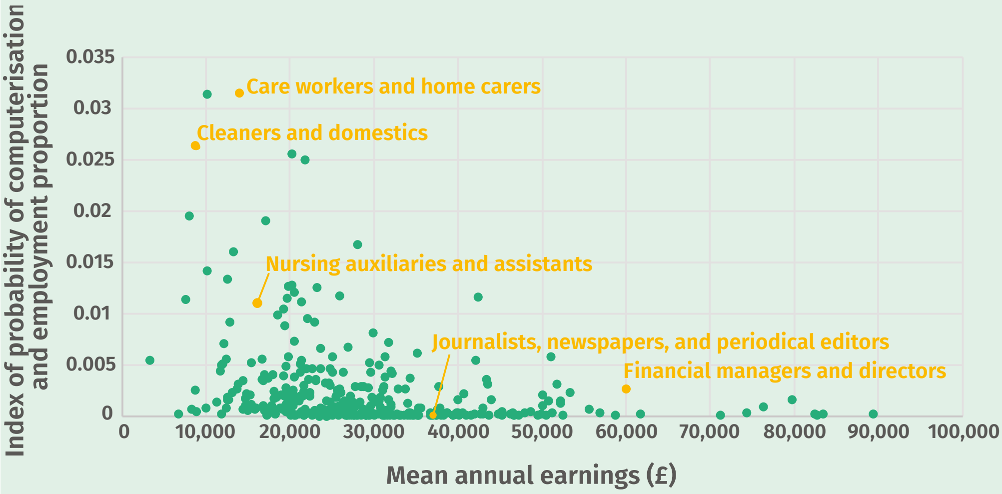

Source: IPPR analysis using ASHE 2017, Quarterly Labour Force Survey (ONS 2017d, 2016-2017 data) and Frey and Osborne (2013). Each data point is a UK minor occupational group. Index is equal to probability of computerisation multiplied by employment proportion. See appendix for methodology.

Both low and high-wage occupations have significant technical automation potential but lower wage jobs in general have higher potential (see figure 1.10).

Again this pattern is amplified by geography. For example, of the 71,000 jobs in accommodation and food services in the East Midlands, 66 per cent have a high potential for automation, and 83 per cent of the workforce do not have a degree-level qualification.

These industries could see a high level of job displacement in the coming years, with workers affected facing an increased chance of falling into long-term unemployment. This is particularly problematic as low-skilled workers are less able to find alternative employment compared to high-skill workers. In retail and wholesale for example, there are over 2.5 million jobs at high risk of automation, and three in four workers who are likely to struggle to adapt to change (see table 1.1).

TABLE 1.1

Wholesale, retail, transport and manufacturing have a large proportion of jobs that are technically susceptible to automation
Industries with high numbers of jobs with the highest technical potential for automation and low levels of qualification among the workforce

Industry	Proportion of jobs with high potential for automation (p>0.7)	Number of jobs with high potential for automation	Proportion of workers without NVQ level 4 qualification (%)
Wholesale, retail, repair of vehicles	63.7	2,638,000	76.3
Transport and storage	57.7	912,000	78.8
Accommodation and food services	64.7	1,093,000	78.0
Manufacturing	48.5	1,453,000	67.6

Source: IPPR analysis using Quarterly Labour Force Survey (ONS 2017d, 2015-2016 data) and Frey and Osborne (2013). Each data point is a minor occupational group. Proportions rounded to nearest percentage point, and job counts rounded to nearest 1000. See appendix for methodology.

Will men lose out most from automation?

PwC analysis has estimated that more men than women are at risk from automation. However, our analysis using a different methodology suggests that a similar number of jobs performed by men and women, making up a greater proportion of jobs that women hold in the UK (46.8 per cent compared to 40.9 per cent of men's jobs), have the technical potential to be automated. Furthermore, women make up a smaller proportion of high-skill occupations that are more likely to be complemented by technology. Automation could therefore increase gender inequality unless women as well as men are able to access new jobs that pay well, or very well. The higher proportion of some ethnic groups (such as Black, Pakistani and Bangladeshi) working in low-skill jobs in the UK, also mean automation risks exacerbating inequality between ethnic groups (Cabinet Office 2017).

The overall pace of automation

The process of adopting new technologies is likely to be slower than the headline figures of of jobs at imminent risk might suggest. This focuses on technical feasibility and doesn't account for economic, social, cultural and political factors that also factor in the decision to automate. Automation is therefore more likely to be a continuous process of change in how we work and the reallocation of employment, rather than an abrupt and imminent elimination of work. Given this, McKinsey modelling that suggests the mid-2050s is the most likely point when 50 per cent of current work activities in the global economy will be automated appears to be a reasonable central scenario for the UK (2017).

However, McKinsey's work suggests significant uncertainty around the timing of automation. The point at which half of existing tasks are automated could occur 20 years earlier or later depending on factors beyond technical feasibility, such as the relative cost of labour and capital. Given the factors shaping the pace of adoption are driven by public policy choices, it suggests the speed and effect of automation will be human-shaped, not technologically determined.

2.
In the absence of policy intervention, the most likely outcome of automation is an increase in inequalities of wealth, income and power

Automation is more likely to accelerate inequalities of wealth and income than create a future of mass joblessness. Without policy intervention, the productivity dividends of automation could create a 'paradox of plenty', in which we produce more, yet it is less equally shared, as the benefits of technological change flow to the owners of capital and the labour market polarises.

A FALLING LABOUR SHARE OF INCOME

One concern is that automation could trigger a falling share of income for labour relative to capital. This could occur in scenarios in which the total number of hours worked declined without a commensurate rise in wages, if wages did not rise as quickly as returns to capital, or in the unlikely scenario in which mass joblessness occurred.

Total hours worked across the economy are most likely to fall without a simultaneous rise in wages in the scenario in which the combination of productivity spillovers and direct effects is negative, and where this impact is large enough to reduce the demand for workers below even the fall in the labour supply induced by the demographic transition. In this case, inequality is likely to rise sharply. In a competitive market, an excess of labour will bring down the equilibrium wage. The lower the wages, the easier it becomes for employers to find tasks where the marginal benefit of employing someone is higher than the marginal cost. Even if technology could do certain activities, employers may still prefer workers simply because they are cheaper. In other words, even in the case where automation destroys more jobs than it creates, it is hard to imagine employers would not find profitable occupations for the newly unemployed, or that workers – depending on the minimum wage policy – would not price themselves into work. This is because comparative advantage, rather than absolute advantage, determines economic activity. Therefore, even if robotics and machines rapidly exceed human capabilities in all tasks and jobs, and these technologies are quickly adopted in the economy, people are still likely to find work in activities where the relative advantage of machines is smallest. The distributional consequences of this scenario would clearly be catastrophic, as employment would pay poorly, possibly at a rate in which living standards would decline for many. Further exacerbating inequality, this scenario would make the capital share of output increase, as the labour share fell.

A falling labour share due to the effects of automation would extend and deepen a long-running trend. Labour's share of national income has been in secular decline in advanced economies, with the average share almost 4 percentage points lower today than in the 1970s (Dao et al 2017). These trends are problematic because falling labour share is associated with increased inequality, as labour income is much more equitably distributed than capital ownership (see figure 2.1). Unless capital ownership is more broadly distributed, a growing share of national income flowing to capital will heighten inequality and dampen the living standards of those – the majority – who rely on their labour as their main income source.

FIGURE 2.1

A declining labour share of national income is associated with rising inequality
Changes in the labour share and in income inequality in OECD countries, 1990s to mid-2000s

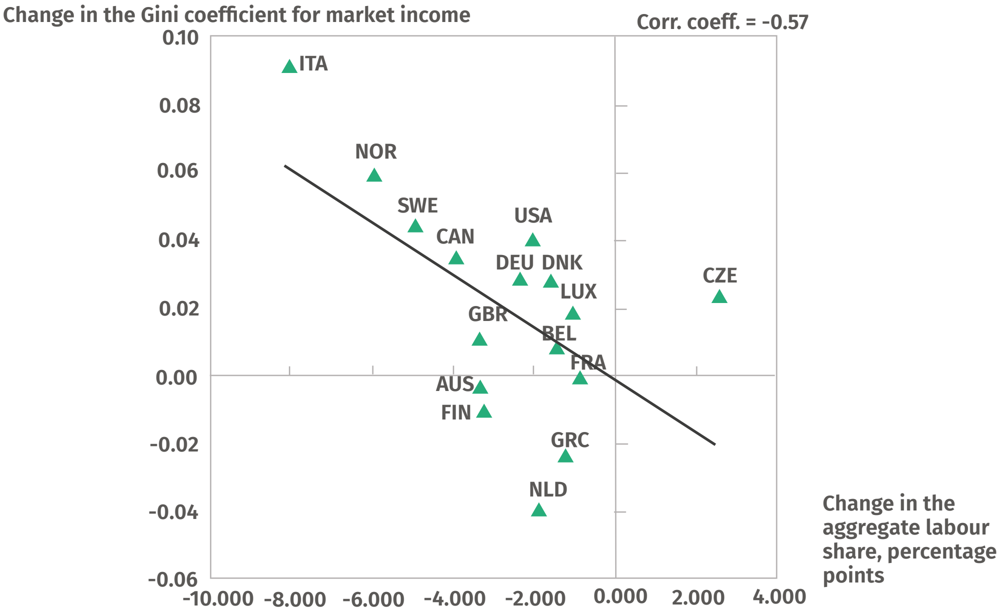

Source: OECD 2015

Technological change is estimated to have caused at least half the decline in the labour share in advanced economies in the last four decades (see figure 2.2). This has been driven by a combination of rapid progress in information and telecommunication technology, and a high share of occupations that could easily be automated. Routine skill-biased automation has taken over many of these tasks, contributing to job polarisation toward high-skilled and low-skilled occupations (ibid.) A fall in the price of investment goods relative to consumption goods, caused by rising capital-labour substitution, has exacerbated this trend. This is expected to accelerate as deploying automating technologies becomes cheaper, suggesting that labour's share could continue to erode (Karabarbounis and Neiman 2014).

