THE IMPACT OF ARTIFICIAL INTELLIGENCE ON WORK

An evidence review prepared for the Royal Society and the British Academy

September 2018

CONTENTS

ACKNOWLEDGMENTS

We would like to thank attendees to workshops on Artificial Intelligence and the future of work organised by the British Academy and the Royal Society, held in London on March 15 and May 9, 2018. Discussions and comments on early outputs from this review provided in these events have been an invaluable source of insight for this project.

We are grateful to members of a steering group from the British Academy and the Royal Society and to members of an external review group for helpful comments and suggestions.

We would also like to thank Andy Haldane, Chief Economist at the Bank of England, Stephen Machin, Professor of Economics at the London School of Economics, Geoff Mulgan, Chief Executive Officer of Nesta, and Richard Susskind, IT Adviser to the Lord Chief Justice of England and Wales, and chairman of the Advisory Board of the Oxford Internet Institute, for participating in interviews conducted by Frontier Economics between March and May 2018. The interviews were used to test the comprehensiveness of the evidence collected, our interpretation of the evidence, and to discuss priorities for public policy.

We extend our gratitude to Sir Richard Blundell, David Ricardo Professor of Political Economy at University College London, who provided expert advice throughout this project and comments on drafts of this report.

EXECUTIVE SUMMARY

Frontier Economics has been commissioned by the Royal Society (RS) and British Academy (BA) to produce a review of evidence on the potential impact of Artificial Intelligence (AI) on work in the UK up to 2030.

In this review, we examine evidence from a wide a range of disciplines to inform the discussion on the impact of AI on work. The report:

- Draws insights from earlier periods of technological change and the more recent evidence on the impact of digital technology
- Presents ***what*** claims have been made about the potential consequences of AI for the future of work, and
- Analyses ***why*** and ***how*** such claims have been made with a focus on which frameworks have been used, and what assumptions (implicit or explicit) any conclusions rest on.

Evidence on the impact of existing technology on work

Technological change is a key driver of economic growth. However, the invention, diffusion and effective use of new technology are in turn likely to be influenced by other factors, including economic conditions, institutions, and social conditions. For example, the adoption of labour-saving technology in the first industrial revolution may have been driven by the specific economic conditions of 18th century. Moreover, sociological research has shown that patterns of work organisation vary across countries even between establishments that use very similar technologies.

It is broadly accepted that the first industrial revolution eventually led to improving standards of living in society and in particular for the working class, but there is evidence to suggest that it took some time for these improvements to materialise.

In the context of the first industrial revolution, technological change was also linked to changes in the nature of work: the mechanisation of textile production involved work moving from artisans' homes to the factories, from rural to urban areas, and from independent work often filling downtime in rural work to full-time, predictable work in a hierarchical structure.

Recent automation does not seem to have led to overall decline in employment levels but there have been income losses for low-educated workers employed in the manufacturing sector. Employment losses in manufacturing have typically been compensated by increasing employment in services, leading to stable or growing overall employment levels. It is not clear from the literature included in this review whether the aggregate figures conceal employment losses for specific demographics and how the new service jobs compare to the manufacturing jobs lost in terms of opportunities for progression, security, quality of working environment.

Evidence on AI from theoretical economic models

Recent theoretical work in economics aims to provide a framework to specifically understand the impact of AI on employment, including both its immediate, first-order effects, and following, second-order effects. The main findings from this literature suggest the following:

- A number of factors can counterbalance initial declines in labour demand due to automation. As automation increases productivity (leading to better or cheaper products), increasing consumer demand, greater investment, and innovation can lead labour demand to rise.
- In the short term, it is not clear whether countervailing effects will be sufficient to offset potential job losses from automation. Even if they are, transitions could be challenging. But new jobs could be generated, in principle, in the same industries where automation is taking place.
- In the long term, as production processes are re-organised, countervailing effects are expected to become stronger and fully compensate the initial decline in work demanded by businesses that have adopted AI to automate production. However, workers who have been directly displaced could experience a fall in their earnings relative to other workers (and potentially in absolute terms). This would increase inequality if the displaced group is mostly composed of low earners.
- The competitiveness of product and labour markets is important to drive better outcomes for workers. The results above typically assume that product and labour markets are competitive, that is, consumers can choose between different products and workers can choose between different employers. This is not always the case, as acknowledged in some of this literature. In particular, lack of competition in labour markets would mean that productivity benefits from automation flow into greater profits rather than into higher wages. There is some concern that digital technology may enable large firms to increase and maintain their market power. There is on-going research on this but evidence is still limited.

These models suggest that automation can lift living standards for all, but this is not necessarily true in the short term, and even in the long term inequality may increase as some workers benefit more than others. This is a departure from earlier economic theory, where technological progress was typically seen to benefit all workers under a broad set of conditions. It is worth noting that the nature of these models is to consider second-order effects of automation driven by price changes (the ‘productivity effects’ and impacts on ‘consumer demand’ mentioned above), but not to consider possible social or institutional changes and the effect these may have, in turn, on economic outcomes.

Evidence on the adoption of AI

This review has identified limited evidence on the adoption and diffusion of AI, though it has been noted that neither AI or other forms of digital technology appear to have led to significant productivity improvements in recent years. However, surveys of business leaders and predictions consider the following drivers of business adoption of AI:

- The profitability of investing in AI: technical feasibility will be a necessary condition, but AI will only be adopted when the expected revenues exceed the costs from adoption;
- The ability to re-shape organisations and processes in the ways needed to take advantage of the new technology;
- Regulatory enablers and constraints: adopting AI in a way that would ensure compliance with relevant regulations may not be possible; vice versa, regulation may enable adoption, for example by setting necessary standards;
- Social attitudes and preferences over the role of machines. In many settings, people may require interaction with a human even where a cheaper automated alternative could be available. This could be the case, for example, in the case of health and social care work.

The role of social factors in influencing the development and adoption of AI has been given less consideration compared to drivers of business adoption. But, insights from sociological research caution against considering technology as an autonomous force driving economic and social change. Research on earlier technological development suggests that social and technical factors are closely linked. For example, the choices around specific technical solutions depend on whether those solutions will be accepted by relevant social groups.

Quantified predictions about potential job losses and gains related to AI

Several studies aim to estimate the proportion of current jobs that could technically be automated in the future. These estimates aim to provide a sense of the scale of potential transformation that could be enabled by technology – specifically, the proportion of current employment for which the immediate, first-order effect of AI adoption could be radical transformation or displacement. This research relies on a small number of assessments of the technical automatability of existing work tasks, coupled with data analysis to investigate which job characteristics (for example, a requirement to interact with customers) are correlated with the assessed automatability. Two of these assessments are based on expert opinion, and a third uses a questionnaire disseminated on a crowdsourcing platform to score work tasks in terms of their suitability for machine learning, against pre-determined criteria.

Studies in this literature find that between 10% and 30% of jobs in the United Kingdom could be automated at some point in the coming decades. However, a large majority of jobs involve performing at least some tasks that have been assessed as not suitable for automation.

The potential impact of automation would depend not only on its scale but also on its distribution across the workforce and its timing. There is broad consensus in this literature that jobs typically performed by low-educated workers are at relatively high risk, except where they involve complex interaction with others. However other studies expect that significant transformation will be required across virtually all jobs, while employment losses will be limited.

There is significant uncertainty around when the potential automation described in this literature would be technically feasible and widely adopted.

Foresight exercises suggest that a number of trends could generate sufficient job creation to compensate for first-order job destruction linked to automation – even in the absence of any job creation from automation. However, the ‘new jobs’ may not have the same characteristics or emerge in the same places as the ‘old’ that are destroyed, and therefore adjustments may be challenging. The literature in this area suggests that occupations that are likely to grow between now and 2030 are disproportionately high-education, although some middle-education occupations are also likely to grow.

Evidence from the application of digital technology in specific settings and industries suggests that this is leading to changes in the organisation and quality of work:

- There are reasons for concern (e.g. potential loss of autonomy for many workers) and for optimism (e.g. automation of hazardous tasks);
- Evidence on the likely balance between the two is limited – outcomes are also likely to depend on the actions of employers, workers, governments, and other stakeholders.

1 INTRODUCTION

Frontier Economics has been commissioned by the Royal Society (RS) and British Academy (BA) to produce a review of evidence on the potential impact of Artificial Intelligence (AI) on work in the near and medium term – up to 2030.

This review is part of a wider programme of work undertaken by the RS and BA, which is producing a rigorous evidence synthesis, to establish areas of uncertainty and consider how evidence from across disciplines can best inform policymaking.

Recent developments in the capabilities of AI systems and their application in robotics have created considerable interest in the potential effects of AI on the economy and society. Recent reports have provided an overview of the likely long-term developments in AI (Stone, Brooks, & Brynjolfsson, 2016; Hall & Pesenti, 2017) and of their potential impact on work (United States National Academy of Sciences, 2017; Executive Office of the United States President, 2016). And, in the United Kingdom, a House of Lords Select Committee on AI was appointed in 2017, to 'consider the economic, ethical and social implications of advances in artificial intelligence, and to make recommendations'. The Committee has reported on its enquiry in April 2018 (House of Lords, 2018).

This report reviews evidence from a wide a range of disciplines to inform the discussion on the impact of AI on work. The report:

- Draws insights from earlier periods of technological change and the more recent evidence on the impact of digital technology
- Presents ***what*** claims have been made about the potential consequences of AI for the future of work, and
- Analyses ***why*** and ***how*** such claims have been made with a focus on which frameworks have been used, and what assumptions (implicit or explicit) any conclusions rest on.

The types of impact this review focusses on are the following:

- What will work in 2030 look like?
  - How much work will be done? How many people will work and how many hours will they work on average?
  - What type of work will be done? Which of the existing occupations will increase, decrease, or be radically transformed?
  - What skills will be required?
  - How will work be organised?
  - How much will workers earn? How will that compare to the cost of goods and services people will need or want to buy?
  - How will all of this, along with broader social changes linked to the use of AI, affect people's wellbeing?
  - How will this vary across places and across personal and socioeconomic characteristics, from formal education levels to gender, age and other characteristics?

- How quickly will changes happen? If change is faster than in the recent past, will this be a one-time adjustment or will change be 'the new normal'?
- If a significant number of workers are displaced, what impact will this have on their earnings and on their wellbeing?

The rest of this report is structured as follows:

- Chapter 2 describes the approach followed to gather the interdisciplinary evidence base this review draws and the characteristics of this evidence.
- Chapter 3 considers technological development from a historical and theoretical perspective providing a high-level framework for considering the development of AI and its likely impacts.
- Chapter 4 examines recent quantitative and qualitative research that focuses on specific predictions and analysis related to ongoing investment in AI.
- Chapter 5 offers a set of conclusions to help structure the ongoing debate on AI.

Box 1. Definitions of AI and related technology used in this report

This review does not adopt a specific definition of AI but rather includes evidence that explicitly refers to AI, and other evidence that was considered to be relevant for the impact of AI on work according to a research protocol - further details on this in Section 2 and Annexe 1 of this report.

- This report also uses the following terms, related to AI:
  - 'Automation', which can be defined as 'the technique, method, or system of operating or controlling a process by highly automatic means, as by electronic devices, reducing human intervention to a minimum'.[1] AI can be thought of as one of the 'techniques, methods or systems' that can enable automation.
  - 'Robots', 'industrial robots', and 'Advanced Industrial Robotics'. Robots, which can be defined as 'machines capable of carrying out a complex series of actions automatically', have been used in industrial processes for some time, and evidence on the effect of their introduction on work (reviewed in Section 3 of this report) could be useful to inform predictions about the potential future effect of AI. Moreover, recent advances in AI are expanding the capabilities of robots.[2]
  - 'Information and Communication Technology' (ICT) refers to computing and communication hardware (e.g. computers, telephones, hard disk drives) and software (e.g. computer programs, mobile applications). This term is used in literature concerning the adoption of ICT in businesses and in the workplace between the late 1980s and early 2000s (reviewed in Section 3). 'Digital technology' is used as a synonym for ICT.

[1] National Academy of Sciences (2017).

[2] See for example House of Commons Science and Technology Committee (2016) for a recent overview of the links between robotics and AI.

- ‘Digitalisation’ refers to the adoption of ICT in work and society at large. In the literature reviewed in this report, the term is more frequent in more recent publications, and therefore may include the adoption of smartphones, the use of algorithms to regulate production processes (see Section 4 of this report), and the use of AI, among other devices and techniques.

2 APPROACH TO THE EVIDENCE BASE

Summary

- This review has collected over 170 relevant English-language documents published since 2000, across a wide range of disciplines.
- The 160 documents include: articles published in peer-reviewed journals, academic manuscripts, and reports published by public sector organisations, international organisations, think-tanks and consultancies since 2000.
- To ensure that the review would focus on the most relevant issues and that it would be completed within the agreed time frame, a short list of 450 documents to be reviewed in detail was selected from the long list of 170. This short list of 50 publications includes:
  - Evidence on historical and recent effects of technology on work, and theoretical frameworks for considering the future impacts of AI
  - Specific quantified predictions relating to future impacts of AI, and
  - Specific non-quantified predictions and analysis of AI impacts.
- The potential impact of AI on employment has drawn considerable attention from researchers in economics, but valuable insights can be gathered from other disciplines, from sociology and history to data science.
- Evidence from the selected publications has been integrated with insights from two workshops organised by the RS and BA, attended by researchers, policymakers and other stakeholders, and by a small number of interviews.

This evidence review draws on a wide range of disciplines to inform the on-going discussion on the impact of AI on work. In collecting, selecting, and reviewing relevant evidence we were guided by the principles for good evidence synthesis for policy developed by the RS (Royal Society, 2018). The review followed a research protocol, agreed with the RS and BA, that included a mapping of relevant topics ('topic map'), drawn in concentric circles from most directly relevant (at the centre of the circle) to least directly relevant (on the periphery). Based on this protocol, contributions from any discipline that would match the topic map would be included in an initial long list of relevant evidence. The map was slightly updated throughout this evidence collection phase to reflect the results of our search for evidence. An updated topic map is presented in Figure 1 overleaf.

The evidence collection phase of the review identified over 170 relevant articles and reports, covering a wide range of disciplines (including, but not limited to: Anthropology, Computer Science, Data science, Economics, History, Sociology, Management studies, and Organisational studies).

Figure 1 Map of relevant topics

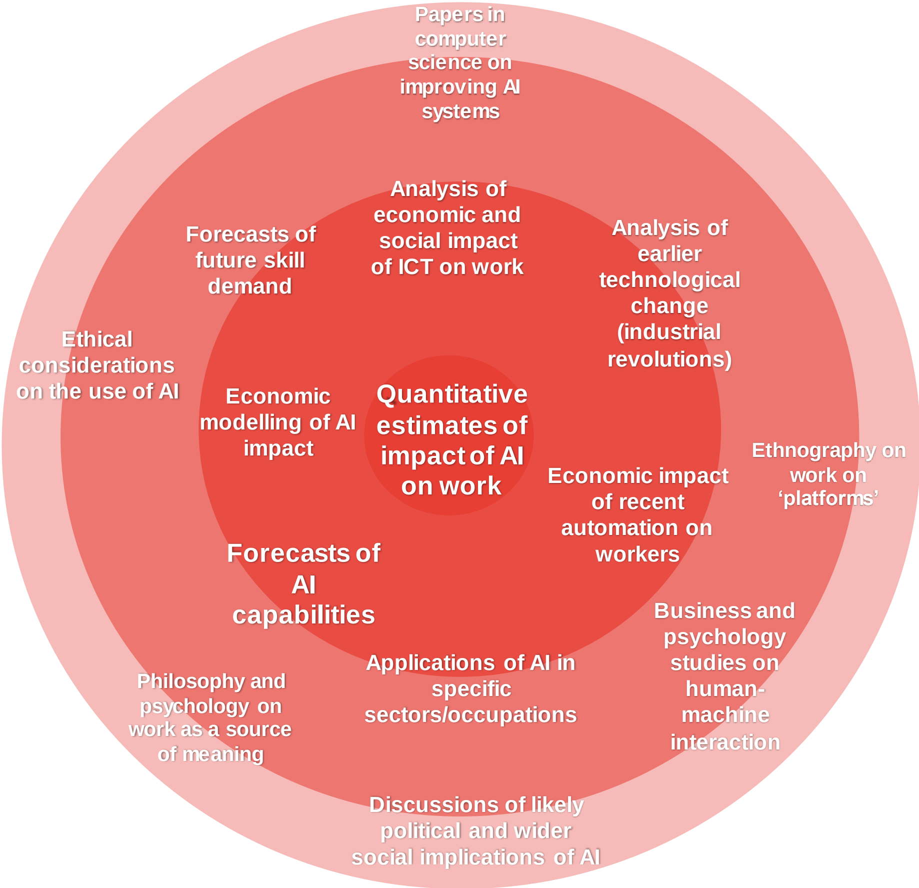

From the long list of over 170 publications, we agreed with the RS and the BA a short list to be reviewed in detail. This short list was then integrated with additional articles following comments from interviewees and further feedback from the RS and BA. The short list includes three types of evidence:

1. Evidence on historical and recent effects of technology on work, and theoretical frameworks for considering AI's future impacts. This literature considers:
   - The impact on the labour market of the first industrial revolution in Great Britain;
   - Adoption of computers and other types of digital technology ('ICT') and recent industrial automation, and;
   - Theoretical models considering the future impact of AI on employment and wages.
2. Specific quantified predictions relating to future impacts of AI, including:
   - Estimates of the proportion of existing jobs that may be lost due to AI-enabled automation ('gross job losses' due to AI); and
   - Estimates of the future composition of employment in terms of occupations and education levels.
3. Specific non-quantified predictions and analysis of AI impacts, including:

- Analysis on the future transformation of existing occupations (e.g. service industries, professional occupations); and
- Identifying the characteristics of new jobs linked to the adoption of AI.

