The Potential Impact of Artificial Intelligence on UK Employment and the Demand for Skills

A report by PwC for the Department for Business, Energy and Industrial Strategy

BEIS Research Report Number: 2021/042

Final report

August 2021

Acknowledgements

This report has been prepared by PwC for the Department of Business, Energy & Industrial Strategy in line with the tender process of January 2019.

The Department would like to thank the project team within PwC, which included John Hawksworth, Yuval Fertig, Eduardo Rodriguez Montemayor, Jonathan Gillham, Hoa Duong, Patricija Petkeviciute, Joe van Rietschoten and Richard Berriman.

PwC was advised by Professor Haile Getinet of Nottingham University Business School. PwC would also like to thank participants in the expert workshop held at their offices in July 2019, as listed in Annex 2 of this report.

Any enquiries regarding this publication should be sent to us at:
enquiries@beis.gov.uk

Contents

Executive Summary

Aim of this study

This research provides estimates of the potential impact on UK employment and skills demand of artificial intelligence (AI) and related technologies such as robotics, drones and autonomous vehicles.

The research began with an expert workshop in July 2019 that assessed the potential automatability of a broad range of occupations over 5, 10 and 20 year time horizons. The results of this assessment were used as training data to estimate an explanatory model for all jobs within the OECD's PIAAC[1] survey database taking account of the tasks involved in each job and other relevant characteristics. The estimated automation probabilities for the UK were then disaggregated further using ONS data and a separate PwC modelling exercise was conducted to estimate potential job creation and impacts on skills demand, using both ONS data and online job advert data from Burning Glass Technologies (BGT).

Compared to previous UK studies on this topic, the research is more comprehensive in covering both potential job displacement and job creation from AI and related technologies, as well as the consequent implications for skills demand.

The research focuses on the potential distributional effects of AI on UK employment. Our net employment impact estimates (for 5, 10 and 20 year time horizons) are broken down, in as much detail as the data allow, by occupation, industry sector, region, education level, age group and gender. We also consider how these estimated employment impacts might affect the future demand for skills not just in the AI area but more widely across the economy.

Given the many uncertainties involved in any such estimates of the future development and economic impact of technological change, the results of this study should be interpreted with appropriate caution. The focus should be on the broad directional trends identified, rather than specific numerical estimates.

Overall economic and employment impact of AI

Past research suggests that AI and related technologies should be significantly positive for productivity and real income levels, boosting UK GDP by up to 10% by 2030.[2] However, other studies[3] have pointed, to varying degrees, to potential displacement of human workers as these technologies are rolled out across the economy over the coming decades.

[1] Programme for the International Assessment of Adult Competencies. Further details available online at: https://www.oecd.org/skills/piaac/

[2] PwC (2017b)

[3] In particular, Frey and Osborne (2013) for the US, Arntz, Gregory and Zieharn (2016) and Nedelkoska and Quintini (2018) for a range of OECD countries, and ONS (2019) for England. All of these studies adopted a broadly similar approach to the present study in relation to estimated job displacement due to automation, but none of them covered job creation. McKinsey Global Institute (2017) provided estimates for both job creation and displacement, but their methodology was not readily comparable to the other studies cited and their estimates did not cover the UK specifically. The same is true of WEF (2018), which projected short term net employment trends from 2018 to 2022.

Our base case estimate is that around 7% of existing UK jobs could face a high (over 70%) probability of automation over the next 5 years[4], rising to around 18% after 10 years and just under 30% after 20 years. This is within the range of estimates from previous studies and draws on views from an expert workshop on the automatability of occupations and detailed analysis of OECD and ONS data on how this is related to the task composition and skills required for different occupations.

But AI will also create many jobs through the boost it gives to productivity and economic growth[5]. Whilst some of these extra jobs will be in areas linked directly to AI and related technologies (e.g. data scientists, robotic engineers or people involved in the design and manufacture of sensors for driverless vehicles and drones), most of the additional employment will not be in high tech areas. Instead, these additional jobs created will mostly be in providing relatively hard-to-automate services (e.g. health and personal care) that are in greater demand due to the additional real incomes and spending arising from higher productivity generated by AI.

The overall net effect on employment is unclear, with the most plausible assumption based on historical trends and past macroeconomic research for the UK[6] being for a broadly neutral long-term effect. We adopt this as our base case assumption, but also test the robustness of our general conclusions on the potential distributional effects of AI on UK employment patterns to alternative scenarios on the aggregate net employment effect.

The uncertainties involved in projecting the economic impact of any major technological change mean the focus should be on the broad patterns of effects found, not the specific quantitative estimates. Estimates are also generally more uncertain over longer term horizons and at greater levels of disaggregation by occupation, industry sector or geography.

Distributional employment impacts of AI

Analysis by industry sector

The health and social care sector had the largest estimated net employment increases from AI over the next 20 years (see Figure 1). This is projected to be an area of high demand due to an ageing population and rising income levels, and also one where AI and robotics are likely to be complementary to human labour in most cases rather than substitutes.

[4] This 5 year estimate of around 7% is similar to the estimate of the probability of automation for existing UK jobs in ONS (2019), although the latter study did not specify a time horizon for its estimates.
[5] This is sometimes referred to as the 'income effect' on jobs from technological innovation, as in Oxford Economics and Cisco (2017) and PwC (2018c).
[6] PwC (2018c). McKinsey Global Institute (2017) suggested a small net positive employment effect for the US and a small negative net impact for Germany and Japan, but did not include detailed estimates for the UK. WEF (2018) suggested a small net positive global jobs impact between 2018 and 2022, but again no specific UK estimates.

Figure 1: Estimated net employment effects of AI on selected industries over 20 years (SIC 1)

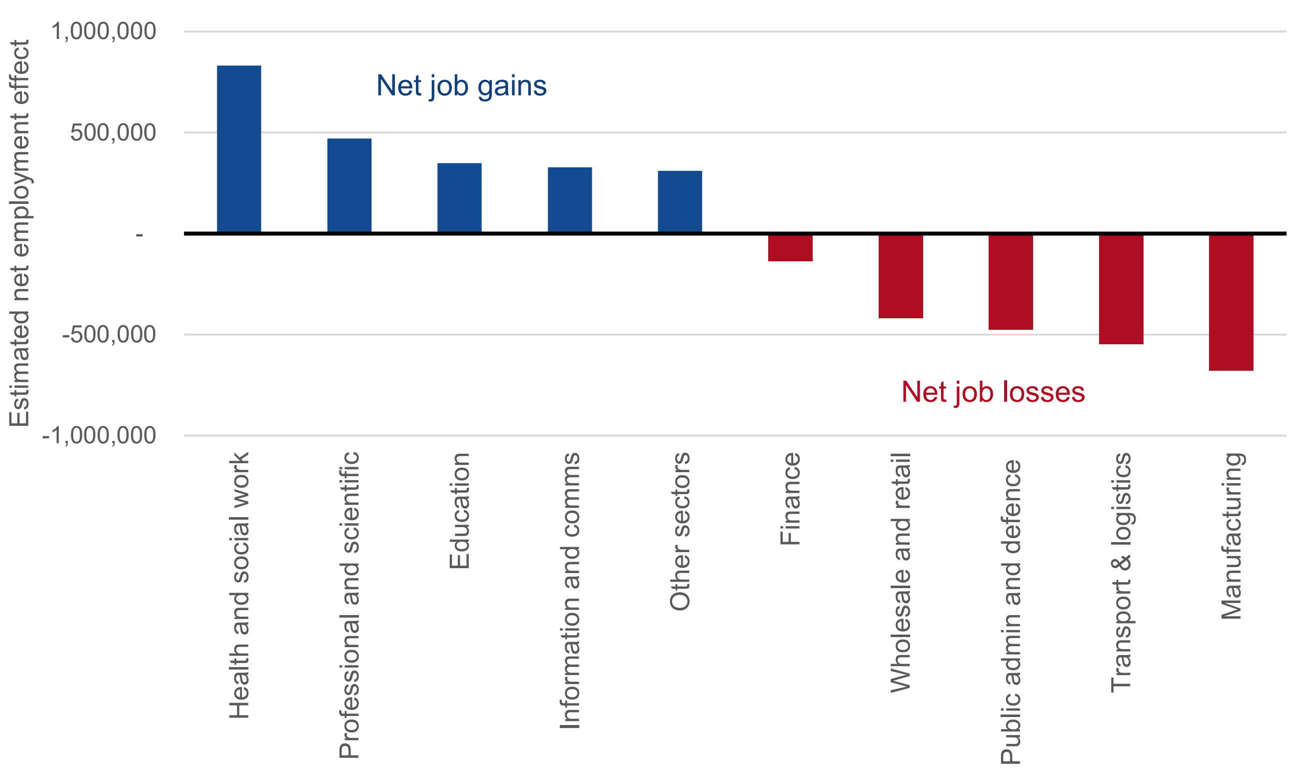

Source: PwC analysis of OECD PIAAC and ONS APS data

Information and communications and other professional, scientific and technical services are also projected to see significant net gains - which will include many highly skilled jobs linked closely to AI and other emerging technologies. By contrast, **significant net employment reductions are projected in wholesale and retail, finance and public administration areas in the short to medium term, all of which are relatively automatable sectors, and in transport in the longer term.**

Analysis by occupation and earnings levels

These sectoral findings reflect the occupational patterns evident from the analysis. Managerial and professional occupations with higher median earnings levels tend to see significantly positive estimated net employment effects (Figure 2).

Figure 2: Estimated net employment effect vs median earnings of occupations (SOC 4) over 20 years

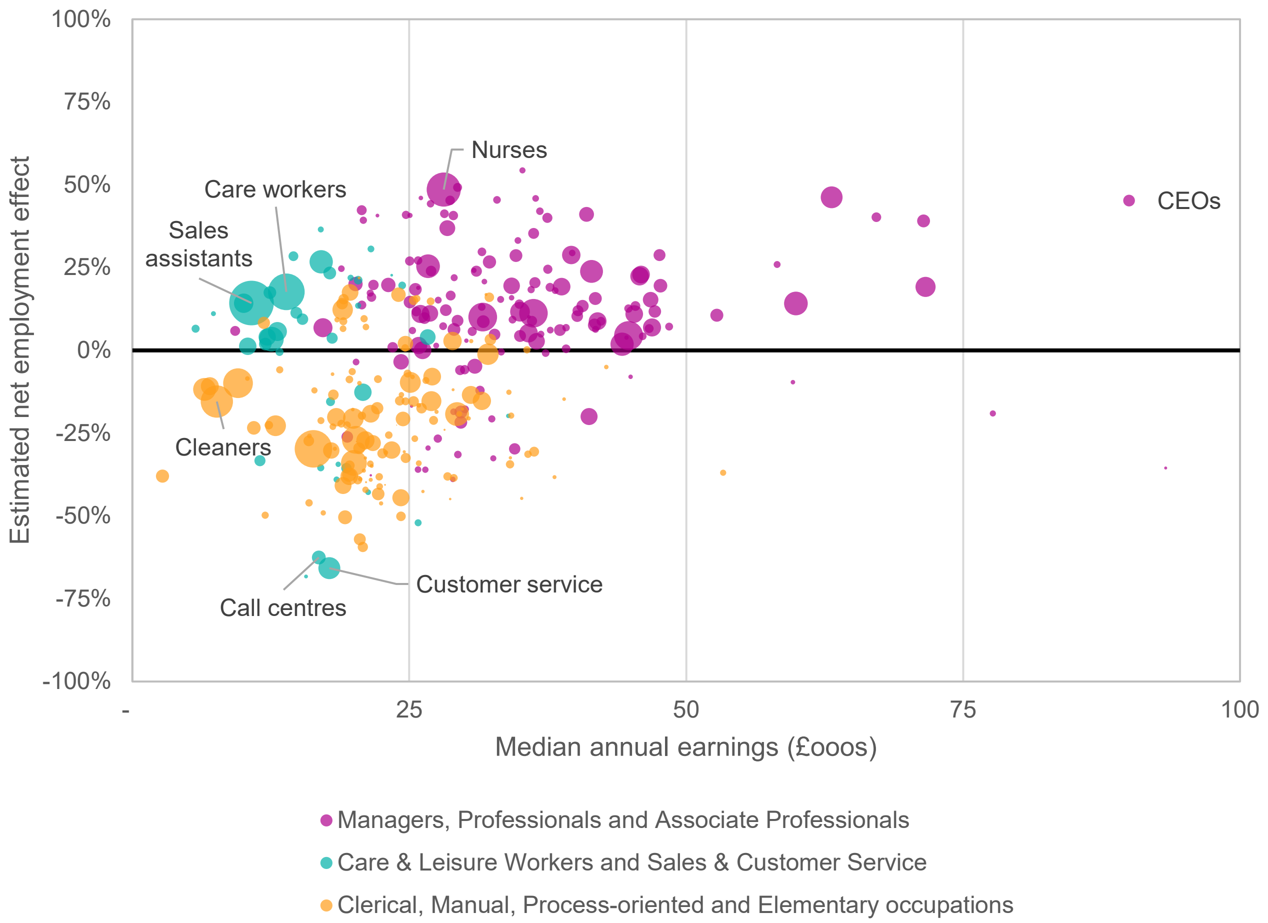

Source: PwC analysis of OECD PIAAC and ONS APS data

Less well-paid clerical and process-oriented roles will experience negative estimated employment effects, although only in the longer term for manual workers such as truck or taxi drivers. This suggests that AI will drive a continuation of skill-biased technological change with the potential to widen existing earnings differentials. We found that the general pattern of our estimates by occupational group was robust to plausible alternative scenarios on the overall balance of net job creation and displacement.

Analysis by education, age and gender

Our conclusion that AI may see a continuation of skill-biased technological change is also supported by the **positive correlation we find between estimated net employment effects of AI and education levels.** This conclusion is also robust to plausible alternative scenarios on the overall scale of job displacement and creation.

Figure 3: Estimated net employment effect of AI by education level over 20 years

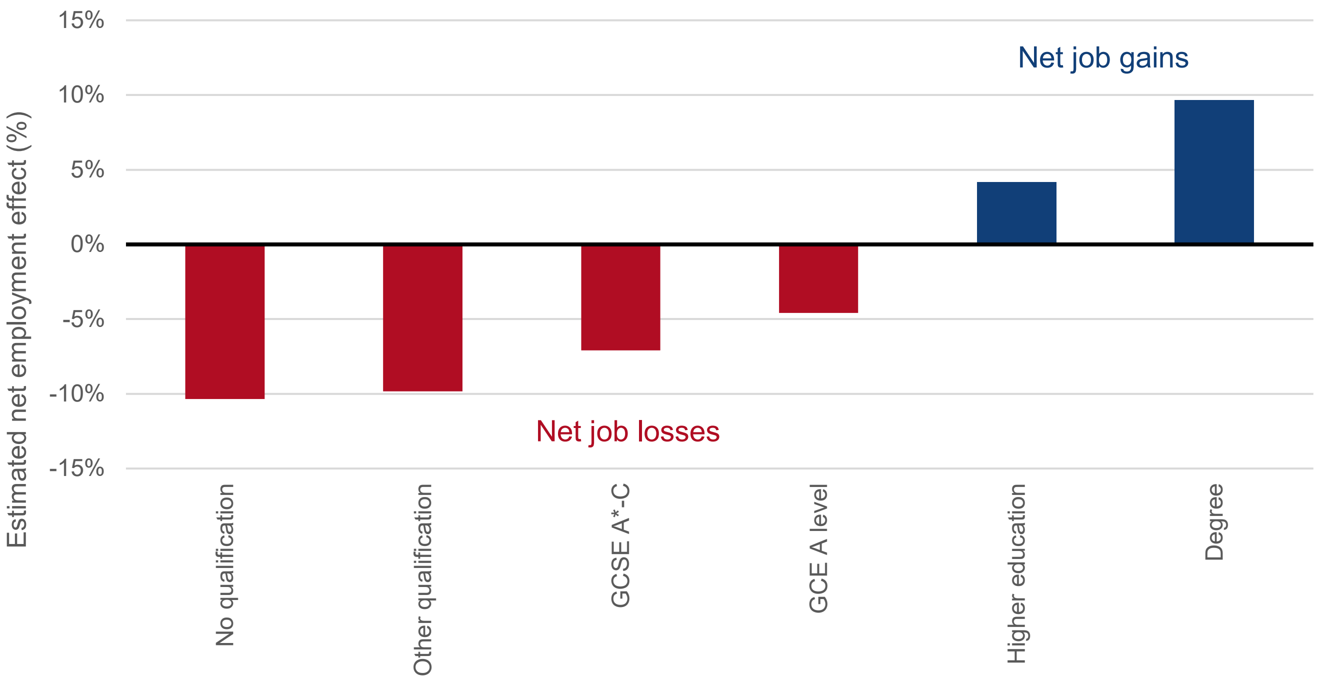

Source: PwC analysis of OECD PIAAC and ONS APS data

By contrast, we find less evidence of material differences in net employment impacts by gender or age group, although there is **some indication that entry-level jobs for younger workers may be more likely to be automated**. However, young workers may also be more adaptable in adjusting to new technologies and digital 'upskilling' will be important for all demographic groups.

Analysis by region

When we break the results down by UK region and sub-region, our base case estimates, as well as plausible alternative scenarios, suggest **somewhat more positive net effects in higher income areas such as London and the South East and more negative net effects in some cities in Northern England and the Midlands** (see Figure 4). But there are also considerable variations within regions, reflecting different occupational mixes across towns and cities.

Figure 4: Estimated net effect of AI by region (NUTS 3) over 20 years (in terms of no. jobs)

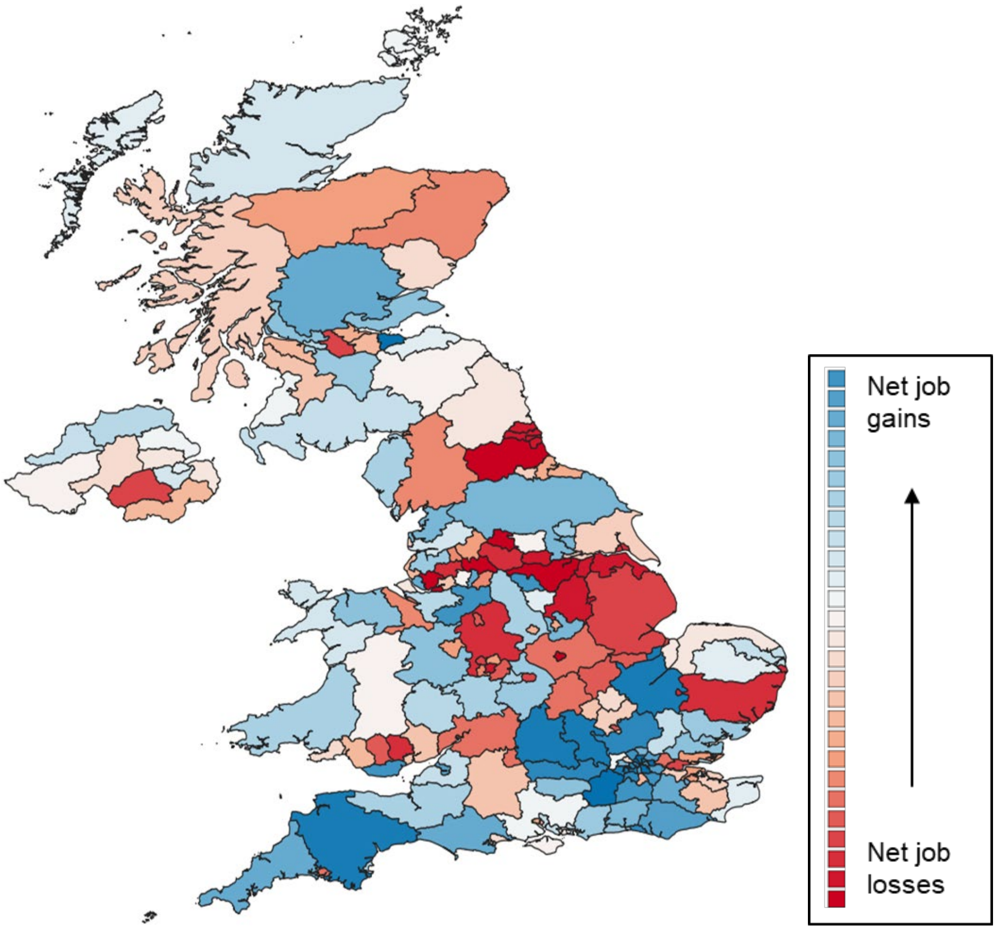

Source: PwC analysis of OECD PIAAC and ONS APS data

Estimated impact on future skills demand and shortages

We combined the employment estimates described above combined with data on the skills required in online job adverts from Burning Glass Technologies (BGT[7]) to projected forward how demand for different types of skills might evolve over the next 20 years. This analysis suggested **particularly strong increases in demand for skills related to health care. We also find smaller but still significant increases in demand for science and research and IT skills**, although the nature of the skills in this latter group will likely change significantly as AI and other digital technologies evolve. Demand for skills relating to administration, finance and customer support may see net decreases over time as the tasks involved are increasingly automated.

Our analysis also shows that **almost all occupations on the current government shortages list[8] are projected to show net employment increases due to AI.**

[7] BGT is a company that collects and analyses large volumes of data from online job adverts in the UK and other countries. PwC purchased access to BGT's UK data from 2012 to mid-2019 for use in this study.
[8] UK's Shortage Occupations List at the time the analysis was completed in October 2019, found here: https://www.gov.uk/guidance/immigration-rules/immigration-rules-appendix-k-shortage-occupation-list

Conclusion

AI and related technologies should not cause mass technological unemployment, but our analysis suggests that they may well lead to significant changes in the structure of employment across occupations, sectors and regions of the UK. The effects may be relatively small over the next five years, but could become more material over the next 10-20 years. These technological changes may also add to income inequalities to the extent that our analysis suggests that they may tend to favour people with higher education and skills levels, who also tend to have higher earnings levels.

1. Introduction

Background to this study

The technical capabilities of artificial intelligence (AI) and related technologies such as robots, drones and autonomous vehicles have progressed substantially in the last decade. New applications of AI are transforming whole sectors of the economy via increased productivity and innovation (e.g. NBER, 2017). As PwC argued in previous research[9], these technologies have the potential to boost the economy significantly, perhaps by as much as 10% of GDP in 2030 in the UK.

At the same time, there are also concerns that these technologies could displace large numbers of human workers from their jobs over the coming decades, leading to mass technological unemployment. Economic research papers such as Frey and Osborne (2013) have added to these concerns, even though the message of these authors was much more nuanced than media headlines about nearly half of US jobs potentially being automated might suggest.

Such concerns have been seen several times before in relation to previous major new technologies from mechanical weaving machines and tractors to digital computers, but they have proved unfounded in the long run as indicated, for example, by the fact that UK employment rates are currently at record highs since the 1860s according to ONS and Bank of England estimates.[10]

However, these past waves of technological progress have been associated with significant disruption to labour markets, with employment shifting from agriculture to manufacturing and then to services. Some occupations have either vanished (e.g. elevator operators) or been radically altered to focus on tasks that human workers can do better than machines (e.g. bank tellers), while new occupations have been created (e.g. computer programmers) and others have greatly expanded in number (e.g. doctors and nurses) as richer economies have increased demand for existing services (e.g. healthcare).

Two key research questions

One key research question is whether AI and related technologies (as defined in Box 1-1 below) will follow this historical pattern of triggering significant structural labour market change without causing mass technological unemployment. While there have been dissenting voices that argue that this time may be different, leading to widespread unemployment[11], the consensus of most recent economic studies[12] of which we are aware has been that history will continue to provide a reasonable guide to the past in this respect. We will get both job displacement and job creation due to AI, but these will broadly balance in the long run once economies have had time to adjust.

[9] PwC (2017b, 2017c)

[10] Chiripanhura and Wolf (2019) gives an analysis from 1861 to 2018. The latest ONS data for Q4 2019 show a higher UK employment rate for 16-64 year olds (76.5%) than at any point in that historical period.

[11] For example, Ford (2015) and Harari (2016).

[12] For example, Acemoglu and Restrepo (2018, 2019), McKinsey Global Institute (2017), World Economic Forum (2018), PwC (2018c).

This is the basis for the working assumption in this study that job creation matches job displacement, though there are clearly many uncertainties around this and we do therefore test the robustness of our key conclusions to alternative scenarios on the net employment impact of AI.

Box 1-1: Definition of AI and related automating technologies used in this study

There is no single definition of AI. Sometimes it is defined in a narrow sense to refer to machine learning. However, in the broader definition used for this and previous PwC studies in this area, 'AI' is used as a collective term for digital systems and machines that can, in at least some ways, sense their environment, and can think, learn and take action in response to what they are sensing and their objectives.[13] Throughout this report we refer for brevity to 'AI', but this should be interpreted in this broader sense as including the related automating technologies shown in the figure below.

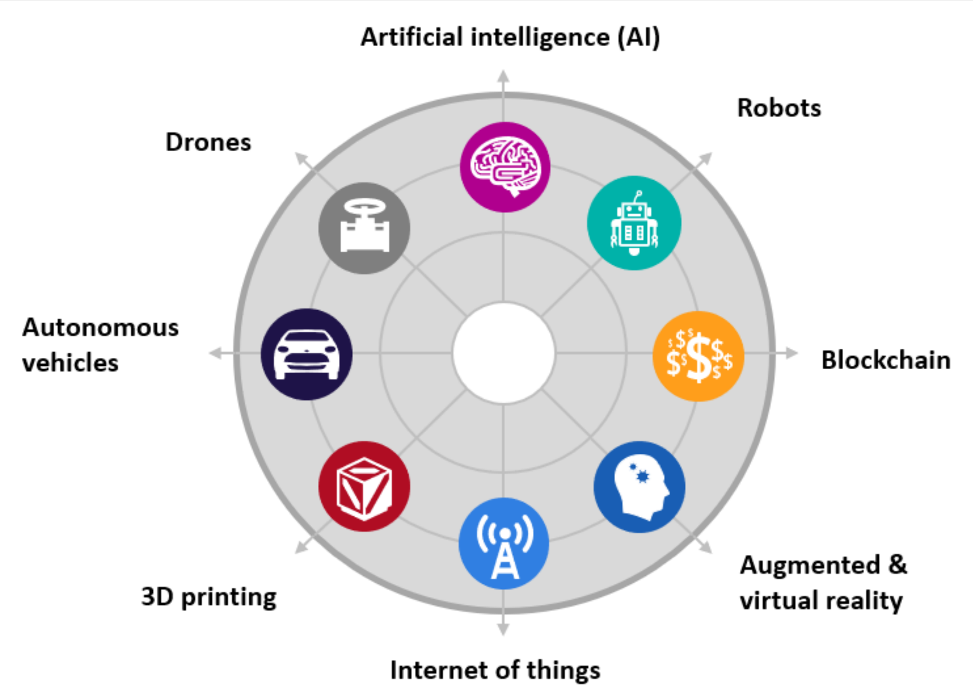

In this broader sense, AI enables computers, robots and other machines (e.g. drones, driverless vehicles) to do things that only people (or teams of people) use to be able to do in the past - via techniques such as machine learning (including deep learning), natural language processing, computer vision and robotic process automation. This is also linked to a range of supporting technologies in areas like augmented and virtual reality (AR/VR), satellite communications, 5G networks and digital sensors used to create an 'internet of things', 3D printing and blockchain.

In another sense of the term, however, we are still focusing on 'narrow AI' rather than 'general purpose AI'. This is because, over the 20 year time horizon adopted for this study, we assume that AI and related technologies will be focused on automating specific tasks or clusters of tasks, and not on replicating all or most of the functions of human beings in the way that robots or androids are often depicted in science fiction books, TV

[13] The progress of AI has gone through three major phases: a shift toward machine learning during the late 1990s and early 2000s; a rise in the popularity of neural networks beginning in the early 2010s (with the progress of image and speech recognition by automatically processing unlabelled data); and the recent growth in reinforcement learning in the past decade (Hao, 2019).

series or films.[14] The latter seems very unlikely to be achievable over the next 20 years - or potentially for some considerable time after that - based on the current state of AI and robotics, so we treat this possibility as falling outside the scope of this study.

Glossary of key technological terms

Artificial intelligence (AI): field of computer science involved with developing systems with ‘intelligence’.

Machine learning (ML): sub-field of AI involving development of systems that use statistical algorithms to ‘learn’ from training data, rather than being pre-programmed as with traditional digital computers.

Robot: a machine capable of carrying out a series of actions automatically. This may be directed by a computer programme based on pre-set instructions as with most industrial robots, or it could learn from its environment using sensors and some form of AI/ML.

Autonomous vehicles: cars, vans, trucks and other motor vehicles that can drive themselves using sensors and AI systems without any input from a human driver. Autonomous drones could be seen as an aerial variant of this, and both could be seen as types of AI-powered robots.