The scale of transfer of income from labour to capital, which could be caused by automation, is potentially very large. McKinsey estimate that about half of all the activities people are paid to do in the world's workforce could potentially be automated by adapting currently demonstrated technologies. This amounts to almost $15 trillion in wages, shifting the flow of income from workers to the owners of their firms (McKinsey 2017) (see figure 2.3). In Europe's five largest economies, the wages associated with technically automatable jobs amount to $1.9 trillion annually, covering 62 million full-time-equivalent employees.

FIGURE 2.2

Technological change and globalisation have driven labour's falling share of income

The drivers of labour's declining share of income in the global economy

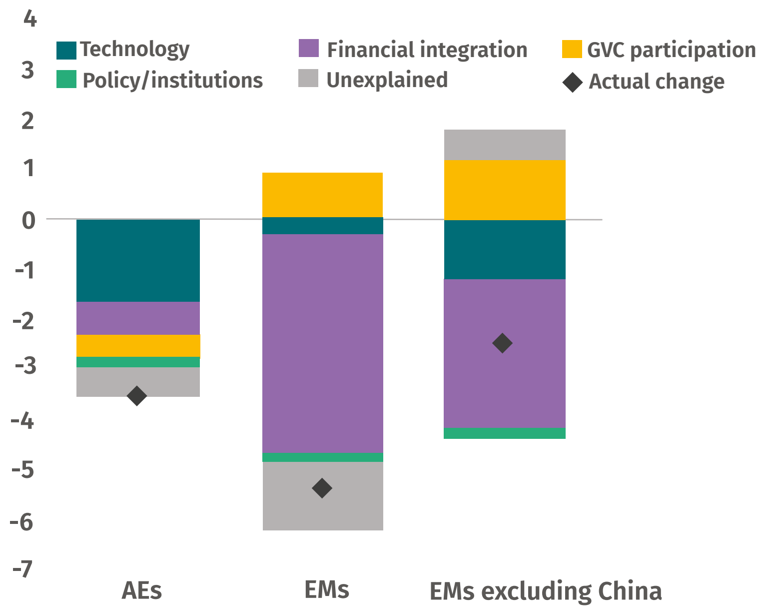

Source: Dao et al 2017

Our analysis finds that the value of earnings associated with UK occupations that could feasibly be automated using current or near-breakthrough technology is £290 billion.[8] This represents 33 per cent of all wages and earnings from labour in the UK economy.[9] In reality, the net value of wages lost through automation would be considerably lower. While increased automation of activities will replace some workers and labour earnings, employment and wages will rise in other areas of the labour market due to higher output and productivity. The final value of the transfer from wages to capital is hard to predict and in large part it will depend on the extent of friction in the labour market: whereby displaced workers are unable to move or retrain to find work, causing an oversupply of labour in certain geographies and sectors which results in high unemployment and suppressed pay. We would expect the final figure to be some part of, but lower than, the full £290 billion.

Automation therefore risks transferring very significant amounts of income from labour to capital, depending on whether automation substitutes or complements labour, and the employment effect of automation at the macro level (Freeman 2015). The extent to which substitution would accelerate inequality depends on how far income can be recirculated throughout the broader economy and how the transition is managed.

The risk is therefore less mass joblessness and more the 'paradox of plenty'. Technological change would make society richer in aggregate. However, capital-biased economic change would create a problem of distribution: those who can provide labour but do not own capital might have inadequate means of making a

8 Authors' analysis using Frey and Osborne's estimates of probability of computerisation by occupation (Frey and Osborne 2013), Quarterly LFS (ONS 2015-2017) and Annual Survey of Hours and Earnings (ONS 2017). See appendix for methodology and interpretation of this figure.

9 Authors' analysis based on OBR 2017a. Total wages and earnings figure excludes employer social contributions and includes labour share of mixed income.

reasonable living as technological change puts pressure on wage income in certain sectors and shifts income to capital owners.

How new technologies affect workers and inequality depends on who owns them and how the benefits are recirculated, whether via wages, tax and benefits, or capital returns. Without policy intervention, the benefits of growing productivity risk flowing disproportionately to capital owners; patterns of ownership therefore become increasingly vital in a more automated economy in which capital has a growing role as a factor in production.

FIGURE 2.3

Technical automation potential is concentrated in countries with the largest populations and/or high wages

Wages associated with technically automatable activities ($trillion in 46 countries, $14.6 trillion in total)

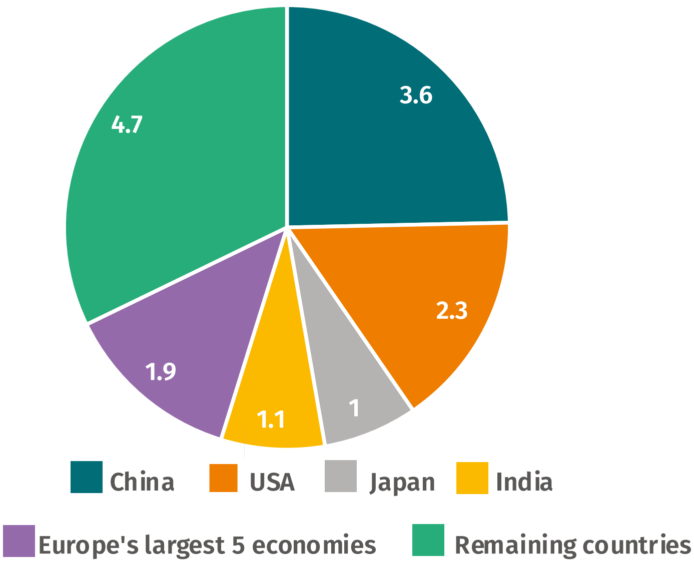

Source: McKinsey 2017

PLATFORMS, AUTOMATION AND INEQUALITY

Even if total wages in the economy do not fall, the capital income share may increase if returns to capital grow faster than returns to labour (wages). One trend making this outcome more likely, and set to deepen, is the set of recent advances in digital technology which have enabled and accelerated the growth of digital 'platform companies', such as Amazon, Google, and Facebook.

The growth of digital platforms reflects and is driven by a deeper trend: the rise of data as a critical driver of growth. The emergence of a set of networked technologies – smartphones, mobile internet, cloud computing, sensor and locative technologies – is enabling the capture of a gigantic and rapidly growing volume of unstructured mass data produced by numerous decentralised sources. The exponential expansion of datasets too large and complex to manipulate or interrogate with standard methods or tools would be of little use in and of itself. The sheer volume, velocity and variety with which networked devices and platforms are generating data is far too large for humans to analyse in a way that can generate useful information and economic value. The second crucial technological development has therefore been the accelerating capability of machine learning technologies and the growth of automated algorithms that can create patterns, information and insight out of datasets, and thus produce value.

A combination of vast datasets and computational power enables the customisation of products and production, improved decision-making by replacing or supporting human decision-making with automated algorithms, and the development of new business models, products and services (Schmidt 2017). AI and machine learning applied to data has also brought about self-improving technology, and enabled both improved and new automating technologies reliant on large quantities of processed data, such as driverless cars. A combination of 'big data' and AI technologies are therefore likely to enable the platforms to automate a greater number of work tasks in future. In turn this will further concentrate returns in the economy to the founders and investors of the leading digital platform companies (Srnicek 2016).

Given the relative labour 'lightness' of digital platforms due to their extensive use of algorithmic management instead of people, the rise in digital monopolies already appears to be driving rising inequality. Indeed, the concentration of economic power among digital platforms has been termed a form of neo-feudalism (Morozov 2016). Furman and Orszag (2015) show that economic 'rents' – payments to factors of production above what is required to keep them in the market – have been rising, particularly since 2000, and have been a central factor in increasing wage inequality and capital's increasing share observed during this period. They attribute this in part to digitalisation and the platform economy. The rise in 'supernormal returns on capital' at firms with limited competition is consequently leading to a rise in economic inequality. Relatedly, Autor and van Reenen found that industries where top firms' share of the market had most increased – where so-called 'superstar firm' had emerged, of which platform companies are the most prominent – had experienced the largest declines in the share of income going to workers on average. In each of these, there has been a growing share of income going to capital and to the most highly paid workers (Autor et al 2017). Digitalisation and the platform economy – enabled by automating technologies – therefore further threaten to increase economic inequality.

INCREASING WAGE INEQUALITY

Even with a stable labour income share, inequality will increase if wage inequality increases. There are strong reasons to believe that automation will lead to this outcome, as the effects of automation are highly unlikely to be even across the skills distribution. In the previous wave of automation, the substitution of routine tasks with computers predominantly replaced middle-skill and middle-wage jobs. By contrast, high-skilled workers were complemented by computers rather than substituted, helping increase their productivity and driving up their wages, leading to wage polarisation (Acemoglu and Autor 2011). For example, using data on the United States, Japan, and nine European countries including the UK from 1980 to 2004, Michaels, Natraj and Van Reenen (2014) found that industries with faster ICT growth shifted demand from middle-educated workers to highly educated workers, consistent with ICT-based employment polarisation. Other studies have found that computer use is also associated with greater inequality of wages within occupations, with greater wage dispersion arising if new skills are costly or difficult to acquire, so that only some workers acquire the skills. This association contributes to wage inequality, accounting for an estimated 45 per cent of the growth in the wage gap between the 90th and 50th percentiles of the entire workforce since 1990 in the United States (Bessen 2016).

Even if AI has the potential to perform a range of tasks across the skills distribution, including the analytical and manual tasks unaffected by past automation, it does not follow that those who currently perform those jobs will be replaced. But analyses to date suggest that jobs at the low-skill end of the labour market have the greatest potential for complete automation, or substitution. Though McKinsey's analysis of the kind of tasks likely to be automated suggests that workers right across the skills distribution will be affected, it also suggests that a much larger proportion of lower-skill tasks are susceptible. Looking at occupations as a whole, Frey and Osborne (2013) find that the next generation of automating technologies are much more likely to be able to complete *all* the tasks of lower-skills jobs, relative to higher skill ones. This outcome is different from previous waves of automation in which mid-skilled jobs were affected more than low-skill jobs, because increasingly new technologies are able to perform manual tasks that previously were not technologically feasible.