Figure 2 below reports the number of publications included in this review by evidence type, topic and relevant disciplines.

Figure 2 Interdisciplinary evidence

Evidence type	Topic	Relevant disciplines	Number of publications
Evidence on recent and earlier effects of technology on work	Industrial revolutions	History	8
	Adoption of ICT and recent industrial automation	Anthropology, economics, sociology	8
Theoretical approaches	Economic modelling	Economics	9
Quantified predictions	Gross job losses	Data science, economics, engineering	8
	Future skill needs	Data science	2
Non-quantified predictions and analysis	Future transformation of existing occupations	Law, management, sociology, economics	3
	Interaction of social and technical change	Sociology	2
	Identifying possible new jobs linked to the adoption of AI	Data science, management, economics	1
	Human-machine interaction and algorithmic management	Social sciences	6
Cross-cutting			3
Total			**50**

History, economics, and mixed approaches combining data science with other disciplines are most frequently represented in this short list. This is a result of the assessment of relevance presented in Figure 1 above, of the emphasis of this review on the relatively short term effects of AI (up to 2030), and of the level of interest the potential impact of AI has drawn in economics research. However, valuable insights can be gathered from other disciplines, including law, sociology, and other social sciences.

In order to keep the scope of this review manageable within the available time frame, this report could not cover in detail a number of potentially relevant topics, including, for example:

- The potential for AI to transform a number of specific sectors (e.g. healthcare) and the ethical concerns around this potential transformation;
- Critiques of social and technological determinism in the sociology of technology literature;

- The characteristics of specific regions, cities, or local areas that may be linked with different adoption of technology and different consequences of that adoption.

Moreover, the relatively rapid nature of the review implies that it can only include a small proportion of the rich literature on historical technological change (in particular, the first Industrial Revolution).

It is worth noting that different disciplines tend to consider a different time frame and focus on different aspects of the adoption of or adaptation to technology, as illustrated in Figure 3 overleaf. Historical research allows us to look at the very long term picture on the impact of technology. Adapting from the discussion in Mokyr, Vickers & Ziebarthm (2016) and other literature, the long term picture includes the following categories:

- **The initial conditions** that determine the evolution of technology, institutions, social norms, culture;
- **The development and adoption of technology**, influenced by the initial conditions above;
- **The 'First order' effects** of adoption: immediate impact on work;
- **The 'Second order' effects**: longer-term impact, including firms fully reorganising production, consumer responses to any effects on the price and quality of goods, social responses, and;
- **The long-term,** including substantial social and institutional changes.

Figure 3 Time frames considered in different types of evidence

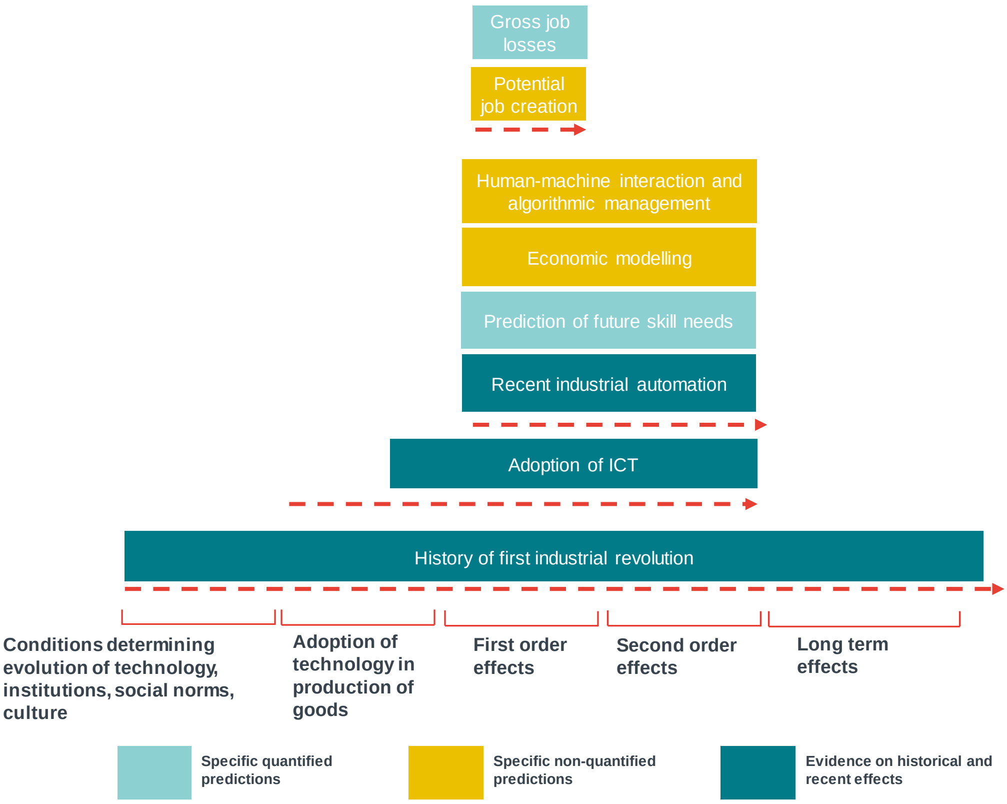

Source: Frontier Economics analysis of selected literature.

The evidence on recent history and theoretical modelling take into account both first and second order effects, with disciplines focusing on different aspects. For example, sociology has a greater emphasis on interaction between employers and workers as groups, while the economic literature in this context typically explains changes in employment and wages as a result of the choices made by employers to maximise profits.

Specific quantified predictions focus mainly on first order effects of technology on the importance of current occupations. Specific qualitative analysis focuses on other types of first order effects, including how occupations may be transformed by technology, and how non-economic factors may shape the adoption of technology.

Further detail on the methodology of the review can be found in Annex 1 of this report.

3 EVIDENCE ON THE IMPACT OF TECHNOLOGY ON WORK

Summary

- Historical analysis of technological changes highlights that while innovation is key to economic growth, technological change is not independent of prior economic, social, cultural conditions. These can hinder or foster innovation, and direct the way in which technology is used.
- In the context of the first industrial revolution, technological change was also linked to changes in the nature of work: the mechanisation of textile production involved work moving from artisans' homes to the factories, and from rural to urban areas. Textile production in factories involved predictable work in a hierarchical structure rather than independent work often filling downtime in rural work.
- The adoption of digital technology in advanced economies between the 1980s and the 2000s has been linked to 'job polarisation': an increase in the shares of job roles requiring low and high levels of formal education, at the expense of those in the middle. This may have been affected by the specific conditions of the period, including trade liberalisation and declining union membership.
- Recent industrial automation has been linked with falling employment and earnings in manufacturing for workers with low-medium levels of formal education. However, overall employment levels in areas affected by industrial automation have often remained constant, as job losses in manufacturing were compensated by job creation in service industries. It is not clear, from the literature reviewed as part of this project, how the services jobs created compare with the manufacturing jobs destroyed in terms of earnings and other characteristics.
- Recent theoretical work in economics models the potential consequences of AI-enabled automation, showing that:
  - A number of countervailing effects can counterbalance initial declines in labour demand due to automation. As automation increases productivity (leading to better or cheaper products), increasing consumer demand, greater investment, and innovation can lead labour demand to rise.
  - In the short term, it is not clear whether countervailing effects will be sufficient to offset potential job losses from automation. Even if they are, transitions could be challenging. But new jobs need not be in different industries compared to those that decline due to automation.
  - In the long term, as production processes are re-organised, countervailing effects are expected to become stronger and fully compensate. However, workers who have been directly displaced could experience a fall in their earnings relative to other workers (and potentially in absolute terms). This would increase inequality if the displaced group is mostly composed of low earners.

- Even if all workers benefit from automation there could still be distributional issues as an increasing proportion of income flows to owners of capital.
- More radical innovation is not necessarily worse for workers as a whole– the greater productivity gains, the greater the potential for increasing demand to offset displacement from automation.
- The way in which consumers will respond to any productivity gains (at a global level) will be important in dictating which new jobs will be created, and where they will be located.
- Workers are more likely to benefit from AI if markets are competitive. There is emerging evidence that the competitiveness of several markets has declined in recent years. Public policy interventions that foster competition, if needed, could help steer markets towards better outcomes for workers.

3.1 Historical accounts of technological change

In this Section, we provide an overview of findings from selected researchers who have written about industrial history prior to the 1980s. A majority of these findings relate to the first industrial revolution, taking place in Great Britain between around 1780 and 1830 (hereinafter, for simplicity, 'the industrial revolution').[3]

Our goal is to draw key messages from the research of scholars who have written about potential lessons from history for the future of work (Mokyr, 2016; Allen, 2017), about lessons from history on the contribution of technology to economic growth (Crafts, 2010), or whose findings appeared to be particularly relevant based on our reading of the researchers above and our engagement with researchers and policymakers (Broadberry & Wallis, 2016; Clark, 2005, 2014; Humphries & Weisdorf, 2015; Humphries, 2016).

The role of non-technological factors in shaping innovation and economic growth

All the articles reviewed here discuss technical change in a wider institutional, social, cultural and economic context.

Crafts (2010) discusses how incentive structures shape the invention, diffusion and effective use of new technologies. In economic history and economics, technological progress is seen as the key driver of economic growth and improvement in living standards. However, whether technological progress does happen and lead to economic growth may be determined by other initial conditions, that may include:

[3] Goldin (2017) argues that we should look for lessons from an earlier period of change – the European Renaissance, a time of increasing availability and use of information technology (paper, the printing press) and potential democratisation of knowledge that was followed by unrest and conflict. The comparison appears to be echoed in Ferguson (2017), but we did not identify other publications that built on this idea.

- **Economic conditions**. For example, the high labour costs and low energy costs of 18th century England may explain why labour-saving technology was developed and adopted, sparking the industrial revolution happened (Allen, 2009). With other economic conditions, different types of technology may have emerged, or the incentive to innovate may have been lower.[4]
- **Institutions**. Institutions that enforced property rules, constraining the power of ruling elites, may have encouraged technological progress, shaping both the invention and diffusion of new technologies. Under these institutions, gains from technological progress would not be extracted by the elites but could benefit wider society more than under other institutions.
- **Social conditions**. For example, it has been argued (Abramovitz & David, 1996) that the specific conditions prevailing in Europe after World War II led to industrial relations or forms of government intervention that could support high investment in technology adoption. Clark (2005) also points to research on individual, cultural, social changes that led to greater investment in knowledge (Galor & Moav, 2002).

Research on economic, institutional and social conditions is central in the on-going debate on what explains the timing and location of the industrial revolution, from which a consensus has not emerged (Clark, 2005). Fouquet & Broadberry (2015) provide new evidence to inform this debate, showing that economic growth was already a feature of history before the industrial revolution, but periods of growth were always followed by periods of decline that largely cancelled the benefits of growth. According to Fouquet & Broadberry (2015), then, advances in economic development in Europe from the 19th century are largely due to a fall in the frequency of decline periods, more than to an increase in the frequency or intensity of growth periods.

Broadberry & Wallis (2016) argue that this is explained by a shift in institutional arrangements, broadly defined, consistent with the theory put forward by North, Wallis & Weingast (2009). In their account, pre-industrial societies were generally based on a set of rules that stipulated the relations between specific groups. These rules were inherently based on identity – they would assign different obligations and privileges to those who belonged to different groups. When economic shocks happened, altering the relative economic ranking of different groups, institutions came under tension and this often led to conflict and economic decline. Only when institutions started including 'impersonal' rules – that is, rules that would apply to everyone regardless of their status or belonging to a particular group, could this cycle be broken. Wallis (2011) mentions the emergence of the contemporary business corporation as an example of how western societies shifted to impersonal rules. The first business corporations appeared in the 16th century as privileges (often including monopolies over a certain activity) granted to a particular group within the dominant elite. Only in the 19th century general incorporation acts 'allowed a corporation to be formed through a simple administrative act that did not require explicit permission of a legislative or government body'.

[4] Note that this is only one of many theories that have been put forward to explain why the industrial revolution took place at a particular point in time and specifically in Great Britain.

The consequences of technological change on living standards

An important question is when living standards for the working class started rising after the industrial revolution. Several recent contributions to the discussion on the impact of AI discuss whether 'this time will be different': that is, whether technological change will be linked with widespread improvement in living standards, without causing substantial unemployment, as has been the case in history so far.[5]

Recent evidence that this increase in living standards may only have happened after 1850, later than had been previously thought (Allen, 2009; Temin & Voth, 2005) may give us pause – even if 'this time' (the time of technological change linked to AI) 'is not different' there may still be reason for concern about how changes in the economy may affect workers, particularly at the low end of the income distribution. However, to the best of our knowledge, the literature provides mixed findings on this issue. It is also worth noting that a lack of improvement in living standards at a time of population growth was a better outcome than what had generally happened in the past in similar situations – that is, a fall in living standards akin to those described in Fouquet & Broadberry (2015).

Mokyr et al. (2016) describes some of the thinking by contemporaries of the industrial revolution on its possible consequences. When it came to effects on living standards:

- Optimists (James Stuart, David Ricardo) nevertheless had concerns around short-term displacement;
- Pessimists did not think technological advances would benefit society, but thought that adopting technology would still be the best choice in a context of international competition (William Mildmay).

Other concerns were expressed around the nature of the new forms of work that were emerging. There were worries that industrial work, with its repetition and organised discipline, would be morally inappropriate, compared to the much greater variability of pre-industrial work. Mokyr also notes that 'while people today worry about the exact opposite phenomenon', the separation between home and work was 'cause of great anxiety'.

The consequences of technological change on working conditions and different demographics

Humphries (2016) and Mokyr et al. (2016) set out how economic, social and cultural changes are deeply linked, and suggest that some of the characteristics of work we may give for granted are in fact relatively new. Technological change came with the move of textile production from the home to the factory, which implied:

- Workers now being monitored and placed in a hierarchical structure;
- A new separation of place of work from home and place of leisure;

[5] See for example Furman (2016), Elliott (2017).

- An increase in the predictability of work, compared to the earlier situation where textile production often complemented (often unpredictable) agricultural work.

Research on more recent periods shows that the organisation of work in the same industry, using similar technology, can be very different in different countries. This has been shown in the context of process automation that emerged in a number of industries (including oil refining and chemical industries) in the 1960s (Gallie, 1978), and of the automotive industry in the 20^{th} century (Jurgens, Malsch & Dohse, 1993).

This suggests that technological change has been accompanied by related changes in the organisation of work. These changes are not entirely determined by intrinsic characteristics of the new technology, but they are also linked to other social and cultural factors

Moreover, Humphries & Weisdorf (2015) points to differences in the impact of these changes across different demographic groups. Work as textile artisans was a source of independence for many married women, even where this work involved relatively limited hours and earnings. Industrialisation and the separation of work and home limited their opportunities to access relatively well remunerated casual work. At the same time, the wages of young unmarried women, who were more likely to be able to accept full-time work in industry, tracked those of young unskilled males.

These considerations are highly relevant for our time. Section 4.4 reviews selected literature on the changing organisation of work through the use of digital 'platforms', and several of the publications we have identified describe the characteristics of work 'crowdsourced' through Amazon Mechanical Turk – a leading service that defines itself as 'operating a marketplace for work that requires human intelligence' (Yin, Gray, Suri & Wortman Vaughan 2016, Gray, Suri, Ali & Kulkarni, 2016).[6] As we will see in Section 4, men have been traditionally over-represented in many of the occupations estimated to be at high risk of automation due to AI. On the other hand, predictions reviewed in Section 4.3 of this report (Bakhshi, Downing, Osborne & Schneider, 2017; McKinsey, 2017) show that many occupations in social care and education, typically dominated by women, are expected to become increasingly important in the UK and other advanced economies.[7]

3.2 Recent technological change and its impact

3.2.1 The impact of digitalisation on the structure of the labour market

A rich literature documents changes in the structure of employment in advanced economies between the 1980s and the 2000s that have been described as 'job polarisation'. 'Job polarisation' refers to a decline in the number of 'middle-education' jobs (typically requiring an upper secondary education or a vocational qualification), along with an increase in the number of 'high-education' jobs

[6] https://www.mturk.com/

[7] For further information on the gender composition of different occupational groups in the UK, see for example Office for National Statistics (2013).

(requiring a higher education qualification) and 'low-education' jobs (requiring limited formal education qualifications). At the same time, there has also been an increase in the inequality of earnings – specifically, an increase in the distance between high earners (the top 10%) and other workers.

There is a broad but not universal consensus in the economics literature that job polarisation is explained in large part by the increasing importance of Information and Communication Technology (ICT).