Robotic process automation (RPA): class of software ‘robots’ that replicate the actions of a human being interacting with the user interfaces of other software systems. Enables the automation of many ‘back office’ functions (e.g. finance, HR) and could be seen as a stepping stone to a more fully developed form of AI.

The second key research question is how large the disruption to labour markets will be from AI and related technologies and what form will this take? This is the main focus of the present study, as applied to the UK over three broad timescales: short term (c.5 years), medium term (c.10 years) and long term (c.20 years).

Past studies

In terms of quantifying the labour market effects of AI and related technologies, our study follows in the research tradition begun in 2013 by Frey and Osborne (‘F&O’ hereafter) from the Oxford Martin School, who made global headlines by estimating that nearly half of jobs in the US could potentially be susceptible to automation by computer-controlled machines using machine learning (ML) and related technologies.[15]

The F&O study attracted the attention of academics, governments and policy think tanks from around the world, who have since devoted substantial effort to studying the potential future pace of automation related to AI/ML and how this may affect employment and other economic variables of interest. Technical Annex 1 provides more details on past research on the

[14] This includes the deeper question as to whether digital systems could ever achieve ‘consciousness’, leading to what some writers (e.g. Vernor Vinge, Ray Kurzweil) have speculated might be a ‘technological singularity’ with potential existential implications for the future of humanity and its relationship with digital systems. We assume that, if this is ever possible, it would only occur well beyond the 20 year time horizon of this study.

[15] In their original paper, Frey and Osborne emphasised that this was a potential upper bound estimate based on technological feasibility alone, rather than a prediction of what they actually expected to happen given the many possible economic, legal and regulatory, political and social barriers to such automation. However, these important caveats were lost in much of the media coverage of their analysis, as discussed further in Frey (2019).

potential economic and employment impacts of AI, but Box 1-2 highlights some key studies that followed in the tradition of F&O.

Box 1-2: Previous estimates of potential job displacement from AI/automation

Estimates of job displacement relating to AI have been diverse. In the seminal study in this field, Frey and Osborne (2013) estimated that up to 47% of existing US jobs could be at high probability of automation (defined as over 70% for this and other studies reviewed in this box). F&O began with an expert workshop of machine learning (ML) specialists assessing whether 70 selected occupations were or were not technologically susceptible to automation through AI/ML. These assessments were used as training data for an ML algorithm to estimate the probability of automation across over 700 occupations included in the US O*NET database, based in particular on the tasks involved in each occupation (on which O*NET included data).

However, two later studies published by the OECD estimated that only around 9-14% of jobs across a range of OECD member countries face a high probability of automation (although many more face moderate risks of automation).[16] The key difference from the F&O study, despite using their results as a starting point, was that these later studies used the OECD's PIAAC database, which allowed analysis of the characteristics of individual jobs, rather than whole occupations as in the O*NET data used by F&O. The authors of these two studies argued that this allowed them to produce more realistic estimates of how many jobs might actually be displaced by automation, as opposed to modified to allow humans to focus on tasks where they have a comparative advantage over digital systems, and vice versa.

Based on the 2016 OECD research by Arntz et al, the ONS (2019) has estimated that the proportion of existing UK jobs at high probability of automation could be as low as 7%. The ONS study is notable in that it used raw data from the UK Annual Population Survey (APS) to produce a much more granular breakdown of UK estimates by occupation, industry sector, geographical area, socio-economic and demographic characteristics than the estimates based on the UK PIAAC survey, which has a much smaller sample size than the APS.

Previous PwC research (2017a, 2018b) has found, however, that while there is value added in using the PIAAC data on individual jobs, this cannot in itself explain much of the difference in the estimates of the proportion of existing jobs at high probability of automation between the F&O and OECD studies. Instead, after replicating both sets of results, we found that most of the difference was due to technical methodological issues. In particular, we found this was mostly due to the lower predictive power of the models used in the OECD studies. This had the effect in the OECD studies of causing the majority of jobs to be classified as 'medium risk' (30-70% probability) rather than 'high risk' (>70% probability) or 'low risk' (<30% probability). This problem also applies to the ONS estimates since they were derived initially from the Arntz et al results, albeit with considerable further analysis to make them more granular.

This previous PwC research used a hybrid version of these past approaches in that we estimated our model using PIAAC data for individual jobs, but were able to come up with a model with significantly higher predictive power. The outcome was that our estimate of the proportion of existing jobs at potential high risk of automation moved back in the

[16] See Arntz, Gregory and Zieharn (2016) and Nedelkoska and Quintini (2018). These studies use the OECD PIAAC database, a detailed survey of adult skills that can be found here: https://www.oecd.org/skills/piaac/

direction of the original F&O study, which makes sense because all these studies took the F&O results as a starting point. For the UK, we estimated that up to around 30%[17] of existing UK jobs could face a high probability of automation for the UK in the next 20 years. However, these studies also emphasise that not all of these jobs may actually be displaced given practical barriers to rolling out AI and related technologies across the economy. They also emphasised that many additional jobs would be created due to the macroeconomic benefits of AI.[18]

One interpretation of the difference between our estimates and those of the OECD/ONS studies is that the latter apply to shorter time horizons, since their methodology only picks up those jobs that F&O estimates to have very high probabilities of automation. These may be low hanging fruit, such as clerical jobs susceptible to robotic process automation, rather than jobs like drivers that will only be automated at scale over much longer timeframes according to our analysis in the present study. However, this is only one possible interpretation, since the OECD and ONS authors do not attempt to put a timescale on their estimates.

While headline estimates vary widely, however, there is more consensus across these studies on the fact that jobs involving higher levels of skills, and typically performed by more educated workers, will face lower potential probabilities of automation in the medium to long term than those requiring lower skill levels. This is a topic that we also consider further in the present study, with broadly consistent general conclusions.

Valued added elements in this study

Despite the extensive past research summarised in Box 1-2 (and other studies mentioned in Technical Annex 1), there are still some important gaps in current research on AI and automation, which we aim to address in this study for the UK.

Firstly, the question that F&O attempted to answer is far from settled, as estimates regarding how many jobs may be displaced by AI have been very wide-ranging (see Box 1-2 above). **This study aims to provide as up-to-date and comprehensive as possible a set of job displacement estimates for the UK, while also recognising the uncertainties surrounding any such estimates. This has included conducting a new expert workshop in July 2019 to update the judgements on the automatability of different occupations from a 2013 workshop that underpinned F&O's study (and, directly or indirectly, all the other studies that followed in that tradition as described in Box 1-2).**

Secondly, by focusing on potential job displacement, many previous studies only give a partial view of the economic impacts of AI and other technologies. As well as displacing jobs, major new technologies also create jobs either directly through enabling new tasks to be undertaken by human workers, or indirectly via the boost they give to productivity and so to real income and spending levels across the economy. Some occupations will wholly or largely disappear over time, but many others will be re-engineered so that some tasks are automated where it is more cost-effective to do so, while human workers focus on other tasks where they retain a comparative advantage over machines, including some brand new tasks enabled by these

[17] This compares to a Bank of England estimate of around 35% in Haldane (2015), which translated Frey and Osborne's US estimates to the UK.
[18] PwC (2018c) provided high level estimates of how many UK jobs might actually be displaced and created by AI, using a simplified version of the methodology in the present study but at a much lower level of granularity.

technologies.[19] **This study focuses on the net employment impact of AI, paying as much attention to job creation as job displacement.**

Thirdly, behind the headline estimates of the total number (or proportion) of jobs created or displaced by automating technologies, there are important distributional impacts. Future jobs will not necessarily be in the same place, in the same occupations and industries, or offering the same pay and conditions as today's jobs. The demand for skills will also be different and workers, supported by government, business and education and training providers, will need to upgrade their skills accordingly to remain employable. **This study therefore focuses primarily on the distributional impact of AI on employment and the implications for skills demand.**

In addition to its holistic and more up-to-date approach to estimating job displacement and job creation, this study also brings some more technical methodological improvements with respect to previous studies. In particular, we use a random forest classification model to estimate the level and distribution of job displacement.[20] This improves on the logistic regression model used by Arntz et al (OECD, 2016), which led in our view to underestimates of the proportion of existing jobs with a potential high probability of automation (as discussed also in Box 2 above).[21]

We have also developed in much more detail the model used to estimate potential job creation from AI in PwC (2018c), as described in depth in Technical Annex 3. This has allowed us to produce estimates of net employment effects on a more consistent basis and at a much greater level of granularity in terms of how these effects vary across sectors, occupations, regions, demographic and socio-economic groups within the UK.

We have also extended previous PwC analysis to look not just at employment but also at the implications for future skills demand. This analysis makes use of Burning Glass Technologies (BGT) data on online job adverts together with our estimates of the net employment effects of AI.

Structure of the report

Section 2 presents an overview of our analytical approach to estimating the potential impact of AI on UK employment (the technical Annexes provide more details).

Section 3 presents the estimated potential net impacts of AI on employment across the UK economy. These results are disaggregated by occupational group, industry, region and socio-economic and demographic group.

---

[19] See, for example, the arguments presented by Erik Brynjolfsson of MIT about the reinvention and re-engineering of jobs: https://sloanreview.mit.edu/article/four-ways-jobs-will-respond-to-automation/. According to the OECD (Marcolin et al. 2016), across its 32 member countries at least 54% of people are employed in routine occupations, and up to one in two jobs are likely to be affected by automation to a certain extent - even if those jobs do not disappear.

[20] This model of job displacement was essentially the same as used in PwC (2018a), which in turn updated an earlier version in PwC (2017a). The main difference was applying the model over different time horizons for this study and then going on to produce more disaggregated estimates for the UK.

[21] This potential underestimation also fed through to the absolute estimates of automation probabilities for the UK by the ONS (2019), given that the latter took Arntz et al's estimates as a starting point. In other respects, however, the ONS study had clear added value in providing more granular estimates for the UK of the relative probabilities of automation across occupations, industries, regions and demographic groups.

Section 4 discussed the implications for skills. It first analyses how occupations in shortage in the UK will evolve in the context of automation before presenting estimates of the effects of AI on the potential future demand for skills.

Section 5 concludes and highlights areas for possible future research.

Additional supporting material in technical annexes including: a review of relevant past studies; discussion of the methodology and results of the expert workshop held in July; further details of the methodologies behind our job displacement and job creation estimates; and the results of the sensitivity and scenario analysis carried out on key assumptions.

2. Analytical Approach

Overview of approach

This research is based on a holistic approach to estimating the potential net employment effect of AI and related technologies for the UK. The approach is built upon estimating two countervailing effects:

1. **Job displacement effect:** we estimate the potential magnitude and distribution of job displacement by occupation and industry.
2. **Job creation effect:** we assume, for reasons discussed further below, that total job creation matches job displacement and then estimate how these additional jobs are allocated across occupations and industries.

The key outputs of this analysis are up-to-date estimates of both job displacement and job creation in terms of the impact on occupations and industries. We then derive estimates of the **net effect** of AI on employment by occupation and industry, which is the difference between job creation and job displacement. We then use the occupational breakdown to infer the distributional impacts by region, socioeconomic and demographic groups, as well as to infer implications for the future demand for skills in the UK.

Our approach is summarised in Figure 2-1 and explained in more detail in Technical Annexes 2, 3 and 4, which also describe the sources of data used for each step. We would also emphasise the uncertainties inherent in any such analysis and the associated caveats to bear in mind when interpreting the results of this study (as discussed further in Box 2-1 below).

Figure 2-1: The Approach

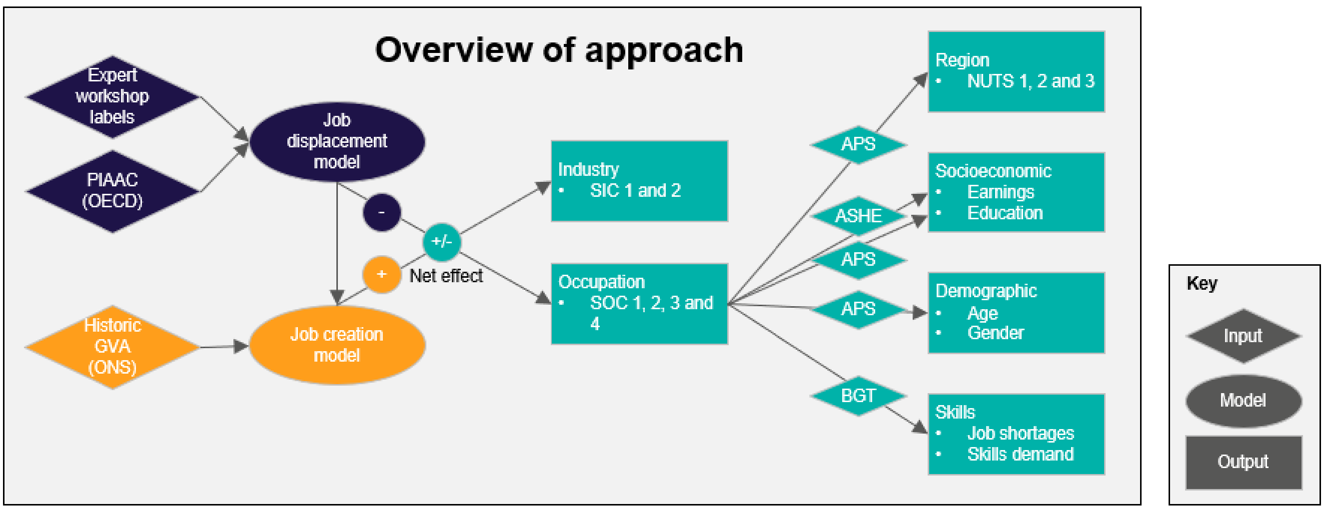

Source: PwC

The first step of the methodology was to update estimates of the automatability of 70 selected occupations that had earlier been labelled in Frey and Osborne (2013) - and subsequently used directly or indirectly as a starting point for many later studies (including by the OECD, the Bank of England and the ONS). For those purposes, an expert workshop was held in July 2019 (as described in detail in Technical Annex 2). Experts from academia, industry and government were asked to consider the task composition of occupations in order to evaluate the 'automatability' of the different occupations. Figure 2-2 shows that the overall proportion of occupations labelled by the majority of experts as automatable over 20 years was broadly consistent with Frey & Osborne's workshop result in 2013 (although they didn't specify a time horizon in their study). However, this disguises the fact that expert opinions on the specific occupations that were deemed automatable has changed. The likelihood of complete automatability of jobs in our base case estimates is based on a simple majority rule of workshop participants - a super-majority voting rule (at least ⅔) is used for sensitivity analysis as presented in Technical Annex 5.

Figure 2-2: Overview of results of the workshop on occupational automatability: Proportion of occupations labelled as likely to be automated at scale by workshop participants

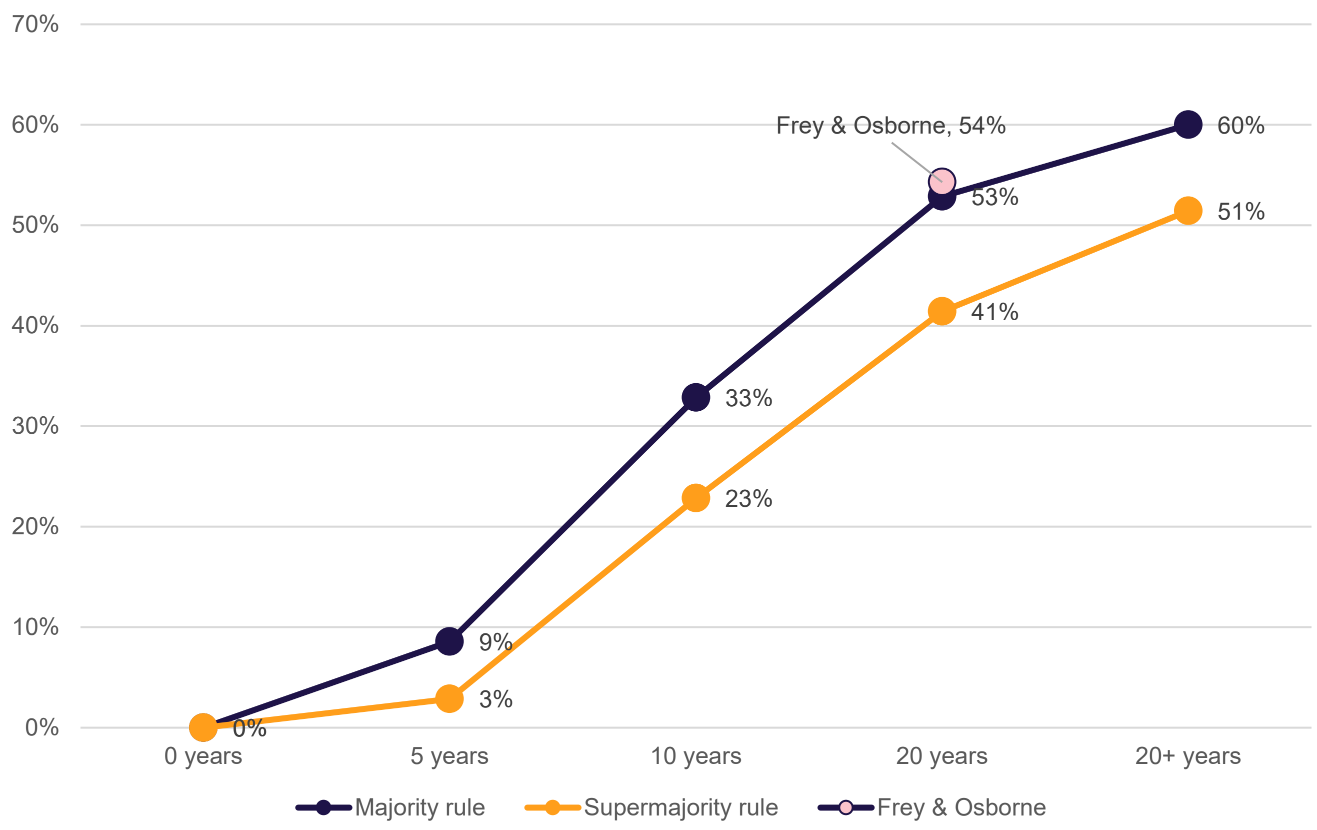

Source: PwC analysis of data from expert workshop, Frey and Osborne (2013)

Further details of the workshop results, as well as details of all of the steps followed to estimate the job displacement and job creation models, are set out in the Technical Annexes 2 and 3. The key steps in the process are as follows:

- Job displacement effect:
    - we used the expert judgements from the workshop as initial training data to estimate, using machine learning techniques, a probabilistic model of automatability based on detailed data on the tasks involved and other characteristics of over 200,000 jobs from the OECD's PIAAC database[22];

---

[22] In using this PIAAC database, our approach follows that of Arntz et al (2016) and Nedelkoska and Quintini (2018), although the technical details of our modelling approach differ as discussed in Technical Annex 3.

  - this model was then used to produce estimates using the UK PIAAC data of the proportion of existing UK jobs that face a high probability of automation (defined here as over 70%, which is in line with the threshold used in past studies) and so may be displaced by AI; we refer to these as job displacement estimates;

- Job creation effect:
  - in our central scenario we assume that total job creation from AI equals total job displacement, which is consistent with past research, as discussed further below (we also consider alternative scenarios in the Annex);
  - we then model how this job creation would be allocated across industries in a way related to past and projected future sectoral growth and across occupations in a way inversely related to estimated automatability; this means that the extra jobs created by AI are mostly allocated to industries and occupations that are projected to see high growth in demand and relatively low automatability, which includes areas like health care and other personal services;

- Net employment effect:
  - we further break down these UK estimates of job displacement and job creation (and so net employment effects) into more disaggregated groups by occupation (to SOC 4 level), by industry sector (to SIC 2 level), by region (to NUTS 3 level) and by socio-economic characteristics (e.g. education, age, gender) using data from the ONS's Annual Population Survey (APS);
  - using data from the ONS's Annual Survey of Hours and Earnings (ASHE), we also estimate how net employment effects may vary with the median earnings levels in different occupations, industries and regions, which gives an indication of the distributional effects of AI;

- Impact on demand for skills
  - online job advert data from Burning Glass Technologies (BGT) is used to estimate how net employment effects by occupation may translate into future trends in the demand for different types of skills (based on skill requirements set out in online job adverts); we also used the BGT data to look at recent trends in the demand for skills linked specifically to AI.

Box 2-1: Some important caveats

We have built on previous research to provide estimates that are as up to date and reliable as possible. But we also recognise that any such projections will be subject to significant margins of uncertainty both about the pace of technological change and the speed with which new technologies are adopted across the economy. Given these uncertainties, all estimates should be interpreted as plausible indications of potential general trends and patterns of impacts - they are not intended to be precise forecasts or predictions.

The more disaggregated estimates are also subject to higher levels of uncertainty since they rely on smaller sample sizes in the data sets used to generate the estimates. Longer term estimates will also generally be more uncertain than shorter term estimates.

There are a range of uncertainties relating to the job displacement effect, starting with the way in which the assessments from the expert workshop are translated into training data. As discussed below, we have used a simple majority voting rule here but also considered the robustness of our qualitative conclusions to using a ‘super-majority’ (2/3rds) voting rule instead. This affects the pace with which some occupations are affected by automation as well as the extent of such automation.

There are also uncertainties relating to the particular type of algorithm used to analyse the OECD PIAAC data and the explanatory variables to include. To improve predictive power of our model, we have used a wider set of variables than studies such as Arntz et al (2016). We consider our approach to be an improvement, but this is a matter of judgement and should be recognised as an additional source of uncertainty around the results of this study.

Uncertainties are particularly large when estimating job creation. We have therefore looked at a range of alternative scenarios depending on whether job creation is significantly higher or lower, over different time horizons, than our base case assumption that job creation matches job displacement. This includes one scenario in which job creation matches job displacement over 20 years, but lags behind so there are net employment losses over 5-10 year time horizons as labour markets take time to adjust to the impact of the new technologies (e.g. as it takes time to reskill workers and for them to move to locations where additional jobs are created).

We also continue to use the UK’s standard occupational classification (SOC) and standard industrial classification (SIC) given this is how currently available data are classified. In practice, entirely new job categories may be created over the next 20 years, but we cannot anticipate this in advance.

Estimates of the overall job displacement effect

In terms of the aggregate effect on employment, we estimate that around 7% of existing UK jobs could be displaced over the next 5 years, rising to around 18% after 10 years and nearly 30% after 20 years. Based on total UK employment of around 32.4 million in 2018, according to the ONS Labour Force Survey, this could be equivalent to around 2.2 million jobs displaced over the next 5 years, just under 6 million jobs displaced over the next 10 years, and around 9 million jobs displaced over the next 20 years.

This is within the range of estimates from previous studies[23] and shows a plausible S-curve time profile that is often observed as new technologies take off slowly at first, gather momentum as they are adopted more widely and then see slower growth as they reach maturity.[24]

[23] Frey and Osborne (2013) estimated that 47% of US jobs might face a high probability of automation and Haldane (2015) estimated this could translate to around 35% of UK jobs using the same approach. But Arntz, Gregory and Zieharn (OECD, 2016) used the PIAAC data to estimate that only around 10% of UK jobs might face a high probability of automation and ONS translated this to a 7% estimate using more detailed APS data. Nedelkoska and Quintini (OECD, 2018) sought to improve on the 2016 estimates for the OECD, resulting in a slightly higher estimate of around 14% of existing jobs at high risk of automation. The reasons for the differences in these results are technical and complex, as discussed in detail in Annexes 1 and 3, but we consider our methodology to be an improvement on these past studies, while building on the many useful insights in all these research papers.

[24] Rogers (1962) Diffusion of innovations

Figure 2-3: Estimated % of existing UK jobs displaced due to AI

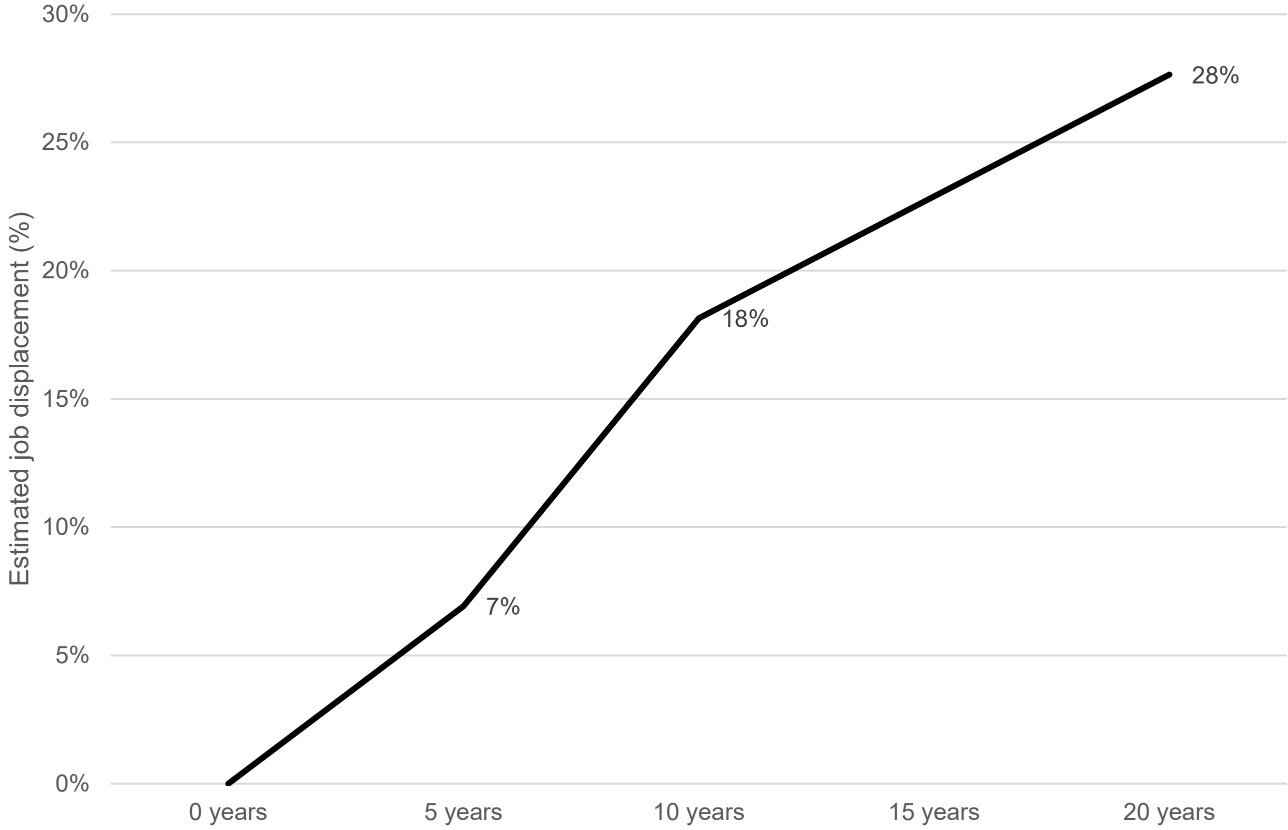

Source: PwC analysis of OECD PIAAC data

When we apply the same model to other OECD countries in the PIAAC dataset, we find that this estimate places the UK broadly in the middle of the range in terms of the estimated proportion of existing jobs that could be displaced by AI after 20 years, as seen in Figure 2-4. The UK is also broadly in the middle of the international range over 5 and 10 year time horizons.