As a result, the low-paid and low-skilled are likely to bear the brunt of adjustment costs as automation reshapes the labour market over time. For example, Deloitte estimate that, in the UK, jobs paying less than £30,000 are five times more susceptible to being automated than jobs paid over £100,000 ((Frey and Osborne 2015). As Daniel Susskind has argued, as automating technologies become increasingly capable both in quantity and productivity, they erode many of the types of tasks in which labour had an advantage over machines. Labour is forced to specialise in a diminishing set of types of tasks, creating winners and losers in the labour market (2017).

If these analyses are correct, it suggests that AI will substitute for relatively lower-skilled jobs and complement higher-skilled jobs. It can therefore be expected that wages for high-skill jobs are likely to increase relative to jobs that require lower skill levels. This is because complementary technology increases the productivity of the worker and therefore the wages which can be commanded. Furthermore, as high-skill workers become more productive, they are able to complete more tasks that previously would have been done by middle-skilled workers, pushing down those workers' relative wage and reallocating mid-skilled labour to low-skill tasks (Acemoglu and Autor 2011). This is not a necessary feature of all automation; in the nineteenth century, new technologies often complemented low-skill workers most, pushing up wages for this group (ibid). But the pattern of the last century, and the anticipated pattern of the coming decades, is for automating technologies to complement higher-skilled workers and substitute for lower-skilled workers.

Even in the case where aggregate employment doesn't fall, it is extremely implausible that automation will affect all groups in society – for instance, defined by geography and education – uniformly. This is especially true in the short term. Even if aggregate employment and wages stay constant (or even increase), those in automated industries will be made worse off unless they also acquire comparable employment opportunities. Indeed, the local effects of automation, whether in the previous wave (Autor and Dorn 2013) or in the early stages of this one (Acemoglu and Restrepo 2016; Acemoglu and Restrepo 2017), tend to be negative, with local employment and/or wages falling. Between those areas where automation happens at speed, and those where it doesn't, inequalities are likely to grow.

More generally, higher-skilled individuals are better equipped to adapt their skills to changing circumstances, and thus of finding ways to complement the capabilities of AI. By contrast, individuals lower down the skills distribution often have skillsets and qualifications particularly well suited to their current work but less adaptable to other kinds, and will thus find it harder to find new work (Burkhardt and Bradford 2017). This dynamic is also likely to exacerbate geographic, gender and ethnic inequalities, as high-skill roles in most sectors are concentrated in London and dominated by men and particular ethnic groups.

Our analysis of automation potential by sex finds that while a similar absolute number of men and women are presently at risk this makes up a greater proportion of all jobs performed by women (46.8 per cent compared to 40.9 per cent of men's jobs). Because occupations with the greatest potential for automation tend to be paid below the average, the direct result of displacement would see the gender pay gap narrow by around 1 percentage point. This would remain the case, however, only if those displaced workers who re-enter the labour market do so at around the new average salary for respective genders.[10] However, as things stands, this is not likely to be the case. Firstly, investment in productivity raising technology is lower in sectors dominated by women, such as care and retail, and therefore automation may not occur; instead women may disproportionately remain in low-paid jobs which employers have chosen not to automate. Rising demand in these sectors, especially care, could also absorb workers, predominantly women, from automating industries. Secondly, complementary technology is likely to raise the wages of some of the highest paid – who are more likely to be men - further, leading to greater wage disparity.

The challenge of automation is therefore likely to be more about redressing the maldistribution of the dividends of technological change than a problem of production in which human labour becomes redundant. Society's objective should be to manage and distribute the gains of higher productivity so that the benefits of automation are fairly and widely shared.

10 Difference in gross hourly pay between men and women as a proportion of men's wages, using 2017 ASHE data. See appendix for methodology.

3.
A managed acceleration of automation is required to reap the full productivity benefits and enable higher wages

Automation has the potential to deliver enormous productivity dividends. McKinsey estimate that widespread adoption of AI, machine learning, and robotics could raise global productivity growth by between 0.8 and 1.4 per cent annually, while the UK's Digitalisation Review estimates that industrial digitalisation could realise efficiency gains of approximately five per cent of GDP per year (McKinsey 2017; HMG 2017a). PwC estimate that AI alone will mean GDP in 2030 is 10.3 per cent higher than it would otherwise have been – the equivalent of an additional £232 billion, or extra spending power per household of £1,800–£2,300 a year if distributed equally (2017b). The gains are driven by productivity and consumption impacts.

If managed well, this could have profoundly positive effects: better, more 'human' work, greater leisure time, improving living standards, more efficient and less environmentally damaging forms of production, and increased non-rivalrous consumption of goods and services.

FIGURE 3.1

Growth will need to be powered by productivity gains in the coming decades to maintain growing GDP per capita growth

Composition of GDP growth for G19 and Nigeria, compound annual growth rate (%)

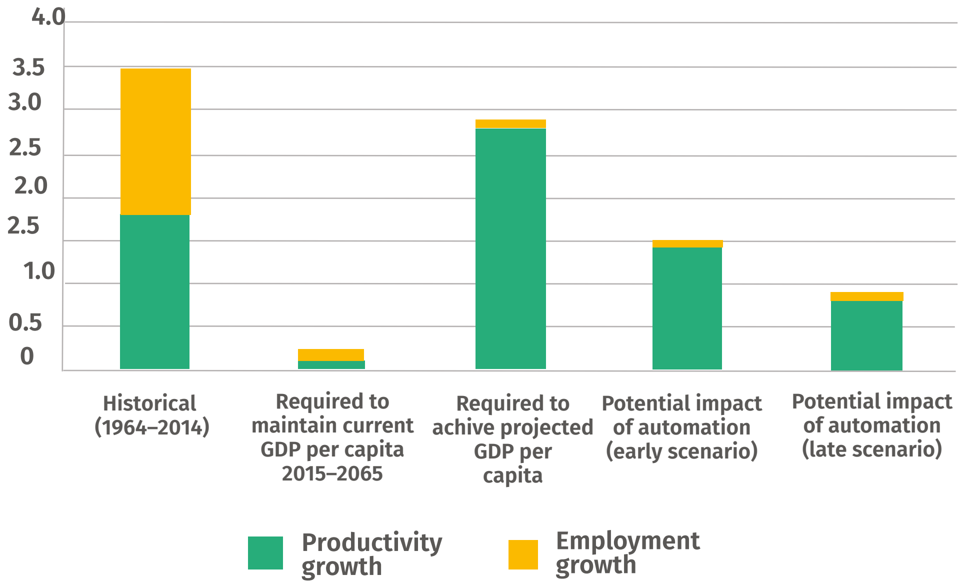

Source: McKinsey 2017

Accelerating the adoption of automating technologies is particularly important to ensure that living standards do not stagnate as a result of demographic change in the coming decades. As the population of advanced economies age, employment growth will inevitably become less important as a source of overall growth. Given that employment growth contributed almost half of all increased output in the major global economies between 1964 and 2014, future prosperity will have to be rooted in productivity advances, not growth in the working age population (see figure 3.1). If productivity performance is not improved, the rate of per capita GDP growth among advanced economies is expected to experience a historic fall. The pace at which we adopt automating technologies is therefore critical to determining our overall economic output in the decades ahead.

Crucially, society has the capacity to shape the pace and impact of automation. While the demonstrated ability to automate specific tasks or occupations is a necessary condition for automation to occur, the interaction of a series of other factors will determine whether it actually does (McKinsey 2017). Five will be particularly important in influencing the pace and extent of automation: the cost of developing and deploying these technologies, the relative cost of capital and labour, the economic benefits of automation beyond labour costs, the balance of economic power between labour and capital, and social and regulatory acceptance.

Crucially, these factors are strongly affected by public policy. Technology is not destiny: its adoption and use are shaped by broader public policy choices, social attitudes and business preferences. The speed and scale at which automation occurs will depend significantly on the collective choices we make. The machine age will be human shaped. A managed acceleration of automation will be vital if is to deliver its productivity dividend and its potential to improve living standards.

WHERE ARE THE ROBOTS?

Accelerating automation is an important policy goal, because, whatever the rhetoric about 'the rise of the robots', the reality is that the UK is adopting automating technologies at a relatively slow pace. We lag behind other EU nations in terms of adopting digital technologies (Jacobs et al 2017). The sale of industrial robots actually fell in 2015 despite growing by 15 per cent worldwide, with fewer robots purchased in the UK than in France, Germany, Spain or Italy. The UK has just 33 robot units for every 10,000 employees, compared with 93 in the US and 213 in Japan (International Federation of Robotics 2016). This of course partly reflects the UK's smaller manufacturing base than in these countries (though it may also help explain it too).

What is more worrying is that overall business spending on ICT, machinery and other equipment has hardly grown in real terms since 2000, while the proportion of UK GDP accounted for by gross fixed capital formation (GFCF) – another measure of business investment – is seven percentage points lower than it was in 1990 (CEJ 2017). In fact, corporate investment in fixed assets (not including construction) fell from 11 per cent of GDP in 1997 to just 8 per cent in 2014 – below the rate of capital depreciation, meaning that the stock of fixed business capital is actually falling (ibid). This should be balanced against the fact that the UK businesses are leading investors in non-R&D intangible assets such as product and service design including digital products, reflecting the size of service sectors in the economy (Haskel and Westlake 2017). Many of these forms of investment are complementary to the skills and infrastructures on which platforms and companies are particularly reliant. Nonetheless, these are not signs of an economy on the cusp of a wave of automation.