The mechanism through which ICT leads to polarisation has been explained using a 'task-based' model of the economy (Autor, Levy & Murnane, 2003, Acemoglu & Autor, 2011). This model considers a starting situation where workers are assigned to jobs as follows: low-educated workers perform mainly routine cognitive tasks and non-routine manual tasks; middle-educated workers perform mainly routine cognitive tasks; high-educated workers perform mainly non-routine cognitive tasks. The model assumes that ICT substitutes for workers in performing routine tasks, while complementing cognitive non-routine tasks. This means that:

- Middle-educated workers, who tend to perform routine cognitive tasks, are displaced.
- Middle-educated workers move towards low-education occupations, dampening wage growth for low-educated workers. This is because high-educated workers have a strong comparative advantage over middle-educated workers in performing non-routine cognitive tasks, and therefore middle-educated workers will not be able to move into high-education occupations.
- Demand for high-educated workers, who are complemented by the technology that can now perform routine cognitive tasks, increases.

This model is consistent with evidence on polarisation described above, where low- and high-education occupations grow at the expense of middle-educated occupations, and high-education earnings grow relative to other workers'.

Empirical work has generated evidence that the increasing role of ICT in production processes does indeed help explain job polarisation. In the UK, Goos & Manning (2004) find evidence of job polarisation in the UK between 1975 and 1999 having the following consequences[8]:

- An increase in high-education jobs, observed across industries;
- A fall in middle-education jobs, due both to a decrease within industry (e.g. machine operatives becoming less important in manufacturing) and to a shift of employment away from manufacturing towards services;
- An increase in low-education jobs, due both to a rise within industry and between industries (again reflecting a shift of economic activity towards services).

Although there is evidence that job polarisation is explained by the impact of ICT on the economy, there are other explanations, not necessarily mutually exclusive:

[8] Goos & Manning (2004) is part of a wider literature testing the hypothesis that ICT explains the job polarisation phenomenon in the UK, US, and other advanced economies (see for example Michaels, Natraj & Van Reenen, 2014 for international evidence on 11 countries within the OECD). A full review of this literature is beyond the scope of this review.

- Krugman (2008) notes that increasing imports from low-wage countries could also be the main explanation for increasing inequality between high-educated and low-educated workers in the United States. These imports may have substituted low- and middle-educated manufacturing work that used to take place in the United States. This would have led to lower earnings growth for low- and middle-educated workers in the United States, now effectively faced with greater competition from workers abroad. Bloom, Draca & Van Reenen (2011) show evidence that competition from Chinese imports encouraged firms in the United States to adopt technology. Trade would have therefore accelerated the technology-driven process of polarisation described above.
- Dwyer (2013) attributes part of the increase in the share of low-educated jobs in the United States to social changes that have led to increasing demand for care work (defined to include teaching, health care, child care, elder care, domestic service, and housekeeping). Population ageing and the ‘professionalisation’ of care work, now exchanged on the market rather than provided unpaid by women in the household, not technology, explain this increasing demand.
- Moreover, Dwyer (2013) notes that other researchers (sociologists, economics, geographers) have generated evidence supporting potential ‘institutionalist explanations’ of job polarisation – changes in institutions that ‘shifted the balance of power between capital and workers’. Research in this area notes that the way in which technological change affects work is mediated by institutions, and that job polarisation has been less pronounced where economic policy has protected middle-wage jobs, maintained unions, and encouraged new industrial development rather than only service-sector growth’. Dwyer (2013) also discusses potential ways in which institutions could upgrade low-paid, low-education care jobs– greater collective action, clearer progression paths between low-education and middle-education jobs defined by public policy or industrial relations.

It is worth noting that the ‘institutionalist’ findings reported in Dwyer (2013) are not entirely inconsistent with economic research on the conditions under which ICT investment has led to productivity growth. An extensive review of this literature (Van Reenen, Bloom, Draca, Kretschmer & Sadun, 2010) concluded that stricter employment protection regulation and, to a lesser extent, stricter product market regulation are linked with a lower impact of ICT on productivity. It is possible that when it comes to institutional arrangements linked to the adoption of ICT, there has been a trade-off between overall productivity growth and job polarisation - both enhanced by a lower degree of employment protection (conversely, a higher degree of labour market flexibility).

In parallel with job polarisation, it has also been observed that high-educated workers have increasingly clustered geographically in the United States and in other advanced economies (e.g. Diamond, 2016). High-educated and other workers have grown apart not only in terms of earnings but also in terms of location. In the future, job growth linked to AI may also be concentrated in different areas compared to those affected by any job decline linked to AI. This could pose significant challenges, particularly given evidence that low-educated workers are less likely to move in response to potential job opportunities compared to high-educated workers (Diamond, 2016 for the US; Manning & Petrongolo, 2017, for the UK)

Box 2. Discussion of empirical strategy: ICT and employment polarisation

- Articles reviewed in this section establish a relative decline in the employment share of middle-education workers and workers performing ‘routine’ cognitive tasks, observed starting around the 1980s. They aim to explain whether this is due mainly to the increasing importance of ICT in the production process or to other factors (including trade liberalisation).
- The articles use statistical techniques to estimate the impact of ICT (specifically, a measure of the importance of ICT in the industry or industry-country pair) on the share of low-, middle-, and high-educated workers in employment or in the overall wage bill.
- A significant concern in this estimation is that any correlation between the high-education share and the importance of ICT capital across industries and countries could be driven by external factors (innovative industries could both employ high-educated workers and use ICT intensively) or reflect a reverse causation link- an impact of high-education workers on investment in ICT.
- To mitigate these concerns, the studies reviewed adopt one of the following strategies:
  - They estimate the relation between current employment composition and past ICT investment. This aims to account for the concern that current employment composition may determine current investment.
  - They use an instrumental variables approach (IV), which broadly includes two steps: i) estimating what proportion of the variation in ICT investment is due to a third factor that is not related to employment composition (an ‘instrument’); ii) estimating the relation between the instrument and the outcome of interest. If the instrument is valid, the estimated coefficient only reflects the causal impact of ICT investment on employment composition.

3.2.2 Recent evidence on the effect of automation on employment and wages

A number of recent publications has estimated the impact of automation on productivity and employment between the early 1990s and the late 2000s. Here ‘automation’ is defined as the adoption of industrial robots[9] (Acemoglu & Restrepo 2017; Graetz & Michaels, 2018; Dauth, Findeisen, Südekum & Wößner, 2017), except for Mann & Püttmann (2017), who use a broader definition.[10]

[9] Defined as ‘multipurpose manipulating industrial robots’, consistent with the International Standards Organisation’s ISO 8373 standard.

[10] Mann & Püttmann (2017) classify all patents granted in the United States between 1976 and 2014 as automation or non-automation related patents, and then link patents to the industries where they are likely to be used. The classification defines automation patents as those which define a ‘device that carries out a

Articles in this literature focus on the United States (Acemoglu and Restrepo 2017; Mann & Püttmann, 2017), Germany (Dauth et al., 2017), and international evidence (Graetz and Michaels, 2018 use data on 17 countries). Mann & Püttmann (2017) consider a longer period than other studies, investigating the impact of automation between 1980 and 2014.

This literature suggests that:

- Automation may **not lead to declining employment** in the areas that are most affected. However, evidence on the impact of automation on wages is less clear.
- Moreover, there is evidence of a **negative impact** of automation specifically on employment and earnings of workers most directly substituted by industrial robots: **low-educated** and (to a lesser extent) middle-educated workers in manufacturing industries.[11]
- Employment and earnings losses for workers directly affected by automation may be more likely in the case of industrial robots compared to other forms of automation.

There is some variation in findings across studies – in particular Acemoglu & Restrepo (2017) find a *negative* impact of robots on both employment and earnings in areas most affected within the United States. This is in contrast to Graetz & Michaels (2018), who find that robot adoption leads to an *increase* in mean hourly wages, and a decrease in the share of hours worked by low-educated workers, and to Dauth et al. (2017), who find that robot adoption in Germany is linked with small *increases* in local employment, as the fall in manufacturing employment is more than compensated by increasing employment in service industries.

Graetz & Michaels (2018) suggests that differences between their findings and Acemoglu & Restrepo (2017) may be explained by the different geographic focus of the two studies. Differences in managerial attitudes and organisational structures (e.g. firms in the United States being 'more aggressive than European counterparts in promoting and rewarding high performing workers and removing underperforming workers', Bloom, Sadun & Van Reenen, 2012) or in institutional arrangements (public policy, the decline of unions being more pronounced in the United States than in many European countries) could explain why robots are linked with a decrease in overall employment in the United States but not elsewhere.

Within literature on the United States, Mann & Püttmann (2017) find automation to be linked with an increase in local employment, in contrast with Acemoglu & Restrepo (2017). The difference between the two studies may be explained by the different measures of automation adopted – industrial robots in Acemoglu & Restrepo (2017), a wider group of automation technology in Mann & Püttmann (2017).

Automation is also linked with shifting employment from manufacturing to services. The positive impact on service manufacturing more than offsets the fall in

---

process independently'. Manual classification of a random sample of 560 patents is used to train a machine learning algorithm that classifies all patents into the automation and non-automation categories.

[11] Manufacturing industries are typically most directly affected by automation – this is true by definition where automation has been defined as the adoption of industrial robots.

manufacturing employment in Germany (Dauth et al., 2017), and a shift from manufacturing to services is also found in Mann & Püttmann (2017)'s work on broader automation in the United States.

Dauth et al. (2017) find that every robot destroys roughly two manufacturing jobs. This implies a total loss of 275,000 manufacturing jobs between 1994 and 2014, which accounts for roughly 23% of decline during those two decades. This loss was more than offset by employment growth in the services sector.

Dauth et al. (2017) are also able to analyse individual-level data to understand how specific workers were affected:

- The adoption of robots is not linked to lay-offs – in fact, workers in local areas more exposed to robots were more likely to stay in their job. Manufacturing employment in areas affected by robots declined between 1994 and 2014 compared to other areas because the former hired fewer additional workers compared to the latter.
- Manufacturing workers in areas more exposed to robots kept their jobs but were faced with lower earnings growth. Workers in areas with the highest robot density (90^{th} percentile of the robot adoption distribution) have earned on average EUR 800 per year less than workers in areas with the lowest robot density (10^{th} percentile of the distribution). This earnings loss is concentrated among low- and medium-educated workers, while high-educated workers experience income gains from robot adoption.

A significant question is then how the 'new' service jobs compare to the 'older' manufacturing jobs in terms of earnings, progression opportunities, required skills, and other work conditions. So far, this question has not been addressed in the literature reviewed in this section.

All the articles reviewed in this section adopt an empirical strategy that aims to identify a causal link from automation to productivity and employment outcomes. We discuss these empirical strategies in Box 3 below.

Box 3. Discussion of empirical strategy: the impact of robots on employment

A positive or negative correlation between automation and employment levels, even where regression analysis allows to control for a number of confounding factors, may still not indicate that the change in employment is *caused* by automation. For example, the correlation may reflect a reverse causal link (more robots are adopted in a given area/industry because employment is growing, not vice versa).

All articles reviewed here adopt an Instrumental Variables strategy (IV). We describe the basic intuition of this approach in Box 2 above. The validity of this strategy depends on the credibility of the chosen instrument – a variable that influences the independent variable of interest (in this case, the adoption of automation) without being related to the outcome of interest (in this case, employment levels and earnings).

The IV strategies adopted in these articles can identify the causal impact of automation if greater 'potential' for automation in the past (e.g. 1980) in a given industry-area pair explains actual automation in a following period (e.g. 1993 to 2007) but is not linked in itself to greater or smaller changes in employment over that following period. Past 'potential' for automation is computed differently in different articles, but essentially is a measure of whether the economic activities performed in each industry-area were amenable to automation, before automation became widespread.

3.3 Theoretical work on the role of AI in shaping future employment

A growing body of recent theoretical work in economics builds on the literature on job polarisation to consider the potential impact on labour of the adoption of AI. The aim of this work is to establish a framework that models how technology that could automate work would feed through the economy. This explicitly sets out the underlying mechanisms and generates non-quantified predictions of the impact of technology on economic growth, employment and earnings. This literature includes Acemoglu & Restrepo (2016, 2017b, 2017c), Aghion, Jones and Jones (2017), Bessen (2017), Caselli & Manning (2017), among others.

These articles typically set out a model of the economy where final goods are produced using capital and labour as inputs, and labour consists of a number of distinct tasks. A new technology (AI) that allows automating a proportion of tasks is assumed to become available.[12] The models analyse how this changes firms' demand for capital and labour.

This is consistent with models used to analyse the recent effect of digitalisation (specifically, 'job polarisation'). This is however a novel approach compared to traditional economic analysis of technology – where technology is seen as augmenting, not substituting, labour, and therefore as generally leading to greater earnings for workers both in the short and the long term. Bessen (2017) and Caselli & Manning (2017) adopt different modelling choices to focus, respectively, on the effect of technology on employment in a specific industry, and on the impact of technology on labour in the long run.

[12] It is worth noting that the articles reviewed in this Section explicitly refer to AI, but the models presented could also be relevant to predict the effects of automation enabled by other types of technology.

It is worth noting that there is some disagreement among economists as to the importance of current and future advances in digital technology for the economy – for example, Brynjolfsson & McAfee (2014) see digital technology as potentially enabling substantial productivity growth, but this view is not shared by Gordon (2016).

The impact of AI on overall employment and average earnings

Researchers in this space suggest that their work provides initial findings in a new area for research, where additional theoretical and empirical work is needed. The articles we reviewed are nevertheless useful to set out the mechanisms through which AI is expected to affect employment, beyond any initial substitution effects that may lead to displacement or reallocation of workers. Specifically, most of the analysis focusses on comparing how an initial 'snapshot' of the economy would compare with snapshots that involve increasing automation in the short term and long term. The adjustment period between the short and long run is not always modelled.

To simplify the analysis, Acemoglu & Restrepo (2018a) first analyse effects of technological change on workers considered as one group:

- In the short run, AI **substitutes labour** in performing specific tasks that are automated. This can lead to job losses (displacement), reduce the share of income that flows to workers (compared to owners of capital) and falling earnings.
- This negative impact on labour is offset by the following countervailing effects:
  - **Productivity effects**: goods and services whose production is increasingly automated become cheaper or increase in quality. This leads consumers to demand more products, either from sectors that are increasingly automated or from other sectors. Consequently, demand for labour increases.
  - Introduction of **new tasks**: as new technology is introduced and demand increases, new activities emerge. In Section 4.3 of this report, we will discuss examples of new activities suggested in the literature. Acemoglu & Restrepo (2016) show that the creation of new tasks explains around half of all employment growth in the United States between 1980 and 2000.
  - In the long term, there is a **capital accumulation** effect. Automation increases the demand for capital, but in the short run capital is fixed – it takes time to produce new machines (if automation requires tangible capital, e.g. robots) and to deploy them. This limits the impact of automation on productivity. In the long run, more machines can be produced and deployed; this drives down the cost of capital, and therefore the cost of production, and productivity effects become stronger.

In the models discussed by Acemoglu & Restrepo (2018a) and by Caselli & Manning (2017), in the short term the substitution effect may dominate, and workers may be worse off. In the long term, the productivity effect is always strong enough to at least offset the substitution effect. Real wages increase compared to the pre-automation situation. Yet as an increasing proportion of economic activity is automated, that is, produced with capital and little labour, the share of income gained by workers relative to owners of capital may fall. This generates potential distributional issues even if everyone is better off than without automation.

In this analysis, the short-term effect of AI depends on the balance between substitution effect, productivity effect and creation of new tasks. Perhaps counterintuitively, more disruptive technologies may have a more benign effect on workers, compared to technologies that are just good enough to substitute workers, but not effective enough to raise productivity (and therefore consumer demand) substantially.

This discussion compares an initial situation without automation with a new scenario where greater automation has been introduced. Korinek & Stiglitz (2017) also consider the transition between these two states. In the transition period, unemployment may rise even if the countervailing effects seen above are strong enough to offset displacement. New jobs may require different skills or could be located in different areas compared to jobs that decline or are transformed as a result of automation. Korinek & Stiglitz (2017) allude to the Great Depression in the US as a case of unemployment partly related to the challenges of transitioning from agriculture to manufacturing. They suggest that government intervention helped enable this structural transformation.

The impact of AI on specific groups of workers

As discussed by Bessen (2017), the direction of consumer responses to productivity effects could matter to determine what happens to workers in specific industries.

- If productivity effects lead to sufficient additional demand for products from automated industries, employment in those industries can increase despite the increasing automation. This has been the case for manufacturing industries, up to a point - until the 1930s for textiles and the 1950s for steel in the United States.
- If instead demand increases more in other industries, workers will have to shift from the automated sector to other parts of the economy. As seen in Section 3.2.2, this seems to be the pattern in the case of automation through industrial robots, which is linked with declining employment in manufacturing, offset at least partially by increases in services.

Bessen’s model implies that faster adoption of automation does not necessarily make job losses in the automating industry more likely. The outcome depends not only on the capabilities and speed of adoption of technology, but also on the response of consumers – though it is worth noting that in a globalised world the choices of consumers in one country may have relatively little effect over patterns in global demand.

This discussion in this section so far considers workers as one homogenous group. But historical evidence shows that technology has different effects on different people. Specifically, if technology substitutes mainly tasks performed by low-educated workers, economic inequality could increase in the short and long term, despite the offsetting mechanisms described above. Acemoglu & Restrepo (2018b) distinguishes two types of workers, ‘low-skilled’ and ‘high-skilled’, and models automation which may substitute workers in tasks performed by either group:[13]

- In the **short term**, automation substituting one group of workers (e.g. high skilled) leads to a **fall in earnings of the displaced group**. Earnings of the other group (e.g. low skilled) could also decline, as the displaced workers compete for jobs with the non-displaced group.
- In the **long term**:
  - **Average wages increase** compared to the no-automation baseline. This is because, consistent with the model for one group of workers, productivity effects offset displacement.
  - In spite of this, the group of **workers directly displaced experiences a fall in earnings relative to the other group**, and potentially even in absolute terms.