Figure 2-4: International comparison of estimated job displacement (over 20 years)

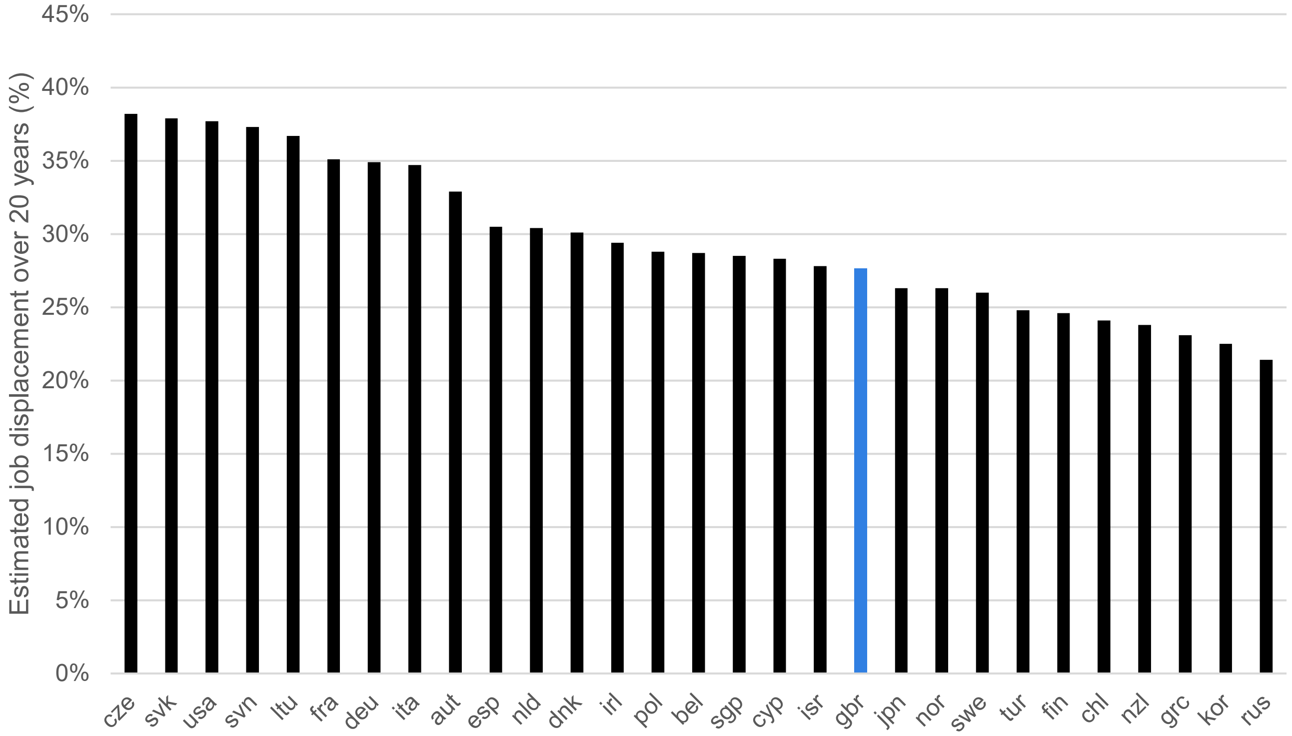

Source: PwC analysis of OECD PIAAC data

Basis for assuming total job creation matches total job displacement in our base case

Since the primary focus of this study is on the potential distributional impact of AI on employment patterns (as presented in Section 3), we do not attempt to independently quantify the total number of jobs that will be created by AI. Instead we assume in our base case scenario that the total number of jobs created by AI is equal to the total number of jobs displaced by AI. The aggregate net effect of AI on UK employment is therefore assumed to be zero. We consider this a plausible main scenario because:

- it is consistent with our previous attempt to quantify the net effect of AI on jobs in the UK in PwC (2018c), which suggested a broadly neutral impact over 20 years drawing also on the estimated macroeconomic impacts of AI using a Computable General Equilibrium (CGE) model of the UK economy in PwC (2017c);
- it places us broadly in the middle of the range of other estimates, some of which suggest a positive net employment effect[25] and others a negative net effect of AI (see, for example, the summary of studies by Winick, 2018, as well as historical estimates of the labour market impacts of industrial robots by Graetz and Michaels, 2016); and
- it is consistent with the fact that previous technological revolutions have not led to mass technological unemployment - indeed the UK employment rate is currently at a historical

[25] Including McKinsey Global Institute (2017) for the US, and World Economic Forum (2018) for short term global employment changes between 2018 and 2022.

high[26] despite rapid rates of digital innovation in recent decades, including personal computers, the internet and worldwide web, smartphones and previous advances in AI and robotics. It is not clear why future waves of digital technology involving further advances in AI should change this general pattern.

The evidence for these latter two points is described in more detail in Annexes 1 and 3. While it is not possible to prove definitively that the net employment effect of AI will be zero (or any other specific figure), it seems to be plausible working assumption for the purposes of this study.

In terms of economic theory, this assumption can be justified by the view that human workers will always retain a comparative (if not an absolute) advantage in some types of tasks, particularly in non-tradable services sectors where softer skills are required. It will therefore make sense for humans to specialise in those tasks, while AI systems and robots focus on tasks where they have a comparative advantage.

On this basis, employment should, perhaps after a period of adjustment due to labour market frictions, return to similar levels in the long term, but focusing on different industry sectors and types of occupations (just as workers have moved from agriculture to manufacturing and then to services over the past 250 years since the start of the Industrial Revolution). This assumes there are still costs to producing AI systems and robots, which need to be traded off against the cost of employing human workers to find the optimal balance across the economy.

Alternative scenarios

We accept, however, that there are many uncertainties as to how far this net zero employment assumption will hold over the whole or part of the 20 year time horizon for this study. To reflect this uncertainty, we also consider three alternative job creation scenarios as well as one alternative job displacement scenario. See Technical Annex 5 for the rationale behind alternative job creation and job displacement scenarios as well as estimates of the net effect of AI on occupations, industries regions and demographic groups under these alternative scenarios.

In general, we find that our qualitative conclusions regarding the redistributive impacts of AI in Section 3 below are robust to these alternative scenarios.

[26] According to the ONS, the latest UK employment rate for 16-64 year olds of 76.5% in Q4 2019 is the highest since comparable records began in 1971. The Bank of England has extended this analysis back to the 1860s, supporting the same conclusion, as discussed in Chiripanhura and Wolf (2019).

3. Net impact of AI on employment across the UK economy

This section summarises the results of our analysis of the potential distributional effects of AI and related technologies on net UK employment across occupations, industry sectors, regions and socio-economic groups. The commentary below also draws on qualitative insights from our expert workshop, as well as from previous research in this area.

The results should be interpreted with appropriate caution given the uncertainties inherent in any such analysis (as discussed in Box 2-1 above), but we highlight below the broad qualitative findings that appear robust to plausible alternative assumptions based on our scenario analysis (as described further in Annex 5).

Key findings:

- Managerial and professional occupations will tend to see positive net employment effects from AI that build up steadily over time.
- Clerical and routine manual roles may see negative net employment effects that start to come through over the next 5-10 years as technologies such as robot process automation roll out across the economy.
- By contrast, negative net employment effects for manual workers such as truck or taxi drivers may only emerge over a longer time horizon (i.e. 20 years) once autonomous vehicles have been introduced at scale across the economy.
- The largest net employment gains over the next 20 years may be seen in the health and social care sector, building up steadily over time. The information and communications sector is also projected to see net gains, but wholesale and retail, finance and public administration may see net losses over the next decade and transport could see net employment declines in the longer term.
- Net employment effects will tend to be more positive for higher paid occupations on average, and also more positive for graduates than for lower educated groups.
- Variations in estimated net employment effects by age and gender are smaller than by education, although younger workers in entry level jobs involving routine tasks may be more vulnerable to automation.
- There will be jobs gains and losses from AI in all regions, but net employment effects may be more positive in London and the South East on average, and less positive in more industrial areas in the Midlands and the North of England due to their different occupational mixes. But there will also be considerable variations within regions, reflecting different occupational mixes across towns and cities.

Net employment effects by occupation

Our model consistently estimates that the **highest net job gains over time will be in professional occupations** (see Figure 3-1), with nearly half of the net increase in jobs for health professionals. The other half of the gains are for scientists, researchers, engineers, technologists, educators, businesspeople, media professionals and civil servants. Most of these occupations are estimated to have around twice as many jobs created as displaced by AI. AI adoption in professional occupations are likely to be largely labour-augmenting, as professionals draw on AI as a tool or platform to perform specific tasks that increase their productivity (e.g. lawyers using AI to read large numbers of past cases to search for precedents and other arguments to use in a current case, or marketeers using machine learning models to improve sales forecasting accuracy).

Figure 3-1: Estimated net employment effects of AI by broad occupational category (SIC 1)

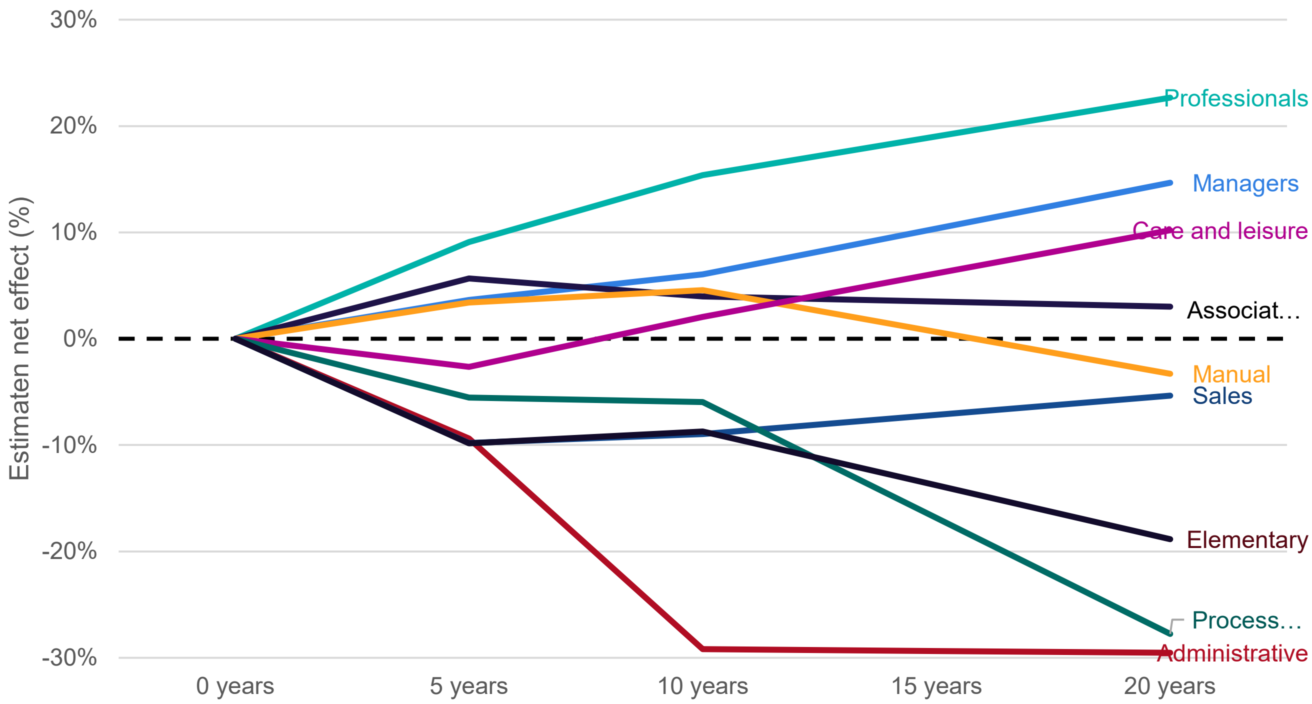

Source: PwC analysis of OECD PIAAC and ONS APS data

Net job creation is also estimated to occur in managerial occupations - for which the tasks involved are difficult to automate. Also, **other occupations requiring 'human touch' such as those related to caring or leisure are expected to increase as well**. Another insight from our expert workshop is that in the service economy, people increasingly value 'experience' and will still want to be served by humans for several reasons - including those related to status (e.g. a hotel concierge or a waiter in a high-class restaurant). In the area of care, humans will continue to transmit more trust when it comes to taking care of relatives. Oxford Economics and Cisco (2017), for example, refer to the phenomenon in which the rollout of AI makes the remaining jobs more human as an "AI paradox".

It is true that AI is becoming increasingly emotionally aware, capable of recognising emotion on faces or detecting depression in voices more accurately than humans. However, these advances will generally augment human labour, for example allowing care workers to focus more on patient care, and enabling retail workers to focus on providing a positive experience for shoppers rather than just stocking shelves, manning check-out tills and providing routine advice.

For some occupational groups, the net effect on employment is expected to be positive and increase steadily over the next 20 years (e.g. professionals and managers). In other cases, the estimates suggest changing patterns over time. As shown in Figure 3-1 this is especially apparent for occupations in sales and customer service (where the highest rate of job displacement is estimated over the next 5 years), administrative occupations (where we estimate particularly high job displacement in 5-10 years), and **manual occupations including taxi drivers, where we expect high rates of job displacement in the 2030s but probably not before then** based on views from our expert workshop and other studies (e.g. PwC, 2018d) on the likely pace of roll-out of driverless vehicles.

The different patterns and timeframes in which AI will impact different occupations can be seen in the light of three overlapping waves of AI (see Table 3-1, which is based on an updated version of the analysis in PwC, 2018b). The technologies that drive these waves have already been developed to varying degrees, but will only be rolled out at scale once they become sufficiently cheap and once legal, regulatory and organisational barriers to adoption have been overcome.

Table 3-1: Waves of artificial intelligence

Wave	Description	Examples of labour displacement	Examples of labour augmentation
Algorithm wave	Automation of simple computational tasks and analysis of structured data in areas like finance, information and communications, and retail, for example in the form of trading algorithms and recommendation algorithms - this is already well underway.	Travel agents (online chatbots) Sales assistants (recommendation algorithms) Financial traders (high-speed trading algorithms)	Economist (machine learning forecasts) Marketers (behavioural analytics) Strategists (dynamic price optimisation)
Augmentation wave	Automation of repeatable tasks such as filling in forms, communicating and exchanging information through dynamic technological support, and statistical analysis of unstructured data in semi-controlled environments such as aerial drones and robots in warehouses – this is also underway, but is likely to come to full maturity in the 2020s.	Personal assistants (virtual assistants) Administrators (unstructured data classification) Postal workers (aerial drones and robots in warehouse)	Doctors (medical diagnosis) Musicians (Human-AI collaboration) Police (face recognition)
Autonomy wave	Automation of physical labour and manual dexterity, and	Taxi, van, bus, train, crane, fork-	Transport and logistics

Wave	Description	Examples of labour displacement	Examples of labour augmentation
	problem solving in dynamic real-world situations that require responsive actions, such as in manufacturing and transport (e.g. driverless vehicles) – these technologies are under development already, but may only come to full maturity on an economy-wide scale in the 2030s.	lift drivers (autonomous vehicles) Waiters, waitresses and chefs ("robo chef")	managers (autonomous vehicles) Programmers (evolutionary algorithms)

Source: Updated analysis based on material in PwC (2018b)

Figure 3-2 decomposes the estimated net employment effect of AI on SOC 3 occupations by the countervailing effects of job creation and job displacement over 20 years. **The upward sloping relationship in the scatter graph shows that, at the occupational level, estimated job creation and job displacement tend to be inversely related to each other.** This is partly due to the replacement effect: workers will switch from occupations involving tasks that are highly automatable to jobs which are harder to automate.

Figure 3-2: Estimated % jobs created and displaced by AI for each SOC 3 occupation over 20 years

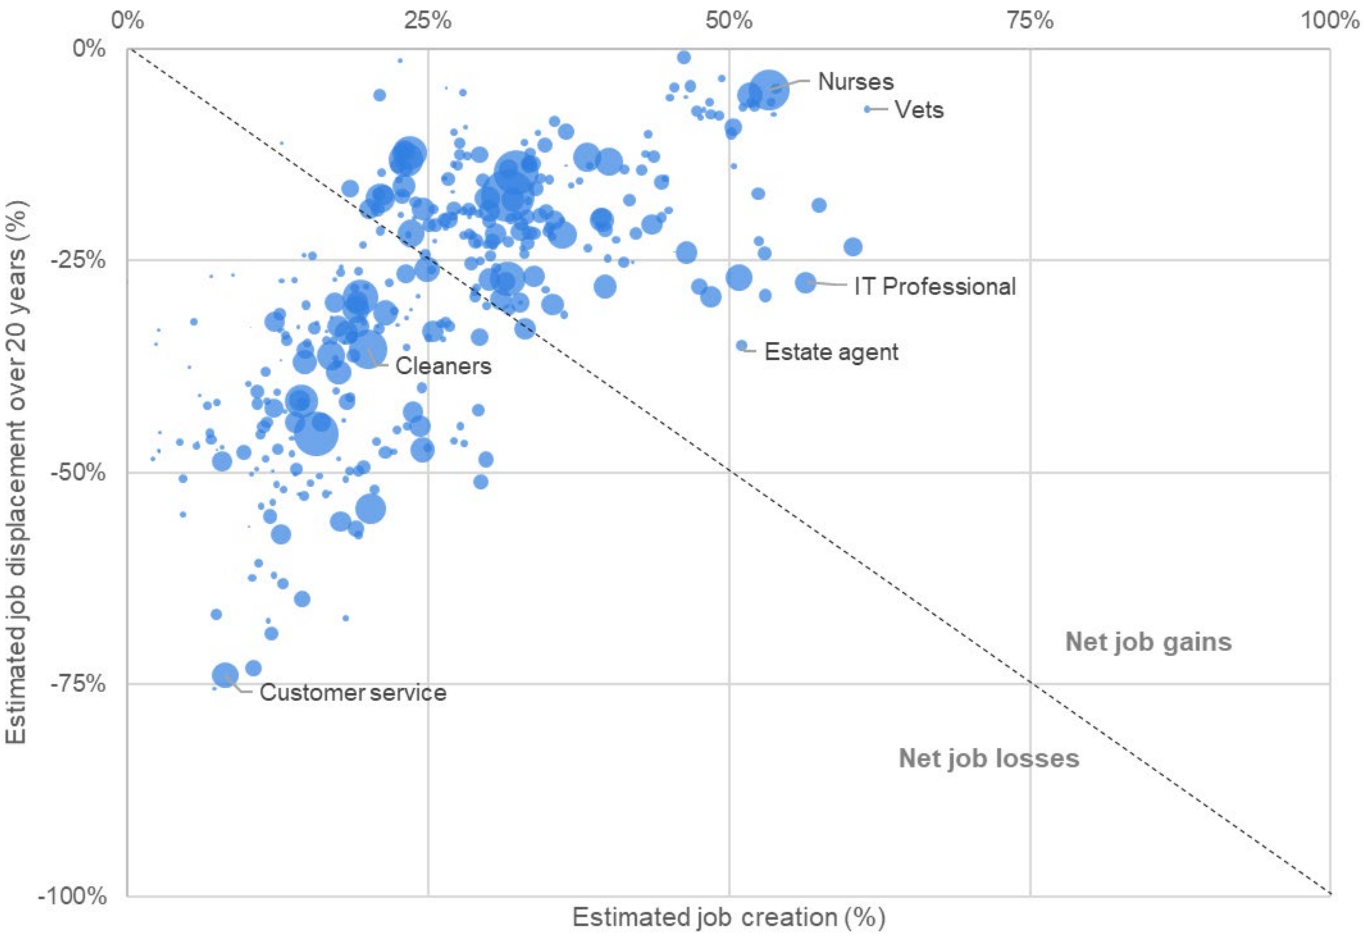

Source: PwC analysis of OECD PIAAC and ONS APS data

Note: This chart plots the rate of job creation and job displacement after 20 years as a result of AI and related automating technologies for each SOC 4 occupation, with the bubble size weighted according to the level of employment in 2018. Subsectors to the right of the 45 degree line are estimated to see net job gains over 20 years and we expect net job losses in sub-sectors to the left of this line.

Occupations in the top right of the graph are estimated to increase the most over the next 20 years in net terms, with employment in some cases estimated to grow by over 50%. **Nurses are the largest of the cluster of occupations expected to see relatively high job creation and low job displacement.** Most of the other occupations that share a similarly positive outlook are also professional occupations such as medical practitioners, vets, solicitors and psychologists. According to our model, the least automatable occupation is chief executive.

The most automatable occupation is estimated to be a telephonist (i.e. switchboard operator), although there are less than 10,000 left in the UK, reflecting the fact that automation has already largely displaced workers in this area. **There are nearly 300,000 people employed in customer service occupations, where nearly three quarters of jobs are estimated to have a high probability of displacement over the next 20 years.**

We also expect broadly similar levels of job creation and job displacement for SOC 3 occupations within the same SOC 1 category. Figure 3-2 shows that broad occupational categories (SOC 1, represented by different colours in the chart), form clusters in the job creation-displacement space. This has important implications for the ease with which workers that have been displaced by AI can find new jobs, because workers displaced in manual occupations, for example, may struggle to find similar jobs - in similar occupations in the same industry or the same occupation in another industry (in what would be a ‘local’ transition to a similar occupation). Evidence shows that workers tend to transition to jobs that are part of the same ‘occupational family’ (OECD, 2019).

Net employment effects by industry sector

As described below, our analysis suggests that the adoption of labour-replacing AI is likely to vary systematically by sector. We find that job creation from AI also varies significantly by industry, mainly due to differences in the estimated growth in demand for different types of goods and services, but also to differences in the extent to which labour-augmenting AI will affect different industry sectors.

Figure 3.3 derives the net effect of AI by broad industry sector (SIC 1) from the absolute number of jobs we estimate will be displaced and created in each sector, assuming the size of the UK labour force remains constant.

Figure 3-3: Estimated net employment effects of AI by broad industry sector (SIC 1)

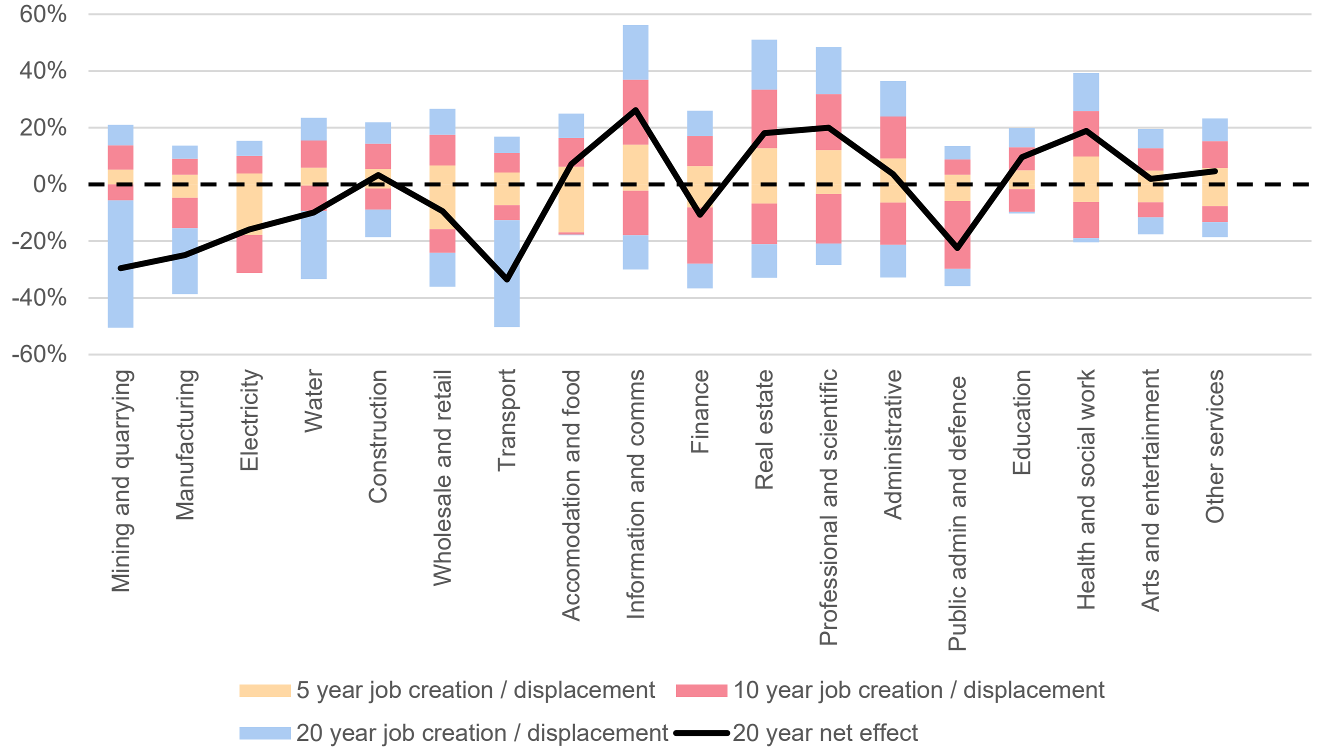

Source: PwC analysis of OECD PIAAC and ONS APS data

Note: This chart plots the estimated job creation and job displacement effects over 5, 10 and 20 years and the net effects of AI on employment over 20 years by industry. The job creation and job displacement effects are represented by stacked bar charts: the yellow portion of the bar above (below) the dashed line corresponds to the 5 year job creation (displacement) effect; the red portions correspond to the 5-10 year job creation and displacement effects; and the blue portions correspond to the 10-20 job creation and displacement year effects. The 20 year net effects are represented by the black line.

The largest estimated net job gains from AI over the next 20 years are in the health and social care sector (Figure 3-3), which is an area of high demand due to an ageing population and the fact that health is a 'superior good', which in standard economic terminology refers to the fact that the demand for health services generally increases as society gets richer (see Box 3-1 for more details about skills gaps in these industries).

As well as a generally higher demand for health services, we expect high AI-related job creation in health as the industry transitions to collaborative and preventative care models and embraces precision medicine. We expect that many of the applications of AI and robotics in healthcare are likely to be complementary to human labour rather than substitutes. However, there are still barriers to overcome to realise the potential of these technologies in health, as highlighted by a 2017 survey in which only 27% of respondents from the UK said they were willing to have a major surgery performed by a robot instead of a doctor - the lowest willingness of any country surveyed, although still not an insignificant proportion of patients (PwC, 2017d).

Information and communications and other professional, scientific and technical services also see significant estimated net employment gains - which will include many highly skilled jobs linked closely to AI and other emerging technologies. Jobs in programming are estimated to increase by 70% as programmers are required to design innovations in AI and related technologies. In Section 4 below we show that many of these high growth industries are already facing skills shortages. For example, according to the official UK Employers Skills Survey (ESS), information technology was last year the sector with the second highest proportion of hard-to-fill vacancies after hospitality.

By contrast, **significant estimated net employment reductions are projected in wholesale and retail, finance and public administration sectors in the short to medium term**, both of which are relatively automatable sectors. Due to the size of the sector, the automating effect of AI in the retail sector may be particularly disruptive to the economy (wholesale and retail constitute more than 12 percent of UK employment).[27] Unlike in health, the uptake of AI in retail is likely to be 'consumer-led',[28] which affects employment in multiple ways. Sometimes technology will increase retail profits without displacing labour, for example augmented reality (AR) can allow customers 'try on' make-up before online purchases and so can reduce returns.[29] However, technology in the wholesale and retail sector may also be labour-replacing, for example fashion buyers must compete with AI systems that can make decisions based on a huge volume of data on individuals' specifications and preferences. While some senior fashion buyers may still be needed to make final judgments based on the recommendations of AI systems, as well as to build relationships with suppliers, the number required may be significantly less than at present and this could apply to a wide range of other retail jobs.

The **significant estimated net employment reductions in transport are only seen in the longer term because the experts from the workshop generally believed that driverless vehicles are only likely to roll out at scale across the economy in the 2030s, rather than the 2020s**, which is consistent with other research on this topic as noted earlier in this section.[30]

More importantly, once AI-directed robots become significantly cheaper due to mass production and become able to dexterously and dynamically interact with and learn from humans and the environment, they won't just be confined to factories and warehouses; service robots may replace humans in hotels, homes and hospitals (PwC, 2015). But some of this may not come to pass over the next two decades according to views from our expert workshop given the economic, legal and regulatory and cultural barriers to adoption of such robots at scale in the service sector.

Box 3-1: UK health and social care workforce of the future

The prospects of automation in the health care sector seem to be expanding, given potential applications of AI to reduce the burden of administrative work, to improve diagnosis and treatment, and to improve the management of resources. Yet, the health care sector remains in general very labour-intensive, and current occupational shortages

[27] Source: https://www.ons.gov.uk/employmentandlabourmarket/peopleinwork/employmentandemployeetypes/datasets/employmentbyindustryemp13

[28] By segments of the population that regularly shop via mobile, are receptive to recommendations and new experiences, open to rent and share products, amenable to new subscription-based business models, and demand next-day delivery (PwC Global Consumer Insights Survey (2018d) suggests that men aged 18 to 34 are more prone to this).

[29] See for example the Sephora Virtual Artist app.

[30] Bootle (2019), for example, contains a detailed discussion from an economist's perspective about why some technology or motor industry predictions of an earlier roll-out of driverless vehicles are likely to prove overly optimistic given the economic, legal and regulatory barriers that need to be overcome here, as well as the normal inertia of consumers in accepting driverless vehicles.

are likely to increase. Vacancies in the healthcare sector have increased during the past few years, according to official data. The vacancy rate has risen from 2 open positions for every 100 employed workers in 2013, to 3.2 in 2018. Social care in particular faces rapid growth in demand for its services in the context of an ageing population and the likelihood of automation is low according to our modelling results (and the intuition that jobs in this sector mostly still require the human touch).

According to estimates from Skills for Care (2018), the adult social care workforce will need to grow to over 2.0m employees by 2030 to maintain the same level of service we have currently, up from 1.6m today - spread across elderly, nursing and specialist care. This workforce is made up of frontline workers (carers in care homes, personal assistants and nurses) as well as managers and senior managers who coordinate care delivery and are integral to the working of the social care system.