ACCELERATING ADOPTION OF AUTOMATING TECHNOLOGIES AMONG NON-FRONTIER FIRMS

Given this, a managed acceleration of automation to boost productivity across the economy should be a key goal of public policy. Maximising the benefits of automation will require improving the rate of diffusion of technologies from the minority of frontier firms to the majority of slow-adopting, low productivity firms in the rest of the economy (Haldane 2017; Andrews et al 2016; Andrews et al 2015). It is estimated that three-quarters of potential productivity improvements related to automation come from the broader adoption of best practices and technologies, as companies catch up with sector leaders. Only a quarter is from technological, operational, and business innovations that go beyond best practices and push the frontier of the world's GDP potential (McKinsey 2015).

Yet the diffusion of technologies to the non-frontier economy has to a considerable extent broken down, slowing the pace of automation and weakening productivity growth (Andrews et al 2015). For example, across the OECD labour productivity in the manufacturing sector at the global frontier increased at an average annual rate of 3.5 per cent over the 2000s, but only 0.5 per cent in non-frontier firms. In the services sector, productivity of frontier firms grew 5 per cent but actually fell by -0.1 per cent in non-frontier firms (ibid).

The gap in productivity performance between the minority of leading firms and the broader, lagging majority of companies is even more pronounced in the UK than other advanced economies (OECD 2017). While 1 per cent of firms have seen average productivity growth of around 6 per cent per year this century, around one-third of UK companies have seen no rise in productivity at all (Haldane 2017).

FIGURE 3.2

Frontier firms have generated rising productivity as non-frontier firms have stalled
Productivity among frontier and non-frontier firms by GVA per person

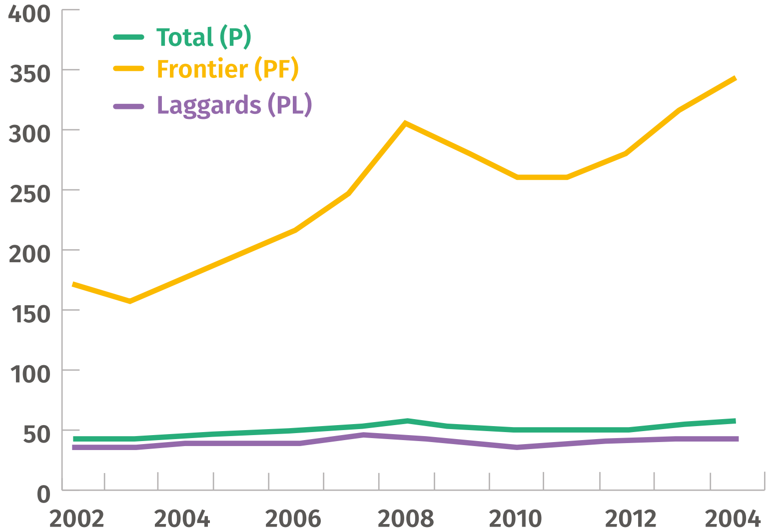

Source: Haldane 2017

Inhibitors of effective technological adoption include legal and economic barriers to the dispersal of innovation, such as IP law and patents (Haldane 2017). Relatedly, the surge of 'superstar firms' in the digital economy, driven by powerful network effects, could be generating monopolies that dominate the market and hinder innovation and the dispersal of its benefits (Autor et al 2017). Sectoral shifts in the economy – with the growth of low-productivity service sectors that are poor technology adopters – also suggest more support for the everyday economy in adopting technologies is required (IPPR 2017). Poor management practices in the UK are also a factor in the long tail of low productivity non-frontier companies, contributing in general, to a lower take-up of new technologies and processes than in other advanced countries (Van Reenen et al 2016). Relatedly, low levels of worker voice in the average British firm mean their knowledge in how best to deploy and use technologies is not properly drawn upon, inhibiting the effective up-take of technology (TUC 2017). The cumulative effect is a significant and growing difference between the productivity of the UK's leading firms, in part driven by technological adoption, and the rest of the economy.

Closing this gap – in part by speeding up the diffusion and adoption of productivity-boosting technologies – could prove transformative. As Andy Haldane, chief economist of the Bank of England, points out, if productivity growth in the second, third and fourth quartiles of the distribution of UK firms' productivity could be boosted to match the productivity growth of the quartile above, it would deliver a boost to aggregate UK productivity of around 13 per cent, taking the UK to within 90–95 per cent of German and French levels (Haldane 2017). The Government's 'Made Smarter' review of industrial digitalisation technologies (IDTs) estimates that the faster adoption of IDTs in particular could improve industrial productivity by more than 25 per cent by 2025, increasing manufacturing sector growth by between 1.5 and 3 per cent per annum, leading to an estimated net gain of 175,000 jobs (HMG 2017a).

Accelerating the pace of automation, as part of the wider adoption of digital technologies across the economy, should therefore become a key goal of public policy. But the goal should be more specific than this. It should be to ensure that automation is managed in a way which enhances the quality of jobs of those working alongside automating technologies, gives workers displaced by automation opportunities for retraining and redeployment elsewhere in the economy, and ensures that the productivity gains derived from automation are shared equitably between capital and labour, and across society as a whole.

In its discussion paper on industrial strategy, the IPPR Commission on Economic Justice has proposed the use of 'missions' as a key means of focusing policy on major challenges facing society (Jacobs et al 2017). Drawing on work by Mariana Mazzucato (2013, 2017), the paper argues that by creating expectations of future demand such missions can stimulate private sector investment, and at the same time can strengthen UK innovation and domestic value chains to meet that demand. It argues that one such mission should be to make the UK the world's most digitally advanced economy by 2040, including through the development and diffusion of automation technologies. In its subsequent Industrial Strategy white paper, the Government has adopted a similar approach, identifying AI and robotics as one of its central 'grand challenges' (HMG 2017b).

Three key strands of policy are needed to fulfil such a mission.

First, **government needs a clear strategy, allied to serious resources, to support frontier innovation in robotics and artificial intelligence, along with other related fields of industrial digitalisation**, such as advanced and additive manufacturing. The UK is already a global market leader in AI, with over 200 SMEs in the field – compared to just 81 in Germany and 50 in France (HMG 2017a). We have particular strengths in some sub-fields, such as in data analytics, virtual and augmented reality and the development of 'collaborative robots' ('cobots') in aerospace. But other countries – including Germany, China and the US – have much more developed strategies to support frontier innovation in this field (Ibid). In its 2017

Autumn Budget the Government provided £75m for R&D in AI, and funding for 450 PhDs (HMT 2017). But as its 'Made Smarter' review pointed out, it does not have a strategic vision, direction or coordination (HMG 2017a). Such a strategy should be led by the public innovation agency Innovate UK, working closely with the leading businesses and academic research centres in the field. The 'Made Smarter' review proposed the creation of a network of digital research centres (in AI, machine learning and data analytics, additive manufacturing, robotics and automation, virtual and augmented reality and the internet of things and connectivity), aiming at upward of £400 million additional R&D investment.

Second, **there needs to be a dedicated programme to accelerate the diffusion of automation, robotics and other digital technologies throughout the economy**. The 'Made Smarter' review described this as a 'national adoption programme' (ibid). In its industrial strategy discussion paper, the IPPR Commission on Economic Justice proposes the establishment of a new body, Productivity UK, to drive firm-level productivity across the economy, including through the more rapid adoption of digital and automaton technologies (Jacobs et al 2017). Its goal would be to help small and medium-sized businesses to understand the productivity-raising potential of new technologies and to accelerate their introduction, through information and advisory services, and grants and loans for investment. This would be done alongside wider management education and consultancy services, and support for skills training. It would work with sectoral bodies in key sectors, with the Scottish, Welsh and Northern Ireland governments, and in England with local enterprise partnerships (LEPs) and combined authorities. It would draw on the work of the new 'Be the Business' initiative of the Productivity Leadership Group.[11]

A particular feature of such a national adoption programme should be an increase in the voice given to workers in the process. If automation and other digital technologies are to bring better, and better-paying, jobs alongside them, it is important the workers are involved in their introduction. To ensure workers' voice is embedded in how Productivity UK operates, we recommend it is governed as a tripartite body between government, businesses and trade unions. Productivity UK should also seek to empower worker voice in how technologies are adopted and used, for example working with trade unions to ensure that employees play a partnership role in the development and deployment of new technologies in the workplace. It should also consider how best to ensure that workers in non-unionised workforces have an effective voice as new technologies are introduced.

Third, **a major national programme should be introduced to equip the UK's workforce with the skills and capabilities needed to complement new technologies.** The Made Smarter review calls for such a programme to re-skill and upskill one million workers over five years (HMG 2017a). But this cannot just be a question of skills supply. As IPPR analysis has shown, the UK's skills problem – in this and other fields – is not just one of inadequate vocational education and training. It is also that British firms exhibit low levels of demand for higher skills (Dromey et al 2017). Around one-third of UK workers are estimated to be over-qualified for their jobs, the highest proportion in the EU (CEDEFOP 2015). IPPR has therefore called for the current apprenticeships levy to be widened in scope, turning it into a 'productivity and skills' levy which firms can use to invest in training and skills utilisation right across their workforces (Dromey et al 2017).

Responding to changing demand for skills in a way that ensures equitable outcomes and improved rates of technological adoption will require radically improving and adapting the skills systems of all advanced economies (The White

11 See https://www.bethebusiness.com/

House 2016; TUC 2017). Lifelong, continuous learning and updating of skills will increasingly necessary in order to support adults to adapt to new technologies and a changing labour market. Demand for technical skills to develop and effectively use advanced technology, as well as 'soft' – cross-functional – skills such as communication, problem-solving and emotional intelligence, is expected to grow. Conversely, those skills associated with tasks that are routine in nature are likely to decline (OECD 2017). In addition to the productivity and skills levy, we would support the proposals which IPPR has made in its comprehensive 'Skills 2030' programme (Dromey et al 2017). These are as follows.