Acemoglu & Restrepo (2018b) assumes that productivity effects affect all workers in the same way, but this need not be the case – responses to productivity improvements could disproportionately raise demand for the products of an industry that employs a particularly high (low) proportion of low (high) skilled workers.

Drivers of the impact of AI

In the baseline models discussed in Acemoglu & Restrepo (2018b) and Caselli & Manning (2017), technological progress leads to increasing earnings for all workers in the long run. However, these articles also discuss how different assumptions may lead to different results.

First, **limited competition** between producers **could lead wages to fall in the long term**, as indicated in Caselli & Manning (2017). Sufficient competition in product markets, and limited market power of employers over their employees are necessary for workers to benefit from productivity gains.

There are indeed concerns that several markets have become less competitive in recent years. Recent research has documented a possible recent increase in profits in the United States, potentially linked to technology (Autor, Dorn, Katz, Patterson & Van Reenen 2017; De Loecker & Eeckhout, 2017). Furman & Seamans (2018) also note that digitalisation includes potential drivers of market concentration: i) the emergence of ‘platform’ markets, which tend to be dominated by one or few firms (e.g. Facebook in social media, Uber in urban transport); ii) the importance of large datasets for the development of machine learning algorithms, which could limit the extent to which market leaders can be challenged by smaller competitors (for example, it could be very difficult to challenge Google’s position in internet search without access to the data on past searches Google holds); iii) the use of algorithms to set prices, which could make collusion between competitors easier to implement and harder to detect. In Section 4.4 of this report, we review

[13] These two groups can be thought of as ‘low-educated’ and ‘high-educated’ as in previous sections of this report. However the discussion that follows here would still be relevant for other definitions of skill levels.

selected evidence on the organisation of work in platform markets that are relevant for this discussion.

The regulation of technology firms and evolution of competition policy to reflect challenges posed by digitalisation are a widely debated topic, beyond the scope of this review. However, the evidence reviewed here suggests that policy measures aiming to reduce the market power of employers or mitigate its effects could help ensure that workers benefit from technological change.

Moreover, AI could lead to **changes in the innovation process**:

- The models discussed so far generally consider a one-time increase in the work tasks that are automated. However, if adopting AI makes it easier to then generate new automation, as considered in Acemoglu & Restrepo (2016), the **impact of AI for workers could be worse**. In this case, the countervailing effects described above (productivity effects, creation of new tasks) still help maintain employment and wages, but it is more likely that fewer people will work and that the share of income flowing to workers will fall in the long run.
- However even this situation does not necessarily imply a jobless future. In Aghion et al. (2017), as automation takes over an increasing proportion of work tasks, industries where production is automated shrink as a share of the economy. This is consistent with the long-term impact of technological progress in agriculture and to a lesser extent in the manufacturing sector, as noted by Bessen (2017). As automated sectors shrink, the share of income that goes to labour remains constant despite increasing automation. Work is concentrated in non-automated sectors, where prices are high relative to goods produced in automated sectors.

Korinek & Stiglitz (2017) devote specific attention to distributional issues. As seen above, in the short term AI could lead inequality to increase, even if all workers benefit to the same extent, because a larger share of income would be received by owners of capital. Korinek & Stiglitz (2017) point out that this could also happen in the long term, if production processes require an input that is in fixed supply (for example, land). Taxing fixed inputs would be a direct way to redistribute income, but that may not be desirable for other forms of capital (e.g. robots). An alternative way to achieve redistribution could be lowering the protection awarded to innovators through Intellectual Property, for example shortening the term of patent protection.

4 PREDICTIONS AND ANALYSIS OF POTENTIAL IMPACTS OF AI

Summary of this section

- There is limited evidence on the factors that may influence the adoption of AI, beyond its technical capabilities. Business adoption is likely to be influenced by economic considerations, regulatory concerns, individual preferences and social norms, and by the need to reorganise production processes to take advantage of AI.
- It is not clear whether adoption of AI is progressing at a particularly fast pace compared to previous innovative technology. There is evidence that the speed of technology adoption has accelerated over time. On the other hand, recent productivity growth in the UK has been sluggish and there is evidence that the diffusion of innovation beyond leading firms may be slower than in the past.
- It has been estimated that between 9% and 47% of jobs in the United States, and 10% to 30% of jobs in the UK could be automated at some point in the coming decades. Differences within these ranges do not appear to be driven by different assessments of the capabilities of machines, but rather are linked to:
  - Different sources of data on the skills needed to perform a particular job;
  - Modelling choices – in particular, assumptions over whether automation would hit all jobs under granular occupation groups at the same time (e.g. all waiters) or hit first only some jobs under that group (e.g. only waiters in fast-food restaurants).
- The literature estimating potential job losses suggests that jobs typically performed by workers with relatively low levels of formal education are at higher risk of automation compared to high-educated jobs.
- There is significant uncertainty around when the potential automation described in this literature would be technically feasible and widely adopted.
- A different assessment of the technical automatability of work tasks has been provided by Brynjolfsson, Mitchell & Rock (2018). According to Brynjolfsson et al. (2018), very few occupations consist primarily of performing automatable tasks, although nearly all occupations include automatable components. As a result, the degree of automatability varies relatively little between occupations. The authors of this article conclude that:
  - AI will affect occupations at all levels of pay and education;
  - However, automation of entire jobs will be much less significant than transformation, which will involve separating automatable tasks so they can be performed by machines, and aggregating non-automatable tasks into new job roles.
- Foresight exercises suggest that occupations that are likely to grow between now and 2030 are disproportionately high-education, although some middle-education occupations are also likely to grow.
- A number of trends other than technological change could generate sufficient job creation to compensate for first-order job destruction linked to

automation. However, the 'new jobs' may not have the same characteristics or emerge in the same places as the 'old' that are destroyed, and therefore adjustments may be challenging.

- There is evidence that the application of digital technology in specific settings and industries is leading to changes in the organisation and quality of work:
  - There are reasons for concern (e.g. potential loss of autonomy for many workers) and for optimism (e.g. automation of hazardous tasks);
  - Evidence on the likely balance between the two is limited – outcomes are also likely to depend on the actions of employers, workers, governments, and other stakeholders.

4.1 From research to innovation and adoption in society

The literature reviewed in Section 3 of this report suggests that the use of AI will be driven not only by technical feasibility, but also by economic, social, and cultural factors. However, this review has identified relatively little discussion of social and cultural factors in the literature relative to AI and work. Indeed, this is one of the points raised in Wajcman (2017), in a review of four recent books dealing with the impact of automation on the future of work.[14] There is work in progress around the public perception of AI[15] and on the factors that shape the interaction between humans and machines, including robots and decisions made on the basis of statistical models, but we have not identified published outputs from this research that discussed the implications of these factors for the future of work.[16]

Wajcman (2017) is critical of the discussion of technology in the debate on the future of work as a 'neutral inevitable force driving these changes [in work]'. She suggests that more attention should be devoted to:

- The factors driving how technology is developed and used. The essay raises concerns around the power of a small elite of leading companies and their staff in shaping the technical systems they design – even as the work of this elite may displace other elites. Wajcman (2017) points to Urry (2016)'s explanation of how social practices are 'constitutive' of technology.
- The way in which technological change is discussed. Urry (2016) argues that visions of the future imply ideas about 'public purposes and the common good', and can have powerful consequences. Wajcman (2017) warns against visions

[14] Wajcman (2017). The books reviewed are: Ford (2015); Susskind, R. and Susskind, D. (2015), Brynjolfsson and McAfee (2014); Urry (2016).

[15] Including the AI Narratives Project launched by the Leverhulme Centre for the Future of Intelligence and the Royal Society.

[16] The programme of the Technology, Mind & Society conference organised by the American Psychological Association in April 2018 includes several research projects that would be relevant for this review, but we have not identified outputs from those projects that could be included in the review. The programme of the conference is available at: https://pages.apa.org/tms/.

of the future that focus on the potential capabilities of technology while turning 'their backs to society'.

Earlier research in social studies of science and technology presents insights for the study of the interaction between technology and work, summarised in Wajcman (2007). This research challenges the notion that 'social' and 'technical' aspects of the implementation of technology are separate:

- Choices between different technical options for embedding technology in production processes, goods, and services can be affected by a range of social factors (for example, will a particular option be accepted by the relevant social group?);
- The use of technology changes relationships between people, the nature and meaning of work. The interaction of social and technical considerations affects the use of technology, and this is linked to social changes (which can then be relevant for the definition and adoption of new technology).
- A consequence of the interaction between the technical and the social is that the final use of a technology cannot be completely predicted by its designers and promoters. This may be one of the factors that makes predicting what, if any, job losses could be linked to AI challenging.

More recently, a growing body of research has focussed on new forms of organisation enabled by digital technology – in particular, 'platforms' such as Airbnb and Uber. Pasquale (2016) draws attention to the role that narratives around platforms may have in influencing the future. The article contrasts two narratives about the role and impact of platforms that would have very different implications for policymaking.

A wider literature in history and economics investigates what may explain which specific innovations are generated and adopted in periods of technological change (for example: Rosenberg, 1969; Mokyr, 1990, as well as research cited in Section 3 on how work organisation varies given the technology used). The importance of this question is highlighted in some of the recent economic theory on AI, but a review of this literature is beyond the scope of this report.

The technical capabilities of AI

An overview of what is known about the existing and likely capabilities of AI in the near future is provided in a number of publications, including but not limited to House of Lords (2018), National Academy of Sciences (2017), and Royal Society (2017). This review focusses on the contributions that have provided quantitative evidence on the workforce implications of their assessment of technical capabilities. High-level insights on technical capabilities from this literature are provided here before providing a detailed review in Section 4.2 below:

- There is a consensus between these publications literature that advancements in machine learning allow AI to perform tasks that cannot readily be codified into a set of explicit instructions. This is a new feature compared to what machines had been capable of up to the recent past (e.g. Autor, 2015).
- However, 'routine' activities are still considered to be most amenable to automation. Within this, there is variation in the literature around the specifics

of which activities can or cannot be considered routine. For example, Frey & Osborne (2017) find that measures of dexterity required to perform an occupation do not predict their estimated probability that the occupation could be automated in the future. However, McKinsey (2017a) and Brynjolfsson & Mitchell (2017) consider the relatively limited existing capabilities of robots in moving and manipulating objects using dexterity as a bottleneck to automation.

- There is also a broad consensus that machines are increasingly capable in a range of cognitive tasks, including retrieving information, recognising patterns, and generating predictions. Several articles recognise that, at least at the moment, it is difficult to explain how or why an AI system has generated a prediction. Brynjolfsson & Mitchell (2017) note that this makes it difficult to automate activities where there is a need not only for an output but also for an explanation of how the output was achieved.
- All publications reviewed here suggest that tasks requiring a high degree of social intelligence are generally not within the capabilities of machines currently or in the near future.

The drivers of business adoption

Compared to the social drivers mentioned by Wajcman (2017), greater attention has been devoted to the considerations that influence business decisions to invest or not to invest in AI, and whether that investment is likely to improve business performance. And indeed, any impact of AI on work is likely to be limited if this technology is not in being adopted by workers and employers.

RSA (2017) suggests four key considerations influencing business decisions to adopt AI:

- **Economic considerations**: the economic benefits and costs of investing in AI and uncertainty around them. The cost of investing in AI may be too high to make the investment profitable for many companies. Moreover, some may find waiting to be the best strategy, given the possibility of further decreases in costs, and the possibility that what is now leading edge may soon become obsolete.
- **Individual preferences and social norms** around what should be done by humans or by machines;
- **Regulatory concerns –** in certain circumstances AI could in principle be used but it may not be clear whether this is consistent with relevant regulation;
- The need for broader **organisational change** and for relevant skills to accommodate and take full advantage of new technology. This is part of what Brynjolfsson, Rock & Syverson (2017) have called ‘complementary investments’, consistent with earlier research on the economic impact of Information and Communication Technology (ICT).[17] Brynjolfsson & Hitt (2003) found that the productivity benefits of US firms’ ICT investments took around seven years to be fully realised.

[17] Including, for example, Bresnahan et al. (2002), Aral et al. (2012).

The sociological literature reviewed earlier in this section also suggests the interaction between different groups (not only employers and employees, but also different groups of workers) as a key factor shaping whether and how AI may be adopted.

Moreover, the use of AI for automated decision making may pose a risk to the fairness of those decisions. Citron & Pasquale (2014) discuss some of the limitations of existing scoring systems based on predictive algorithms, credit scoring in particular. They recommend regulatory oversight to ensure the transparency, accuracy, accountability, and fairness of scoring systems.

McKinsey (2018) analyse targeted applications of AI to business challenges (‘use cases’) from their own practice and other sources, to understand where AI offers the most incremental value over and above alternative techniques. They find that the greatest potential for AI to create value lies in use cases where established analytical techniques (e.g. regression) can already be used. The potential additional value generated by AI is assessed to be particularly large in the travel, transport and logistics, retail, automotive and ‘high tech’ industries. This is driven by applications such as personalised marketing, improved forecasting for supply-chain management, and credit risk measurement. The authors note that some of these applications may need to overcome a number of challenges to generate value, including avoiding or limiting the risks outlined by Citron & Pasquale (2014).

Economic considerations will be guided by a number of factors, including the characteristics of the workforce. As seen in Section 3.1, there is an on-going debate on whether relatively high wages in 18th century Great Britain generated particularly strong incentives to invent and adopt labour-saving technology. There is also evidence that age and education of the workforce have been important in determining the adoption of ICT and of automation technology.

Acemoglu & Restrepo (2018c) show that geographical areas in the United States characterised by greater ageing have adopted industrial robots and other automation technologies to a greater extent.[18] This is because industrial robots have the potential to automate activities such as welding, or loading and packaging, which are typically performed by middle-aged workers. An ageing population makes middle-aged workers relatively scarce and therefore more costly to businesses. Therefore ageing increases firms’ incentive to adopt this particular type of labour-saving technology.

The characteristics of the workforce in terms of education have also been linked to the invention and adoption of digital technology. The increasing availability of high-educated workers in advanced economies in the second half of the 20th century may explain in part the invention and adoption of ICT, as ICT have been shown to complement the skills of high-educated workers (as discussed in Section 3.2).

Acemoglu & Autor (2011) note that large increases in the proportion of university-educated US workers between 1940 and 1980 have been accompanied and followed by a rise in their relative wages (a rise in the ‘graduate premium’). If technology had stayed the same, one might have expected to observe a fall in the graduate premium (as more workers are available to perform the same amount of

[18] Ageing here is measured as an increase in the ratio of older (aged 55 and older) to middle-aged workers (aged between 36 and 55).

work). Instead, the adoption of ICT and related transformations in the organisation of production processes may have led to increasing demand for high-educated workers, more than compensating the effect of increasing supply.

The UK has also experienced a large rise in the proportion of university graduates, starting later than in the US. At the same time, Blundell, Green & Jin (2018) show that the graduate premium remained constant before falling slightly in more recent years (for people born after 1980). Blundell et al. (2018) suggest that differences compared to the US (constant rather than rising graduate premium) are explained by the fact that the UK was a technological follower. In the UK, the increase in university graduation rates in the UK led firms to adopt existing ICT technology and related decentralised organisational structures. This in turn increased the demand for high-educated workers, particularly in management positions.

It is also worth noting that what is technically feasible for one organisation may not be for others. As noted in Brynjolfsson & Mitchell (2017), among others, applications of machine learning typically require large training datasets to be effective. There will be wide variation between organisations in whether their activities can generate such data, and in whether the necessary investments to collect the data have been made.

McKinsey (2017a, 2017c) model different patterns of AI business adoption in their work estimating the proportion of jobs that could be automated in the future, reviewed in Section 4.2 below.

The current state of business adoption

A notable initiative tracking how AI is being used is the AI Index, part of the Stanford 100 Year Study.[19] The AI Index shows increasing research outputs related to AI (a nine-fold increase in the number of Computer Science papers tagged with the keyword ‘Artificial Intelligence’ in the Scopus database of academic papers); increasing early-stage funding in the United States for private companies developing AI systems (attracting over US$3bn in 2016 compared to just over US$500m in 2010); increasing shares of on-line job postings requiring AI skills in the United States (a five-fold increase in October 2017 compared to January 2013), in Canada (an eleven-fold increase), and in the UK (a nine-fold increase). This suggests rapidly increasing interest in AI, not only in research but also among businesses.

In the UK, a RSA/YouGov poll performed in 2017 finds that 14 percent of business leaders are currently investing in AI or robotics, or plan to in the near future (RSA, 2017). The UK adoption rate implied by this poll does not appear significantly lower than international adoption rates as estimated, for example, in McKinsey (2017b), where 9-12% of business leaders across 10 advanced economies reported having adopted AI. RSA (2017) also discuss data from the International Federation of Robotics showing that in 2015 the UK had 10 robot units per million hours worked, compared with 131 in the United States and 133 in Germany. This suggests lagging adoption of AI in the UK, but the difference may be explained in large part by differences in the industrial composition of these economies.