The need could be greater. If the workforce demand grows at the same rate as the 75+ year old population, without any workforce efficiency savings, the UK could need closer to 2.5m employees by 2030. Although there may be some scope for automation, our modelling suggests there will be limits to this as the human touch will still be important in many health and social care roles.

At the same time, if use of AI and related technologies boosts productivity in the economy as a whole, this will feed through into strong GDP growth and so stronger tax revenues, providing the funding for expansion of the NHS and social care workforce. This is the key macroeconomic mechanism behind the rapid job creation in this sector in previous PwC research using CGE modelling techniques (2017c, 2018c). It is also the driving force behind the net employment gains projected in this sector in the present study.

Net employment effects by region

Regional variations in estimated employment effects from AI are not as large as between occupations or industries - the net effect of AI on regions varies by about +/-10% at most, compared to more than 50% across occupations. This is because all regions have a broad range of occupations, albeit with some variations in occupational mix that drive differences in estimated automatability and job creation by region in our model.

As Figures 3-4 and 3-5 show, **we do find some tendency for more positive net employment effects from AI and related technologies in higher income areas such as London and the South East and more negative net employment effects in some towns in Northern England and the Midlands. These variations are also often more marked within rather than between broad (NUTS 1) regions**, as illustrated by the two NUTS 3 level maps in Figure 3-4.[31]

[31] However, as noted above, results are also less reliable at greater levels of regional (or other) disaggregation, so we would caution against focusing on specific estimates for sub-regions as opposed to broad patterns of results across the country.

Figure 3-4: Estimated net employment effects of AI for NUTS 3 regions over 20 years

a) % net employment effect *b) net employment effect (no. jobs)*

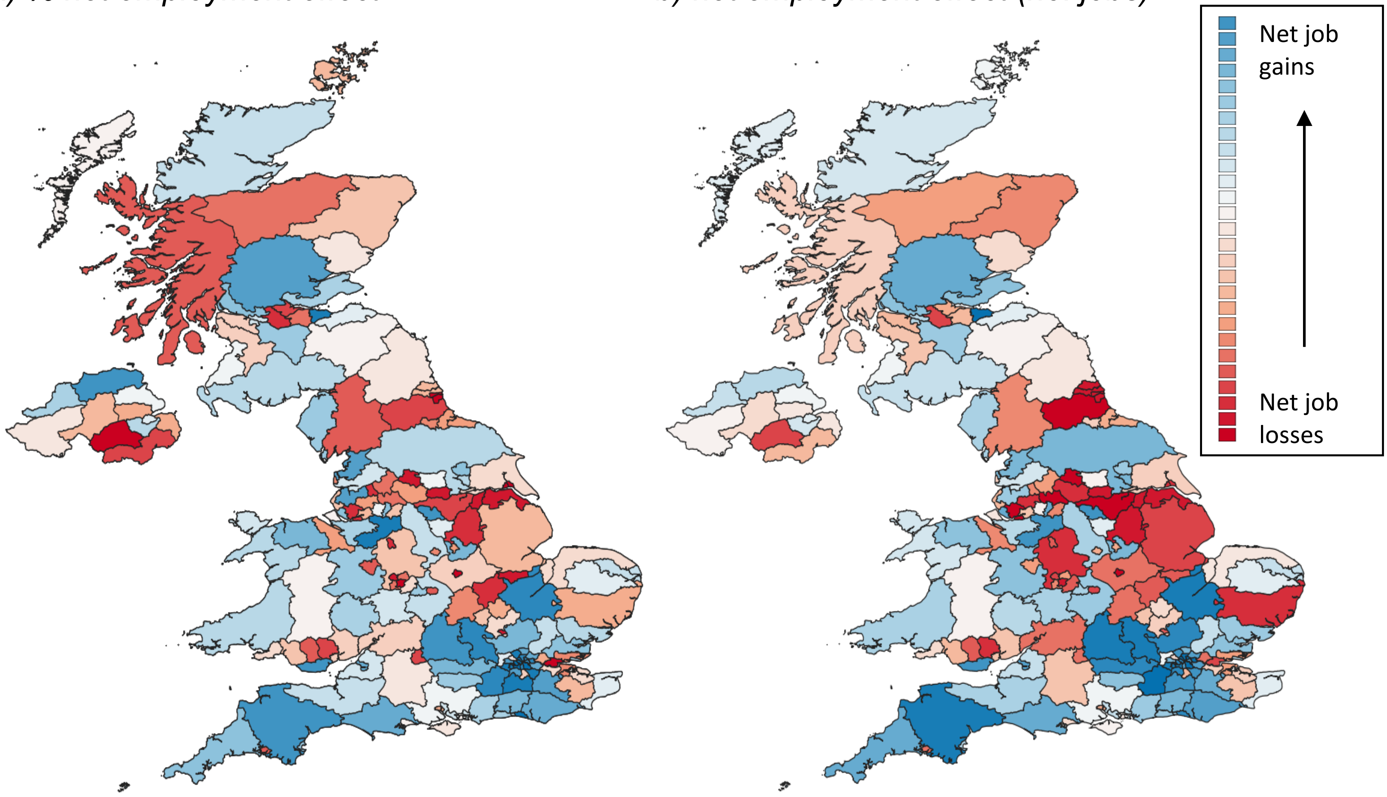

Source: PwC analysis of OECD PIAAC and ONS data

Note: Regions in the maps are shaded by the estimated net employment effect over 20 years, in a) % terms and b) job numbers. High positive net employment effects are estimated in regions shaded dark blue. Regions shaded in red are estimated to see net job losses as a result of AI over 20 years.

Figure 3-5: Estimated % of jobs created and displaced by AI for NUTS 1 regions in 5, 10 and 20 years

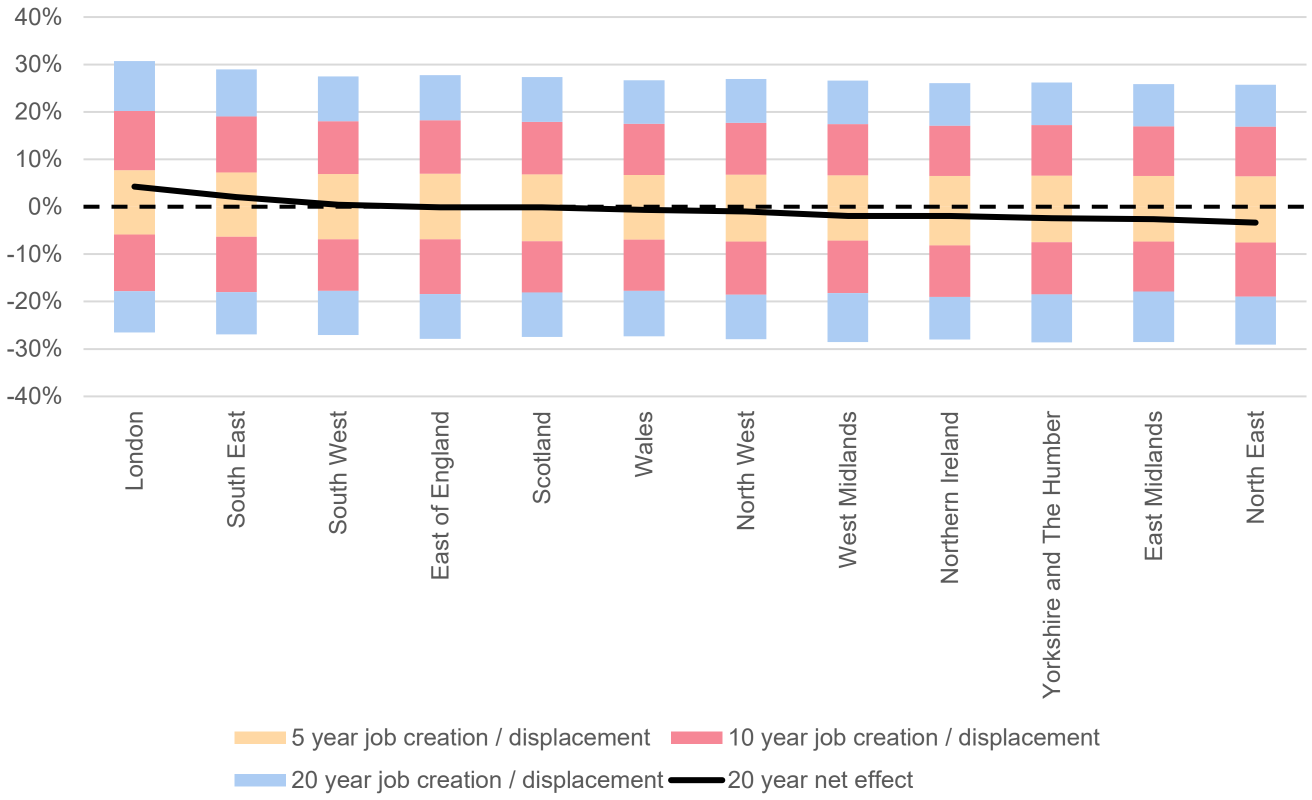

Source: PwC analysis of OECD PIAAC and ONS APS data

Note: This chart plots the estimated job creation and job displacement effects over 5, 10 and 20 years and the net effects of AI on employment over 20 years by region. The job creation and job displacement effects are represented by stacked bar charts: the yellow portion of the bar above (below) the dashed line corresponds to the 5 year job creation (displacement) effect; the red portions correspond to the 5-10 year job creation and displacement effects; and the blue portions correspond to the 10-20 job creation and displacement year effects. The 20 year net effects are represented by the black line.

We have estimated the potential employment impact of AI on UK regions using the above estimates on the impact of AI by SOC 4 occupation and a matrix containing the occupational mix of each region obtained from the ONS's latest Annual Population Survey (APS). We have not taken into account other factors that may vary systematically by region, such as skills or access to re-training programmes. These other factors are likely to increase differences in the estimated employment impact that AI will have on regions.[32]

[32] Another factor we have not modelled is that is the cost-effectiveness of automation may vary by region. Robo-chefs, for example, are likely to be rolled out in London before they come to the Scottish Highlands. This is partly due to the nature of existing restaurants; London has a higher concentration of chain stores such as McDonald's which have the profits to invest in automating technologies. It is also a function of wages; there is more incentive to automate waiters if they cost more. There is also likely more appetite from consumers for automating technologies in London, who expect tech-enabled services and there is a large young population willing to pay for novel experiences. As a general point, certain people, regions and countries may adapt faster and move ahead, which could exacerbate different types of inequalities (Acemoglu, 1998).

Given that regional differences in the impact of AI are a product of the occupational makeup of regions, an important factor determining the net effect by region is occupational diversity.

As most large regions are occupationally diverse, with a core workforce of nurses, teachers, care workers, sales and retail assistants etc., the net effect of AI on regions even at the NUTS 3 level only varies by about +/-10%. Nevertheless, there are some striking patterns.

Net job losses are focused in the North England and the Midlands. In fact, out of the 10 NUTS 3 regions with the most negative percentage net effect, 8 of these are towns or cities in the North and Midlands (see Table 3-2). These urban areas are characterised by high estimated job displacement and low estimated job creation. At a superficial level, this can be explained by broad occupational categories. For example, of the workers in Sunderland, only 5%, 12.5% and 9% are managers, professionals or associate professionals, respectively, compared with the national average of 10%, 20% and 14%. This is generally also true for other regions in which we estimate a negative net employment effect from AI.

Table 3-2: NUTS 3 regions with the most positive and negative estimated % net effect of AI on employment after 20 years

NUTS 3 regions with the most positive estimated % net effect of AI on employment after 20 years	**NUTS 3 regions with the most negative estimated % net effect of AI on employment after 20 years**
Camden and City of London	Leicester
Westminster	Thurrock
Wandsworth	Sunderland
Kensington & Chelsea and Hammersmith & Fulham	Armagh City, Banbridge and Craigavon
Lewisham and Southwark	Sandwell
Brighton and Hove	Kingston upon Hull, City of
Bromley	Bradford
Cheshire East	Wolverhampton
Barnet	Peterborough
Hounslow and Richmond upon Thames	Wakefield

Source: PwC estimates based on analysis of OECD PIAAC and ONS data

However, there are important differences in the occupational mix of these towns and cities at a more granular level. For example, according to the APS, 8% of workers in Bradford are van or taxi drivers. In Leicester and Wakefield, there is a high proportion of people (9% and 6.5% respectively) employed in elementary storage occupations. In Sunderland, there is a relatively high number of people working in call centres, customer service occupations and doing government and other administrative jobs. The speed with which AI-related automation could take hold in these towns is therefore likely to vary according to the specific occupations upon which they are most reliant.

The net employment gains projected in London and the South can largely be explained by the greater prevalence of managers and professionals in these regions, who are more likely to experience AI as augmenting their capabilities rather than displacing their jobs. Of the NUTS 3 regions with the highest estimated net effect of AI, only Cheshire East and Brighton and Hove are not in Greater London. London in particular is characterised by a relatively high proportion of CEOs, directors in finance and sales and marketing, programmers and other IT occupations, and other professionals with relatively low estimated rates of job displacement.

To get an accurate picture as to how AI may affect regional inequality, however, it is necessary to look beyond North-South or East-West divides. Our estimates suggest that the more significant differences will play out between urban and rural areas, and even more so between smaller and larger urban areas as smaller towns and cities face greater displacement than large cities as economic growth and tech talent concentrate in bigger, prosperous urban areas.[33]

Moreover, the regional differences in the estimated net employment effect of AI could tend to worsen income inequality. Figure 3-6 shows a **significant positive relationship between the net employment effect of AI by region and the median earnings in that region: in other words, poorer areas face a greater threat from automation.** Note that this relationship still holds, albeit somewhat less strongly, if the three outliers in the top right corner are excluded from the analysis.

Figure 3-6: Median earnings vs % estimated net employment effect of AI by region

Source: PwC analysis of OECD PIAAC and ONS data from APS and ASHE

Note: Each circle represents a NUTS 3 region and the size of the circle is weighted by the level of employment in the region. The circles are coloured by whether the region is in the (i) Midlands and North England, (ii) South England and East of England, or (iii) Scotland, Wales and Northern Ireland (see legend)

---

[33] According to the OECD (2018), regions with smaller risk of automation are characterised by a larger share of workers with tertiary education, a larger proportion of jobs in services, and are highly urbanised.

According to a recent OECD study, place-based policies may grow in importance in the light of growing public discontent with the economic, social and political status quo in many regions (OECD, 2019) and projected technological, demographic and environmental megatrends. The impact of automation could be one such trend for the UK, based on our regional analysis, the broad patterns from which appear robust to alternative job creation scenarios based on the further analysis summarised in Technical Annex 5.

Net employment effects by socio-economic and demographic group

Our finding that AI has differential impacts on occupations has important implications for people as well as places. As with the regional analysis, we have used the latest detailed (SOC 4) occupational breakdown from the APS to estimate the potential employment impact of AI for selected socio-economic and demographic groups. Specifically, we have made estimates by earnings level, education attainment, age and gender.

Analysis by earnings levels - will skill-biased technological change continue?

As Figure 3-7 below indicates, managerial and professional occupations with higher median earnings levels tend to see significantly positive estimated net employment effects. This relationship is broadly similar to the analysis for NUTS 3 regions by earnings level in Figure 3-6 above.

Figure 3-7: Estimated net employment effects of AI by occupation and median earnings level over 20 years

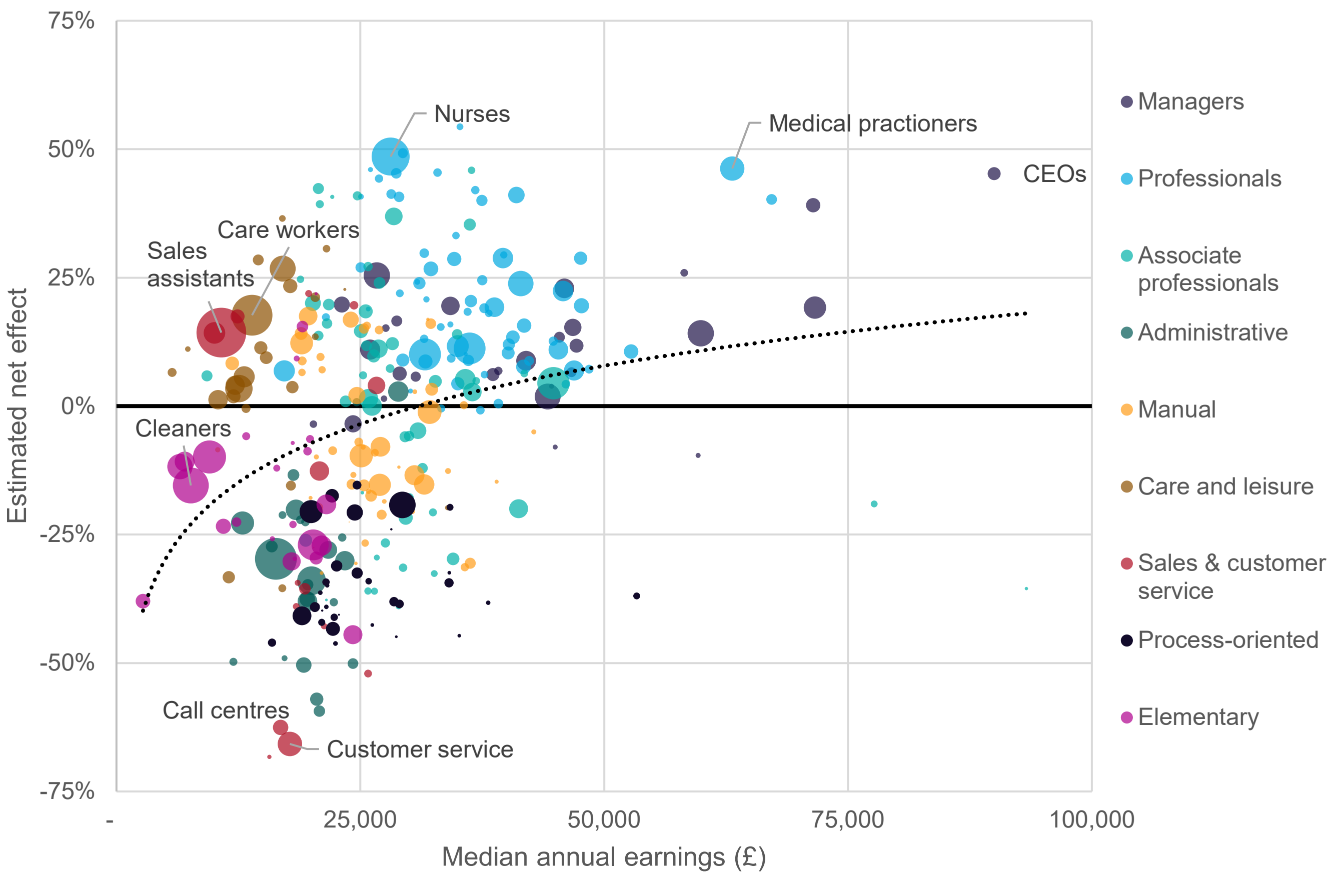

Source: PwC analysis of OECD PIAAC and ONS data from APS and ASHE

Note: Each circle represents a SOC 4 occupational category. SOC 1 occupational categories are denoted by the colour of the circle (see legend) and the size is weighted by the level of employment. For each occupational category, the chart plots the estimated net effect of AI over 20 years on the y-axis and the median annual earnings on the x-axis. Occupations towards the top right have a positive estimated net employment effect and high earnings.

Less well-paid clerical and process-oriented roles see negative estimated net employment effects from AI, although only in the longer term for manual workers such as drivers. Having said that, there are still many relatively lower paid jobs (e.g. care workers) that may see rising net labour demand due to the positive macroeconomic effects of AI on real income levels and so the ability to pay for these services (either directly or via taxation for public services like the NHS). This could also raise wage levels in these occupations, depending on labour supply constraints, although quantifying this effect on wage levels within an occupation is beyond the scope of this study.

Our analysis suggests that the roll-out of AI and related technologies across the economy could see a continuation of skill-biased technological change with the potential to widen existing earnings differentials. Skill-biased technological change, which mainly enhances the earnings potential of skilled professionals, seems based on past research to have been a common feature of the US, the UK and other advanced economies since the 1980s (Berman et al., 2018). Past research suggests that this has led to a polarisation of the labour market, with more highly paid jobs being created for the most skilled, whose work is enhanced by technology (Acemoglu, 1998). At the same time, lower quality and relatively poorly paid jobs emerged for the less skilled (often in low-value services sectors with relatively stagnant real wages).

Over the same period, many middle-level jobs have been displaced by automation and/or offshoring (itself facilitated by new information and communications technologies as well as the opening up of China, India and other emerging economies since the early 1990s). See, for example, the discussion in Mortensen and Pissarides (1999). Similar patterns have been found in other developed countries.[34] How such trends evolve in the future will depend in part on how AI is used across industries in ‘labour-replacing’ and ‘labour-augmenting’ tasks.

How do our results compare with these past studies? Although there is a generally positive relationship between the estimated net employment effect by occupation and median earnings, some of the lowest paid jobs (e.g. in social care) are not the most automatable according to our estimates. The occupations that we estimate will see the biggest net employment losses generally have a median wage of around £20,000. The main reason for this apparent anomaly is known as ‘Moravec’s paradox’. Writing in the 1980s, Moravec explains that "it is comparatively easy to make computers exhibit adult level performance on intelligence tests or playing checkers, and difficult or impossible to give them the skills of a one-year-old when it comes to perception and mobility" (Movarec, 1988). In other words, AI is good at many things that we find hard, and bad at many things that we find easy. The lowest paid jobs, which includes bar staff, waiters, cleaners, kitchen and catering assistants, involve skills such as perception and mobility that many people, including young and relatively uneducated people, are able to do, but which remain hard for robots and seem likely to remain so for some time.

[34] For job polarisation in Britain, see Goos, M., & Manning, A. (2007). A cross-country analysis is provided by Michaels et al. (2014).

In addition to our analysis for this study, PwC (2019) has also recently published a survey of attitudes to digital disruption for a representative sample of around 2,000 workers across the UK (see Box 3-2). One of the notable findings was that higher earners were generally more positive about the impact of automation on their employment and earning prospects than lower paid workers, which is consistent with the estimated impacts in the present study.

Box 3-2: How do workers view the challenge of automation?

The jobs in which displaced workers move to, and the way in which workers respond to digital disruption more generally, will also depend on their attitudes towards AI and digital upskilling. In general, the UK population is becoming more aware of the ongoing transformations. According to PwC's 2019 Future of Government survey for the UK, almost a third (30%) of workers think that their job will be significantly impacted by automation in the next 10 years. Almost half (46%) think that their job will change mostly or completely in 10 years due to technology. More respondents see risks (40%) than opportunities (21%) arising from automation, with 38% feeling concerned and 26% feeling anxious. This cautious view of automation across the UK contrasts with global attitudes, where 50% of respondents in a recent PwC global survey saw more opportunities than risks.

According to the UK survey results, there are also notable differences across occupational and socio-economic groups, as described below.

Manual workers in the UK are more likely to think that automation presents more risks than opportunities (48% manual workers vs 42% senior manager and above), but they are also less likely to be willing to take action themselves to reskill if affected by automation (e.g. 45% would not take evening classes vs. 34% overall).

Manual workers are also more likely to think the government should be banning certain automation processes to prevent job losses (74% vs 66% overall), and they're more likely to be in favour of introducing quotas to businesses to ensure a percentage of human jobs remain in place (81% vs 76% overall).

Low earners in the UK, defined here as those on a gross annual income of up to £20,000, are more likely to see risks than opportunities from automation (47% vs 40% overall) and are more likely to look to the government for financial support (47% think the government should definitely help vs 40% overall). They're also more likely to think the government should subsidise training courses (49% vs 45%), introduce quotas for human jobs (80% vs 76%) and ban certain automated processes (70% vs 66%).

In comparison, UK workers considered to have a mid-to-higher level income, defined here as between £40,000 and £60,000 a year, are more likely to say that automation is unlikely to make their job obsolete in the next 10 years (44% vs 38% overall). Similarly, they are more likely to say that their job will mostly be the same in 10 years' time (46% vs 39% overall).

High earners in the UK, defined here as those earning £60,000 or more a year, are more likely to think that automation presents opportunities compared to other income groups (28% vs 21% overall). They're also more likely to be curious (26% vs. 20% overall) and excited (17% vs. 10% overall) about the prospect of up to 30% of jobs being automated by 2030. These differences by income group seem to be directionally consistent with our modelling results for this study, as shown in Figure 3-7.

More men (40%) than women (33%) in the UK think that their job will be significantly changed or made obsolete by automation. This is directionally consistent with the results of this study looking 20 years ahead (see Figure 3-10 below), although not over shorter time horizons where clerical workers, where women are currently more heavily represented, could be most at risk of automation from AI and related technologies such as robotic process automation. But our modelling does not suggest large differences by gender in net employment effects (as compared, in particular, to differences by education levels).

Those aged between 16 and 24 are more likely to believe that their job will be made obsolete or significantly change due to automation in the next 10 years (59% vs 36% overall). But they are generally also more willing to take steps to address this (e.g. 74% are prepared to take an evening class vs. 66% overall, and 56% are prepared to accept a lower salary than their previous job vs 45% overall).

There are still, though, some awareness gaps: young people in cities, and black and minority ethnic groups, appear to be more excited about the prospect of automation and learning new skills than the average UK citizen. In contrast, those in low paying jobs seem, on the basis of this survey, to be less ready and willing to reskill.

This is only one survey, so a broader range of attitudinal evidence would be useful to test the robustness of these findings. However, if skill-based interventions are to be successful it is important that policymakers are alive to the different attitudes to digital disruption that are prevalent in different parts of the working population.

Analysis by educational levels

Our conclusion that AI may see a continuation of skill-biased technological change is also supported by the positive correlation we find between estimated net employment effects of AI and education levels (Figure 3-8). As with our occupational and regional results, this finding seems robust to plausible alternative scenarios for aggregate job displacement and job creation (see Technical Appendix 5 for details).

Figure 3-8: Estimated net employment effects of AI by education levels

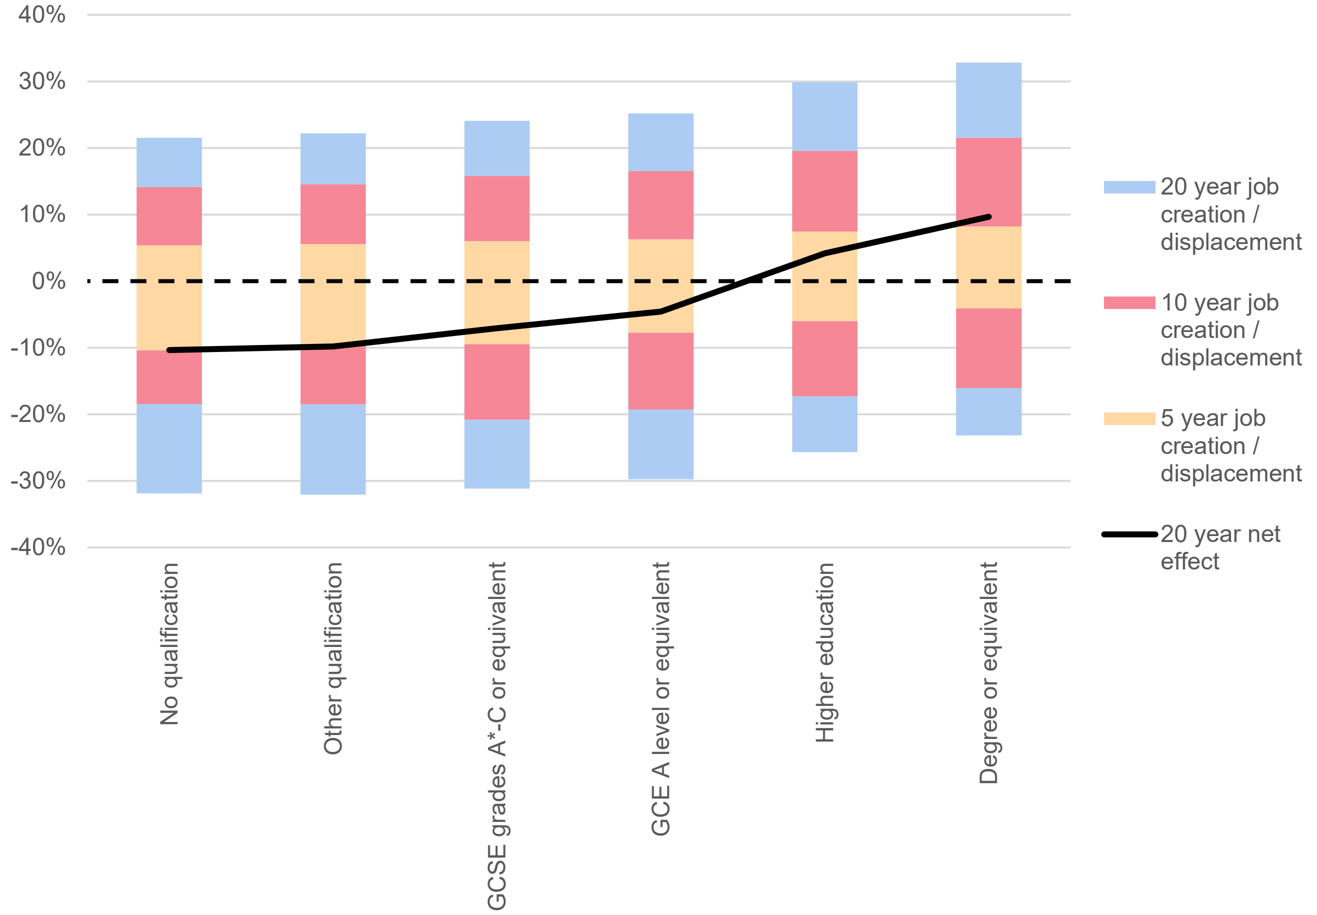

Source: PwC analysis of OECD PIAAC and ONS APS data

Note: This chart plots the estimated job creation and job displacement effects over 5, 10 and 20 years and the net effects of AI on employment over 20 years by education level. The job creation and job displacement effects are represented by stacked bar charts: the yellow portion of the bar above (below) the dashed line corresponds to the 5 year job creation (displacement) effect; the red portions correspond to the 5-10 year job creation and displacement effects; and the blue portions correspond to the 10-20 job creation and displacement year effects. The 20 year net effects are represented by the black line.