- **Provide a 'Personal Training Credit' to support low-paid and low-skilled individuals to invest in their training and career.** This credit would focus resources on those who need support the most, giving people control over their training and careers, and would help close the participation gap. The credit would be worth up to up to £700 a year for adults with low qualifications who are either in low-paid work or who are unemployed. In addition to focusing the Personal Training Credit on low-paid, low-skilled workers, Government could choose to prioritise workers in industries which are more vulnerable to job displacement as a result of automation.
- **Introduce a 'Personal Retraining Allowance', worth £2,000, for workers who are made redundant and lack an NVQ level 3 to invest in upskilling.** This could be funded by reducing the tax free entitlement on redundancy payments to the statutory maximum payment.
- **Develop strong sectoral and local institutions to drive skills policy that reflects the regional composition of current and expected employment and industry**. As part of its industrial strategy, the Government should use sector deals to build new institutions to improve the quality of training, and drive skills utilisation and workplace performance. These institutions would be responsible for identifying and articulating demand in their sector, designing standards, training content and career pathways, overseeing awarding bodies, investing levy underspend and boosting job quality.
- **Establish a cross-government framework to identify and monitor industries in transition as part of the government's new industrial strategy.** This should be targeted at those industries with both a high number of jobs with the potential to be automated, and a high proportion of workers with lower-level qualifications who are more at risk of slipping into long-term unemployment.

4.
An Authority for the Ethical Use of Robotics and Artificial Intelligence should be established to regulate the use of automating technologies

The widespread use of increasingly powerful AI and robotics is likely to generate a range of ethical, legal and regulatory challenges. The range of questions automation provokes is broad and profound (UK Parliament 2016). These include:

- establishing where responsibility lies in the operation of autonomous technologies, particularly where they fail
- concerns around privacy
- ensuring human safety
- the extent and necessity of human control of autonomous systems
- the treatment and rights of non-human intelligence
- application of lethal non-human directed violence
- the lack of scrutiny or understanding of technologies, such as machine learning, and the risk that technologies reproduce social inequalities and biases
- the question of how autonomous technologies should choose between 'competing bads' in situations where a bad outcome for humans is unavoidable
- the implications of AI surpassing human intelligence

There are many others, and more that we cannot yet properly foresee but are likely to arise (Dellot and Wallace-Stephens 2017).

The risk is that the ethical norms and regulatory architecture that respond to these issues will be shaped not by democratic debate and decision but by leading technology companies. This would reflect Lawrence Lessig's claim that 'code is law' (Lessig 2000). In other words, code – and technologies in general – create binding restrictions on the types of behaviour they allow, with legal and ethical frameworks often adapting to emerging technologies in the interests of the developers, rather than their use being shaped to reflect societal norms and preferences.

There are serious problems with this approach. It risks granting excessive and undemocratic influence to a small number of private companies to shape the rules governing the use of emerging technologies, and the ceding of the power that resides in that control. It also risks stifling innovation, as market leaders shape the regulatory environment to best support their business models. The lack of neutrality in shaping the rules governing use of AI and robotics would also likely act as a barrier to public trust, and inhibit the emergence of a clear and transparent regulatory environment.

A better solution is for society to proactively develop the legal, ethical, professional and behavioural framework that would govern the development and use of AI and robotics. This should include developing the institutional mechanisms by which those norms are established and enforced. **To this end, we propose the establishment of an 'Authority for the Ethical use of Robotics and Artificial Intelligence' to provide guidelines and recommend regulatory frameworks for the use and governance of these technologies**. The authority should be modelled on the Human Fertilisation and Embryology Authority (HFEA), which regulates the creation and use of human embryos, issuing regulatory and ethical guidelines to encourage but also ensure the safe use of embryonic technologies.

The proposed authority would play a similar role for robotics and AI. It would provide the technical, ethical and regulatory expertise needed to support relevant public actors to ensure a timely and well-informed public response to the new opportunities and challenges arising from automation. Its members should have a mix of scientific expertise, philosophical, legal and ethical backgrounds, and lay views, as well as specific technical and methodological expertise relevant to the technologies reviewed.

The goal would be to embed the public interest through the Authority defining fair and ethical regulatory guidelines and legal frameworks for the governance and use of automating technologies. The authority should consider and make recommendations relating to the ethical use of autonomous systems including standards and regulation, as well as how best to design and programme ethical autonomous agents. It should seek to answer the questions, in Adam Greenfield's words, of: "What type of deal do we wish to pursue with the technologies, what role do we want to have within these economies, and how do we ultimately occupy ourselves?" (2017).

There is a growing recognition of the need for a regulatory institution to oversee robotics and AI development and use. For example, the Department for Transport's *Pathway to Driverless Cars* report stated that, if driverless cars are to become widespread, regulatory action would be needed to tackle the ethical issues they would create. Similarly, the RSA have recently argued that a new ethical framework for the use of technologies should guide policymakers in regulating automation (Dellot and Wallace-Stephens 2017), Nuffield Foundation are establishing a 'Convention on Data Ethics' (Nuffield Foundation 2017) while the British Standards Institute have issued a *Guide to the ethical design and application of robots and robotic systems* (2016). The House of Commons Science and Technology Committee report into the use of algorithms in decision-making also stressed the need for more effective oversight (2017), while n its 2017 Autumn Budget, the Government committed to the creation of a new Centre for Data Ethics and Innovation, "to enable and ensure safe, ethical and ground-breaking innovation in AI and data-driven technologies" (HMT 2017).

An Authority for the Ethical Use of Robotics and AI would provide a UK counterpart to international institutions which are already seeking to shape the fair and ethical use of technologies. The EU is currently establishing a 'European Agency of Robotics and AI' to provide technical, ethical and regulatory expertise to inform regulatory responses to technological developments (European Parliament 2017). Similarly, the US, Japan, China and South Korea have or are establishing agencies responsible for safely integrating robotic and AI technologies (Calo 2014).

There are many areas where this authority could offer guidance. How should liability be framed and how should the legal framework around the use of autonomous technologies adapt? How should companies respond to the growing use of AI? For example, DeepMind has established a unit to examine the ethical and societal questions raised by AI (Hern 2017). An authority could examine whether such innovations should be rolled out more widely. Similarly, it could work with the Health and Safety Executive, trade unions and business to update workplace laws to protect workers and improve health and safety given the increasing use of autonomous technologies in the workplace.

On algorithmic transparency, it could require companies to supply the rationale behind any decision taken with the aid of AI that can have a substantive impact on human lives and ensure AI decision-making processes are made comprehensible for ordinary people. It could also trial new regulatory initiatives. For example, the data scientist Cathy O'Neil recently suggested creating a National Algorithmic Safety Board, modelled after the National Transportation Safety Board. The latter investigates plane crashes and suggests improvements; the former could act in a similar fashion if and when algorithms 'crash' (O'Neil 2017). These are just some of the regulatory initiatives and guidelines an authority could set out in relation to the challenges of automation. The vital need is to acknowledge that these questions must be addressed by society through collective debate and decision, and not simply left to the developers of the technologies themselves.

5. New models of capital ownership are needed to ensure automation broadens prosperity rather than concentrates wealth

Automation could deepen existing inequalities, disempowering labour, concentrating reward and reinforcing hierarchy. Alternatively, it could underpin an economy where technological change provides people with more resources and power, with automation enabling what Roberto Unger calls “deep freedom”: the “empowerment of the ordinary person - a raising up of ordinary life to a higher plane of intensity, scope and capability” (Unger 2013). Crucially, there is no technologically-determined future, whether dystopian, emancipatory or simply mundane. Technologies are technical artefacts whose effects are powerfully shaped by the broader social, technical, and economic systems they are embedded in. If we want the growing capability of machines to deliver for all, we must build the economic institutions that support a future of shared plenty.

A critical determinant of the distributional effect of automation is the underlying pattern of ownership of economic assets. Unequal patterns of capital ownership are likely to act a fundamental driver of inequality. As the economist Richard Freeman puts it: “As long as the relative advantage of machines varies, there will be work for humans. The problem is that the owners of the machines will receive the vast bulk of the benefits of the technological progress… who owns the robots rules the world” (2015). This is because if capital increasingly dominates the economy at the expense of labour then unequal levels of ownership will accelerate inequality as more and more income flows to unequally owned capital relative to labour’s share. Compounding capital’s growing share of income, labour income risks becoming more unequal due to a skills-based polarisation of the labour market. Automation therefore risks creating a ‘paradox of plenty’: society would be far richer in aggregate, but for many individuals and communities, technological change could reinforce inequalities of power and reward. We therefore need to solve a problem of distribution more than one of production when responding to automation.

This would not be the case if capital ownership were more broadly distributed. If ownership were widespread, then, as capital became more important, this would drive broad-based improvements in living standards as people supplemented their labour income with growing capital income. More equal patterns of ownership would be a force for economic convergence.

However, capital ownership in the UK is highly concentrated. The wealthiest 10 per cent own almost half the private wealth in the UK, and the top 1 per cent own 14 per cent. By contrast, the poorest 50 per cent own just 9 per cent of private wealth (ONS 2015). As well as inequality between families, who owns wealth is sharply filtered by region, generation, gender, ethnicity and class (Roberts and Lawrence 2017). Ownership of financial assets is even more unequally distributed than other forms of wealth. Indeed, outside of housing wealth – which is itself becoming more concentrated, particularly by generation – most people's wealth is relatively small and few have ownership of the productive assets of the economy. They lack both control and a right to share in the benefits of ownership in the form of economic return. For example, the wealthiest 10 per cent own almost 70 per cent of the UK's financial wealth, including almost four-fifths of shares (ONS 2015). While the median financial wealth of the richest 10 per cent is £153,900, for the least wealthy half of households, it is just £400 (ibid). Moreover, those with an income of over £1 million a year receive a fifth of their incomes from capital income in the form dividends, interest and property income, compared to less than 5 per cent for those earning between £20,000 and £30,000, and virtually nothing for the poorest households (ONS 2017). Since the wealthy receive a greater proportion of their income in capital returns than the rest of the population, if capital's share increases, and returns to capital outpace returns to labour, inequality necessarily increases.