[19] https://aiindex.org/

It is not immediately obvious whether an adoption rate of 14 percent suggests that the adoption of AI is moving at a particularly fast pace compared to previous innovative technologies. There is evidence suggesting lags between invention and adoption of technology have been narrowing over time: according to Comin & Hobijn (2010), the average lag was 14 years for personal computers, compared to over 60 years for the telephone. If this trend has continued, we might expect many of the recent advances in AI to be widely adopted by 2030, assuming that the drivers of AI adoption are broadly consistent with those of earlier technology.

However, any acceleration in the pace of technology invention and adoption is not yet evident in productivity statistics. As widely documented (e.g. Haldane, 2017), recent productivity growth in the UK and other advanced economies has been weak, judged by the standards of the last 50 years. Brynjolfsson et al. (2017) propose four potential ways to explain why productivity growth has been sluggish even as digital technology is increasingly pervasive and capable:

- **Misplaced optimism**: digital technology can provide benefits in terms of new products and services but does not have the potential to substantially raise productivity;
- **Measurement errors**: there have been productivity gains, but they are not evident from official statistics because much of the activity related to the digital economy is hard to measure. Many investments are intangible (for example, they relate to machine learning processes rather than to computer hardware or other forms of machinery) and many new services are provided to consumers at low or no cost.
- **Distribution issues**: digital technology is leading to productivity gains, but these are concentrated among a small group of leading firms. Andrews, Criscuolo & Gal (2015) show that in the OECD sluggish productivity growth is not the story for all firms. The most productive firms have experienced robust productivity growth, but there has been an increasing gap between this group and all others.[20]
- **Implementation lags**: productivity gains do not materialise immediately when a new technology is adopted. Production processes need to be reorganised to take advantage of the new technology, particularly if the technology can be widely applied.

Brynjolfsson et al. (2017) present evidence suggesting that implementation lags are the most likely explanation, although measurement errors and distribution issues could also be part of the story. In other words, they suggest that the need for organisational change, one of the drivers of business adoption mentioned above, is a significant factor that can delay the impact of AI on productivity and on the workforce.

4.2 Estimates of potential displacement linked to AI

A number of recent publications, starting with Frey and Osborne (2017, hereinafter 'FO')[21], has aimed to estimate which components of existing work could technically be performed by machines in the future. This amounts to estimating for which workers the immediate ('first order') effects of AI could involve substantial transformation of the tasks performed or job losses, given the technical capabilities of machines. This review includes articles published by PWC (2017, 2018 - hereinafter 'PWC'), McKinsey (2017a, 2017c – hereinafter, 'McKinsey'), and the OECD (Arntz et al. 2017, Nedelkoska and Quintini, 2018 – hereinafter 'AGZ' and 'NQ' respectively), as well as Brynjolfsson, Mitchell & Rock (2018 – hereinafter, 'BMR').[22]

There is broad agreement in this literature that many of the existing jobs (around 50% in the US and UK) have significant potential for automation. However, there is no consensus on what proportion of jobs could be fully automated or radically transformed. The estimated proportion of current jobs that could technically be automated in the future varies from 9% (AGZ) to 47% (FO) for the United States, and from 10% (AGZ) to 30% (PWC) for the UK (note that FO provide estimates for the US only). BMR find that only in very few cases it is possible to automate nearly all work tasks making up an occupation. They suggest that automation of entire jobs will be much less significant than transformation.

There is also uncertainty on the time frame within which the automation described in this literature might be feasible. AGZ, PWC, and NQ all rely on the expert judgement collected by FO. FO asked a group of machine learning researchers convened in a workshop held at the Oxford University Engineering Sciences Department the following question: 'Can the tasks of this job be sufficiently specified, conditional on the availability of big data, to be performed by state of the art computer-controlled equipment?'. FO interpret their analysis as assessing 'potential job automatability over some unspecified number of years'.

FO, AGZ, PWC, McKinsey and NQ are all in broad agreement on the types of jobs that are most exposed to automation:

- Jobs that require relatively **low levels of formal education** (food preparation, machine operators in manufacturing, personal service occupations, administrative support workers). Jobs with higher education requirements (professional, scientific and technical occupations, business and management occupations) are instead consistently found among those with the lowest estimated automation probabilities
- Occupations that **do not involve relatively complex interaction** with other people – in particular, influencing or persuading others, assisting and caring for others, training others, and managing other people's work. It is not clear whether face to face interaction per se is linked with lower automation

[21] First published as a working paper by the Oxford Martin Programme for Technology and Employment in 2013: https://www.oxfordmartin.ox.ac.uk/downloads/academic/future-of-employment.pdf

[22] McKinsey (2017b) also provides estimates of potential employment growth to 2030. These estimates are reviewed in Section 4.3of this report.

probability – sales occupations are typically assessed at high automation probability.

- Occupations **that involve manual skills** – in the language of FO (O*NET data), higher proficiency in 'finger dexterity' (for example, 'Perform open-heart surgery with surgical instruments') and 'manual dexterity'. This is inconsistent with the expectation – discussed in FO and Autor (2015) among others, that manual tasks involving dexterity would be relatively difficult to automate. However, within physical activities, McKinsey suggest that those performed in 'unpredictable environments' are significantly more likely to be automated than those in 'predictable environments'.

Given these patterns, occupations at high risk of automation are concentrated in

- The **agriculture, manufacturing, postal and courier services, land transport, and food services** industries;
- Jobs typically held by **younger workers**: NQ show that workers aged under 20 are over-represented in many occupations at high risk of automation (e.g. sales, 'elementary occupations' such as cleaners). But workers under 20 are also at higher risk of automation than older workers who perform the same occupation. This is because, even within occupations, workers under 20 often perform roles that involve manual tasks (linked with greater automatability) and do not involve providing advice, supervising others, performing research (tasks linked with lower automatability).

BMR, instead, report relatively little differences between occupations in their degree of automatability. This is because, in their assessment, nearly all occupations include a minority of tasks that have a high degree of automatability. Therefore, **all occupations could be transformed** by AI to some extent, but **very few could be completely automated**. As a result, there are no systematic differences between low- and high-paid occupations in their suitability for automation.[23] Their assessment of the susceptibility of tasks to automation is consistent with Brynjolfsson & Mitchell (2017), reported in Section 4.1 above.

Studies in this literature typically acknowledge that technical feasibility is only one of the determinants of automation: as noted in Section 4.1, tasks may not be automated even once it is technically feasible to do so – automation may not be profitable, consistent with regulation, or consistent with consumer preferences and with managers' and workers' attitudes. It is also worth bearing in mind that on the other hand production processes could be re-organised so as to make automation more (or less) technically feasible. This is a point raised in Susskind & Susskind (2015) about professional occupations: although professional occupations may involve a significant amount of non-automatable activities (e.g. empathy), it is not obvious that consumers value these activities specifically over the outcome (e.g. accurate legal advice).

[23] Brynjolfsson et al. (2018) do not discuss how their assessment of automatability varies with the education requirements of each occupation. However as high-educated occupations are typically high-paid, it is likely that their assessment of automatability does not vary with education levels.

The number of jobs 'at risk' of automation

The proportion of jobs estimated to be at 'high risk' of automation varies considerably across studies. Figure 4 below reports estimates for the United States.[24] The proportion of jobs estimated to be at high risk of automation varies between 9% (AGZ) and 47% (FO). Across all studies included in the figure, 'high risk' has been defined as an estimated probability of automation of 70% or higher. BMR do not provide a quantitative estimate of the proportion of jobs that could be automated, and therefore is not included in the figure. McKinsey sets 2030 as the time frame for estimated automation, while PWC predicts into the mid-2030s, and FO, AGZ and NQ do not provide an estimated timing.

Figure 4 Estimated proportion of employment at high automation probability

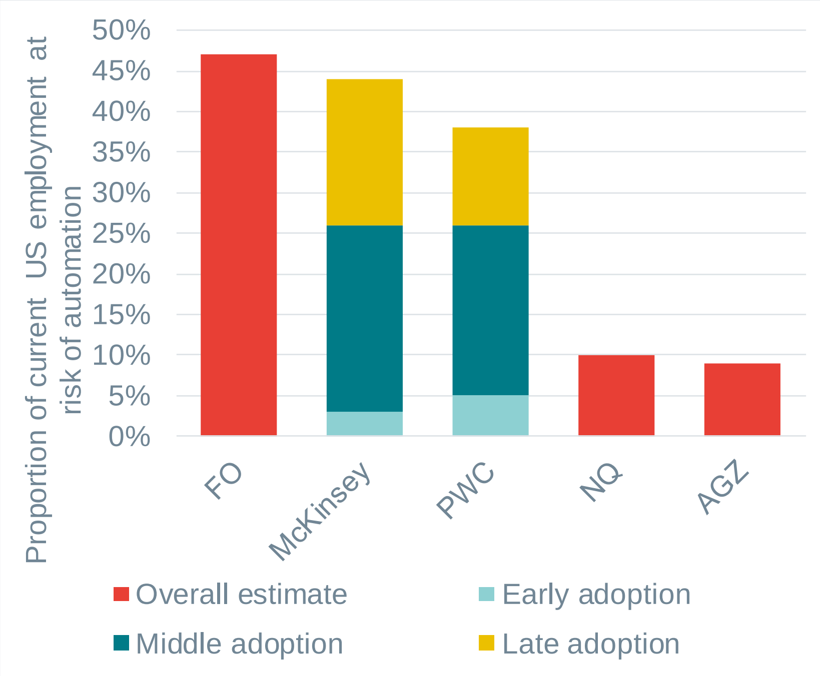

Source: Artnz et al. (2017), Frey and Osborne (2017), McKinsey (2017), Nedelkoska and Quintini (2018), PWC (2017).

Note: The figure focusses on the US as FO do not provide estimates of automation probabilities for other countries, and McKinsey do not provide estimates comparable to FO, NQ and AGZ for the UK.

McKinsey and PWC aim to provide additional detail on the likely timing of potential automation. The three portions of the bars in Figure 4 above should be interpreted as different adoption scenarios in the case of McKinsey, and different stages of automation in the case of PWC:

- For McKinsey, the preferred estimate is given by the 'midpoint adoption' scenario (26% employment automated by 2030).[25] The 'early adoption' and 'late adoption' figures (just under 5% and 44% respectively) are alternative

[24] This is the only geography considered in FO, while other studies extended their scope to other countries including the UK.

[25] Note: McKinsey estimates are provided as proportions of Full-Time Equivalent (FTE) working hours rather than proportions of jobs.

estimates under different assumptions about the development and adoption of AI – further detail on these assumptions is provided below.

- PWC's overall estimate of automation potential by the mid-2030s is given by the overall height of the bar (38%). Here 'early', 'middle', and 'late' adoption represent potential automation by the early 2020s, the late 2020s, and the mid-2030s respectively.

Estimates for the UK, where available (PWC, NQ and AGZ), are very similar to US estimates in the OECD analysis (NQ and AGZ), but lower according to PWC (just over 30% by mid-2030s).

It is worth noting, when assessing the size of potential impacts, that in every quarter over the last 15 years, around 3% of workers become unemployed or inactive, and that around 2% of workers move to a different job within the quarter.[26]

It could be useful to compare these estimates to data on existing or planned adoption of AI. RSA (2017), reports UK business leaders' estimates of the proportion of jobs in their organisation that could be fully automated within 10 years. Over 70% of respondents estimated that full automation could be achieved for at least some jobs in their organisation, but under 15% estimated the proportion of automatable jobs to be over 30%. The study does not report how many within this 15% also reported to be adopting AI in their organisation.

Drivers of variation in estimated jobs at risk

All reports in this literature use the two following inputs:

- A judgement on what can or cannot be automated.
- Data on what activities are performed in different jobs, or what capabilities are required to perform a job ('job requirement data').[27]

In all studies except for BMR, the judgement of automatability is provided by experts for a subset of all possible standardised job categories ('occupations').[28] In the FO study, experts classified 70 occupations (drawn from a total of 702 possible occupations). Out of the 70, 37 occupations (53%) were assessed as automatable. The expert judgements and job requirement data are then combined, using machine learning methods, to generate predictions of automatability ('automation probability') for each occupation – including those for which an expert judgement was not provided. This also allows investigating which characteristics of individual work activities or of occupations are linked with greater automation probability.

BMR, instead, develop a rubric to assess the automatability of work tasks, and then crowdsource the assessment.[29] The rubric asks respondents to grade each task against a set of 21 criteria which BMR judge to be important drivers of

[26] Source: Office for National Statistics, Labour Force Survey estimates, available at: https://www.ons.gov.uk/employmentandlabourmarket/peopleinwork/employmentandemployeetypes/datasets/labourforcesurveyflowsestimatesx02

[27] This includes the O*NET data, collected by the United States Bureau of Labor Statistics, used in FO, and the OECD Survey of Adult Skills (PIAAC), collected by the OECD and used in AGZ, NQ and PWC.

[28] As defined by the Standard Occupational Classification system.

[29] They use an on-line platform called CrowdFlower, now known as Figure Eight, which recruits people to perform 'human intelligence tasks' – in this case, assessing work tasks using the rubric provided by the authors.

automatability. This methodological difference is likely to explain why BMR provide different results compared to other studies in this section on the variation in automatability across occupations.

However, differences in estimated automatability remain even among studies that rely on the same expert judgement over what machines can or cannot do. FO, AGZ, NQ, and PWC all rely on the same expert judgement, provided as part of the FO study. These four studies should therefore be the most comparable, and the following discussion focusses on them before investigating what may drive differences with McKinsey estimates. FO, AGZ, NQ and PWC all agree that only a minority of workers (33% for the US in FO, 26% across OECD countries in the NQ sample) are in roles with low (30% or less) automation probability. There is greater disagreement over what proportions of workers fall in the medium (30 to 70%) and high (70% or higher) ranges of estimated automation probability, as shown in Figure 5 below. Figure 5 focusses on FO and NQ as these studies report a split of current employment into low, medium, and high probability of automation.

Figure 5 Proportion of workers by estimated automation probability according to FO and NQ

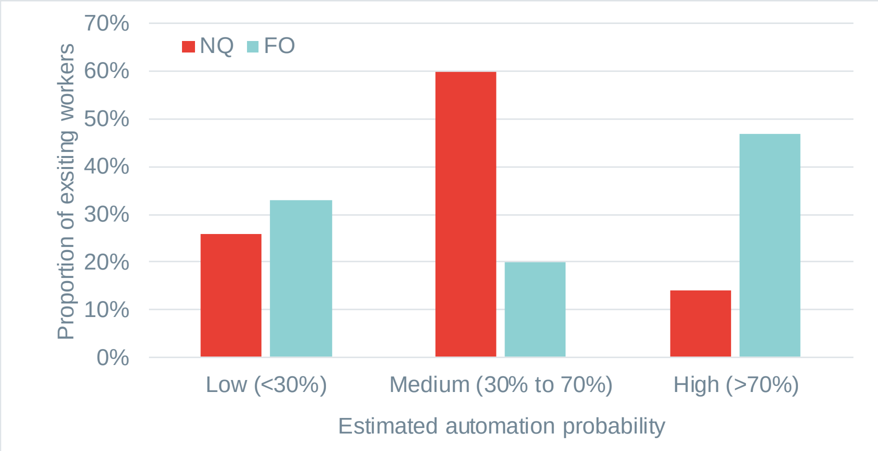

Source: Frey and Osborne (2017), Nedelkoska and Quintini (2018)

Note: NQ figures are averages across OECD countries in the study sample; FO figures concern US employment. It was not possible to present US figures from NQ based on publicly available data.

Due to differences across studies in data and machine learning algorithms used, it is not possible to disentangle this variation into specific drivers. However, AGZ and NQ explain the difference between their estimates and FO as follows:

- FO treat all jobs within the same occupation code (for example, 29-1111, ‘Registered Nurses’, or 47-2012, ‘Brickmasons and Blockmasons’) as identical and therefore having the same probability of being automated. This approach can be interpreted as assuming that if some jobs within an occupation can be automated, all other jobs will follow.
- AGZ, and PWC and NQ following AGZ, allow instead for different jobs within the same occupation code to involve different activities and therefore to have different automation risk. For example, some Registered Nurses may be required to solve complex problems more often than others.

Discussion of results presented in PWC and NQ suggest that estimates of the proportion of jobs at ‘high risk’ of automation are sensitive to modelling choices.

PWC shows that the distribution of estimated automation probabilities depends on the number of worker and job characteristics ('features') included in the machine learning algorithm that predicts automatability. When the number of features increases, the proportion of jobs in the 'low risk' and 'high risk' areas increases, getting closer to the shape of the FO 'occupation-based approach'. Therefore, moving from feature set 1 to feature set 3 implies very different conclusions about the proportion of jobs at high (>=70%) risk of automation. PWC do not describe which features are included or excluded in each of the sets.

Moreover, NQ show how estimated automation probabilities change when alternative sources of job requirement data are used for Germany and for the UK. The proportion of workers at 'high risk' changes significantly (33% for Germany and 3% for the UK using alternative data, compared to 18% and 12% respective baseline estimate). Average automation probabilities and characteristics of jobs most at risk are unaffected by the choice of dataset.

Overall, despite this discussion, all studies reviewed here provide relatively little description of the sensitivity of their results to specific modelling choices – in particular, the composition of the feature set.