We do still find some areas of significant potential job displacement by AI for graduates, since some highly skilled tasks (e.g. medical diagnosis or analysing stock market trends) may also be relatively amenable to AI. But, if anything, our estimates of relative net job creation for graduates may be under-stated as they may also be more adaptable in developing new skills to make themselves complementary to AI rather than substitutable by it. This appears to have been true for earlier waves of digital innovation that have been associated with increased job polarisation (e.g. getting a laptop with internet access have made professionals more productive by eliminating routine tasks such as going to a research library to get data, or making data analysis much quicker and more productive). This could tend to further widen the gap between the positive net effect for those with higher educational qualifications and those with lower qualifications.

Analysis by age group

We have not found large variations in estimated net employment effects by age group, except for somewhat higher potential job displacement for younger workers (aged 16-24). We estimate that the jobs performed by people aged 16-24 could see net job losses of around 6% over 20 years (Figure 3-9).

Workers aged 16-19 are concentrated in certain occupational groups, notably sales and retail assistants (20% of workers aged 16-19), kitchen and catering (12% of workers aged 16-19) and waiters and waitresses (10% of workers aged 16-19). Although not all of these face high automation probabilities, younger workers are, for understandable reasons, not generally yet in managerial or professional roles for which we estimate net positive impacts from AI.

From a policy perspective, a potential negative net effect of AI on young workers may not be a significant cause for concern, as most of the jobs held by people aged 16-24 are only temporary. Moreover, due to Moravec's paradox, as discussed above, the basic skills that are required by the jobs that many young people do - perception, mobility etc. - are relatively hard to automate. So there should still be entry level jobs for young workers, even if not all of these are the same as today.

Figure 3-9: Estimated net employment effect of AI for different age groups

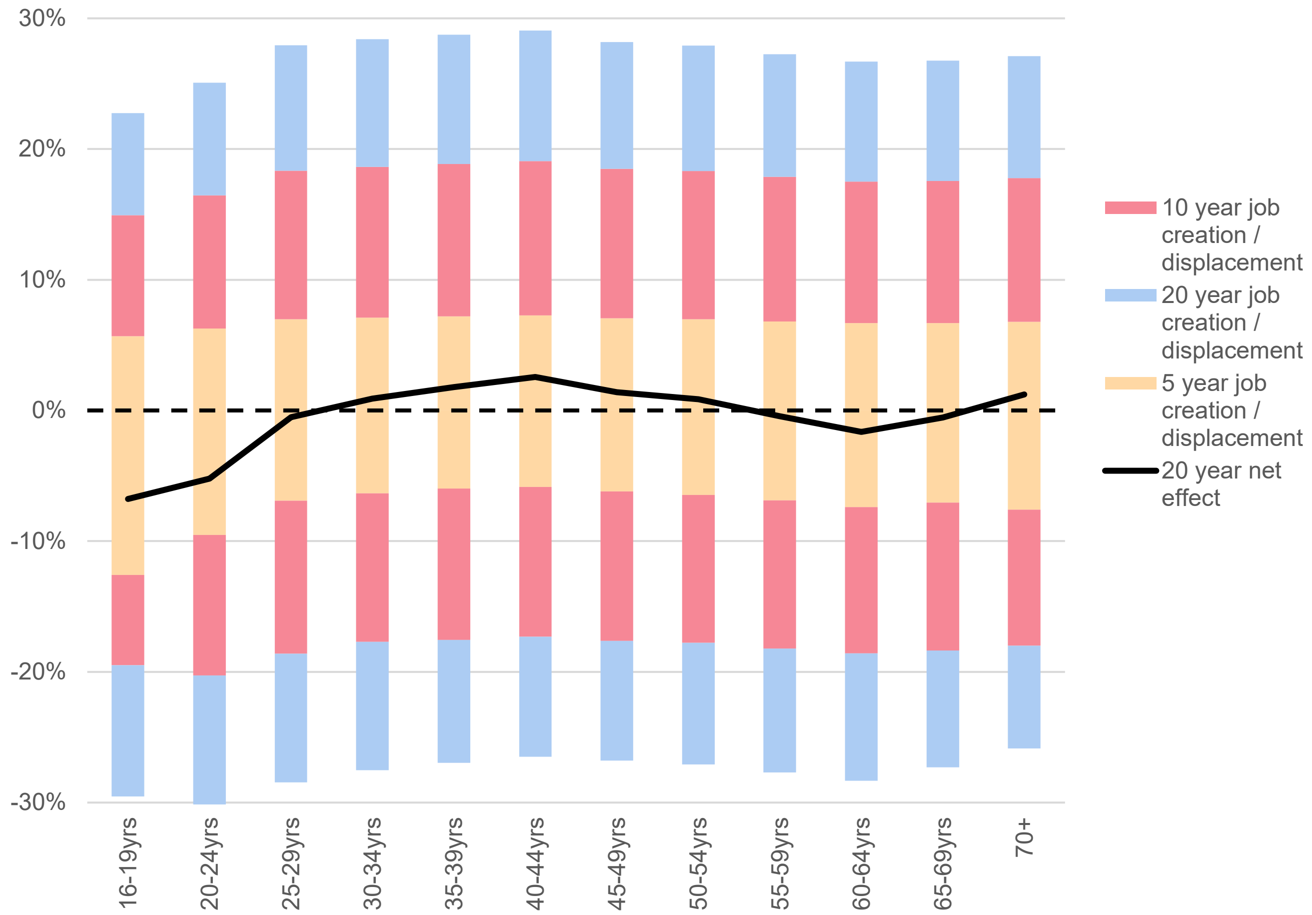

Source: PwC analysis of OECD PIAAC and ONS APS data

Note: This chart plots the estimated job creation and job displacement effects over 5, 10 and 20 years and the net effects of AI on employment over 20 years by age group. The job creation and job displacement effects are represented by stacked bar charts: the yellow portion of the bar above (below) the dashed line corresponds to the 5 year job creation (displacement) effect; the red portions correspond to the 5-10 year job creation and displacement effects; and the blue portions correspond to the 10-20 year job creation and displacement effects. The 20 year net effects are represented by the black line.

Analysis by gender

We also do not find any very material differences in potential net employment effects by gender (see Figure 3-10), although the results seem slightly more negative for occupations (e.g. clerical workers) with a higher current proportion of female workers over the next 5-10 years, but slightly more negative for occupations with a higher proportion of male workers (e.g. drivers) over a 20 year time horizon.

Figure 3-10: Estimated net employment effect of AI over 5, 10 and 20 years by gender

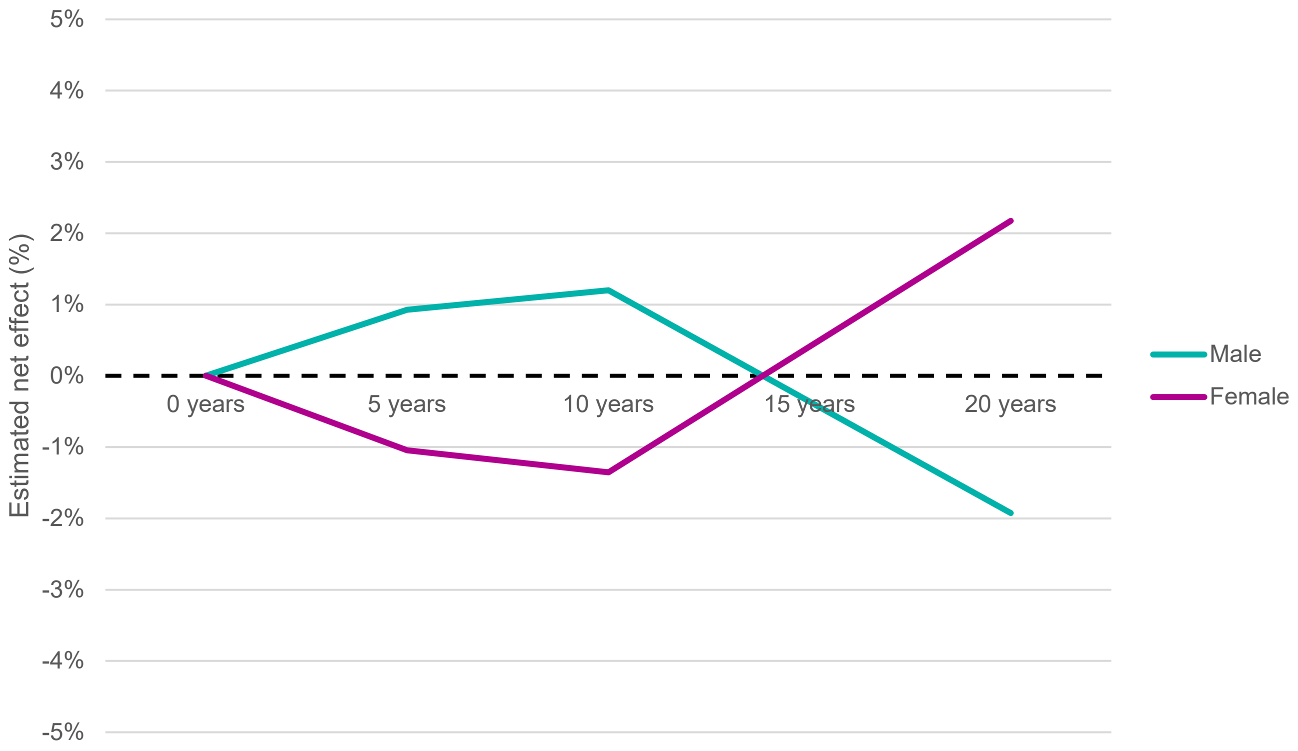

Source: PwC analysis of OECD PIAAC and ONS data from APS and ASHE

Note: The chart above plots the estimated net effect of AI over time by gender. We have estimated the effect over 5, 10 and 20 years. Since we have assumed in this scenario that the total net effect of AI is zero, it follows that a positive net effect of AI in one gender implies a negative net effect in the other.

That the difference between the impact on men and women is estimated to be so small might seem somewhat surprising given that the current distribution of employment by occupation is significantly related to gender. Figure 3-11 shows that administrative occupations and care and leisure are dominated by women at present, while tradespeople and machine operatives are generally men. Nonetheless, it appears that these occupational differences broadly offset each other in terms of the overall estimated net effect of AI on employment by gender, albeit with some smaller differences in the projected time profile of these impacts as Figure 3-10 above illustrates.

Figure 3-11: Proportion of male and female workers in each broad occupational category (SOC 1)

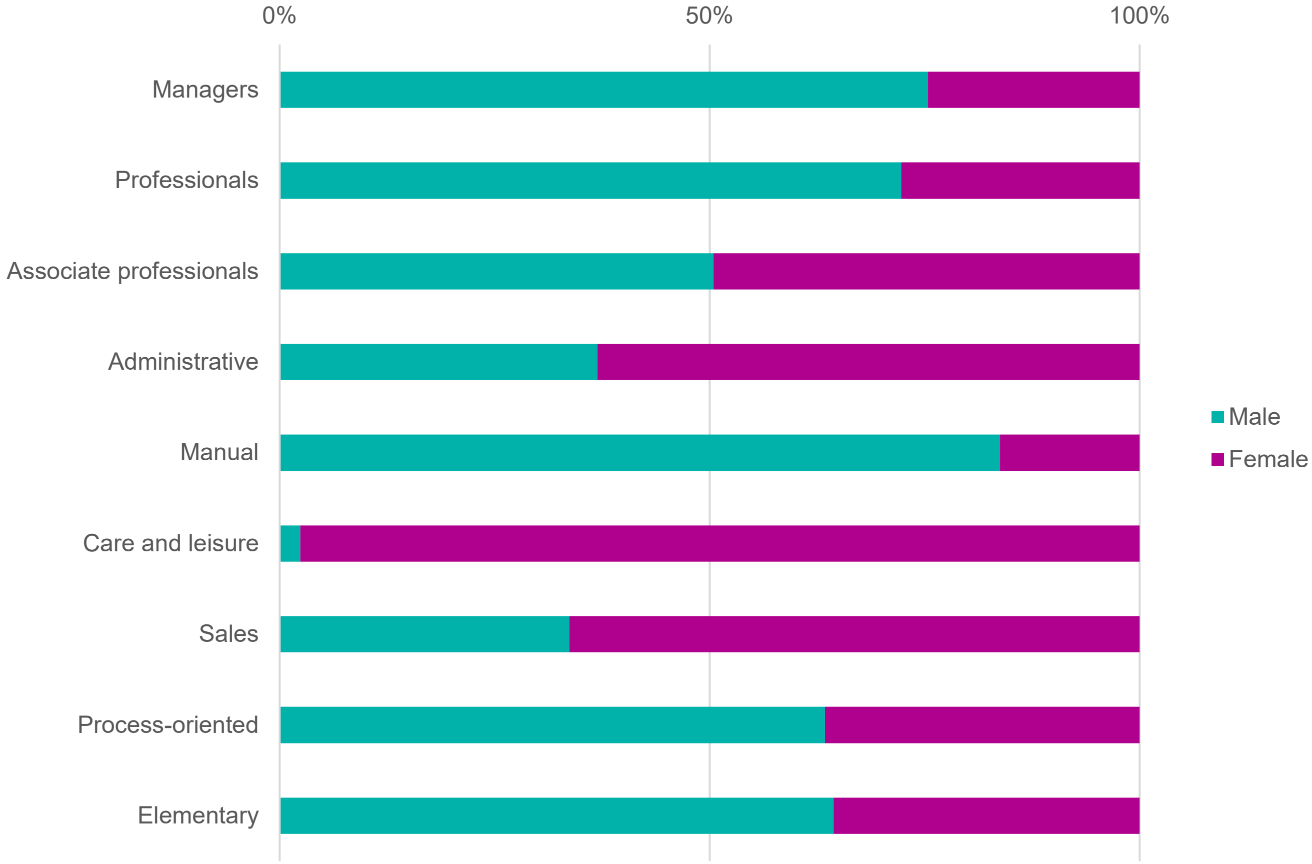

Source: PwC analysis of ONS data

Conclusion

Our analysis in this section suggests some significant potential variations in the net employment effect of AI across occupations, industry sectors, regions, educational groups and earnings levels. In some respects, these relative impacts also vary over time. We find less evidence of significant variations by age or gender. In the next section, we consider what these trends - if they come to pass - may mean for the future demand for skills.

4. Implications for skills

Key findings

- Almost all occupations on the current government shortages list are projected to show net employment increases due to AI, particularly nurses and medical practitioners.
- Estimates show a strong increase in demand for skills related to health care, particularly for medical support and basic patient care. In general, jobs in this sector will see low impact of automation – including care workers, which are not part of the official list of skilled-worker shortages but demand for whom is increasing rapidly with an ageing population.
- The demand for science and research and IT skills is also growing fast, although at lower rates.
- The demand for IT skills is more heterogeneous: some IT skills are clearly expected to decline over the years such as document management analysis. Skills related to some programming languages and those related to augmented reality and artificial intelligence are projected to increase over time.
- Skills relating to administration, finance and customer support may see net decreases over time as the tasks involved are increasingly automated.
- Our qualitative analysis suggests that adaptability and resilience will become key personal features needed to navigate the future of work.

Introduction

Our analysis in the previous two sections shows that AI is expected to displace some jobs by accelerating the automation of work tasks. When combining job displacement estimates with job creation scenarios by occupation and sector, the net employment effects will be positive for certain occupations (e.g. professional, managerial) and negative for others (e.g. occupations with more routine tasks). This section will address two additional research questions:

1. What are the implications of the estimated net employment effects of AI and related technologies for those occupations that are identified as in shortage in the UK?
2. What are the implications of the employment effects of AI for the future demand for skills in general?

The first question is relevant because the UK is already facing important skill shortages in a series of occupations. Our estimates of the net effect of AI on occupations suggest that these shortages may tend to increase further unless appropriate action is taken to develop the relevant skills of current and future workers.

The second question is addressed quantitatively by combining our net employment effect estimates (as described in Section 3 above) with data from online job adverts on the most common and important skills used by different occupations. This captures a “compositional” effect: the skills that will be in more or less demand as a consequence of the varying growth rates of different occupations in the context of technological change (measured in this analysis by our estimates of job displacement and job creation related to AI).

One limitation of this quantitative analysis is that it is assuming that the skills composition of specific occupations remains fixed over time.[35] Since AI and other automating technologies are likely to affect the set of tasks performed by different occupations in the future (as they start taking over tasks previously done by humans), this section also includes a qualitative discussion about how a "task" effect can reshape the skills composition of specific jobs and occupations.

The 'compositional' and 'task' effects define the expertise (or domain knowledge) needed by people. Moreover, relevant skills are also being determined by new work practices that are fuelled by technology. Given that an increasing number of people may in future need to adapt to careers and work practices that are becoming more flexible and subject to change, this section also includes a discussion about 'transversal' skills – i.e. those that cut across different occupations and industries. These include, but are in no way limited to: soft skills for more collaborative environments; resilience and adaptability for a more dynamic labour market; and leadership for organisations facing more dynamic environments.

Implications of AI for occupations with current skill shortages

The impact of AI on the rise and decline of different occupations is likely to create skills gaps if people do not re-skill for the jobs most likely to increase in the future. Skills gaps arise not only when people leave education and enter the workforce, but also each time they change jobs.[36] Past international evidence tells us that among displaced workers who are re-employed within a year, between 20% and 70% change occupation or industry. Of this group, roughly a quarter experience a major change in the skills required to do their jobs (OECD, 2012).

The UK is already facing considerable skills gaps. The OECD report 'Getting Skills Right' (2019) stated that about 40% of British workers are either overqualified or under-qualified for their job – with many working in a field that is different to the one in which they studied. A recent study by the Industrial Strategy Council (2019) finds that, by 2030, 7 million additional workers could be under-skilled for their job requirements. Moreover, many firms find it difficult to fill positions. According to the Employer Skills Survey, 20% of establishments showed vacancies in 2017 (vs. 14% in 2011), many of which are "hard to fill because of skills shortages".[37]

Importantly, the UK government tracks significant skill shortages in many types of occupation[38] from health professionals to specialists in science, technology and engineering. The net employment effects estimated in Section 3 give an indication as to whether such shortages might be expected to widen in the future due to the impact of AI and related technologies. As

[35] This means that, by only relying on the compositional effect, the estimated demand for skills becomes more uncertain as we extend the time horizon from 5 years to 10 and 20 years.
[36] This is called the 'sectoral shift hypothesis' in the literature (see Lilien, 1982; Petrongolo and Pissarides, 2001).
[37] This amounts to 209,000 'skill shortage vacancies' experienced to some degree by 6% of all employers - with some variability across sectors. According to the Government Office for Science, skills shortages vary according to sector and occupation. The gas, electricity and water industries have the highest proportion of jobs classed as hard to fill due to skills shortages (36%), followed by 34% in construction and 30% in manufacturing. By occupation, the proportion is highest in skilled trades (43%), machine operators (33%) and in professional occupations (32%) (Vivian et al., 2016: p151).
[38] We used the UK's Shortage Occupation List at the time when the analysis was carried out in October 2019, found here https://www.gov.uk/guidance/immigration-rules/immigration-rules-appendix-k-shortage-occupation-list

discussed earlier in this report, these estimates combine both the direct impact and the indirect macroeconomic impact from a boost to productivity growth and so to average real income and spending levels.

Many healthcare-related professions, for instance, are already part of the government's list of occupations in shortage and demand is expected to grow - particularly in the context of ageing populations. Since these occupations are labour-intensive, they are not expected to be highly automatable and more workers will be needed to cover the shortages.[39] Official estimates suggest that the NHS will need 5,000 extra nurses every year - three times the figure it currently recruits annually. The net employment effect estimates of our model show that nurses constitute the occupation that will be growing the most in absolute terms since the effect is positive and large (Figure 4-1). Medical practitioners are also at the top of the list. Figure A4-4 in Annex 4 presents the Top 20 'growing' occupations, as determined by the magnitude of the impact of AI on net employment in 20 years, and the skills that they require.[40]

Figure 4-1: Net employment effect of AI over 20 years for skilled occupations in shortage

[39] Not all occupations with growing projected demand are part of the list of occupations in shortage. The net employment effect for sales and retail assistants is positive and large, but the main skills that are needed for such jobs are to perform tasks related to basic customer service or merchandising.
[40] The most frequent skills demanded from medical practitioners, for instance, include: surgery; mental and behavioural health specialties; teaching; emergency and intensive care; and general medicine (according to Burning Glass Technologies data).

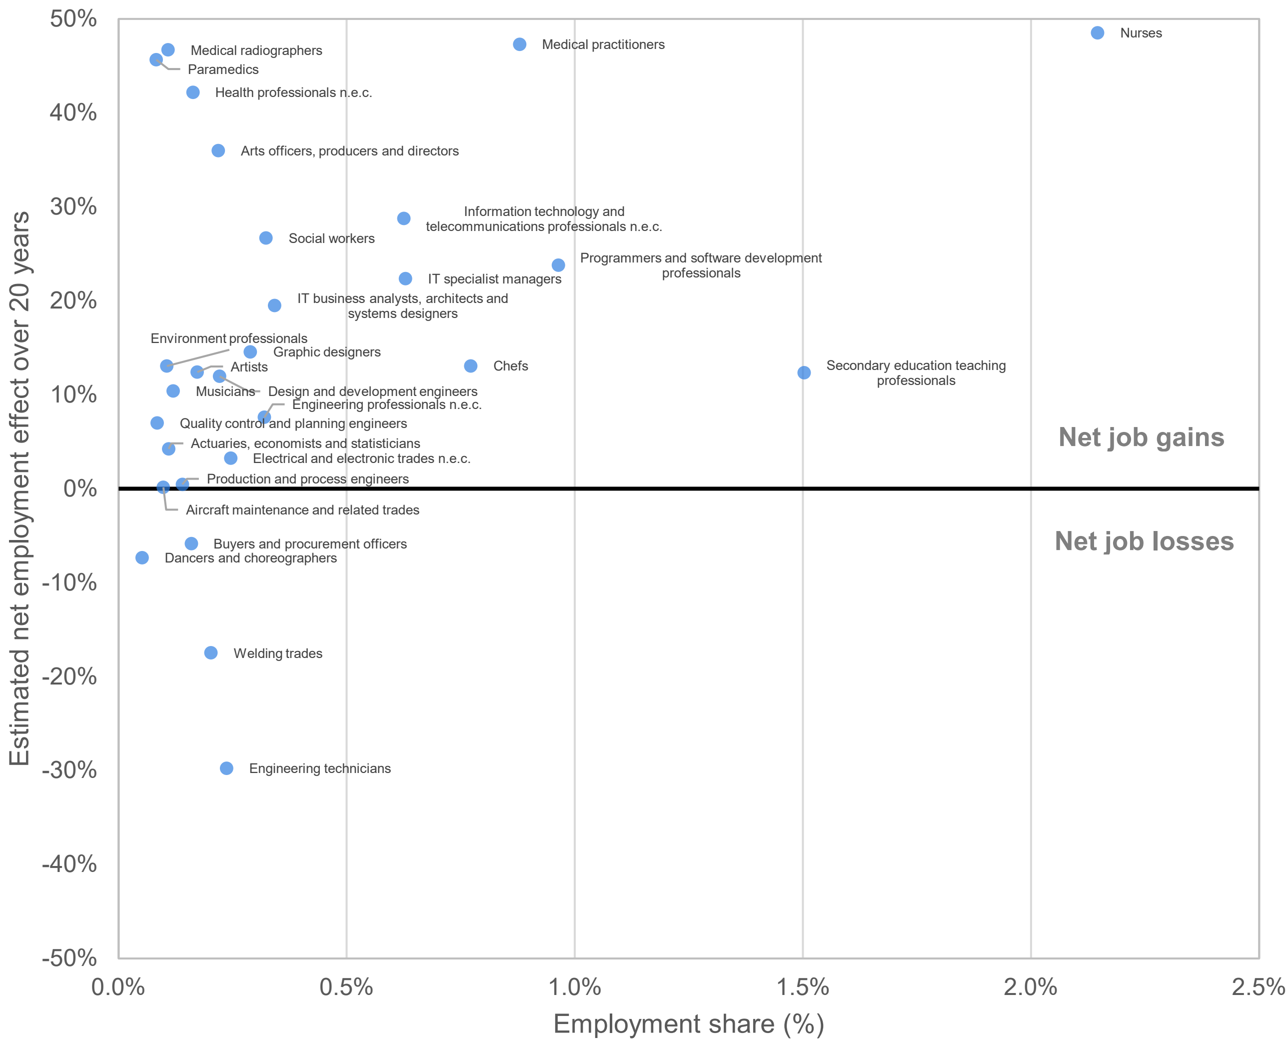

Source: PwC analysis of OECD PIAAC, ONS and BGT data

The work of chief executives, senior managers and other high-skilled professionals working in business strategy or managing people also faces low risk of automation and shortages will continue to exist. Many IT-related occupations featuring in the list of occupations in shortage show a strong net effect in employment in our model, including IT business analysts, architects and systems designers.[41] There are other occupations that, while not being large employers, will experience large growth relative to current levels of employment: notably chefs but also social workers (child development; advanced patient care; mental and behavioural health specialties) and even arts officers, producers and directors with the continued strength of the UK's creative industries (Figure 4-1).[42]

Whether other occupations that show a large positive net employment effect will also become part of the list of occupations in shortage will depend on the skills needed. Some of the growing occupations do not necessarily require specialised skills; the hospitality sector shows large shortages but is not necessarily facing a problem of upskilling. It will also depend on how labour supply evolves, which will be driven by current educational enrolment in relevant fields

[41] Some of these occupations do not appear in Figure A4.4 in the Annex (they are not part of the Top 20) but still show a strong net effect in absolute employment (with a positive net effect of more than 20 thousand jobs).
[42] Film, TV, radio, photography, music, advertising, and digital creative industries are all part of this thriving sector.

and also by immigration - the health care sector is one that relies heavily on professionals from other countries migrating to the UK (12% of the workforce by some estimates).[43]

Are occupational shortages likely to widen in the future with the progress of AI and automation?

The net employment estimates of our model suggest that shortages are likely to continue increasing – particularly if workers losing their jobs to automation do not get reskilling relevant to the growing occupations.

Of the occupations that are currently in shortage in the UK, the majority are not highly automatable (see Figure 4.2, the highlighted bubbles). If anything, AI and related technologies could 'augment' the work of those occupations, rather than replacing workers. Even in 20 years, the displacement effect is expected to be moderate - compared to the whole universe of occupations, those occupations in shortage largely see a positive net effect of AI on jobs. There are only very few occupations featuring on the UK Shortage Occupation List that are estimated to see a negative net impact of AI (i.e. more jobs are displaced than created). In those cases, AI and automation could alleviate shortages.

For some occupations, the estimated net effect is positive mainly because job displacement is expected to be low, for instance: secondary education teaching professionals. By contrast, the demand for health professionals is driven by a combination of low job displacement and large job creation according to our estimates (see Figure 4-2). For the latter, it will prove more difficult to close the existing skills gaps.