Yet the traditional mechanisms to distribute capital are becoming weaker. Share ownership by individuals has fallen from 53 per cent of the total in 1963 to just under 12 per cent in 2014, while UK pension funds own less than 10 per cent of the FTSE (ONS 2015). Employee share ownership schemes are not widely spread and are generally regressive (Lawrence and Mason 2017). Home ownership is falling, with sharp regional and intergenerational divides (D'arcy and Gardiner 2017). The UK therefore has both a highly unequal distribution of capital and increasingly weak mechanisms for broadening capital ownership (Roberts and Lawrence 2017).

To make sure that automation enriches all of society, we therefore need to broaden and democratise capital ownership. Expanding the distribution of capital and pluralising models of ownership would help democratise who has a claim on the economic dividends of automation, both in terms of control and benefit. It would broaden who has a claim on the common wealth. Broadening society's claim to the dividends of automation would also recognise that technological progress is enabled by a collective inheritance, the social accumulation of knowledge and investment (Alperovitz 2016). We will need new models of ownership to institutionalise that acknowledgement. Broadening ownership would help ensure technological change drives rising living standards across the board, sustain aggregate demand, and dampen the drive to inequality due to the dynamic of rising returns to capital relative to labour widening income and wealth disparities.

A major goal of public policy should therefore be to seek to disperse and pluralise ownership in society, building institutions where the wealth generated by technological change is more widely shared. This should go hand in hand with efforts to recirculate the productivity gains of automation through wage bargaining, by improving workforce skills and through taxation of the profits of automating companies.

A number of measures could more widely distribute capital ownership. Three in particular have the potential to extend ownership at scale, and are currently being examined in detail by the IPPR Commission on Economic Justice.

1. ESTABLISH A CITIZENS' WEALTH FUND

A Citizens' Wealth Fund should be established to invest in company shares and other productive assets on behalf of the public.[12] At its core, the Fund would be a mechanism to transform national private and corporate wealth into public wealth, providing an institution for the collective ownership of capital. By owning wealth in common, and distributing the proceeds for the benefit of citizens, the Fund would act as a force for economic convergence by sharing returns to capital. To the extent that automation improves productivity and corporate profitability,

12 For further details, see Roberts and Lawrence 2017.

the whole of society would share in the benefit through the fund's ownership of a broad range of assets. The fund's investment mandate including ethical investment requirements should be established by Parliament, but be managed independently from government in pursuit of its mandate, with an independent Board and management agency. A proportion of its annual income should be returned to the public (its owners), either through individual transfers such as a universal capital dividend, or collectively via budget transfers or community investments of various kinds (Cummine 2016). The pay-out from the fund should be capped to ensure the fund is permanent and able to grow to a sufficient size.

The fund could be capitalised by a mix of capital receipts transfers and new taxes or revenue streams. Sources for capitalisation could include transferring of certain public assets into the fund, proceeds from planned asset sales (according to the OBR, planned asset sales – primarily the winding down of financial assets – are expected to raise around £54 billion between 2017/18 and 2022/23), the hypothecation of particular tax streams such as reformed and new wealth taxes into the fund, the transferring of new sources of revenue such as through spectrum sales or new models of ownership of common data, or even through public borrowing (Lansley 2017). One innovative option would be to introduce a 'scrip tax'. A scrip tax is a tax on corporate profits paid by firms issuing equity to government instead of cash. It transforms a stream of payments in the form of tax into an asset that produces returns, with part of the government's claim on corporate profit paid as a form of stock. While a scrip tax would moderately dilute shareholder wealth, it would not reduce the working capital of businesses and would broaden ownership. Given equity raised by a scrip tax would accumulate in value, it could be set at a low rate and still create a substantial stake for the Fund over time.[13] Corporation tax is expected to raise £276 billion between 2018/19 to 2022/3 (2017). Applying a scrip tax in the UK as part of or in addition to corporation tax has the potential to generate significant amounts of equity. It would therefore enable the Fund to steadily accumulate assets over time, ensuring a social claim on returns to capital.

The fund would also expand due to real returns. The majority of sovereign wealth funds consistently deliver a real rate of return of CPI plus 4.5 per cent (Roberts and Lawrence 2017). Given this, with effective capitalisation and stewardship, the fund should continue to expand substantially over time, acting as a vehicle for the accumulation of assets on behalf of the public to ensure everyone in the UK had a stake in the economy, a collective say in its direction, and benefited from the gains of automation.

2. THE EXPANSION OF EMPLOYEE OWNERSHIP TRUSTS (EOTS)

EOTs are share capital funds held in trust for the benefit of all employees in a company. Established in the UK in 2014, they enable a business owner to be exempted from capital gains tax if he or she sells a minimum 51 per cent stake in the company to the trust. The trust therefore effectively gains a controlling interest in the company on behalf of the employees, and can pay out dividends to them. The effect is to invert the traditional company ownership hierarchy: whereas capital normally hires labour, in an EOT-owned company the employees hire the capital. Expanding this model has the potential to transform firm ownership among small and medium-sized companies, and potentially larger ones as well: it could be achieved with a set of reforms to the tax incentives for firm owners and investors (for further details, see Lawrence and Mason 2017).

13 For example, a form of a scrip tax was the mechanism that capitalised Sweden's 'wage-earner funds' which were at the core of the innovative and economically successful 'Meidner Plan' (Furåker 2015).

3. INTRODUCING COMPULSORY PROFIT SHARING IN FIRMS ABOVE A CERTAIN SIZE

In France, profit sharing has been mandatory since the 1960s for firms with more than 50 employees, with profits shared being exempt from employers' national insurance contributions and employees' income tax if the profit share is democratically agreed within the firm (ETUI 2017). Profit sharing enables employees to benefit from the profitability of their company directly and reflects the fact that a company is a partnership between capital and labour, in which capital's right to profit is not absolute. A similar tax-incentivised scheme for firms of 50 employees or more should be introduced in the UK. This should be accompanied by efforts to expand cooperatives and employee-owned firms.

4. REDUCING WORKING TIME AS PRODUCTIVITY IMPROVES

One further means of spreading the benefits of automation should also be on the agenda: a reduction in working time (Coote 2014). Automation presents an opportunity to reconsider the purpose of work and how we allocate working time. In particular, we believe that the advancing capability of machines should enable us to work better, but less.[14]

As productivity improves, society always has the option of taking the benefits not just as higher income, but as more time free from work. Historically, this has been a vital component of rising living standards, but over recent years the secular decline in average working hours has stalled (ONS 2017b). While some individuals, mostly on above-average incomes, have been able to take higher earnings in the form of shorter hours, this is difficult for most. Consumption patterns are socially determined and compared, and most people would find it difficult to consume less than they do now if their neighbours and friends were not also doing so. It is for this reason that a general reduction in working hours is unlikely to come about simply as a result of individual choices. We will need to make such decisions collectively, as a society. One option would be a general reduction in the 'normal' working week. Another would be a gradual increase in the number of public holidays.

The Commission on Economic Justice is considering the levers that could bring about a reduction in working time. However it might be done, this should surely be one of the key issues which society debates as it contemplates a more automated future alongside new models of ownership. Only if everyone is able to make a claim on the wealth generated by technological progress can we create a future of shared plenty, where prosperity is underpinned by justice.

14 Work-related stress, depression or anxiety accounts for 40 per cent of work-related ill health and 49 per cent of working days lost, in 2016/17 (HSE 2017), suggesting a reduction in working hours would have a wide range of benefits.

Appendix: Methodology

In this paper we present estimates of the proportion of jobs which could feasibly be automated within different industries and within different regions. We also present the wages associated with occupations that could be automated, and the gender pay gap in the scenario in which they are automated. All calculations are based on the following sources of data:

- **Frey and Osborne (2013).** We use the authors' dataset of probability of 'computerisation' by US occupation code (SOC 2010). 'Probability of computerisation' refers to the likelihood that the occupation could feasibly be automated using technology that currently exists or 'near-term technological breakthroughs'. The authors' analysis considers the probability of an occupation being entirely computerised, or substituted, rather than complemented by partial automation. The analysis does not specify a timeline as factors other than technical potential, such as economic and political developments, are not considered.

  To use Frey and Osborne's data, we map US occupation codes onto UK four-digit standard occupational classification (SOC) codes, using official crosswalks between the two coding systems and ISCO-08, the international occupation classification. Where there are more US codes than ISCO codes, we create a weighted average using US Census data. For US codes that do not correspond to UK codes via ISCO, we examine the constituent tasks of the occupation to manually match them to UK codes. For a very small number of codes, probabilities of computerisation are unavailable; we assume these are equal to zero for the wage analysis.
- **The Quarterly Labour Force Survey (ONS 2017d)**. We derive counts of employment by four-digit occupation code using the quarterly LFS. Regional and industry analysis is based on two years of data, from Jan – March 2015 to Oct – Dec 2016[15]. Calculations for the wage analysis as well as the total proportion of jobs at risk are based on four quarters of data, April – June 2016 to Jan – March 2017.
- **Annual Survey of Hours and Earnings (ONS 2017c)**. We source average annual and hourly pay from the 2017 ASHE data set, provided by sex and occupation. We use the value for gross pay.

For the regional and industry analysis, we use job counts in the Labour Force Survey multiplied by probabilities of computerisation to estimate the proportion of jobs with potential for automation within each group of interest.[16] To assess the wages associated with these jobs, we use the mean gross annual income for each detailed occupational group and multiply by the number of jobs in each group vulnerable to computerisation. For the very small proportion of jobs for which we do not have probabilities of computerisation, we assume zero probability.

Wage totals are grossed using factors based on the difference between our estimates for total earnings in the economy and ONS estimates (OBR 2017a). We

15 This analysis was carried out in early 2017 for Dromey et al (2017).

16 This approach follows that taken by the Bank of England (2015). An alternative approach that other researchers including Frey and Osborne have taken is to identify a cut-off for 'high-potential' jobs, such as probability higher than or equal to 0.7, and to estimate the number of jobs and wages associated with these occupations, however, selecting a cut-off could be seen as arbitrary. Some results may differ from other researchers' due to this methodological choice.

treat employee and self-employed earnings differently, as ASHE data pertains to wages for employees, and has been shown to over-estimate self-employed earnings. Our baseline for employee wages is the OBR estimate for 2016/17 of wages and earnings (not including employer social contributions). Our baseline for self-employed earnings is equal to the OBR's estimate for mixed income multiplied by the labour share of GDP in 2016/17.