Compared to FO, AGZ, PWC and NQ, McKinsey relies on a different assessment of automatability. McKinsey do not provide a comparison of their assessment to FO, and provide relatively little technical detail on how human judgement and machine learning were combined to generate their estimates of technical automation potential. However, the types of activities linked with greater automation potential in the McKinsey study are broadly consistent with findings in other studies, as discussed in further detail below.

Differences in automation risk across countries

AGZ, PWC, McKinsey and NQ all provide estimated automation probabilities for several countries, with noticeable differences across studies:

- In McKinsey, estimates vary relatively little across countries (43% for the UK, 51% for China, and 56% for Japan);
- There is greater variation in NQ, where the proportion of jobs at high risk of automation varies from under 5% (Norway, Finland) to over 20% (Slovak Republic), and in PWC, where the proportion varies from 22% (South Korea, Finland) to 44% (Slovak Republic)
- AGZ estimate little international variation but focus on a narrower group – a subset of wealthier OECD countries.

It is not clear what drivers the difference between McKinsey, on one side, and NQ and PWC, on the other.

In NQ and PWC, differences across countries are driven by:

- Differences in industrial structure: countries with larger manufacturing industries, for example, are estimated to be more exposed to automation;
- Differences in the occupational composition of each industry: for example, manufacturing firms in South Korea may employ a smaller proportion of workers in occupations at high risk of automation compared to Italy.

- Differences in the types of activities performed by workers within the same occupations and industries. For example, professionals in all industries may work in teams more often in some countries, and therefore may have a greater requirement for interpersonal skills.

NQ report that differences in occupational composition, controlling for industry composition, are the most important of these three components in explaining differences across countries in automation risk.

However, there are large differences between PWC and NQ for specific countries:

- Greece and Japan are among the countries with highest automation risk in NQ, and among the countries with lowest risk in PWC;
- The United States is one of the countries with lowest automation risk in NQ, and among the countries with highest risk in PWC.

4.3 Estimates of potential job creation linked to AI and future skills requirements

The first-order effects of AI on work may not only include substituting workers in existing tasks, but also generating new tasks for workers to perform. Acemoglu & Restrepo (2016) note that new tasks have emerged in previous periods of technological change, and that they account for around half of new employment creation in the United States between 1980 and 2000. The economic framework presented in Section 3.3 signals that the creation of new tasks (and new products) linked to AI is a key determinant of the impact of AI in the short term. Acemoglu & Restrepo (2018a) suggests that job creation from productivity effects (higher demand for goods and services whose production has been automated) would not be sufficient in its own to offset potential job losses from automation.

Section 4.3.1 below reviews relevant articles on potential job creation linked to AI.

Moreover, Bakhshi, Downing, Osborne & Schneider (2017) and McKinsey (2017c) have considered the potential impact of AI on employment levels within the context of other trends. We review these studies in Section 4.3.2.

4.3.1 Potential job creation

Compared to the automation effects of AI adoption, this review has identified significantly fewer contributions investigating:

- What types of jobs may emerge or grow in importance as a result of the adoption of AI;
- What types of jobs are more likely to be transformed (and are likely to become more productive) as a result of AI, rather than be replaced.

This is perhaps unsurprising, as identifying new occupations and the way in which existing occupations will change is a significantly more demanding exercise, that would require collecting evidence and making assumptions on the outcomes of innovation processes, future demand for goods and services, and the organisations of production processes.

This review has not identified any publications estimating the net impact of AI on employment, wages or other aspects of work. By ‘net impact’ we mean taking into account both potential job losses due to displacement, and potential job gains due to innovation related to AI and to the potential positive impact of AI on consumer demand.

However, several contributions have pointed out that adopting AI will lead to new problems that have not been solved before, some of which may lead to the emergence of new occupational categories. On the face of it, many, but not all of the new potential occupations appear to be appear best-suited to high-educated’ workers:

- Wilson, Daugherty & Morini-Bianzino (2017) suggest that ‘trainers’ (workers performing tasks useful to train AI systems), ‘explainers’ (workers interpreting the outputs generated by AI systems so organisations using AI systems can be accountable internally and to others), and ‘sustainers’ (workers monitoring the work of AI systems to prevent and mitigate any unintended consequences) may all be jobs of the future.
- Eurofound (2017a) suggests that advances in industrial robotics could generate employment in the provision of robotics support services to manufacturing firms, as well as in the manufacturing of robots. Roles in these areas would include programmers and specialists in robot maintenance. These occupations would not be entirely new, but they involve new combinations of skills.
- Francois (2018), commenting on Aghion et al. (2017, reviewed in detail in Section 3.3) notes the need for people to be involved in shaping the ‘values’ of AI – even in the case AI systems were to reach a very advanced stage of development, where a large majority of work tasks could be automated.
- More in general, Agrawal, Gans & Goldfarb (2017) frame AI as a ‘prediction machine’ and note that its use will raise demand for tasks that complement prediction. One such group of tasks they point to is ‘judgement’ – the assessment of the value of potential outcomes, which can be combined with predictions about what the potential outcomes are to generate decisions over what course of action should be taken.

4.3.2 Estimates of future skills requirements

Along with articles described above, which point to specific sources of potential job creation related to AI, others have aimed to predict future levels and composition of employment, taking into account a number of possible future trends including the adoption of AI in the workplace. This Section focusses on two reports that have given specific consideration to the potential impact of technology: Bakhshi et al. (2017) and McKinsey (2017c).

Nesta and Oxford Martin School: ‘The Future of Skills’

Bakhshi et al. (2017) adopt an innovative approach to translate expert judgement into predictions about a detailed set of occupations. Their analysis generates predictions for the probability that the share and the total number of people employed in each occupation will increase, decrease, or remain constant by 2030.

They provide separate estimates, generated using the same approach but judgement by a different group of experts, for the United States and the UK.

In a nutshell, their approach consists of the following steps:

- Eliciting from experts predictions about the future of 30 detailed occupation groups (defined at a detailed level of Standard Occupational Codes), up to 2030; [30]
- Using these predictions, along with detailed information about each occupation group, to train a machine learning algorithm that would generate predictions for all detailed occupation groups.

Using this approach, the study estimates that around 20% of UK workers are employed in occupations that are likely to shrink as a proportion of the workforce, while 10% are employed in occupations that are likely to grow.

Overall, their findings suggest declining employment in many low- and middle-educated occupations, and in some professional occupations (such as financial specialists). They also suggest that food preparation, hospitality occupations, and elementary services (including couriers, cleaners, waiters) are likely to grow in importance.

On the one hand, this is broadly consistent with concerns that the future may continue to provide declining opportunities for employment in the middle of the education distribution, and that technology (and other trends) may erode employment in occupation groups that had previously been ‘protected’:

- At the high end of the education distribution, Financial specialists (including fraud examiners, risk management specialists, financial quantitative analysts, and investment underwriters), and to a lesser extent, Health technologists and technicians (e.g. radiologic technicians);[31]
- At the low end of the education distribution, Assemblers and Fabricators, Food processing workers, and other hospitality occupations, among others.

Moreover, many of the occupation groups expected to grow are found at opposite extremes of education requirements:

- Some typically requiring university education (e.g. creative, digital design, engineering, architecture, medical, education occupations);
- Others, in typically low-education services (e.g. animal care and service workers, personal appearance workers)

However, occupations expected to grow also include some middle-education occupations – ‘Construction Trades Workers’ in the US, Sports and Fitness occupations, Electrical and Electronic Trades in the UK). The authors’ interpretation of the findings also suggests that some of the low-education occupations expected to grow (e.g. hospitality occupations) could become

[30] At 4-digit for the UK and 6-digit level for the US, using the most recent available SOC classification: https://www.ons.gov.uk/methodology/classificationsandstandards/standardoccupationalclassificationsoc/soc2010.

[31] Respectively, O*NET-SOC code 13-2099 and 29-2099: https://www.onetonline.org/link/summary/

increasingly better (paid), perhaps in light of recent evidence on the role of new 'craft' activities in the economy (e.g. microbreweries).[32]

Box 4. Detail and discussion of method: The Future of Skills

Bakhshi et al. (2017) use the following method to generate predictions on the likelihood that existing occupations will experience increasing or decreasing demand up to 2030:

- Experts provided their predictions within two workshops (one for each country) hosting 12 to 13 experts from academia, government, industry and the social sector.
- Experts were provided with factsheets for each occupation and invited to consider the potential impact of the following trends: technological change (including progress in AI); globalisation; demographic change; environmental sustainability; urbanisation; increasing inequality; political uncertainty.
- Experts provided directional predictions and an assessment of their (un)certainty on a scale of 0 (not certain at all) to 9 (completely certain).
- The first set of 10 occupations was chosen randomly, the second and third set were chosen so they would be maximally informative for the machine learning algorithm.
- The machine learning algorithm uses predictions about 30 occupation groups generated from expert workshops in combination with detailed information on the 'content' of each of the 30 occupation groups from the US O*NET database (technically, scores against 120 'features' for each occupation group) as a training set to generate predictions about the remaining occupation groups.
- Once predictions have been generated, the study investigates which work features are predictive of changes in future demand.

Centre for Cities (2018) uses data on the location of UK workers to understand how the predictions in Bakhshi et al. (2017) may vary across cities. The study finds that the proportion of workers in occupations likely to shrink varies from 13% (Oxford, Cambridge) to 29% (Mansfield, Sunderland and Wakefield). Cities where this proportion is higher are also likely to benefit the least from the new job creation predicted by Bakhshi et al. (2017). These are also disproportionately cities that already experience relatively poor economic performance (low productivity). Comparing predictions with data on the rise and decline of different occupations over the 20th century, the study argues that the predicted scale of change is not unprecedented. However, the unequal distribution of predicted growth and decline, combined with existing disparities, poses significant challenges.

McKinsey Global Institute: Jobs Lost, Jobs Gained

McKinsey (2017c) models not only the potential displacement effects of automation enabled by AI, but also potential job creation related to six distinct trends. These trends include the development and deployment of new technology. This does not

[32] Felk & Mengedoth (2014).

explicitly include any effect of AI. However, job creation related to this trend is estimated by projecting forward recent employment growth in IT roles (e.g. computer scientists, IT administrators). It is not clear whether this level of growth would or would not be possible in the absence of development and deployment of AI - but if one were to model the impact of AI deployment and development separately, alongside this trendline scenario, there would be a need to worry about potential double-counting.

In any case, the report provides an indication of the potential size of job creation that may counteract (along with other effects specifically linked to AI) possible displacement linked to technology. The five trends considered along with the development and deployment of new technology are: rising consumption in emerging economies; population ageing; investment in infrastructure and buildings; investment in renewable energy, energy efficiency, and climate adaptation; 'marketisation' of previously unpaid domestic work particularly in developing economies.

For consumption and ageing, related employment growth is modelled extrapolating forward recent growth as in the case of technology ('trendline scenario'). For other drivers, the trendline scenario is accompanied by a 'step-up scenario' in which it is assumed that the relevant trend will accelerate compared to the recent past.

Overall, the study suggests that the trends described above can be linked with significant employment growth, at least in the same ballpark as the potential employment losses linked to automation. Job losses from automation are modelled according to McKinsey's 'midpoint scenario' for adoption of AI, described in Section 4.2 of this report. However, employment losses and gains are different in terms of occupations and location, which implies potentially challenging transitions. For example, employment in office support occupations is expected to shrink in advanced economies and grow in developing economies (e.g. China, India, Mexico).

4.4 Evidence on changes in relationships between workers, employers, and technology

Evidence presented in Sections 4.1 to 4.3 was largely related to the number and types of jobs that are expected to grow or decline in the future, while AI is increasingly adopted in production processes. In this Section, we report evidence on how digital technology is affecting different areas of economic activity, and on the likely impact of AI on aspects of job quality other than earnings.[33] The areas of economic activity mentioned in this Section are not necessarily those where AI has had or is expected to have the greatest impact.

[33] For a definition of job quality please see Section 4.4.2 below.

4.4.1 Case studies of automation in specific settings

Service industries and technical occupations

Lacity and Willcocks (2016) study how 'Robotic Process Automation' (defined as 'tools and platforms that deal with structured data, rules-based processes, and deterministic outcomes') has been adopted in 11 services industries.[34] The study relies on surveys of outsourcing professionals and interviews with 48 people, including service automation adopters, software providers, and management consultants. This article does not set out to provide a representative view of how service automation is being implemented, but rather to investigate why firms are adopting service automation, what outcomes they are achieving, and what makes service automation beneficial to the firm.

All the cases of automation considered in this article involve integration between human and machine capabilities, or automation leading to refocus and restructure human tasks without leading to significant displacement.

In one case, an advantage of automation was increasing the firm's ability to scale up and down routine interactions with customers to accommodate different phases in the product life cycle. Automated routine interactions could increase automatically when required, for example to support the launch of a new product, while workers kept focussing on more complex cases of interactions with customers. The idea of changing market conditions influencing the interaction between workers and machines is also relevant in Wilson et al. (2017) and Brynjolfsson & Mitchell (2017): when a situation requires new tasks being performed (as in the case of the launch of a new product), these tasks may be first performed by workers, and then devolved to machines once sufficient data to train machine learning algorithms have been collected.

Other examples of automation provided in Lacity and Willcocks (2016) involve a redesign of workers' responsibilities towards less routine tasks. One of these cases is the automated production of corporate earnings reports, adopted by the Associated Press (AP) in 2014. Reporters who produce the corporate earnings reports before this was automated were not laid off. Rather, these reporters re-focussed on other stories that could not be produced through an automated process. It is possible that this outcome was driven, or at least influenced, by the characteristics of employer-employee relations in this industry. The authors note that reporters employed by the AP are highly unionised but does not comment on the role of unions in shaping the outcome of this automation process.

While this picture of service automation does not provide examples of workers being displaced as a result, it is worth noting that:

- The automation considered here is likely applicable to a narrower set of tasks compared to those that could be performed by AI systems. It is not clear where more advanced stages of automation may in fact lead to lay-offs.
- The restructuring of occupations described here implies that skills such as creativity and problem-solving are increasingly valuable. While restructuring of

[34] Including Healthcare, Energy, Telecommunications, Media, Financial and accounting services, and Transportation.

this type may not lead to mass layoffs, workers may be required to adapt to significantly different tasks, and the transition may be challenging for some.

The arguments in Liew (2018) and Recht and Bryan (2017) around the use of AI in radiology echo some of Lacity and Willcocks' findings about the potential for integration between workers and machines in other occupations and industries. Liew (2018) describes five key areas of radiology where AI can already be applied given existing technical capabilities, and argues that radiologists should take advantage of these capabilities to deliver imaging at lower cost, and leverage their ability to make clinical decisions based on imaging data.

Manufacturing

Eurofound (2017a) reports evidence from existing literature and interviews with experts on the adoption of robotics in manufacturing industries in Europe. The report focusses on 'Advanced Industrial Robotics' (AIR), defined as 'robots working within industrial environments, which are equipped with advanced functionalities (for example, sensing potential collisions and halting or performing a programme motion with a small offset) allowing them to operate in less structured contexts.

The ability of AIR to operate in less structured contexts implies greater flexibility of production, including the potential to introduce greater variation (personalisation) and to adapt the scale of production to changes in customer demands more quickly. In terms of the economic framework presented in section 3.3, businesses aim to use AIR not only to achieve greater efficiency (lower prices) but also better quality, and to introduce new products. Consistent with that economic framework, experts participating in this study expect these improvements to lead to greater demand for manufacturing products, countervailing the negative impact of automation on the demand for workers in manufacturing.

As robots can take on an increasing range of tasks, the role of workers is expected to shift away from manual activities (including hazardous tasks) and towards the supervision of, and integration with, automation. The roles of production line supervisors and operators, production managers, and forklift operators are expected to change significantly and require upskilling. This could be challenging, as the new roles are likely to involve mixes of skills that are not often combined in existing occupations (note that the recombination of existing tasks into new occupations is a key theme in Brynjolfsson et al., 2018, as seen in Section 4.2). However, this study also suggests that a significant proportion of the required upskilling could be achieved through on-the-job training.

The ageing of the industrial workforce is mentioned as a significant driver of robot adoption, consistent with the insights in Acemoglu & Restrepo (2018b) seen in section 3.2. Social and regulatory factors are not mentioned as important drivers, but the study notes that defining who should be liable for any work-related accidents caused by robots could be problematic. The legal issues around this could be complex due to the variety of commercial arrangements between robot producers and users.

The study also discusses potential implications of AIR for the quality of manufacturing jobs. These findings are reported in section 4.4.2 below.

Digital platforms

The emergence of digital ‘platforms’ (e.g. Uber, Deliveroo, Lyft), connecting consumers and workers through relatively little intervention of an intermediary, has been widely discussed for its implications on work, market competition (e.g. Evans & Schmalensee, 2014), and other outcomes (e.g. traffic congestion). These new forms of work are enabled by digital technology - not necessarily by AI specifically. However, increasing use of digital technology, including AI, could make some of the aspects of these forms of work increasingly relevant:

- Work being organised largely through algorithms, rather than decisions of a larger number of human managers. As a result, workers interacting increasingly with machines (and customers) and less often with co-workers;
- The division of some types of work into smaller tasks that can be sourced from a large pool of potential workers (‘crowdsourced’), as in the case of Amazon Mechanical Turk.[35]

Lee, Kusbit, Metsky & Dabbish (2018) and Rosenblat & Stark (2016) investigate how drivers who use ride-sharing applications (Uber, Lyft) in the United States respond to ‘algorithmic management’: the automated mechanisms that ride-sharing platforms use to direct their behaviour. Both studies analyse posts by drivers in dedicated online forums and performed semi-structured interviews with drivers (21 and seven, respectively) and passengers (12 in Lee et al.). This research suggests that:

- Drivers valued the flexibility provided by the platforms;
- However, drivers felt relatively little control over several aspects of their work, which they linked to:
  - Specific explicit policies around what should or should not be done to use the app, and sudden changes in those policies: for example, the inability to refuse a ride request from a potential passenger based on its destination, or to request payment to return items left behind by passengers;
  - Mechanisms or design choices that influence their behaviour, more or less explicit (from the use of ‘surge’ pricing, to messages around likely future demand);
  - The need to maintain a high rating from passengers to keep driving through the platform.