Developing skills strategies to cover such occupational shortages is likely to be a long-term endeavour if it largely relies on domestic workers. Given that the skills required by occupations in shortage might be substantially different from the skills that the current UK workforce has, the time it takes to close skills gaps through training of local people may be increasing.[44] Many workers would need to pursue what would basically be a 'second career' - which is not easy since until recently workers have tended to transition to jobs that are part of the same occupational family.[45]

Figure 4-2: Estimated job displacement and creation effects of AI by occupation (occupations in shortage in colour) over 10 and 20 years

a) *10 years*

[43] Source: https://www.ons.gov.uk/peoplepopulationandcommunity/populationandmigration/internationalmigration/articles/internationalmigrationandthehealthcareworkforce/2019-08-15

[44] Institutions like the World Economic Forum (WEF) and the OECD continue developing tools to measure the 'distance' between the skills required by jobs and occupations, often by creating clusters of skills that are in the same 'neighbourhood'. NESTA in the UK has created a whole typology of skills based on their similarities.

[45] See OECD (2019).

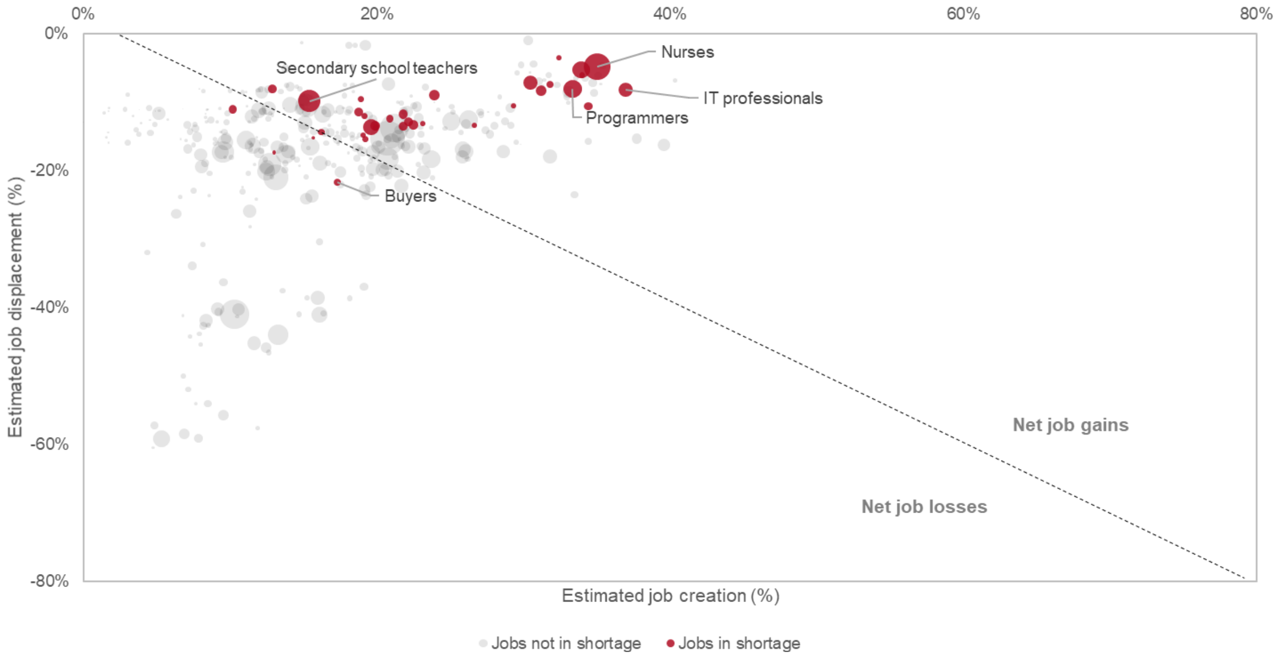

b) 20 years

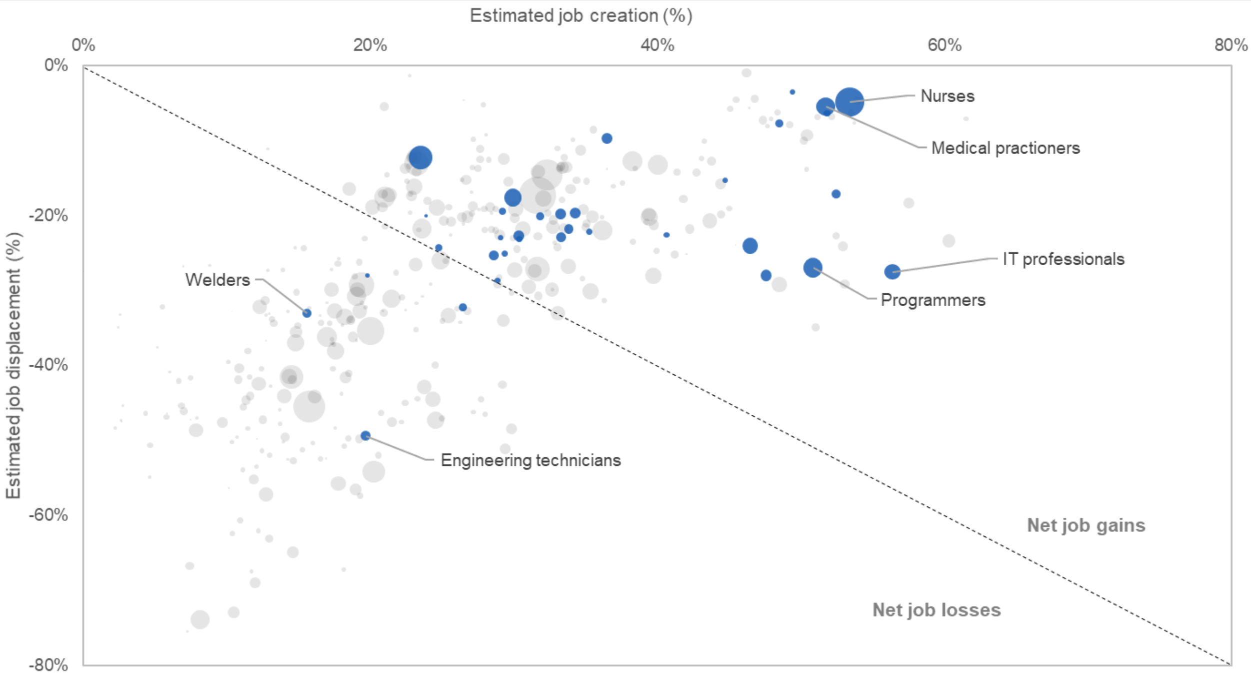

Source: PwC analysis of OECD PIAAC, ONS and BGT data

A certain degree of misalignment of skills is normal over the economic cycle as workers and firms adjust to the evolving demands of the economy. Yet, if skills gaps and mismatches become persistent in the context of technological change, they can limit the ability of firms to innovate and adopt new technologies, while also impeding the reallocation of labour from less to more productive activities.[46] So what are the skills that will be most in demand in the future?

[46] See the evidence presented by Haskel & Martin (1993a, 1993b); Allen and van der Velden (2001); McGowan and Andrews (2015b).

Estimating the future demand for skills

Our model is able to estimate the potential increase (or decrease) in the demand for different skills by combining the net employment changes by occupation derived from the jobs displacement and job creation estimates (in absolute numbers) with data that identifies the skills that are used the most by every occupation. Similar to the net employment analysis of Section 3, estimations are obtained for a timeframe of 5, 10 and 20 years. We use data on job ads from Burning Glass Technologies (BGT)[47] to obtain more up-to-date information about the skills required by different occupations. The methodology we have used to estimate the future demand for skills is described in detail in Technical Annex 4.

Figure 4-3: Overview of BGT job ads data

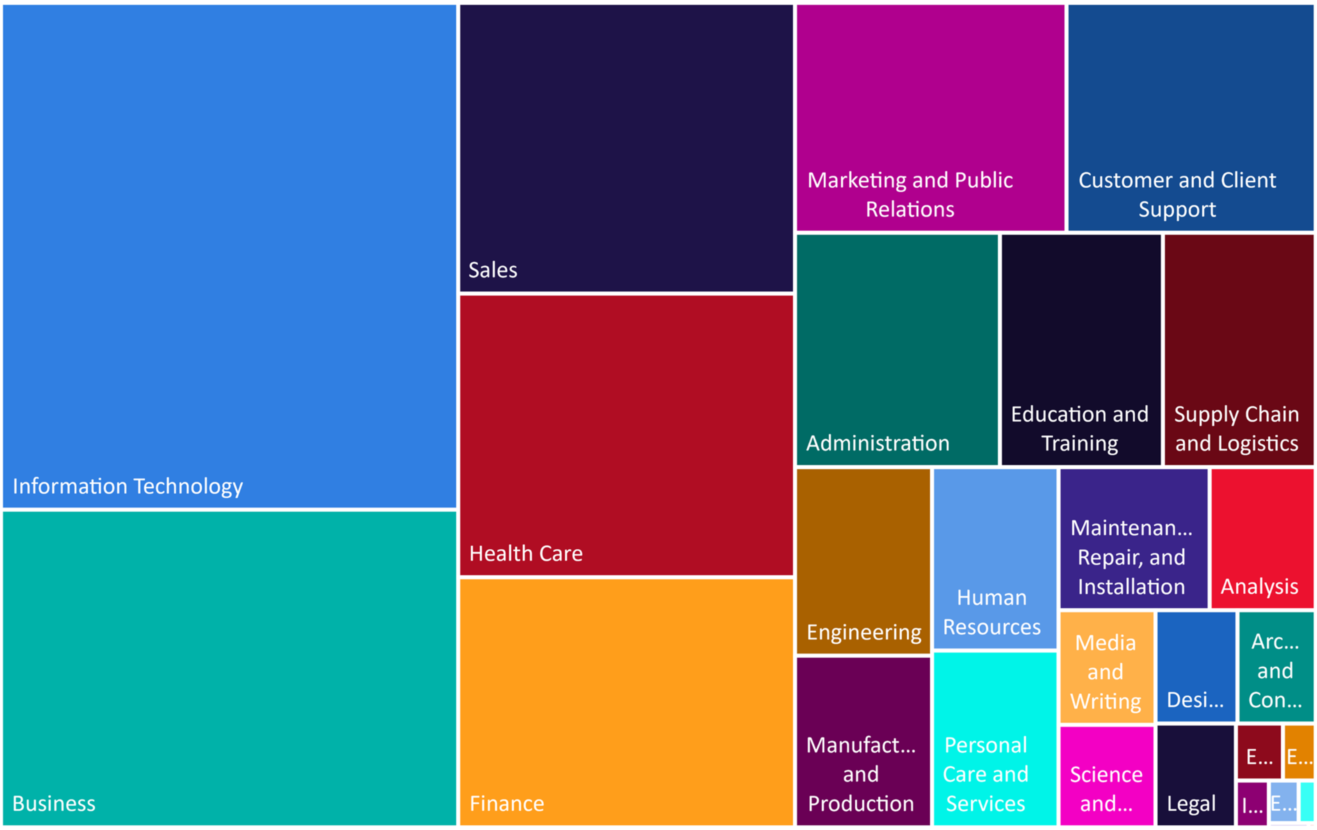

Source: PwC analysis of BGT data

The more disaggregated taxonomy of skills is what BGT calls 'skills clusters' and there is data that shows how frequently such skills are mentioned in job ads targeting specific occupations.

[47] The Burning Glass Technologies (BGT) database is described in Technical Annex 4. BGT collects data on active job postings from thousands of web-pages on a daily basis. For each job posting, in addition to extracting job title, salary, education and experience requirements, Burning Glass identifies keywords from free text job descriptions. The full job descriptions are not available. Keywords include: skills, personal competences and knowledge required by employers.

The majority of skills captured by such skills clusters refer to what the database calls specialised skills. Examples of skills clusters include: 'business intelligence', 'data analysis', 'graphic and visual design' or even broader domain specialisations such as 'cardiology'.[48] BGT then has a more aggregated classification of skills. It groups the more than 500 skills clusters into 28 groups. The groups that are in most demand in the UK today are 'Business', 'Information Technology', 'Health Care' and 'Sales' (Figure 4-3).

The future demand for skills at the level of the economy is obtained through a skills ranking that takes into account how often the different skills clusters are used across all occupations (i.e. the proportion of times that the skill is found in job ads). Figure 4-4 shows the estimated change in demand for the 28 skills groups defined in the BGT database. It shows the estimated net effect in 10 years (i.e. taking into account the net effects of job displacement and job creation over a 10-year horizon) and in 20 years.[49]

Figure 4-4: Net employment effect on demand for skills (aggregated at skills group level)

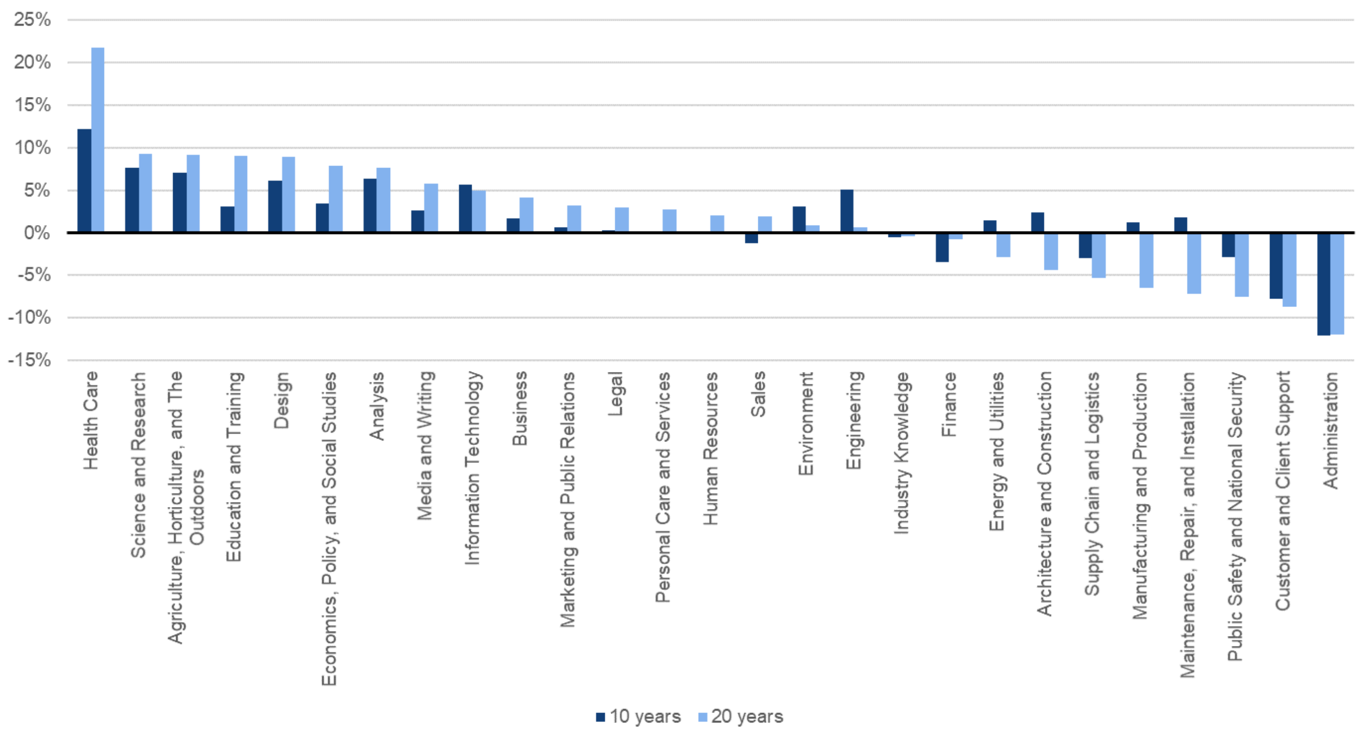

Source: PwC analysis of OECD PIAAC, ONS and BGT data

The main finding is that Healthcare skills are projected to see the largest increase in demand over both 10 and 20 years, followed by Science and Research skills. Conversely, Administration skills, which we might expect to be relatively automatable, are projected to see a fall in demand over 10 and 20 years.

Different skills in a given sector can show a wide dispersion in net demand. Within each of the 28 skills groups shown in Figure 4-4 there are several skills clusters, some in more demand

[48] There are some transversal skills captured such as communication or collaboration skills that the database calls baseline skills, though these are not presented by occupation.
[49] The final impact on skills is called the 'net effect' because it is based on the difference between the job displacement and creation estimations of sections 2 and 3.

than others: for instance, within the Business skills group, the skill cluster ‘project management’ will see a larger positive net effect than ‘customer and client support’ - the latter likely to start declining sometime in the future. Figure 4-5 shows the dispersion of the estimated net effect across the skills clusters contained in each skills group.

Figure 4-5: Dispersion of net effects for skills clusters within each skills group (over 20 years)[50]

[50] Few data points for skills clusters were removed from the figure as they were considered as outliers (i.e. with net effect values that are substantially different to the rest of skill clusters).

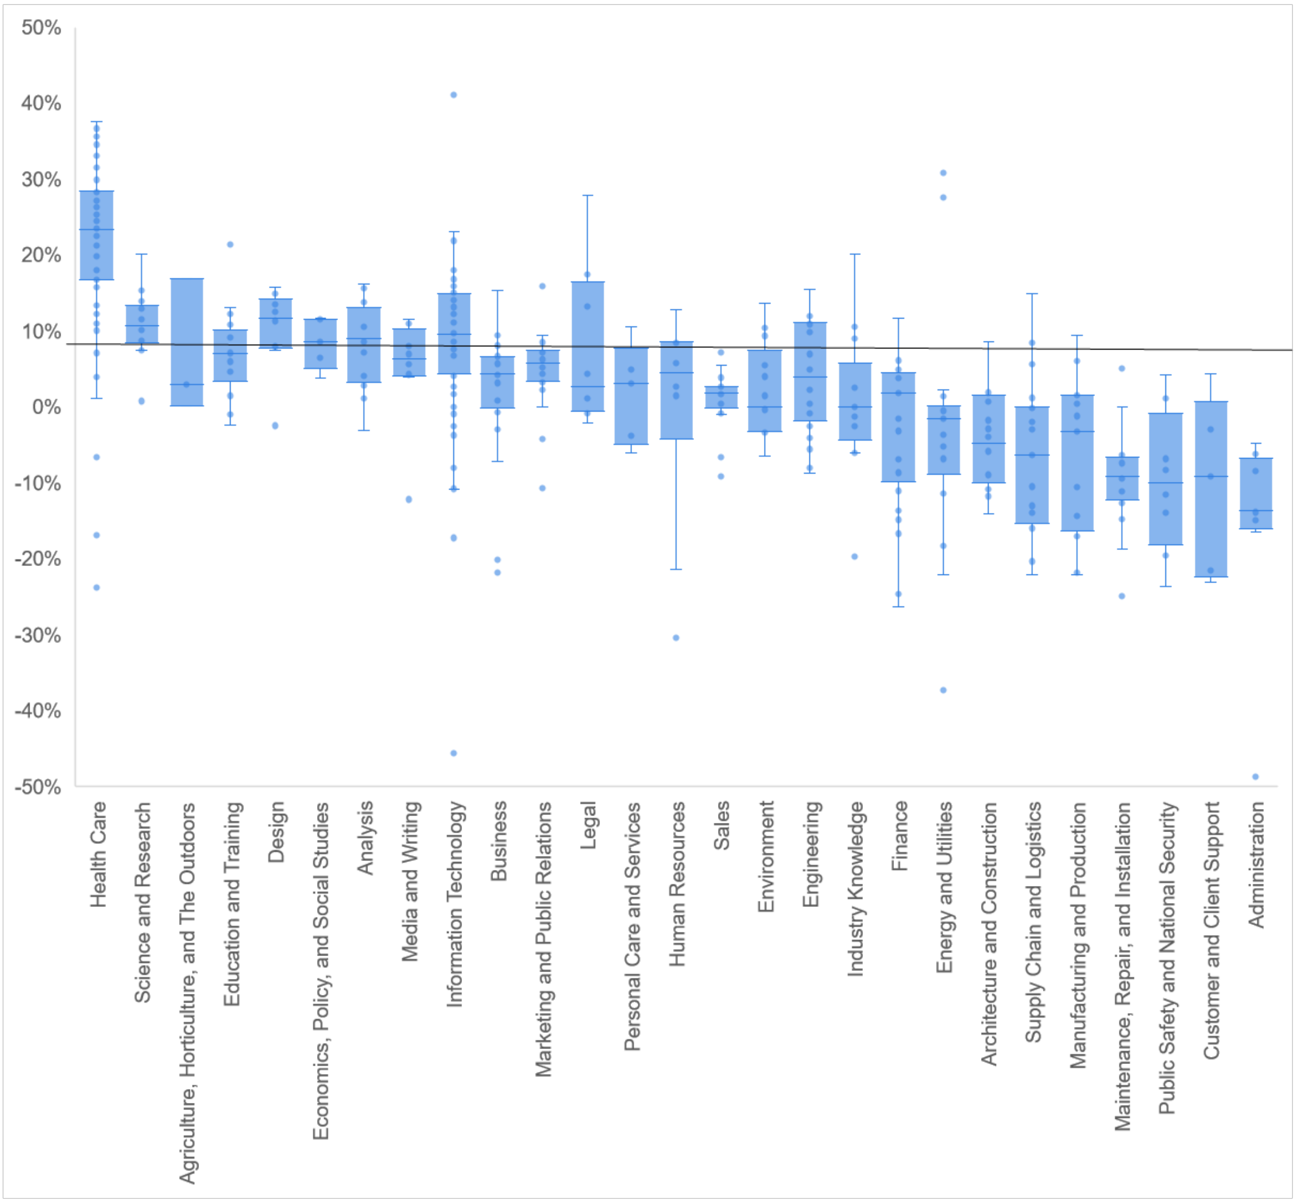

Source: PwC analysis of OECD PIAAC, ONS and BGT data

For the Healthcare skills group, the dispersion of effects (the size of the box plot in Figure 4-5) is not as large as for other groups – which means that the net effect is somewhat homogeneous across all the skills clusters (from medical support skills to oncology specialisation). Almost all skills clusters (even the one with the 'minimum' value in the graph) are expected to see an increase in demand.[51] The impact on skills in the Administration group (e.g. of the clerical type) are also homogeneous (low dispersion) but running in the opposite direction: demand is largely expected to decrease over time.

There are other skills groups that show *wide dispersion*. Skills clusters in the Business group are on average expected to increase, but there are some skills clusters that will decline substantially (as indicated by the 'minimum' value) such as technical assistance or order management. Skills clusters related to the Customer and Client Support group are expected to decline on average, but there is large heterogeneity as you might expect given that some such roles may still require the 'human touch', particularly when dealing with non-routine cases that

[51] Idem.

may remain difficult to automate. But the number of staff needed to deal with such cases will be much less than at present in this area, so reducing the overall demand for skills in this field, but potentially increasing the incomes of those still working in the field since they will be playing a more value-adding role in tackling difficult cases than the average customer support worker today.

The net effect in the demand for skills for all these skills groups is a combination of the separate impacts from job displacement and job creation due to AI. Figure 4-6 decomposes for the 28 skills groups the net effect into these two elements. Healthcare skills already represent a large fraction of employment and will across the board continue increasing. For the skills groups Analysis and Information Technology, the net effect is positive given a large job creation effect but they also will see a high degree of job automation (i.e. some of the skills clusters contained within each of them may decline, for example as some forms of specific IT expertise become redundant or automatable over time). In another example, Science and Research and Education and Training show a similarly positive estimated net skills demand effect in 20 years. Yet, the former shows more job displacement but also more job creation - while the latter is more stable (with low disruption as automation is estimated to be relatively low given the labour-intensive nature of those jobs and job creation relatively moderate).

Figure 4-6: Decomposition of job displacement and creation effect by skills group

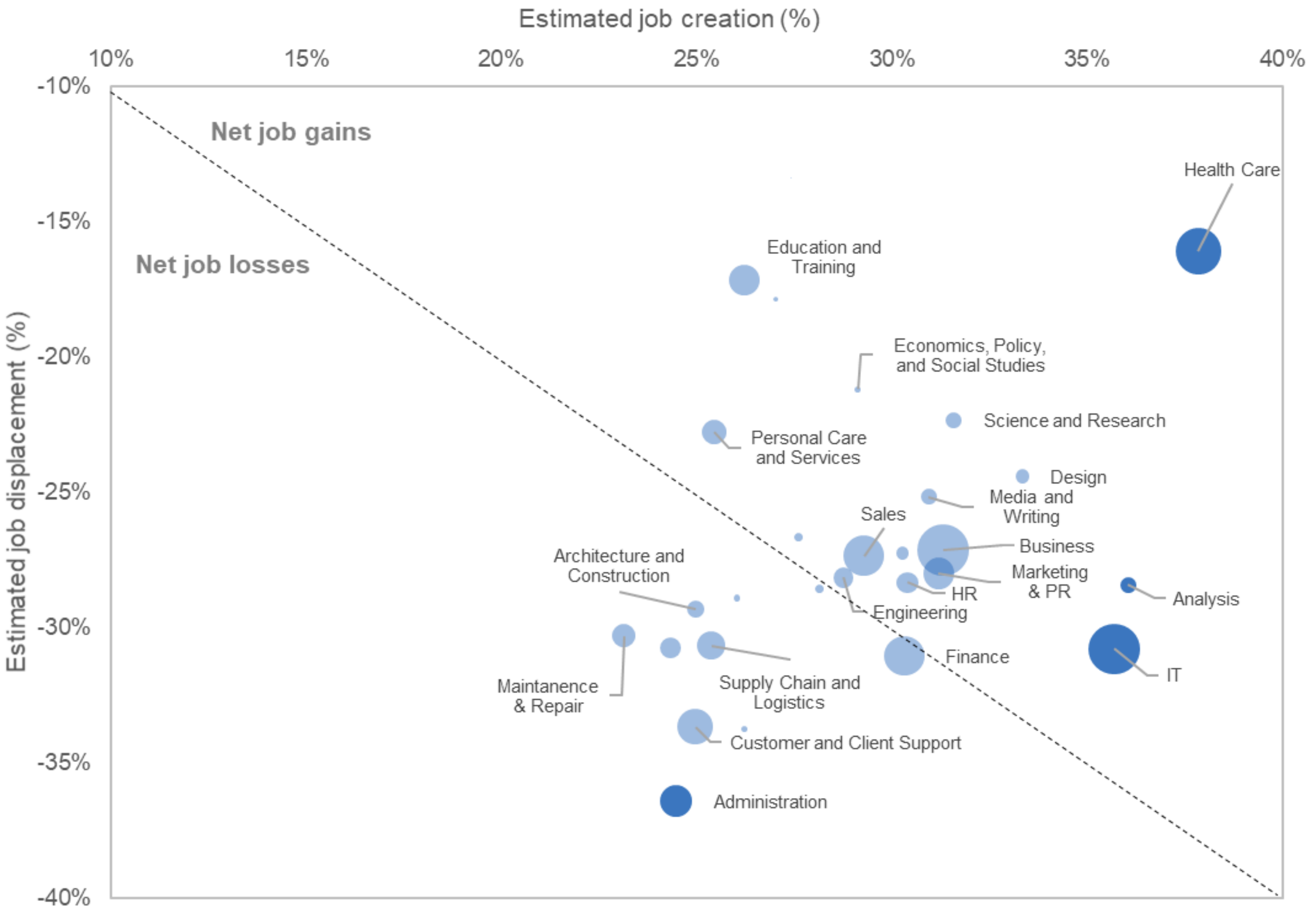

Source: PwC analysis of OECD PIAAC, ONS and BGT data

To better understand the situation of each skills group, we can look separately at estimated job displacement and job creation effects for each of the skills clusters contained within them. Figure 4-7 shows an example of this analysis for the Healthcare skills group. Demand is projected to be high for the majority of skills clusters within this group - and some will face a low risk of automatability such as injury treatment or paediatrics. There are only a few skills clusters in this group that are projected to see a decline in demand due to automation - e.g. managing ‘medical records’ or ‘health information’.