The gender pay gap is presented as the difference in gross hourly pay for male and female earners as a proportion of male hourly pay, in a baseline scenario of 2016/17 data and a counterfactual scenario in which occupations have been automated in proportion to their probability of computerisation. Our estimate of the gender pay gap in the baseline scenario differs slightly to the official estimates due to grossing by occupation, and because we use gross hourly pay including rather than excluding overtime.

References

Acemoglu D and Restrepo P (2016) 'The race between machine and man: Implications of technology for growth, factor shares and employment', *NBER*. http://www.nber.org/papers/w22252

Acemoglu D and Restrepo P (2017) 'Robots and jobs: Evidence from US labor markets', MIT Working Paper. https://economics.mit.edu/files/12763

Alperovitz G (2016) *Technological inheritance and the case for a basic income*. http://www.garalperovitz.com/2016/12/technological-inheritance-case-basic-income/

Andrews D, Criscuolo C and Gal P (2016) 'The global productivity slowdown, technology divergence and public policy: A firm level perspective', OECD. https://www.oecd.org/global-forum-productivity/events/GP_Slowdown_Technology_Divergence_and_Public_Policy_Final_after_conference_26_July.pdf

Andrews D, Criscuolo C and Gal P (2015) 'Frontier firms, technology diffusion and public policy: Micro evidence from OECD countries', OECD. https://www.oecd.org/eco/growth/Frontier-Firms-Technology-Diffusion-and-Public-Policy-Micro-Evidence-from-OECD-Countries.pdf

Arntz M, Gregory T and Zierahn U (2016) 'The risk of automation for jobs in OECD countries: A comparative analysis', OECD social, employment and migration working papers no. 189. http://www.oecd-ilibrary.org/socialissues-migration-health/the-risk-of-automation-for-jobs-in-oecd-countries_5jlz9h56dvq7-en

Autor D (2015a) 'Why are there still so many jobs? The history and future of workplace automation', *Journal of Economic Perspectives*, 29:3, pp. 3–30. https://economics.mit.edu/files/10865

Autor D (2015b) 'Polanyi's Paradox and the shape of employment growth', Federal Reserve Bank of St. Louis: Economic Policy Proceedings, *Re-evaluating Labor Market Dynamics*, pp. 129–177. https://economics.mit.edu/files/9835

Autor D and Salomons A (2017) 'Does productivity growth threaten employment? "Robocalypse now"?', presentation at European Central Bank Annual Conference, Sintra, Portugal 27 June 2017. https://www.ecbforum.eu/uploads/originals/2017/speakers/presentations/Autor-Salomons-ECB-Productivity-Talk-v2.0.pdf

Autor D, Dorn D, Katz L, Patterson C and Van Reenen J (2017) 'The fall of the labor share and the rise of superstar firms', MIT working paper. https://economics.mit.edu/files/12979

Autor D and Dorn D (2013) 'The growth of low-skill service jobs and the polarization of the US labor market', *American Economic Review*, 103:5, pp. 1553–1597. http://dx.doi.org/10.1257/aer.103.5.1553

Autor D and Price B (2013) 'The changing task composition of the US labor market: An update of Autor, Levy, Murnane (2003)', MIT Working Paper. https://economics.mit.edu/files/11600

Autor D and Acemoglu D (2011) 'Skills, tasks and technologies: Implications for employment and earnings', *Handbook of Labor Economics*. https://economics.mit.edu/files/5571

Avent R (2016) *The wealth of humans: Work and its absence in the twenty-first century*, Allen Lane

Baumol W (1967) 'Macroeconomics of unbalanced growth: The anatomy of urban crisis', *The American Economic Review*, 57:3, pp. 415-426

Berg A, Buffie E and Zanna L (2016) 'Robots, growth, and inequality', IMF, *Finance & Development*, 53:3, http://www.imf.org/external/pubs/ft/fandd/2016/09/berg.htm

Bessen J (2017) 'Automation and jobs: When technology boosts employment', Boston University Research Paper. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2935003

Bick A, Fuchs-Schündeln and Lagakos D (2015) *How do average hours worked vary with development? Cross-country evidence and implications*. https://www.cesifo-group.de/dms/ifodoc/docs/Akad_Conf/CFP_CONF/CFP_CONF_2015/isd15-Wohlrabe/Papers/isd15_Bick.pdf

Boppart T and Krusell P (2016) *Labor supply in the past, present, and future: A balanced-growth perspective*. http://perseus.iies.su.se/~pkrus/working_papers/pastpersentandfuture.pdf

British Standards Institute (2016) *Robots and robotic devices. Guide to the ethical design and application of robots and robotic systems*. https://shop.bsigroup.com/ProductDetail/?pid=000000000030320089

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

Burkhardt R and Bradford C (2017) 'Addressing the accelerating labor market dislocation from digitalization', Brookings Institute. https://www.brookings.edu/research/addressing-the-accelerating-labor-market-dislocation-from-digitalization/

Cabinet Office (2017) *Race disparity audit*. https://www.gov.uk/government/publications/race-disparity-audit

Calo R (2014) 'The case for a Federal Robotics Commission', Brookings Institute. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2529151

CEDEFOP (2015) *Annual Report 2015*, http://www.cedefop.europa.eu/en/publications-and-resources/publications/4145

CitiBank (2016) *Technology at work v2.0: The future is not what it used to be*. http://www.oxfordmartin.ox.ac.uk/downloads/reports/Citi_GPS_Technology_Work_2.pdf

Clarke S and D'Arcy C (2016) *Low Pay Britain*, Resolution Foundation. http://www.resolutionfoundation.org/app/uploads/2016/10/Low-Pay-Britain-2016.pdf

Commission on Economic Justice (CEJ) (2017) *Time for change: A new vision for the British economy. The Interim Report of the IPPR Commission on Economic Justice*, IPPR. https://www.ippr.org/research/publications/cej-time-for-change

Coote A (2014) *10 reasons for a shorter working week*, New Economics Foundation. http://neweconomics.org/2014/07/10-reasons-for-a-shorter-working-week/

Corlett A and Gardiner L (2015) *Looking through the hourglass: Hollowing out of the UK jobs market pre- and post-crisis*, Resolution Foundation. http://www.resolutionfoundation.org/app/uploads/2015/03/Polarisation-full-slide-pack.pdf

Cummine A (2016) *Citizens' wealth: Why (and how) sovereign funds should be managed by the people for the people*, Yale University Press

Dao M, Das M, Koczan Z and Lian W (2017) 'Drivers of declining labour share of income', IMF. https://blogs.imf.org/2017/04/12/drivers-of-declining-labor-share-of-income/

D'Arcy C and Gardiner L (2017) *The generation of wealth: asset accumulation across and within cohorts*, Resolution Foundation, http://www.resolutionfoundation.org/app/uploads/2017/06/Wealth.pdf

Dellot B and Wallace-Stephens F (2017) *The Age of Automation: Artificial intelligence, robotics and the future of low-skilled work*, RSA. https://www.thersa.org/globalassets/pdfs/reports/rsa_the-age-of-automation-report.pdf

Dromey J, McNeil C and Roberts C (2017) *Another lost decade? Building a skills system for the economy of the 2030s*, IPPR. https://www.ippr.org/research/publications/skills-2030-another-lost-decade

European Parliament (2017) *Recommendations to the Commission on Civil Law Rules on Robotics*. http://www.europarl.europa.eu/sides/getDoc.do?pubRef=-//EP//TEXT+REPORT+A8-2017-0005+0+DOC+XML+V0//EN

European Trade Union Institute (2017) *Financial participation in Europe: Overview of similarities, differences and trends*. https://www.worker-participation.eu/National-Industrial-Relations/Across-Europe/Financial-Participation

Executive Office of the President (2016) *Artificial intelligence, automation, and the economy*, https://www.whitehouse.gov/sites/whitehouse.gov/files/images/EMBARGOED%20AI%20Economy%20Report.pdf

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

Freeman R (2015) *Who owns the robots rules the world*, IZA World of Labor. https://wol.iza.org/articles/who-owns-the-robots-rules-the-world/long

Frey C and Osborne M (2015) *From brawn to brains: The impact of technology on jobs in the UK*, Deloitte. https://www2.deloitte.com/content/dam/Deloitte/uk/Documents/Growth/deloitte-uk-insights-from-brawns-to-brain.pdf

Frey C and Osborne M (2013) *The future of employment: How susceptible are jobs to computerisation*. http://www.oxfordmartin.ox.ac.uk/downloads/academic/The_Future_of_Employment.pdf

Furåker B (2015) 'The Swedish wage-earner funds and economic democracy: Is there anything to be learned from them?', *European Trade Union Institute*, 22:1, pp. 121-132. http://journals.sagepub.com/doi/abs/10.1177/1024258915619310

Furman J and Orszag P (2015) 'A Firm-Level Perspective on the Role of Rents in the Rise in Inequality'. http://gabriel-zucman.eu/files/teaching/FurmanOrszag15.pdf

Goos M and Manning A (2007) 'Lousy and lovely jobs: The rising polarisation of work in Britain', *The Review of Economics and Statistics*, 89:1, pp. 118-133. http://www.mitpressjournals.org/doi/abs/10.1162/rest.89.1.118?journalCode=rest

Goos M, Manning A and Salomons A (2010) 'Explaining job polarisation: Routine-biased technological C=change and Offshoring'. http://personal.lse.ac.uk/manning/work/ExplainingJobPolarization.pdf

Greenfield A (2017) *Radical Technologies: The Design of Everyday Life*, Verso

Haldane A (2017) 'Productivity puzzles', speech by Andrew Haldane, chief economist, Bank of England, London School of Economics, 20 March 2017. http://www.bis.org/review/r170322b.pdf

Haldane A (2015) 'Labour's Share', Speech by Andrew Haldane, chief economist, Bank of England, Trades Union Congress, London 12 November 2015. http://www.bankofengland.co.uk/publications/Documents/speeches/2015/speech864.pdf