Rosenblat & Stark (2016) conclude that despite Uber’s limited explicit managerial role, its automated systems have considerable power to control how workers behave, and that workers have limited scope to influence these systems.

[35] An important issue concerns the legal status of those who work through digital platforms – whether self-employed, employees, or somewhere in between. In the UK, this issue has been recently analysed in depth by Matthew Taylor’s ‘Review of Modern Working Practices’. This review does not take a view on this debate – we simply refer to people offering their work on digital platforms as ‘workers’ and describe their relation to the platform as ‘users’, although we recognise that other terms may be more appropriate.

4.4.2 Effects on the quality of the working environment

International organisations have defined a number of frameworks that aim to define and support the measurement of 'job quality'. These frameworks include, among others, those developed by the International Labour Organisation (ILO, 2012), the United Nations Economic Commission for Europe (UNECE, 2015), the European Foundation for the Improvement of Living and Working Conditions (Eurofound, 2013), and the OECD (Cazes, Hijzen & Saint-Martin, 2016). A common thread across these frameworks is that they consider earnings as one element among a set of factors that determine how work affects different aspects of people's well-being. For example, the OECD Job Quality Framework suggests that measuring the quality of jobs held in a country and/or by a group of workers should include:

- Earnings, including both average earnings and earnings inequality;
- Labour market security, including the risk of falling into unemployment or very low pay, as well as the degree of insurance against unemployment and the probability of getting out of very low pay;
- Quality of the working environment including:
  - Factors that have a negative impact on quality: time pressure and physical health risk factors;
  - Factors that have a positive impact on quality: work autonomy and learning opportunities, and good workplace relationships.[36]

The deployment of AI could have an impact not only on earnings and labour market security, as discussed throughout this report, but also on the quality of the working environment. This is noted in RSA (2017), and it is also consistent with the changes in working environment observed during the first industrial revolution (as discussed in section 3.1) and with the emergence of digital platforms in several markets.

This review has not identified predictions of the potential impact of AI on job quality. RSA (2017), a recent review of existing literature on automation (Eurofound, 2017a, summarised earlier in this section) and Eurofound (2017b) suggest that the use of AI could lead to positive or negative changes in the quality of the working environment:

- The automation of routine tasks, potentially allowing workers to refocus on non-routine activities and on the supervision of automated work could lead to greater autonomy and learning opportunities for some roles (as in the case of production line supervisors delineated in Eurofound, 2017a). The proliferation in the workplace of machines, typically equipped with sensors, could make workers subject to increasing degree of monitoring, with potential negative impacts on their autonomy.
- The increasing use of robotics in manufacturing could reduce the need for workers to perform dangerous and physically demanding tasks. However, as advances in industrial robotics allow greater interaction between robots and

[36] In Cazes et al. (2016), 'quality of the working environment' captures 'non-economic aspects of employment'.

humans, new risks for health and safety may emerge, at least in an initial period of adoption.

- The impact of AI on working relationships is likely to be complex and to change from role to role. On the one hand, automation may lead increasing numbers of workers to interacting primarily with machines rather than with co-workers. On the other, if work that involves interaction with others resists automation (as suggested for some types of interaction by several of the studies reviewed in section 4.2), an increasing proportion of work tasks will involve interacting with others.

A broader literature has considered the role of digital technology in shaping the quality of work. A recent contribution in this area is Evans & Kitchin (2018), which examines how the use of extensive information management (or ‘big data’) systems affects worker satisfaction, drawing on an ethnography of a retail store in Ireland. In the store, a number of digital devices guide and generate continuously fine-grained data on workers’ activities. This changed working practices in at least two ways:

- A greater emphasis on measurable aspects of work as indicators of performance and drivers of pay (to the expense of ‘symbolic work’, including the amount and characteristics of interactions with customers);
- A move from direct monitoring based on observation from a supervisor to continuous monitoring based on data.

The changes led to worker dissatisfaction, related to the undervaluation of symbolic work, and a lack of understanding of the operational parameters of the organisation as a whole, leading to worker disengagement. The authors note that redesigning how big data systems are used could remove or mitigate some of these issues.

RSA (2017) also notes that AI could be used to automate recruitment processes. Here caution will be needed to ensure that AI, if used, does not entrench existing biases (or introduce other forms of bias). This is a possible risk in using past recruitment decisions to train machine learning algorithms that could be applied to screening future applications. As noted in Citron & Pasquale (2014) in reference to credit scoring, limiting the fairness issues related to the use of automated decision making may require regulatory oversight. Yet there is potential to improve recruitment decisions, reducing rather than magnifying existing bias against groups of applicants. Kahneman (2018) and Thaler (2015) note that in many cases where humans are required to make decisions, they are outperformed by even simple statistical models – and recruitment may be one of these cases (Thaler, 2015). Managers’ reluctance to accept giving up decision to machines could be a significant factor delaying the adoption of further automation.

5 CONCLUSIONS

This review has drawn on literature from a range of disciplines to inform on-going discussions about the potential impact of AI on work.

Evidence from economics, history, sociology and other disciplines shows that the impact of AI is likely to be influenced not only by technology but also by cultural, economic, social factors. Indeed, the way in which work is organised can change significantly across borders even where the same technology is used in production.

There is evidence that digital technology and automation have already significantly affected work over and above the role of trade liberalisation. This evidence shows that the use of digital technology in work is linked with increasing polarisation of work between jobs mainly performed by workers with low levels of formal education ('low-educated') and jobs performed by high-educated workers.

Employment overall has not declined, but there have been shifts from manufacturing to services, across geographical areas, and income losses for displaced low-educated workers. Individual losses from displacement related to automation have not yet been estimated but a broader literature suggests that these losses can be significant and persistent.

There is significant uncertainty around the technical potential for AI to automate tasks currently performed in the workplace. Several studies suggest that AI could affect at least a significant minority of existing jobs, and that occupations performed by low-educated workers more likely to be affected compared to those performed by high-educated workers.

If AI is widely adopted and this leads to automation, economic modelling suggests that potential job losses in the short term will be compensated to an extent by countervailing mechanisms through which increasing productivity leads to greater demand for labour. This may nevertheless lead to significant increases in inequality, particularly if employers have significant market power.

We have identified the following gaps in evidence base:

- There is limited evidence on how AI is being used now and on how workers' tasks have changed where this has happened.
- There has been relatively little discussion of how existing institutions, policies, social responses are shaping and are likely to shape the evolution of AI and its adoption.
- We have identified little consideration of how international trade, mobility of capital and of AI researchers are shaping the development of AI and therefore its potential impact on work.

REFERENCES

Abramovitz, M. & David, P.A (1996). Convergence and Delayed Catch-Up: Productivity Leadership and the Waning of American Exceptionalism. In Landau, R., Taylor, T. & Wright, G. (eds.), *The Mosaic of Economic Growth*. Stanford, CA: Stanford University Press.

Acemoglu, D., & Autor, D. (2011). Skills, tasks and technologies: Implications for employment and earnings. In Handbook of labor economics (Vol. 4, pp. 1043-1171). Elsevier.

Acemoglu, D. & Restrepo, P. (2016). *The Race Between Man and Machine: Implications of Technology for Growth, Factor Shares and Employment* (NBER Working Paper No. 22252). Cambridge, MA: National Bureau of Economic Research

Acemoglu, D. & Restrepo, P. (2017). *Robots and Jobs: Evidence from US Labor Markets* (NBER Working Paper No. 23285). Cambridge, MA: National Bureau of Economic Research

Acemoglu, D. & Restrepo, P. (2018a). *Artificial Intelligence, Automation and Work* (NBER Working Paper No. 24196). Cambridge, MA: National Bureau of Economic Research

Acemoglu, D. & Restrepo, P. (2018b). Low-Skill and High-Skill Automation. *Journal of Human Capital* 12(2), 204-232.

Acemoglu, D. & Restrepo, P. (2018c). Demographics and Automation (NBER Working Paper No. 24196). Cambridge, MA: National Bureau of Economic Research

Aghion, P., Jones, B. and Jones, C. (2017). *Artificial Intelligence and Economic Growth* (NBER Working Paper No. 23928). Cambridge, MA: National Bureau of Economic Research

Agrawal, A., Gans, J.S. & Goldfarb, A. (2018). *Prediction, Judgment, and Complexity* (NBER Working Paper No. 24243). Cambridge, MA: National Bureau of Economic Research

Allen, R.C. (2009). Engels' pause: Technical change, capital accumulation, and inequality in the British industrial revolution. *Explorations in Economic History* 46(4), 418-435

Allen, R.C. (2017). Lessons from history for the future of work. *Nature* 550, 321–324. Retrieved from: https://www.nature.com/news/lessons-from-history-for-the-future-of-work-1.22825

Andrews, D., Criscuolo, C. and Gal, P.N. (2015). *Frontier Firms, Technology Diffusion and Public Policy: Micro Evidence from OECD Countries* (OECD Productivity Working Papers November 2015, No. 02). Retrieved from: https://www.oecd-ilibrary.org/economics/frontier-firms-technology-diffusion-and-public-policy_5jrql2q2jj7b-en

Aral, S., Brynjolfsson, E. & Wu, L. (2012). Three-Way Complementarities: Performance Pay, HR Analytics and Information Technology. *Management Science*, 58(5), 913-931.

Arntz, M., Gregory, T. & Ziehran, U. *The Risk of Automation for Jobs in OECD Countries* (OECD Social, Employment and Migration Working Papers No. 189). Retrieved from: https://www.keepeek.com//Digital-Asset-Management/oecd/social-issues-migration-health/the-risk-of-automation-for-jobs-in-oecd-countries_5jlz9h56dvq7-en#page1

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

Autor, D., Levy, F. and Murnane, R.J. (2003). The Skill Content of Recent Technological Change: An Empirical Exploration. *The Quarterly Journal of Economics* 118(4), 1279-1333.

Autor, D., Dorn, D., Katz, L.F., Patterson, C. & Van Reenen, J. (2017). The Fall of the Labor Share and the Rise of Superstar Firms. Cambridge, MA: Massachusetts Institute of Technology. Retrieved from https://economics.mit.edu/files/12979

Bakhshi, H., Downing, J.M., Osborne, M.A & Schneider, P. (2017). *The Future of Skills: Employment in 2030*. Report prepared by Nesta and Oxford Martin School. Retrieved from: https://www.nesta.org.uk/sites/default/files/the_future_of_skills_employment_in_2030_0.pdf

Ball, K. (2010). Workplace Surveillance: an Overview. *Labor History* 51(1), 87-106

Bessen, J. (2017). *Automation and Jobs: When technology boosts employment* (Boston University School of Law, Law and Economics Research Paper No. 17-09). Boston, MA: Boston University School of Law.

Bloom, N., Draca, M., & Van Reenen, J. (2016). Trade Induced Technical Change? The Impact of Chinese Imports on Technology, Jobs and Plant Survival. *The Review of Economic Studies* 83(1), 87-117

Bloom, N., Sadun, R. & Van Reenen, J. (2012). Americans Do IT Better: US Multinationals and the Productivity Miracle. *American Economic Review*, 102(1), 167-201

Blundell, R., Green, D. & Jin, W. (2018). The UK Education Expansion and Technological Change. Retrieved from http://www.ucl.ac.uk/~uctp39a/BGJ_Jan_22_2018.pdf

Bresnahan, T., Brynjolfsson, E., & Hitt, L. (2002). Information Technology, Workplace Organization, and the Demand for Skilled Labor: Firm-Level Evidence. *Quarterly Journal of Economics*, 117(1), 339-376.

Brynjolfsson, E. & McAfee, A. (2014). The Second Machine Age: Work, Progress and Prosperity in a Time of Brilliant Technologies. WW Norton & Company.

Broadberry, S. & Wallis, J. (2016). *Growing, Shrinking and Long Run Economic Performance: Historical Perspectives on Economic Development* (NBER Working Paper No. 23343). Cambridge, MA: National Bureau of Economic Research

Brynjolfsson, E. & Hitt, L. (2003). Computing Productivity: Firm-Level Evidence. *The Review of Economics and Statistics* 85(4), 793-808

Brynjolfsson, E. & Mitchell, T. (2017). What Can Machine Learning Do? Workforce Implications. *Science* 358 (6370), 1530-1534

Brynjolfsson, E., Rock, D. & Syverson, C. (2017). *Artificial Intelligence and the Modern Productivity Paradox: A Clash of Expectations and Statistics* (NBER Working Paper No. 24001). Cambridge, MA: National Bureau of Economic Research

Brynjolfsson, E., Mitchell, T. & Rock, D. (2018). What Can Machines Learn and What Does It Mean for Occupations and the Economy?. *AEA Papers and Proceedings* 108: 43-47

Caselli, F. & Manning, A. (2017). Robot Arithmetic: Can New Technology Harm All Workers or the Average Worker?. Centre for Economic Performance Discussion Paper No. 1497. London, UK: London School of Economics. Retrieved from: http://cep.lse.ac.uk/pubs/download/dp1497.pdf

Cazes, S., Hijzen, A. & Saint-Martin, A. (2015). *Measuring and Assessing Job Quality: The OECD Job Quality Framework* (OECD Social, Employment and Migration Working Papers, No. 174). Paris: OECD Publishing

Centre for Cities (2018). *Cities Outlook 2018*. London, UK: Centre for Cities. . Retrieved from http://www.centreforcities.org/wp-content/uploads/2018/01/18-01-12-Final-Full-Cities-Outlook-2018.pdf

Citron, D. K., & Pasquale, F. (2014). The scored society: due process for automated predictions. *Washington Law Review*, 89, 1

Clark, G. (2005). The Condition of the Working Class in England, 1209-2004. *Journal of Political Economy* 113(6), 1307-1340

Clark, G. (2014). The Industrial Revolution. In Aghion, P. & Durlauf, S.N. (Eds.), *Handbook of Economic Growth*, Vol.2, (pp.217-262). Elsevier

Comin, D. & Hobijn, B. (2010). An Exploration of Technology Diffusion. *American Economic Review* 100, 2031–2059

Crafts, N. (2010). *The Contribution of New Technology to Economic Growth: Lessons from Economic History* (CAGE Online Working Paper Series 01, Competitive Advantage in the Global Economy). Retrieved from: https://warwick.ac.uk/fac/soc/economics/research/centres/cage/manage/publications/01.2010_crafts.pdf

Crafts, N. & O'Rourke, K.H. (2014). Twentieth Century Growth. In Aghion, P. and Durlauf, S.N. (eds.), *Handbook of Economic Growth,* Vol.2 (pp.263-346). Elsevier

Dauth, W., Findeisen, S., Südekum, J. & Wößner, N.. (2017). German Robots – The Impact of Industrial Robots on Workers (Institute for Employment Research of the Federal Employment Agency Working Paper 30/2017). Retrieved from http://doku.iab.de/discussionpapers/2017/dp3017.pdf

De Loecker, J. & Eeckhout, J. (2017). The Rise of Market Power and the Macroeconomic Implications. (NBER Working Paper No. 23687). Cambridge, MA: National Bureau of Economic Research

Diamond, R. (2016). The determinants and welfare implications of us workers' diverging location choices by skill: 1980-2000. *American Economic Review*, 106(3), 479-524

Dwyer, R.E. (2013). The Care Economy? Gender, Economic Restructuring, and Job Polarization in the U.S. Labor Market. *American Sociological Review* 78(3)

Elliott, S.W. (2017). Artificial intelligence and the future of work and skills: will this time be different?. Blog post Retrieved from https://www.oecd-forum.org/users/69561-stuart-w-elliott/posts/21601-artificial-intelligence-and-the-future-of-work-and-skills-will-this-time-be-different

Eurofound (2012). *Trends in Job Quality in Europe*. Luxembourg: Publications Office of the European Union.

Eurofound (2017a). Advanced industrial robotics: Taking human-robot collaboration to the next level. Retrieved from https://www.eurofound.europa.eu/sites/default/files/wpfomeef18003.pdf

Eurofound (2017b). Automation of Work Literature Review. Retrieved from https://www.eurofound.europa.eu/sites/default/files/wpef17039.pdf

Evans, D. & Schmalensee, R. (2014). The Antitrust Analysis of Multisided Platform Businesses. In Blair, R.D. & Sokol, D.D. (eds.). *The Oxford Handbook of International Antitrust Economics,* Volume 1. Oxford University Press

Evans, L. & Kitchin, R. (2018). A smart place to work? Big data systems, labour, control and modern retail stores. New Technology, Work and Employment 33(1), 44-57

Executive Office of the United States President (2016). *Artificial Intelligence, Automation, and the Economy*. Washington, DC. Retrieved from: https://obamawhitehouse.archives.gov/sites/whitehouse.gov/files/documents/Artificial-Intelligence-Automation-Economy.PDF

Felk, J. & Mengedoth, J. (2014). Bottoms Up. Retrieved from https://www.richmondfed.org/publications/research/econ_focus/2014/q4/feature2

Ferguson, N. (2017). The Square and the Tower: Networks, Hierarchies and the Struggle for Global Power. Allen Lane.