Figure 4-7: Decomposition of job displacement and creation effect by skill clusters within the ‘health care’ skills group (over 20 years)

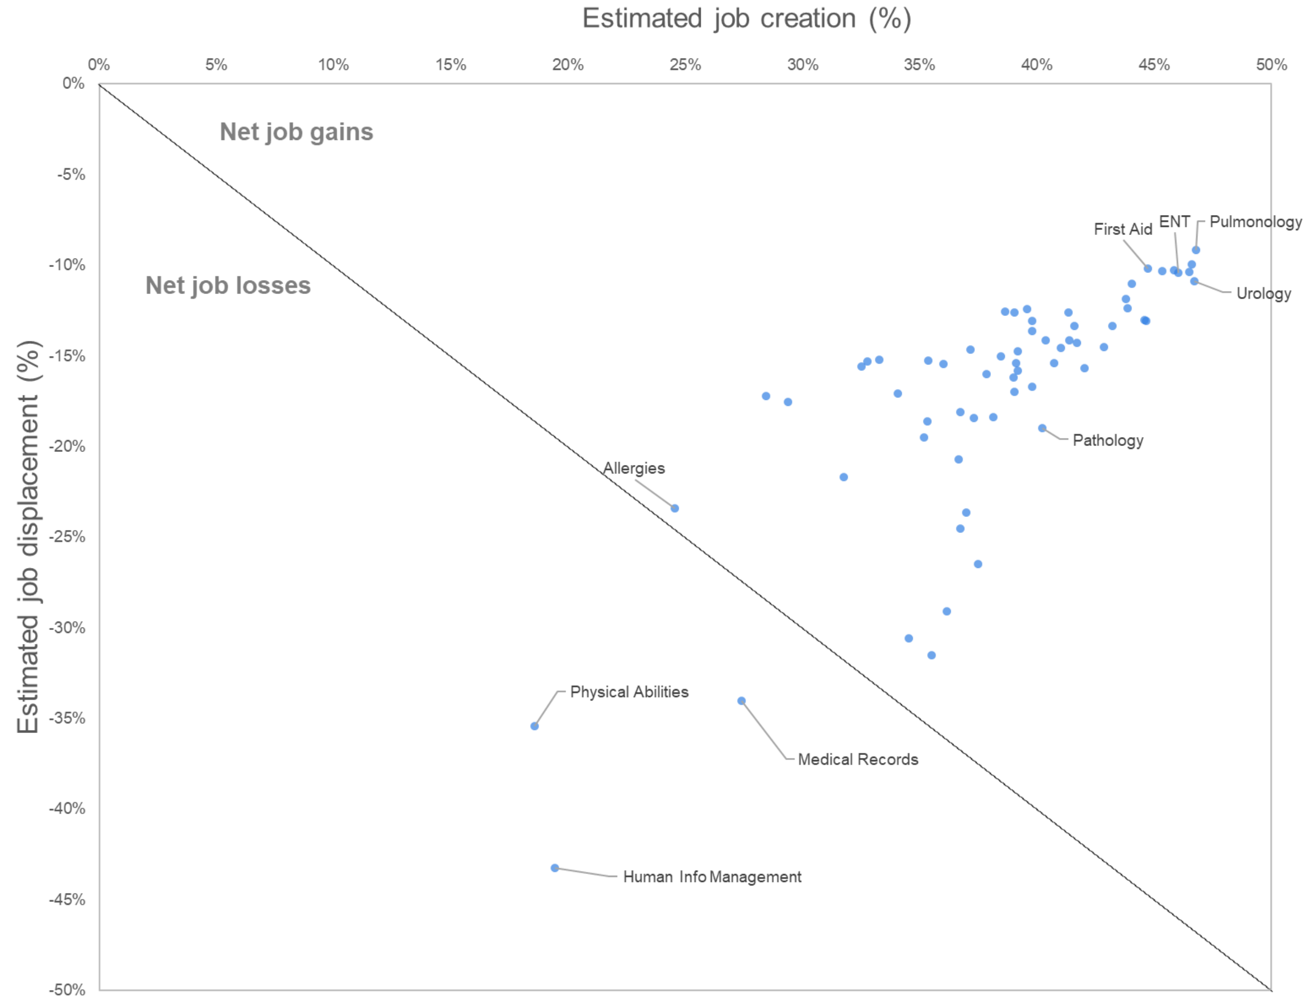

Source: PwC analysis of OECD PIAAC, ONS and BGT data

The picture is more heterogeneous for ‘Business’ skills. Skills such as order management are expected to decline given the expected job displacement of occupations using such skills. By contrast, the occupations using business consulting or knowledge management face lower degrees of automatability - while job creation is expected to remain high. People management is also expected to remain a highly demand skill in the future - this trend is consistent with an increasing number of case studies that show the value of these people management skills for the modern economy (as discussed below).

For the IT skills group, the estimated net demand effect is also largely positive across the different skills clusters contained within this broader group. Many IT-related skills have been trending between 2012 and 2018 according to all BGT job ads (Table 4-1). These trends are corroborated by other sources. In the last decade, many new occupational titles in the UK have been clustered around using new technologies (e.g. web developer or database administrator).[52] According to a 2019 PwC global upskilling survey, 43 percent of UK respondents are planning to learn new skills to better understand or use technology.

Table 4-1: Fast growing skills clusters

Rank	Skill cluster group	Skill Cluster	Rise in Demand (2012-18 %)
1	Information technology	Scripting Languages	283%
2	Health care	Obstetrics And Gynecology (OBGYN)	244%
3	Information technology	Version Control	236%
4	Information technology	Cloud Solutions	183%
5	Marketing and public relations	Social Media	173%
6	Health care	Advanced Patient Care	161%
7	Health care	Emergency And Intensive Care	159%
8	Health care	General Medicine	152%
9	Health care	Surgery	147%
10	Health care	Mental And Behavioral Health Specialties	142%
11	Health care	Mental Health Diseases And Disorders	140%
12	Business	Key Performance Indicators	137%
13	Health care	Social Work	136%
14	Information technology	Software Quality Assurance	136%
15	Human resources	Employee Training	131%
16	Personal care and services	Child Care	123%
17	Health care	Medical Support	120%
18	Customer and client support	Advanced Customer Service	120%
19	Health care	Basic Patient Care	119%
20	Maintenance, repair, and installation	Vehicle Repair And Maintenance	111%

Source: PwC analysis of BGT online job advert data

Despite the positive trend in demand in general, IT skills are quite diverse and many skills clusters will face large job displacement (Figure 4-8 decomposes the estimated job creation and job displacement for each skills cluster within the IT skills group). For instance, many occupations using Oracle[53] are expected to decline, though this is more than compensated by other occupations using Oracle seeing rising demand. Some IT skills are clearly expected to decline over the years such as document management analysis. Skills related to some

[52] New jobs are being created to exploit the new human-machine symbiosis. For instance, according to BGT data, advertised positions in the UK involving scripting languages have increased 283 percent between 2012 and 2018.
[53] Oracle is a proprietary multi-model database management system produced and marketed by Oracle Corporation. It is commonly used for running online transaction processing, data warehousing and mixed database workloads.

programming languages and those related to augmented reality and artificial intelligence are largely expected to increase over time - with very low rates of automatability.

Figure 4-8: Decomposition of job displacement and creation effect by skill clusters in 'IT' skills group (in 20 years)

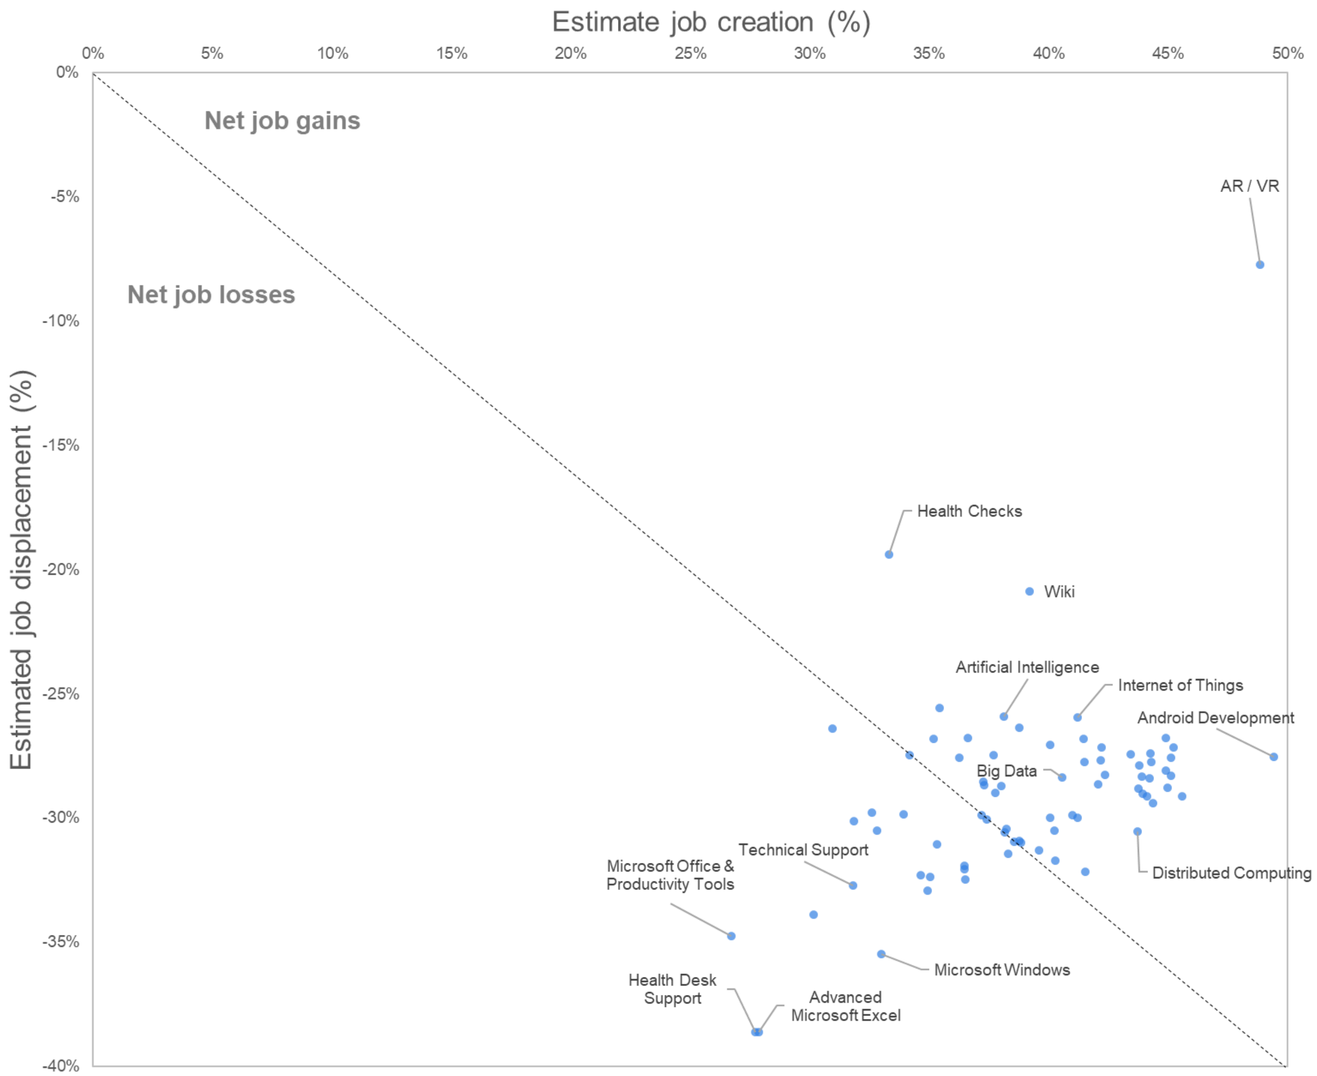

Source: PwC analysis of OECD PIAAC, ONS and BGT data

Demand for AI skills

The rise of AI will demand employees with a certain set of skills:

- AI specialist skills: e.g. AI-related fundamental research, engineering and applications, as well as data science and computational thinking (see examples in table below)
- Skills needed to leverage AI and collaborate with machines (including through AI-human teams on the factory floor and quality control).

Figure 4-9: The demand for AI-related skills

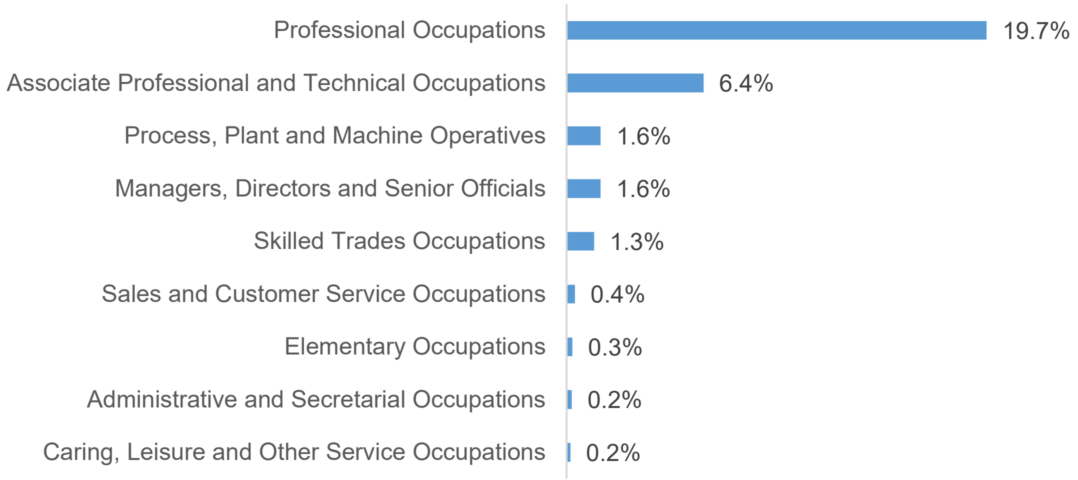

Source: PwC analysis of BGT data

Currently, it is mainly professional occupations that use AI skills (Figure 4-9).[54] In terms of sectors, demand has been the strongest in the ICT industry – the sector of Finance also stands out as having relatively high demand for AI skills, which seems plausible given their value in financial trading in particular. While the trends in Figure 4-10 are somewhat variable by sector, the overall trend is for increasing demand for AI skills across the economy and this only seems likely to increase further based on our wider analysis in this study.

Figure 4-10: Evolution of AI-related skills demand by sector

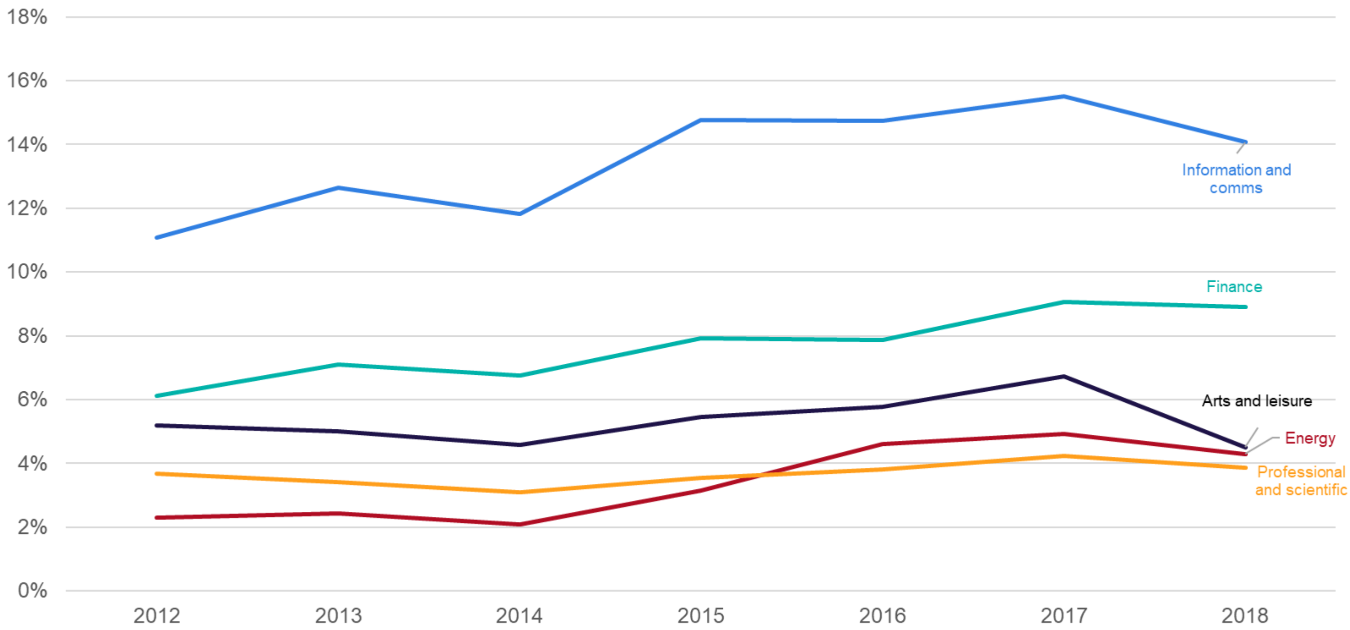

Source: PwC analysis of BGT data

[54] Figure 4-9 shows the demand for AI-related skills using a broad definition of IT-related skills clusters (e.g. scripting languages or distributed computing), analysis-related skills clusters (e.g. data science or natural language processing), and engineering/robotics.

In response to this rising demand, the UK government has a clear objective to boost AI-specific skills (HM Government, 2019) The initiative and programmes in this area of skills are diverse: from publicly-sponsored PhD places[55] to a global Turing Fellowship programme to both attract and retain the best research talent in AI from around the world to the UK. The private sector is also taking actions to support AI skills, including working with universities (e.g. to develop industry-funded Masters programmes in AI at leading universities across the UK) and working with government and universities to assess the potential role for new Masters conversion courses in AI-related expertise, for skilled graduates in other disciplines.

Possible future evolution of the task composition of jobs

As mentioned at the start of this section, an important caveat to the skills demand estimates presented above is that they are only really capturing the 'compositional' effect: i.e. the changes in skills directly coming from projected changes in the occupational mix in the economy due to AI and related technologies (as derived from the estimates of Section 3). Yet, the task composition of jobs and occupations is likely to change significantly – as has been the case in past waves of technological disruption.[56]

Current quantitative work on AI and automation, including estimates in this study, still has limited capacity to project whether and how the skills required by each individual occupation will evolve in the future (e.g. it is still uncertain if the tasks performed by nurses in 20 years' time will look substantially different to what they do now). With the BGT database, we are able to build an occupation/skills matrix that gives a snapshot for a given period of time. We use data on the latest year available (which is 2018) to estimate the demand for skills. The database actually also has data on job ads starting from 2012. We were able to produce a similar matrix using historical data, but no clear patterns or trends emerged when comparing against the main 2018 matrix and we have concerns about the robustness of any such analysis. This means that making predictions about how the 'skill/task composition' of occupations will evolve into the future is difficult using this kind of data.

More generally, the literature exploring the impacts of AI and other automating technologies on the task- and skill- composition of jobs is still in its infancy. Estimates of the task effect have been very limited in the literature because there is currently no micro data looking at how specific AI applications substitute or complement work - thus changing the task-composition of occupations. In fact, until recently, empirical work had been limited in its ability to study directly whether skill requirements in the workplace have been rising, and whether these changes have been related to technological change.[57] Since the influential work by Autor, Levy and Murnane (2003), a task-based approach to skills has gained track, which has made it possible to analyse occupational skill requirements directly by conceptualising work as a series of tasks (the OECD PIAAC survey used in the present study has been constructed following these principles, by asking workers what types of tasks they perform at work). Yet, projections into the future remain limited.

[55] The AI-relevant studentships will be distributed via the current EPSRC call for Centres in Doctoral Training, which will be supported with £100m investment from government.

[56] The history of automation and technological change in the 19th and 20th centuries has been one of task (re)generation, whereby the task content of production has been expanded as a result of new or a broader range of tasks emerging. See Acemoglu and Restrepo, 2019.

[57] Early studies in the area of skill-biased technological change used "traditional" skill measures to assess the skill level of employees, such as the proportion of production workers/non-production workers or blue-collar/white-collar workers. But these classifications use divisions according to occupational groups that are of limited use in determining skill requirements and how they evolve.

How might the task composition of jobs evolve in the future with the increasing reach of AI and related forms of automation?

Given the usual taxonomy classifying tasks into cognitive vs. manual and routine vs. non-routine, a clear trend is that occupations have become more complex since 1970. There has been a sharp increase in non-routine cognitive tasks such as doing research, planning or selling, and a pronounced decline in manual and cognitive routine tasks such as double-entry bookkeeping and machine feeding. The changes in skill requirements have been most pronounced in rapidly computerising occupations (Spitz-Oener, 2006). Changes in occupational content have been estimated to have accounted for around half of the recent 'educational upgrading' in employment.[58] But what next?

Using granular data for tasks, the EU's CEDEFOP[59] specialists have presented evidence of further increases in the complexity of human work with the use of the most recent automating technologies. The increasing sophistication of jobs will surely demand more skills and education (consistent with our broad findings in Section 3 above). Studies also suggest that workers with higher education and specialised skills will need constant retraining.[60]

Recent research[61] linking advances in different categories of AI to different types of human abilities will help to understand how tasks for different types of jobs and occupations may evolve in the future. Open source projects are starting to track the rapid progress of AI in performing tasks at human-like levels of capability in domains including voice recognition, translation, visual image recognition and others.[62] Obtaining a satisfactory understanding of the manner in which technological progress affects labour and skill demand, and its impact on productivity growth, is dependent on whether the 'reengineering' of tasks acts as a countervailing force to the displacement effect.

The rise of AI will demand employees with a certain set of skills that will enable them to have basic knowledge of AI's value and its relevance to big data. In addition to AI specialist skills (e.g. engineering and applications) and AI-savvy functional specialists, many more people will need the skills to work alongside AI (the human-machine symbiosis). As many as 120 million workers in the world's 12 largest economies may need to be retrained or reskilled as a result of Artificial Intelligence (AI) just in the next three years.[63] Not all of the skills required to apply AI are technical; with the automation of more tasks, people will stop doing 'transactional' tasks and concentrate on areas where humans are expected to have an advantage. These 'interaction' tasks involve working in diverse, inter-disciplinary teams, with strong communication and empathy skills and problem-solving capabilities.[64] People will need to work better with each other, since AI's need for data and human expertise crosses functions and business lines.

[58] Goldin & Katz (1996, 1998) provide a historical perspective on the link between technology and skill demand.
[59] CEDEFOP refers to the European Centre for the Development of Vocational Training.
[60] See McGuinness et al. (2019).
[61] See, for instance, recent work by NYU's Robert Seamans linking advances in different categories of AI to different types of human abilities.
[62] See, for instance, the Electronic Frontier Foundation (EFF) platform - a collaborative pilot project, which tracks progress on task-specific AI performance metrics across a variety of separate artificial intelligence categories, such as abstract strategy games and image recognition.
[63] See IBM (2019).
[64] The concepts of 'transactional' vs 'interaction' tasks have been used widely by the McKinsey Global Institute in recent studies.

Work practices and transversal skills

The combination of the 'compositional' and 'task' effects will continue reshaping the demand for specialised skills - those related to specific occupations and/or domain knowledge (such as geriatrics or supply chain management). Furthermore, people will also need to develop the so-called transversal skills that are 'durable', such as adaptation and resilience, to navigate a future of work that is expected to be constantly in flux (MIT Media Lab, 2019). People are living longer and are less likely to stay in the same job throughout their working lives. PwC's own research into the Workforce of the Future found that 60% of survey respondents think that "few people will have stable, long-term employment in the future." These skills are important to adapt to more regular transitions to new jobs and also to the rise of self-employment, including in the technology-powered 'gig economy'.[65]

In short, instead of simply promoting the skills needed today, many experts argue instead for promoting resilience and adaptation ('learning how to learn') because the process of upskilling is expected to become a more regular need for the majority of people.[66] Digital skills are certainly transversal and need to be widely developed across society. Almost all workers, of any occupation or industry, will need to increase their understanding and digital awareness to function successfully. Yet, these skills are not necessarily durable: learning a given software normally has a short shelf life – indeed the average lifespan of a tech skill is roughly 18 months according to one recent estimate.[67]

Transversal skills also include those soft skills that are relevant for more collaborative environments. Skills such as empathy, building relationships, and collaboration provide the tools for the rich and versatile coordination which underpins a productive workplace in the modern economies of developed countries[68] – the subtleties of which computers have yet to master (machines are generally poor at sensing the situation around them). Research shows that many organisations in developed countries are rapidly changing attitudes and increasingly valuing such skills.[69] Communications skills usually come at the top of the list (Figure 4-11 shows the most demanded 'baseline' skills in the job ads tracked by BGT database).

[65] According to the Resolution Foundation, two-thirds of the growth in employment since 2008 has been in 'atypical' roles such as self-employment and agency work (Clarke and Cominetti, 2019). Research commissioned by BEIS estimates that 4.4% of the UK population worked in the gig economy in 2018 (GOV.UK, 2018) and according to research from the Trade Union Congress (2019), the size of the gig economy has more than doubled in the three years to 2019.
[66] See Maree, 2018.
[67] See Fisher, 2019.
[68] Evidence for developed countries shows that cognitive ability accounts for only a small portion of the variation in earnings of workers (see, for instance, Heckman and Kautz, 2012).
[69] The IBV study, "The Enterprise Guide to Closing the Skills Gap", found that in 2016, executives ranked technical core capabilities for STEM and basic computer and software/application skills as the top two most critical skills for employees. Two years later, however, the top two skills sought were transversal skills - willingness to be flexible, agile, and adaptable to change and time management skills and ability to prioritise. This matters for organisations in complex environments where the classic gains from specialisation are eclipsed by the need to adapt flexibly to changing circumstances (see NESTA / Oxford Martin School 2019).

Figure 4-11: Top 'baseline' skills featured in BGT

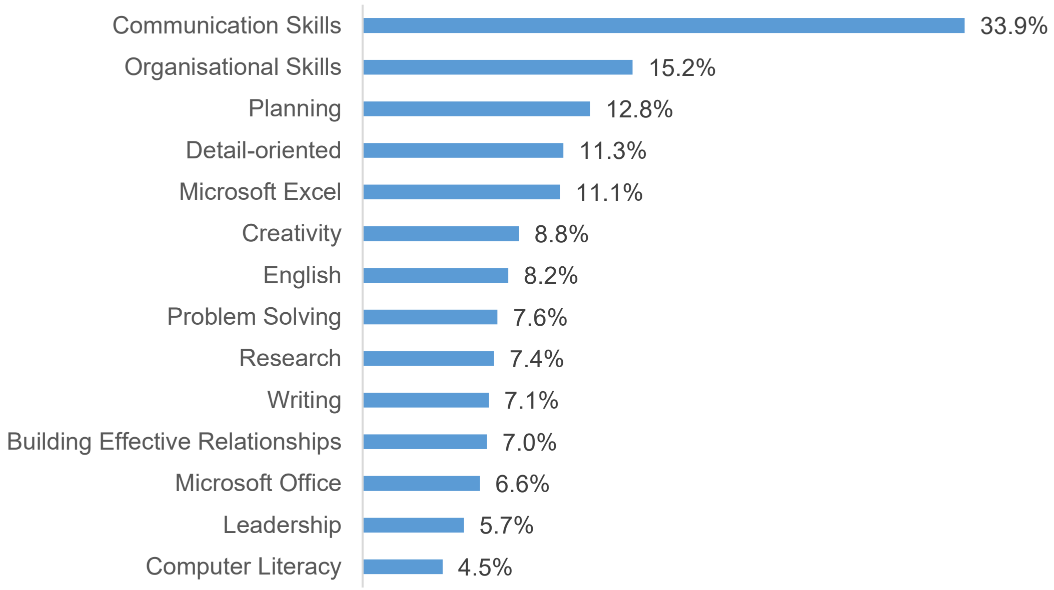

Source: PwC analysis of BGT data

Economist Lawrence Katz memorably titles workers who virtuously combine technical and interpersonal tasks as "the new artisans".[70] These are the types of workers that are likely to thrive in an economy increasingly characterised by open innovation, creative work and diverse agile teams.[71] Many authors suggest avoiding the myopia trap of excessive focus on competence and domain knowledge - rather than the vision, external networking and setting organisational direction that characterise successful leadership.[72] Excessively narrow specialisation may handicap innovation and leadership, which in today's interdependent world requires a global network and mindset.[73]

Summary of skills analysis

Our analysis indicates how the demand for skills may evolve over the next 5, 10 and 20 years due to occupational mix changes associated with the impact of AI. We find that demand for skills groups such as health care and business skills are likely to increase, while demand for other skills groups such as administration decline.

Whilst some of the job creation will be in areas linked directly to AI and related technologies (e.g. data scientists, robotic engineers or people involved in the design and manufacture of sensors for driverless vehicles and drones), most will not. As was discussed also in the expert workshop convened as part of this study, tasks like managing people and relationships will continue to be in high demand. Furthermore, jobs will be created to provide relatively hard-to-automate services that are in greater demand due to the additional real incomes arising from

[70] See Autor (2015).
[71] See Deming (2015) and (OECD, 2001)
[72] See, for instance, Pucik et al. (2017), who also summarise the research on leadership transitions by Herminia Ibarra from INSEAD and London Business School.
[73] See Hanushek et al. (2016) for a discussion on general vs vocational skills.

higher productivity generated by AI.[74] These include health care services but also other personal services and the artistic and creative industries.[75]

The challenge is that many of the skills likely to see high future demand growth are already in shortage: we find that many of the occupations in the official list of shortages will see a net increase in demand according to our estimates of net employment effects in 10 and 20 years. This points to the need for determined action on training and retraining of current and future workers to address these potentially growing skills gaps.