Haskel J and Westlake S (2017) *Capitalism without Capital: The Rise of the Intangible Economy*, Princeton University Press

Health and Safety Executive (2017) *Work-related Stress, Depression or Anxiety Statistics in Great Britain 2017*. http://www.hse.gov.uk/statistics/causdis/stress/stress.pdf?pdf=stress

Hern A (2017) 'DeepMind announces ethics group to focus on problems of AI', *The Guardian*, 4 October 2017. https://www.theguardian.com/technology/2017/oct/04/google-deepmind-ai-artificial-intelligence-ethics-group-problems

HMG (2017a) *Made smarter*. https://www.gov.uk/government/uploads/system/uploads/attachment_data/file/655570/20171027_MadeSmarter_FINAL_DIGITAL.pdf

HMG (2017b) *Industrial strategy: Building a Britain fit for the future*. https://www.gov.uk/government/uploads/system/uploads/attachment_data/file/662508/industrial-strategy-white-paper.pdf

HMT (2017) *Autumn Budget 2017*. https://www.gov.uk/government/uploads/system/uploads/attachment_data/file/661480/autumn_budget_2017_web.pdf

International Federation of Robotics (2016) *Executive Summary World Robotics 2016 Industrial Robots*. https://ifr.org/img/uploads/Executive_Summary_WR_Industrial_Robots_20161.pdf

Jacobs M, Hatfield I, King L, Raikes L and Stirling A (2017) Industrial strategy: Steering structural change in the UK economy, IPPR. https://ippr.org/research/publications/industrial-strategy-steering-structural-change-in-the-uk-economy

Karabarbounis L and Neiman B (2014) 'The Global Decline of the Labor Share', *Quarterly Journal of Economics*, 129:1, pp.61-103. https://www.dropbox.com/s/t3ber11ic9xy9dx/labor_share.pdf?dl=1

Lansley S (2017) 'Reversing the inequality spiral: Citizens' Wealth Fund', *IPPR Progressive Review*, http://onlinelibrary.wiley.com/doi/10.1111/newe.12050/abstract

Lawrence M and Mason N (2017) *Alternative models of ownership* (title TBC), IPPR

Lawrence M (2016) *Future Proof: Britain in the 2020s*, IPPR. www.ippr.org/publications/future-proof-britain-in-the-2020s

Lessig L (2000) 'Code is Law', *Harvard Magazine*

Mazzucato M (2017) *Mission-Oriented Innovation Policy: Challenges and Opportunities*, UCL Institute for Innovation and Public Purpose. https://www.ucl.ac.uk/bartlett/public-purpose/publications/2017/oct/mission-oriented-innovation-policy-challenges-and-opportunities

Mazzucato M (2013) *The Entrepreneurial State: Debunking Public vs. Private Sector Myths*, Anthem Press

Mazzolari F and Ragusa G (2013) 'Spillovers from High-Skill consumption to Low-Skill Labor Markets', *The Review of Economics and Statistics*, 95:1, pp. 74-86. http://www.mitpressjournals.org/doi/pdf/10.1162/REST_a_00234

McKinsey Global Institute (2017) *Harnessing automation for a future that works*. http://www.mckinsey.com/global-themes/digital-disruption/harnessing-automation-for-a-future-that-works

McKinsey Global Institute (2015) *Unlocking the potential of the Internet of Things*. http://www.mckinsey.com/businessfunctions/business-technology/our-insights/the-internet-of-things-the-value-of-digitizing-thephysical-world

McKinsey Global Institute (2011) *Growth and renewal in the United States: Retooling America's economic engine*. http://www.mckinsey.com/~/media/McKinsey/Global%20Themes/Americas/Growth%20and%20renewal%20in%20the%20US/MGI_growth_and_renewal_in_the_US_full_report.ashx

Michaels G, Natraj A, and Van Reenen J (2014) 'Has ICT polarized skill demand? Evidence from eleven countries over 25 years', *Review of Economics and Statistics*, 96:1, pp. 60-77. http://eprints.lse.ac.uk/46830/1/Michaels_Natraj_VanReenen_Has-ICT-polarized-skill-demand_2014.pdf

Morozov E (2016) 'Tech titans are privatising our data', *The Guardian*. https://www.theguardian.com/commentisfree/2016/apr/24/the-new-feudalism-siliconvalley-overlords-advertising-necessary-evil

Nuffield Foundation (2017) 'Response to the 2017 budget announcement on Centre for Data Ethics and Innovation'. http://www.nuffieldfoundation.org/news/response-2017-budget-announcement-centre-data-ethics-and-innovation

O'Connor S (2016) 'When your boss is an algorithm', *The Financial Times*, 7 September 2016. https://www.ft.com/content/88fdc58e-754f-11e6-b60a-de4532d5ea35

O'Neil C (2016) *Weapons of math destruction: How Big Data increases inequality and threatens democracy*, Allen Lane

OBR (2017a) *Economic and Fiscal Outlook*. http://cdn.budgetresponsibility.org.uk/Nov2017EFOwebversion-2.pdf

OBR (2017b) 'November 2017 Economic and Fiscal Outlook – Supplementary economy tables', dataset. http://budgetresponsibility.org.uk/efo/economic-fiscal-outlook-november-2017/

OECD (2017) *OECD Employment Outlook 2017*. http://www.keepeek.com/Digital-Asset-Management/oecd/employment/oecd-employment-outlook-2017_empl_outlook-2017-en#.WhaojFVl_cs#page1

OECD (2015) *The Labour Share in G20 Economies*. https://www.oecd.org/g20/topics/employment-and-social-policy/The-Labour-Share-in-G20-Economies.pdf

ONS (2017a) 'Personal Incomes Statistics 2014-15', Office of National Statistics. https://www.gov.uk/government/uploads/system/uploads/attachment_data/file/608854/National_Statistics_T3_1_to_T3_11_publication_2014_15_revised.pdf

ONS (2017b) *Average actual weekly hours of work for full-time workers (seasonally adjusted)*. https://www.ons.gov.uk/employmentandlabourmarket/peopleinwork/earningsandworkinghours/timeseries/ybuy/lms

ONS (2017c) 'Occupation (4 digit SOC) - ASHE: Table 14', Annual Survey of Hours and Earnings, 2017. https://www.ons.gov.uk/employmentandlabourmarket/peopleinwork/earningsandworkinghours/datasets/occupation4digitsoc2010ashetable14

ONS (2017d) Social Survey Division and Northern Ireland Statistics and Research Agency. Central Survey Unit, (2015-2017) Quarterly Labour Force Survey, Quarters January–March 2015 to Jan–March, 2017, UK Data Archive

ONS (2015) *Wealth in Great Britain Wave 4: 2012 to 2014*. https://www.ons.gov.uk/peoplepopulationandcommunity/personalandhouseholdfinances/incomeandwealth/compendium/wealthingreatbritainwave4/2012to2014

PwC (2017a) *Will robots steal our jobs? The potential impact of automation on the UK and other major economies*. https://www.pwc.co.uk/economic-services/ukeo/pwcukeo-section-4-automation-march-2017-v2.pdf

PwC (2017b) *The economic impact of artificial intelligence on the UK economy*.

Polanyi M (1966) *The Tacit Dimension*, University of Chicago Press

Roberts C and Lawrence M (2017) *Wealth in the 21st century: dimensions and drivers*, IPPR. https://www.ippr.org/files/2017-10/cej-wealth-in-the-21st-century-october-2017.pdf

Schmidt F (2017) *Digital labour markets in the platform economy mapping the political challenges of crowd work and gig work*, Friedrich Ebert Stiftung. http://library.fes.de/pdffiles/

Spitz-Oener A (2006) 'Technical change, job tasks, and rising educational demands: Looking outside the wage structure', *Journal of Labor Economics*, 24: 2. Pp. 235-270. http://www.jstor.org/stable/10.1086/499972?seq=1#page_scan_tab_contents

Srnicek N (2016) *Platform Capitalism*, Polity

Srnicek N and Williams A (2015) *Inventing the Future: Postcapitalism and a World Without Work*, Verso

Susskind D (2017) 'A model of technological unemployment', University of Oxford, Department of Economics Discussion Paper. https://www.economics.ox.ac.uk/materials/papers/15126/819-susskind-a-model-of-technological-unemployment-july-2017.pdf

Susskind D and Susskind R (2015) *The future of the professions: How technology will transform the work of human experts*, Oxford University Press

Total Economy Database (2017) *Output, Labor, and Labor Productivity, 1950-2017* data series. https://www.conference-board.org/data/economydatabase/index.cfm?id=27762

TUC (2017) *Shaping our digital future*. https://www.tuc.org.uk/sites/default/files/Shaping-our-digital-future.pdf

Unger R (2013) 'Deep freedom: Why the Left should abandon equality', *IPPR Progressive Review*. https://www.ippr.org/juncture/deep-freedom-why-the-left-should-abandon-equality

Van Reenen J, Bloom N and Sadun R (2017) 'Management as a technology?', *NBER*. http://www.nber.org/papers/w22327

UK Commission for Skills and Jobs (2014) *The future of work: jobs and skills in 2030*. https://www.gov.uk/government/publications/jobs-and-skills-in-2030

UK Parliament (2017) *Algorithms in decision-making – publications*. http://data.parliament.uk/writtenevidence/committeeevidence.svc/evidencedocument/science-and-technology-committee/algorithms-in-decisionmaking/oral/73859.html

UK Parliament (2016) *Robotics and artificial intelligence inquiry – publications*. House of Commons Science and Technology Committee. https://www.parliament.uk/business/committees/committees-a-z/commons-select/science-and-technology-committee/inquiries/parliament-2015/robotics-and-artificial-intelligence-inquiry-15-16/publications/

World Economic Forum (2015) *Deep Shift: Technology Tipping Points and Societal Impact*. http://www3.weforum.org/docs/WEF_GAC15_Technological_Tipping_Points_report_2015.pdf