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

Fouquet, R. & Broadberry (2015). S. Seven Centuries of Economic Growth and Decline. *Journal of Economic Perspectives* 29(4), 227-244

Francois, P. (2018). *Discussion of: Artificial Intelligence and the Future of Growth*. In Agrawal, A.K., Gans, J. & Goldfarb, A. (Eds.), *The Economics of Artificial Intelligence: An Agenda*. Forthcoming from University of Chicago Press. Retrieved from: http://www.nber.org/chapters/c14028.pdf

Frey, C.B. & Osborne, M. (2017). The future of employment: how susceptible are jobs to computerisation?. *Technological Forecasting and Social Change*, 114, 254-280

Furman, J. (2016, July 7). Is This Time Different? The Opportunities and Challenges of Artificial Intelligence. Remarks at *AI Now: The Social and Economic Implications of Artificial Intelligence Technologies in the Near Term*, New York University. Retrieved from https://obamawhitehouse.archives.gov/sites/default/files/page/files/20160707_cea_ai_furman.pdf

Galor, O. & Moav, O. (2002). Natural Selection and the Origin of Economic Growth. *Quarterly Journal of Economics* 117(4), 1133-91

Gallie, D. (1978). *In Search of the New Working Class*. Cambridge, UK: Cambridge University Press

Goos, M. & Manning, A. (2004). Lousy and Lovely Jobs: the Rising Polarization of Work in Britain. *The Review of Economics and Statistics*, 89(1), 118-133

Gordon, R.J. (2016). Perspectives on *The Rise and Fall of American Economic Growth*. *American Economic Review: Papers & Proceedings* 106(5), 1-7.

Graetz, G. & Michaels, G. (2017). Is Modern Technology Responsible for Jobless Recoveries?. *American Economic Review*, 107(5), 168–73.

Graetz, G. & Michaels, G. (2018). Robots at Work. forthcoming in *The Review of Economics and Statistics*. Retrieved from: http://personal.lse.ac.uk/michaels/Graetz_Michaels_Robots.pdf

Gray, M., Suri, S., Ali, S.S. & Kulkarni, D. (2016). The Crowd is a Collaborative Network. Proceedings of the 19th ACM Conference on Computer-Supported Cooperative Work & Social Computing, 134-147

Gregory, T. et al (2016). Racing With or Against the Machine? Evidence from Europe. ZEW Discussion Paper No. 16-053. Retrieved from: https://ub-madoc.bib.uni-mannheim.de/41403/1/dp16053.pdf

Haldane, A. (2017). Speech given at the London School of Economics. Retrieved from: https://www.bankofengland.co.uk/-/media/boe/files/speech/2017/productivity-puzzles.pdf

Hall, W. and Pesenti, J. (2017). Growing the Artificial Intelligence Industry in the UK. Report prepared for the Department for Business, Energy and Industrial Strategy and the Department for Digital, Culture, Media and Sport. Retrieved from: https://assets.publishing.service.gov.uk/government/uploads/system/uploads/attachment_data/file/652097/Growing_the_artificial_intelligence_industry_in_the_UK.pdf

House of Commons Science and Technology Committee (2016). Robotics and artificial intelligence. Fifth report of Session 2016-17. Retrieved from: https://publications.parliament.uk/pa/cm201617/cmselect/cmsctech/145/145.pdf

House of Lords Select Committee on Artificial Intelligence (2018). AI in the UK: ready, willing and able? Report of Session 2017-19. Retrieved from: https://publications.parliament.uk/pa/ld201719/ldselect/ldai/100/100.pdf

Humphries, J. & Weisdorf, J. (2015). The Wages of Women in England, 1260-1850. *The Journal of Economic History*, 75(2), 405-447

Humphries, J. (2016). Spinning the Industrial Revolution. University of Oxford Discussion Papers in Economic and Social History No. 145. Oxford, England: University of Oxford.

International Labour Organization (2012). *Decent Work Indicators: Concepts and Definitions* (ILO Manual). Geneva: International Labour Organization.

Jurgens, U., Malsch, T. & Dohse, K. (1993). *Breaking from Taylorism. Changing Forms of Work in the Automobile Industry*. Cambridge, UK: Cambridge University Press

Kahneman, D. (2018). Commentary on Camerer (2018), Artificial Intelligence and Behavioral Economics, Chapter in forthcoming NBER book The Economics of Artificial Intelligence: An Agenda. Retrieved from http://www.nber.org/chapters/c14016.pdf

Korinek, A. & Stiglitz, J. (2017). Artificial Intelligence and Its Implications for Income Distribution and Unemployment. Chapter in forthcoming NBER book *The Economics of Artificial Intelligence: An Agenda*. Retrieved from: http://www.nber.org/chapters/c14018.pdf

Krugman, P. (2008). Trade and Wages Reconsidered (Brookings Papers on Economic Activity, 2008. No. 1), 103-154. Washington, DC: Brookings Institution.

Lacity, M.C. & Willcocks, L.P. (2016). A new approach to automating services. *MIT Sloan Management Review*, 58(1), 41

Lee, M,K., Kusbit, D., Metsky, E. & Dabbish, L. (2018). Working with Machines: The Impact of Algorithmic and Data-Driven Management on Human Workers. *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems*, 603-1612

Liew, C. (2018). The future of radiology augmented with Artificial Intelligence: A strategy for success. *European Journal of Radiology* 102, 152-156

Mann, K. & Püttmann, L. (2017). Benign Effects of Automation: New Evidence from Patent Texts. Retrieved from: https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2959584

Manning, A., & Petrongolo, B. (2017). How local are labor markets? Evidence from a spatial job search model. *American Economic Review*, 107(10), 2877-2907

McKinsey Global Institute (2017a). A Future that Works: Automation, Employment, and Productivity. . Retrieved from: https://www.mckinsey.com/~/media/McKinsey/Global%20Themes/Digital%20Disruption/Harnessing%20automation%20for%20a%20future%20that%20works/MGI-A-future-that-works-Executive-summary.ashx

McKinsey Global Institute (2017b). Artificial Intelligence: the Next Digital Frontier? Discussion Paper. Retrieved from:https://www.mckinsey.com/~/media/McKinsey/Industries/Advanced%20Electronics/Our%20Insights/How%20artificial%20intelligence%20can%20deliver%20real%20value%20to%20companies/MGI-Artificial-Intelligence-Discussion-paper.ashx

McKinsey Global Institute (2017c). Jobs Lost, Jobs Gained: Workforce Transitions in a Time of Automation. Retrieved from: https://www.mckinsey.com/~/media/McKinsey/Global%20Themes/Future%20of%20Organizations/What%20the%20future%20of%20work%20will%20mean%20for%20jobs%20skills%20and%20wages/MGI-Jobs-Lost-Jobs-Gained-Report-December-6-2017.ashx

McKinsey Global Institute (2018). Notes from the AI Frontier – Insights from Hundreds of Use Cases. Discussion Paper. Retrieved from: https://www.mckinsey.com/~/media/McKinsey/Featured%20Insights/Artificial%20Intelligence/Notes%20from%20the%20AI%20frontier%20Applications%20and%20value%20of%20deep%20learning/MGI_Notes-from-AI-Frontier_Discussion-paper.ashx

Michaels, G., Natraj, A. & Van Reenen, J. (2014). Has ICT Polarized Skill Demand? Evidence from Eleven Countries over Twenty-Five Years. *The Review of Economics and Statistics* 96(1), 60-77

Mokyr, J. (1990). *The Lever of Riches*. Oxford, UK: Oxford University Press

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

National Academy of Sciences (2017). Information Technology and the U.S. Workforce: Where Are We and Where Do We Go from Here?. Washington, DC: The National Academies Press. Retrieved from: https://www.nap.edu/catalog/24649/information-technology-and-the-us-workforce-where-are-we-and

Nedelkoska, L. & Quintini, G. (2018). Automation, skills use and training (OECD Social, Employment and Migration Working Papers, No. 202). Paris, France: OECD. Retrieved from: https://read.oecd-ilibrary.org/employment/automation-skills-use-and-training_2e2f4eea-en#page1

North, D., Wallis, J.J. & Weingast, B.R. (2009). *Violence and Social Orders: A Conceptual Framework for Interpreting Recorded Human History*, Cambridge: Cambridge University Press

Office for National Statistics (2013). Women in the labour market. Retrieved from https://www.ons.gov.uk/employmentandlabourmarket/peopleinwork/employmentandemployeetypes/articles/womeninthelabourmarket/2013-09-25

Pasquale, F. (2016). Two narratives of platform capitalism. *Yale Law & Policy Review*, 35

http://beyondthetechrevolution.com/blog/second-machine-age-or-fifth-technological-revolution/PWC (2017). UK Economic Outlook March 2017 – Will robots steal our jobs?. PWC Report. Retrieved from: https://www.pwc.co.uk/economic-services/ukeo/pwcukeo-section-4-automation-march-2017-v2.pdf

PWC (2018). Will robots really steal our jobs?. PWC Report. Retrieved from: https://www.pwc.co.uk/economic-services/assets/international-impact-of-automation-feb-2018.pdf

Recht, M. & Bryan, N. (2017). Artificial Intelligence: Threat or Boon to Radiologists?. *Journal of the American College of Radiology* 14(11), 1476-1480

Rosenberg, N. (1969). The Direction of Technological Change: Inducement Mechanisms and Focusing Devices. *Economic Development and Cultural Change* 18(1), 1-24

Rosenblat, A. & Stark, L. (2016). Algorithmic Labor and Information Asymmetries: A Case Study of Uber's Drivers. *International Journal of Communication* 10, 3758-3784

Royal Society (2017). Machine learning: the power and promise of computers that learn by example. Retrieved from https://royalsociety.org/~/media/policy/projects/machine-learning/publications/machine-learning-report.pdf

Royal Society and Academy for Medical Sciences (2018). Evidence synthesis for policy – A statement of principles. Retrieved from: https://royalsociety.org/~/media/policy/projects/evidence-synthesis/evidence-synthesis-statement-principles.pdf

RSA (2017). *The Age of Automation: Artificial Intelligence, robotics and the future of low-skilled work*. Retrieved from https://www.thersa.org/globalassets/pdfs/reports/rsa_the-age-of-automation-report.pdf

Stone, P., Brooks, R., Brynolfsson, E., Calo, R., Etzioni, O., Hager, G., Hirschberg, J., Kalyanakrishnan, S., Kamar, E., Kraus, S., Leyton-Brown, K., Parkes, D., Press, W., Saxenian, A., Shah, J., Tambe, M. &Teller, A. Artificial Intelligence and Life in 2030. One Hundred Year Study on Artificial Intelligence: Report of the 2015-2016 Study Panel. Stanford, CA: Stanford University. Retrieved from: https://ai100.stanford.edu/sites/default/files/ai100report10032016fnl_singles.pdf

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

Temin, P. & Voth, H.J. (2005). Credit Rationing and Crowding Out During the Industrial Revolution: Evidence from Hoare's Bank, 1702-1862. *Explorations in Economic History* 42(3), 325-348

Thaler, R. (2015). Who's Afraid of Artificial Intelligence?. Response posted at https://www.edge.org/response-detail/26083

UNECE (2015*). Statistical Framework for Measuring Quality of Employment*. New York, NY and Geneva: United Nations Economic Commission for Europe

UK Commission for Employment and Skills (2014). The Future of Work: Jobs and skills in 2030. Evidence Report 84. Retrieved from: https://assets.publishing.service.gov.uk/government/uploads/system/uploads/attachment_data/file/303334/er84-the-future-of-work-evidence-report.pdf

Urry, J. (2016). What Is the Future?. Polity Press

Van Reenen, J., Bloom, N., Draca, M., Kretschmer, T. & Sadun, R. (2010). The Economic Impact of ICT (SMART report N.2007/020). Retrieved from https://warwick.ac.uk/fac/soc/economics/staff/mdraca/cstudytheeconomicimpactofictlondonschoolofeconomics.pdf

Wajcman, J. (2006). New connections: social studies of science and technology and studies of work. *Work, employment and society* 20(4), 773-786

Wajcman, J. (2017). Automation: is it really different this time?. *British Journal of Sociology* 68(1), 119-127

Wallis, J. J. (2011). Institutions, organizations, impersonality, and interests: The dynamics of institutions. *Journal of Economic Behavior & Organization*, 79(1-2), 48-64

Wilson, J.H., Daugherty, P.R. & Morini-Bianzino, N. (2017). The Jobs that Artificial Intelligence Will Create. *MIT Sloan Management Review*. Retrieved from: https://sloanreview.mit.edu/article/will-ai-create-as-many-jobs-as-it-eliminates/

Yin, M., Gray, M.L., Suri, S. & Wortman Vaughan, J. (2016). The Communication Network within the Crowd. *Proceedings of the 25th International Conference on World Wide Web*, 1293-1303

ANNEX 1: FURTHER DETAIL ON METHODOLOGY

This review consisted of three key steps:

1. Identifying potentially relevant evidence;
2. Selecting relevant evidence;
3. Analysing the selected evidence.

The review focussed on English language publications, published since 2000. We considered as potentially relevant all publications – regardless of their discipline - that:

1. Provided quantitative estimates of the impact of AI on job destruction, job creation, net levels of employment, or earnings.
2. Provided quantitative and qualitative information that could be used to generate, improve or assess quantitative estimates of the impact of AI.

We were also guided by the questions listed in Figure 6 overleaf, agreed with the Royal Society (RS) and the British Academy (BA)in early stages of this review.

In this phase of the project we aimed to cast a wide net, only excluding publications that are clearly far from the two objectives above. We also set out below further steps to ensure the transparency of our selection process.

To identify the academic literature, we used a combination of searches on key databases (JSTOR, Google Scholar, ScienceDirect) and 'snowballing' (searching for publications that have cited or have been cited by relevant contributions identified previously), as well as building on any existing review publications and following suggestions provided by RS and BA, attendees to workshops held on March 9th and May 15th 2018, and researches and policymakers interviewed as part of this project.

Publications considered to be relevant were recorded in a standardised format in a spreadsheet table, including basic bibliographic information and abstract, where available ('the long list'). The long list included 167 publications, from which we agreed with the RS and BA a short list of 42 to be reviewed in detail.

Due to the relatively short time frame of this project, we restricted our scope to articles published from 2000, and we have not been able to include books among the publications reviewed in detail. Where a book that would be close to the 'core' of the topic map presented above was identified, we aimed to reflect as much as possible the author's view by relying on alternative sources, including:

- Academic articles or working papers published by the author on the topic;
- Other articles published by the author;
- Discussion of the author's findings and views in other researchers' publications;
- Interviews with researchers and policymakers, and events convened by the RS and BA as part of their on-going programme of work on Artificial Intelligence.

For example, we were not able to review Susskind and Susskind's (2105) book 'The Future of the Professions'. However, we aimed to incorporate findings from this work through the following: an interview with Richard Susskind; including articles published by the authors in our long list of relevant publications; attending a workshop organised by the RS and the BA, held in London on March 15th, 2018, where Richard Susskind gave a keynote speech.

Figure 6 High-level questions and specific issues discussed in the literature

High-level question	Specific issues	Selected drivers
What is AI be capable of doing and how will that change in the next 5-10 years?	What tasks will AI be capable of performing more effectively than humans?	Technological: current capabilities of AI and likely speed of improvement
		Cultural, economic, political, and social factors: determining what research is funded, where, and what research is most effective
How will AI be used to produce goods and services? **How will the use of AI change the productivity of workers and capital**?	What tasks will AI perform? Which of these will be tasks currently performed by workers? Which worker tasks or uses of capital will be complemented by AI?	What organisations have the necessary 'absorptive capacity' to adopt new technology for commercial ends?[37]
		What ethical considerations may inform the application of AI?
	What are the effects of AI adoption in specific occupations or industries on the rest of the economy?	How would increasing interaction with machines affect human decision-making (particularly by business leaders)?
	Will AI adoption drive major social changes?	What social changes should be expected, given the history of previous periods of technological change?
	Will the use of AI have any effects on research and innovation?	What innovations will generate new demand for workers, if any?
	How quickly will any changes happen?	How will the interaction between business leaders and workers shape the use of AI?
How will any additional income generated by AI be spent?	What additional goods and services will be demanded by those experiencing greater earnings due to AI?	Economic, cultural, social, psychological and other factors influencing consumer choices
Who will work, how many hours will they work for, under what conditions, and for what wage?	How fast will changes in employer requirements be compared to changes in the supply of labour? How will labour demand and supply interact?	Will employers have greater market power? Will there be changes in collective bargaining?

Source: Frontier Economics analysis based on preliminary review of the literature.

[37] Defined as 'The ability of a firm to acquire relevant external knowledge, transform existing practices, and exploit these new capabilities for commercial ends' (Zahra, S.A., et al. (2002). Absorptive Capacity: A Review, Reconceptualization, and Extension. *Academy of Management Review* 27(2), 185–203).