[74] This is referred to as the 'Baumol' effect in the academic literature.
[75] The number of people working in the areas of sport and fitness has increased by 120 per cent to 189,000. The same applies to the artistic and literary professions, which have doubled and now account for more than 400,000 jobs (Littlewood, 2012).

5. Conclusions

Past research by PwC and others[76] suggests that AI and related technologies such as robots, drones and driverless vehicles should be significantly positive for productivity and real income levels. Net effects on employment are unclear, with the most plausible assumption[77] in our judgement being for a broadly neutral long-term employment effect as we assume in this study. But we would expect significant structural shifts in UK employment patterns due to AI and related technologies over the next two decades, as in past waves of technological advance.

How does this study add value to previous research?

The present study therefore focuses on estimating the potential distributional impacts of AI on different occupations, industries and regions and on different socio-economic and demographic groups. In doing this, we sought to build on and extend past studies in a number of ways.

First, we held an expert workshop in July 2019 to update the assessment of the potential automatability of selected occupations that formed the starting point for the analysis of Frey and Osborne (2013) and many subsequent studies that relied, directly or indirectly, on their results. We also sought to identify how these judgements on automatability would vary over three broad time horizons: 5, 10 and 20 years.

Second, we used these updated expert judgements to update the estimates of potential job displacement for the UK we made in previous PwC studies (2017a, 2018b and 2018c). These used a proprietary PwC machine learning-based model that we found to have greater predictive power than the models used in previous studies using the same OECD PIAAC data such as Arntz et al. (2016) and Nedelkoska and Quintini (2018).

Third, we used more detailed ONS data from the APS and ASHE to produce more granular estimates for how job displacement might vary across occupations, industry sectors, regions and socio-economic and demographic groups. To a degree this mirrors the analysis by ONS (2019) for England, but it differs in that the latter was based on the 2016 analysis by Arntz et al, which suggested a relatively low proportion of jobs were at high risk of automation. This carried through to the ONS estimate that only 7.4% of existing jobs face a high probability of automation (>70%). This broadly matches our estimate over a 5 year time horizon (7%), but is significantly lower than our estimates over 10 year (18%) and 20 year (28%) time horizons. We consider that our estimates can therefore provide a longer term set of more granular job displacement estimates for the UK, as compared to what we interpret as shorter term estimates for England only from the ONS.

Fourth, unlike all of the studies mentioned above[78], we have also looked at net employment effects taking into account job creation arising from the increased real incomes arising from increased adoption of AI across the economy. For the reasons discussed earlier in the report and in past research by PwC and others, we assume that aggregate job creation equals aggregate job displacement so as to focus on potential distributional effects. We then developed a novel methodology for allocating the additional jobs created across industry

[76] For example, PwC (2017b and 2017c) and McKinsey (2017).
[77] Based in particular on our earlier UK estimates in PwC (2018c).
[78] Except PwC (2018c), but that adopted a less sophisticated and much less granular approach to estimating job creation from AI.

sectors and occupations. Based on this allocation, we then looked at how different regions, demographic and socio-economic groups might be affected based on variations in their occupational mix at a granular level (SOC 4). To our knowledge, this is the first study that has derived granular estimates of the potential net employment impact of AI for the UK in this way.

Key findings from the analysis

In some respects, the analysis reinforces the findings from previous research by PwC and others. This reflected the fact that our expert workshop, although entirely independent of the original Frey and Osborne exercise in 2013, produced broadly similar assessments of the potential automatability of different occupations in the majority of cases. However, we were able to produce more refined views as to how these effects might build up over time, with relatively low levels of job displacement over the next five years, but larger impacts expected by the 2030s.

More specific findings of note are that:

- Managerial and professional occupations will tend to see positive net employment effects from AI that build up steadily over time.
- Clerical and routine manual roles may see negative net employment effects that start to come through over the next 5-10 years as technologies such as robot process automation roll out across the economy.
- By contrast, negative net employment effects for manual workers such as truck or taxi drivers may only emerge over a longer time horizon (i.e. 20 years) once autonomous vehicles have been introduced at scale across the economy.
- The largest net employment gains over the next 20 years may be seen in the health and social care sector, building up steadily over time, and this could also be an area seeing rising skills shortages. The information and communications sector is also projected to see net gains, but wholesale and retail, finance and public administration may see net losses over the next decade. The transport sector could see significant net employment declines in the longer term looking out to the 2030s.
- Net employment effects will tend to be more positive for higher paid occupations on average, and also more positive for graduates than for lower educated groups.
- Variations in estimated net employment effects by age and gender are smaller than by education, although younger workers in entry level jobs involving routine tasks may be more vulnerable to automation.
- There will be jobs gains and losses from AI in all regions, but net employment effects may be more positive in London and the South East on average, and less positive in more industrial areas in the Midlands and the North of England due to their different occupational mixes. But there will also be considerable variations within regions, reflecting different occupational mixes across towns and cities.

These findings reinforce the view that AI and related technologies could tend to widen existing education/skill-related income inequalities and geographical inequalities over time. These general conclusions proved robust to considering a range of different scenarios for job creation and displacement, even though the precise numerical estimates are subject to many uncertainties.

This suggests the need for a policy agenda that is aimed at maximising the economic and social benefits of AI, but also helping people to adjust to this through appropriate actions by government, business and other stakeholders. It is beyond the scope of this study to explore specific policy options, but the general agenda could include, for example: promotion of digital upskilling by both government and business; other forms of education, training and retraining focused on human skills that are complementary to AI; and making and facilitating additional growth-enabling investment in regions and cities that could otherwise see potential negative employment effects from AI. These are consistent with the broad direction of current government policy building on the Industrial Strategy.

Possible areas for future research

Our analysis has also highlighted a number of possible areas for future data collection and research beyond the scope of the current study.

On data collection, it would be useful to have a more up-to-date version of the OECD PIAAC survey for the UK with a significantly larger sample size, including a regional variable. This would facilitate direct analysis of the potential impact of new technologies on UK workers and jobs by region without needing to map across to other surveys such as the APS.

On research, there is a need to better understand and map the potential career transitions that workers will need to make in future in response to the labour market disruptions linked to AI and other new technologies. This will involve more detailed data collection and analysis on the proximity of different types of jobs as indicated, for example, by the frequency of current worker transitions between them as well as their characteristics in terms of skill requirements (and perhaps geographical location). This can be used as the basis for further modelling as to how many workers will need to develop wholly new skill sets in order to make such transitions, as opposed to making more incremental changes to move to jobs in near proximity to their current occupations (in terms of the skills required). Such analysis would help to guide policy development by highlighting areas where workers may need particular support in retraining for new careers.

Another research angle worthy of further exploration would be to start with expert assessments of a wide range of different tasks and then use this as a basis for assessing the automatability of different occupations. This has been done for the US[79] but not yet, as far as we are aware, in a comprehensive way for the UK. This would enable an assessment of whether this method produces similar or different job displacement estimates to the approach adopted in this and many previous studies in the tradition based on Frey and Osborne, which assesses the automatability of occupations based on their task composition. There may be data limitations on adopting this approach for the UK at present, but conducting an updated version of the UK PIAAC survey with a much larger sample size could help to make this feasible.

Finally, as discussed in Section 4 above, there is a need for more research on how to quantify potential future trends in the skills composition of particular occupations at a granular (SOC4) level, starting with accumulating more and higher quality historical data on these trends. This

[79] For example, Brynjolfsson and Mitchell (2017).

would allow better estimates to be made of how the demand for skills may evolve in future in response to AI and other technological developments.

References

Acemoglu, D. (1998). Why Do New Technologies Complement Skills? Directed Technical Change and Wage Inequality, Quarterly Journal of Economics, Volume 113, Issue 4, November 1998, Pages 1055–1089.

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

Acemoglu, D. and Restrepo, P. (2017). Robots and Jobs: Evidence from US Labor Markets.

Acemoglu, D. and Restrepo, P. (2018). Artificial Intelligence, Automation and Work, NBER Working Paper No. 24196.

Acemoglu, D. and Restrepo. P. (2019). Automation and New Tasks: How Technology Displaces and Reinstates Labor, Journal of Economic Perspectives, 33 (2): Vol. 33, No.2, pp. 3-30.

Aghion, P. and Howitt, P. (1992). A Model of Growth Through Creative Destruction, Econometrica, vol 60, issue 2, 323-351.

Agrawal, A. Gans. J. and Goldfarb, A. (2018), Prediction Machines: The Simple Economics of Artificial Intelligence, Harvard Business Review Press.

Allen, J. and van der Velden, R. (2001). Educational mismatches versus skill mismatches: effects on wages, job satisfaction, and on-the-job search. Oxford Economic Papers, 53(3), pp.434-452.

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

Autor, D. H., Levy, F., & 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. and Salomons, A. (2018). Is Automation Labor Share-Displacing? Productivity Growth, Employment, and the Labor Share. Brookings Papers on Economic Activity, 2018(1), pp.1-87.

Arntz, M., T. Gregory and U. Zierahn (2016), The Risk of Automation for Jobs in OECD Countries: A Comparative Analysis, OECD Social, Employment and Migration Working Papers, No. 189, OECD.

Berman, E.; Bound, J. & Machin, S. (1998). Implications of Skill-Biased Technological Change: International Evidence, The Quarterly Journal of Economics, Volume 113, Issue 4, November 1998, Pages 1245–1279.

BEIS (2017), Industrial Strategy: building a Britain fit for the future. Available here: https://www.gov.uk/government/publications/industrial-strategy-building-a-britain-fit-for-the-future

Bootle, R. (2019). The AI Economy: Work, Wealth and Welfare in the Robot Age, Nicholas Brealey Publishing.

Brynjolfsson, E. and Mitchell, T. (2017). What can machine learning do? Workforce implications. Science, 358(6370), pp.1530-1534.

Chiripanhura, B. and N. Wolf (2019), Long-term trends in UK employment: 1861 to 2018, ONS Economic Review, April 2019: https://www.ons.gov.uk/economy/nationalaccounts/uksectoraccounts/compendium/economicreview/april2019/longtermtrendsinukemployment1861to2018

Clarke, S. and Cominetti, N. (2019). Setting the record straight: How record employment has changed the UK, available online at: https://www.resolutionfoundation.org/publications/setting-the-record-straight-how-record-employment-has-changed-the-uk/

Computer Weekly (2019). UK artificial intelligence investment reaches record levels. [online] ComputerWeekly.com. Available at: https://www.computerweekly.com/news/252470371/UK-artificial-intelligence-investment-reaches-record-levels.

Cunha, F. and Heckman, J. (2007). The Technology of Skill Formation. American Economic Review, 97(2), pp.31-47.

Deming, D. J. (2015). The growing importance of social skills. NBER Working Paper No. 21473, August 2015, National Bureau of Economic Research.

Dworkin, J.D. (2019). Network-driven differences in mobility and optimal transitions among automatable jobs, Royal Society Open Science, available from: https://royalsocietypublishing.org/doi/10.1098/rsos.182124

Elliott, S. (2017). Computers and the Future of Skill Demand. OECD, Centre for Educational Research and Innovation.

EU JRC (2019). The Changing Nature of Work: Skills in the Digital Age. Available via the following webpage: https://ec.europa.eu/jrc/en/facts4eufuture/changing-nature-work-skills-digital-age

Felten, E.; Raj, M. and Seamans, R. (2019). The Variable Impact of Artificial Intelligence on Labor: The Role of Complementary Skills and Technologies, NYU Stern School of Business, unpublished manuscript, June 4, 2019

Financial Times (2019). The five UK sectors grappling with acute labour shortages. [online] Ft.com. Available at: https://www.ft.com/content/36baacce-ddd0-11e8-9f04-38d397e6661c.

Fisher, A. (2019). These are the most important skills you need to be successful in the modern workplace. [online] World Economic Forum. Available at: https://www.weforum.org/agenda/2019/05/soft-skills-are-hard-to-measure-and-in-demand-can-they-be-taught

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

Frey, C.B. and Osborne, M. (2013). The future of employment: How susceptible are jobs to computerisation?, Oxford Martin School working paper, available at: https://www.oxfordmartin.ox.ac.uk/downloads/academic/The_Future_of_Employment.pdf

Frey, C.B. (2019), The Technology Trap, Princeton University Press.

Goldin, C. and Katz, L. (1996). Technology, Skill, and the Wage Structure: Insights from the Past, American Economic Review 86: 252–257.

Goldin, C. and Katz, L. (1998). The Origins of Technology-Skill Complementarity, Quarterly Journal of Economics 113: 693–732.

Goos, M., & Manning, A. (2007). Lousy and lovely jobs: The rising polarization of work in Britain. The Review of Economics and Statistics, 89(1), 118-133.

Graetz, G. and Michaels, G. (2015), Robots at Work, CEP Discussion Paper No 1335.

GOV.UK (2018). Gig economy research. [online] GOV.UK. Available at: https://www.gov.uk/government/publications/gig-economy-research.

GOV.UK (2019). Leading experts appointed to AI Council to supercharge the UK's artificial intelligence sector. [online] GOV.UK. Available at: https://www.gov.uk/government/news/leading-experts-appointed-to-ai-council-to-supercharge-the-uks-artificial-intelligence-sector [Accessed 28 Oct. 2019].

Hanushek, E., Schwerdt, G., Woessmann, L. and Zhang, L. (2016). General Education, Vocational Education, and Labor-Market Outcomes over the Lifecycle. Journal of Human Resources, 52(1), pp.48-87.

Haskel, J. and Martin, C. (1993). Do Skill Shortages Reduce Productivity? Theory and Evidence from the United Kingdom. The Economic Journal, 103(417), p.386.

Hao, K. (2019). We analyzed 16,625 papers to figure out where AI is headed next. [online] MIT Technology Review. Available at: https://www.technologyreview.com/s/612768/we-analyzed-16625-papers-to-figure-out-where-ai-is-headed-next/.

Harari, Y.N. (2016), Homo Deus: A Brief History of Tomorrow, Harvill Secker.

Heckman, J. and Kautz, T. (2012). Hard Evidence on Soft Skills.

Heifetz, R. (1994). Leadership without easy answers. Cambridge, MA: Belknap Press of Harvard University Press.

HM Government (2018). Industrial Strategy: Artificial Intelligence Sector Deal. [online] Crown. Available at: https://assets.publishing.service.gov.uk/government/uploads/system/uploads/attachment_data/file/702810/180425_BEIS_AI_Sector_Deal__4_.pdf.

Industrial Strategy Council (2019), UK Skills Mismatch in 2030, Research Paper October 2019.

Lilien, D. (1982). Sectoral Shifts and Cyclical Unemployment. Journal of Political Economy, 90(4), 777-793. Retrieved from http://www.jstor.org/stable/1831352.

Littlewood, M. (2012). Robots are on the rise but never underestimate the human factor. The Times. [online] Available at: https://www.thetimes.co.uk/article/the-robots-are-coming-let-s-welcome-them-pp29z8vz2.

Manacorda, M. and Petrongolo, B. (1999). Skill Mismatch and Unemployment in OECD Countries, Economica, 66(262), pp.181-207.

Maree, K. (2018). Psychology of Career Adaptability, Employability and Resilience. New York: Springer.

McGowan, M. A. and Andrews, D. (2015). Skill mismatch and public policy in OECD countries.

McGuinness, S., Pouliakas, K. and Redmond, P. (2019). Skills-Displacing Technological Change and Its Impact on Jobs: Challenging Technological Alarmism? [online] Available at: https://www.iza.org/publications/dp/12541/skills-displacing-technological-change-and-its-impact-on-jobs-challenging-technological-alarmism.

McKinsey Global Institute (2017), Jobs Lost, Jobs Gained: Workforce Transitions in a Time of Automation, December 2017.

Moravec, H. (1988), Mind Children.

Mortensen, D. T. & Pissarides, C. A. (1999). Unemployment responses to ‘skill-biased’ technology shocks: The role of labour market policy. Economic Journal, 109(455), 242–265.

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

MIT Media Lab (2019). Project Overview ‹ AI, Automation, Labor, and Cities: How to map the future of work – MIT Media Lab. [online] MIT Media Lab. Available at: https://www.media.mit.edu/projects/future-of-work-ai-automation-labor/overview/.

NBER (2017). Economics of Artificial Intelligence. [online] Available at: https://conference.nber.org/conferences/2017/AIf17/summary.html.

Nedelkoska, L. and G. Quintini (2018), "Automation, skills use and training", OECD Social, Employment and Migration Working Papers, No. 202, OECD Publishing, Paris.

Neal, D. (1995). Industry-specific human capital: Evidence from displaced workers. Journal of labor Economics, 13(4), 653-677.

Nickell, S. and Bell, B. (1995). The Collapse in Demand For The Unskilled and Unemployment Across the OECD. Oxford Review of Economic Policy, 11(1), pp.40-62.

OECD (2001). Competencies for the Knowledge Economy. [online] Available at: http://www.oecd.org/innovation/research/1842070.pdf.

OECD, (2012). Better Skills, Better Jobs, Better Lives: a Strategic Approach to Skills Policies. OECD Publishing, Paris.

OECD Science, Technology and Industry Scoreboard (2017): The digital transformation, OECD Publishing, Paris, http://dx.doi.org/10.1787/9789264268821-en.

OECD (2018), Job Creation and Local Economic Development 2018: Preparing for the Future of Work, OECD Publishing, Paris. https://doi.org/10.1787/9789264305342-en

OECD (2019a), Getting Skills Right: Future-Ready Adult Learning Systems, February 2019, http://www.oecd.org/publications/getting-skills-right-future-ready-adult-learning-systems-9789264311756-en.htm

OECD (2019b), OECD Regional Outlook 2019: Leveraging Megatrends for Cities and Rural Areas, OECD Publishing, Paris, https://doi.org/10.1787/9789264312838-en.

ONS (2019). The probability of automation in England: 2011 and 2017, https://www.ons.gov.uk/employmentandlabourmarket/peopleinwork/employmentandemployeetypes/articles/theprobabilityofautomationinengland/2011and2017

Oxford Economics and Cisco (2017). The A.I. Paradox: How Robots Will Make Work More Human.

Petrongolo, B. and Pissarides, C. (2001). Looking into the Black Box: A Survey of the Matching Function. Journal of Economic Literature, 39(2), pp.390-431.

Poletaev, M. and Robinson, C. (2008). Human Capital Specificity: Evidence from the Dictionary of Occupational Titles and Displaced Worker Surveys, 1984–2000. Journal of Labor Economics, 26(3), pp.387-420.

Pucik, V., Evans, P., Björkman, I., & Morris, S. 2017. The global challenge: International human resource management (3rd ed). Chicago: Chicago Business Press.

PwC (2015) Service robots: The next big productivity platform.

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

PwC (2017b), Sizing the Prize, available here: https://www.pwc.com/gx/en/issues/data-and-analytics/publications/artificial-intelligence-study.html

PwC (2017c). The economic impact of artificial intelligence on the UK economy. Available at: https://www.pwc.co.uk/economic-services/assets/ai-uk-report-v2.pdf.

PwC (2017d) What doctor? Why AI and robotics will define New Health. Available via this webpage: https://www.pwc.com/gx/en/industries/healthcare/publications/new-health-a-vision-for-sustainability.html

PwC (2018a) Skies without limits: Drones - taking the UK's economy to new heights, available here: https://www.pwc.co.uk/issues/intelligent-digital/the-impact-of-drones-on-the-uk-economy.html

PwC (2018b). Will robots really steal our jobs? An international analysis of the potential long term impact of automation. [online] Available at: https://www.pwc.co.uk/economic-services/assets/international-impact-of-automation-feb-2018.pdf

PwC (2018c), What will be the net impact of AI and related technologies on jobs in the UK, UK Economic Outlook, July 2018. Available here: https://www.pwc.co.uk/economic-services/ukeo/ukeo-july18-net-impact-ai-uk-jobs.pdf

PwC Global Consumer Insights Survey (2018d): Artificial intelligence. Available here: https://www.pwc.com/gx/en/industries/consumer-markets/consumer-insights-survey/2018/survey/artificial-intelligence.html

PwC Global CEO Survey (2019a), available here: https://www.pwc.com/gx/en/ceo-agenda/ceosurvey/2019/gx.html

PwC Future of Government (2019b), available here: https://www.pwc.co.uk/industries/government-public-sector/insights/the-future-of-government.html

Rhodes, C. (2019). Industrial strategy. UK Parliament. Available at: https://researchbriefings.parliament.uk/ResearchBriefing/Summary/CBP-7682#fullreport.

Secretary of State for Health and Social Care (2018). The Government response to the Health and Social Care Select Committee Second Report of Session 2017-19, 'The Nursing Workforce.

Spitz-Oener, A (2006), Technical Change, Job Tasks, and Rising Educational Demands: Looking outside the Wage Structure, Journal of Labor Economics 24(2): 235-270.

Trade Union Congress (2019). UK's gig economy workforce has doubled since 2016, TUC and FEPS-backed research shows. [online] Tuc.org.uk. Available at: https://www.tuc.org.uk/news/uk%E2%80%99s-gig-economy-workforce-has-doubled-2016-tuc-and-feps-backed-research-shows

Nesta (2018), Classifying Occupations According to Their Skill Requirements in Job Advertisements.

Vivian et al. (2016) UK Employer Skills Survey 2015: UK Results, Evidence Report 97, UKCES.

World Economic Forum (2018): The Future of Jobs Report 2018. Available at: http://www3.weforum.org/docs/WEF_Future_of_Jobs_2018.pdf

World Bank (2019). The Changing Nature of Work. Available at: http://documents.worldbank.org/curated/en/816281518818814423/pdf/2019-WDR-Report.pdf.

Winick, E. (2018). Every study we could find on what automation will do to jobs, in one chart. MIT Technology Review. Available at: https://www.technologyreview.com/s/610005/every-study-we-could-find-on-what-automation-will-do-to-jobs-in-one-chart/

Technical Annexes

Technical Annex 1. Past research on economic impact of AI

In this annex we provide some further discussion of key studies on the economic impact of AI to supplement the discussion in the main text.

The general conclusion of this literature is that further development and adoption of AI and related technologies should bring significant economic benefits. Although some tasks will certainly be automated, technology in general is expected to boost the economy in the UK and elsewhere through the productivity gains from using 'labour-augmenting' AI applications, robots and other automating technologies.

Past research by PwC (2017c) has estimated that UK GDP could be up to around 10% higher in 2030 as a result of increased investment in AI. This reflects a combination of productivity gains for existing activities and consumption-side product enhancements and new firm entry stimulating demand.[80] The potential effects of AI on the economy are far-reaching beyond its effects on productivity growth and innovation, also affecting international trade and patterns of competition.[81]

Robots have indeed been found in past studies to be positive for businesses and the economy in many contexts.[82] Also they do not necessarily reduce human employment: in fact, robot density per worker in 2018 was the highest in three countries with relatively strong employment rates: Germany, Korea, and Singapore (World Bank, 2019). With this kind of evidence in mind, it is not so surprising that the cross-party House of Commons Business, Energy and Industrial Strategy Select Committee suggested recently that the real danger for the UK economy and for future jobs is not that we have too many robots in the workplace, but that we have too few.[83]

A consensus is building that AI will be the definitive technology of the next few decades. Some economists believe that it could disrupt the economy on a similar scale to the internet (e.g. Bootle (2019), or Agrawal, Gans and Goldfarb, (2018)). This view is shared by many CEOs worldwide, the majority of whom surveyed in PwC's 22nd Annual Global CEO Survey (2019) agreed with the statement that "AI will have a larger impact than the internet".

For these reasons, AI has become a central element of the UK government's Industrial Strategy - identified as one of the 'grand challenges' facing the country, along with an ageing society, the environment and 'mobility' (BEIS, 2017). The 'AI Sector Deal', a billion-pound joint government and industry deal, aims at putting the nation at the forefront of emerging

[80] See PwC report: The economic impact of artificial intelligence on the UK economy, 2017. Other studies by McKinsey (2017), World Economic Forum (2018) and, for China, PwC (2018d) also find the potential for positive net employment effects from AI and related technologies due to the boost to the size of their economies and so to labour demand, offsetting any initial job displacement effects. But the debate continues on this issue.
[81] Comprehensive research on the link between AI and the economy appears in 'The Economics of Artificial Intelligence', an NBER handbook edited by Ajay Agarwal, Joshua Gans, and Avi Goldfarb.
[82] According to a cross-country study of 17 countries, robotics has added an estimated 0.4 percentage points of annual GDP growth between 1993 and 2007 on average - accounting for about one-tenth of GDP growth during this time period (see Graetz and Michaels 2015). Acemoglu and Restrepo (2017) find a negative net employment effect for the US, but Michel and Bevins (2017) argue that these effects are relatively small and should not be generalised to automation more generally.
[83] https://www.theguardian.com/business/2019/sep/18/uk-economy-has-too-few-robots-warn-mps. The Resolution Foundation also echoes the need for more robots in its report "Robot wars: automation and the labour market", 2016.

technologies - including in the services sector through the Next Generation Services Industrial Strategy Challenge (HM Government, 2018). Leaders from business, academia and data privacy organisations have joined forces to help boost growth and ethical use of AI in the UK in businesses and organisations (GOV.UK, 2019).

The UK's interest in AI is translating into investment. Figures compiled by Tech Nation show that UK AI companies secured a record £800m of investment in the first six months of 2019 - which is already greater than the whole of 2018 and has risen six-fold since 2014, pushing the UK into third place globally behind the US and China. UK start-ups have raised almost double the amount achieved by their competitors in France, Germany and the rest of Europe last year. The UK is second after the US in the number of AI companies in operation, but 89% of businesses have 50 or fewer employees (Computer Weekly, 2019). The GovTech Fund also seeks to increase the uptake of AI by the government for more efficient public services.

The potential 'technological dividend' created by innovative, productivity-enhancing AI applications does not come without challenges - an important one being the potentially redistributive impact. Certain people, regions and countries will adapt faster and move ahead, which could exacerbate different types of inequalities (Acemoglu, 1998).[84] AI could help some people to thrive, whilst others could find their skills outdated or even entirely displaced by machines - which are increasingly capable in a range of cognitive tasks, including retrieving information, recognising patterns and generating predictions. Skill-biased technological change, mainly enhancing the work of skilled professionals, seems to have been a common feature since the 1980s (Berman et al., 1998). This has led to a polarisation of the labour market, with more highly paid jobs being created for the most skilled, whose work is enhanced by technology (Acemoglu, 1998); but at the same time, more lower quality and relatively poorly paid jobs emerging for the less skilled (in low-value services with stagnant real wages).[85] At the same time, many middle-level jobs are being displaced by automation and/or offshoring (itself facilitated by new information and communications technologies as well as the opening up of China, India and other emerging economies since the early 1990s – see Mortensen and Pissarides, (1999)).

How such trends evolve into the future will depend on the types of tasks that AI automates and the occupations and industries in which AI is used to 'augment' human labour.

What AI is good at and bad at is therefore a key question that drives research in this area. For example, using their taxonomy of tasks Brynjolfsson and Mitchell (2017) point out that because AI systems are often seen as 'black boxes' it is 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. More importantly, the fact that AI has largely been about lowering the cost of prediction (through machine learning) might suggest a developmental trajectory more like what we have seen in the past (e.g. when computation became cheaper): a gradual sequence of sector-specific and skill-specific disruptions without significant economy-wide effect.[86] In its current state-of-the-art uses, AI is said to be "narrow" – designed to accomplish a specific problem-solving or reasoning task (or cluster of related tasks). There are large uncertainties about the timeframe for the 'generalisation' of AI to be able to address multiple tasks in the way that humans can

[84] Also, in past PwC research, we estimated that there will be significant variations in net employment impacts by broad industry sector (SIC 1).
[85] Similar patterns have been found in other developed countries (Goos and Manning, 2007; Michaels et al., 2014).
[86] See work by Robert Seamans and also the work by Agrawal, Gans and Goldfarb (2018 and forthcoming).

routinely do.[87] There are also continuing challenges substituting machines for workers in tasks requiring adaptability, common sense (tacit knowledge), social intelligence, and creativity.[88]

[87] In an Artificial General Intelligence scenario, autonomous machines would become capable of general intelligent action, like a human being, including generalising and learning across different cognitive functions (see OECD 2017).
[88] See GTCI (2017), Chapter 6.

Technical Annex 2. Expert workshop

An expert workshop on the labour market impact of AI was hosted by PwC at their More London offices on Thursday 18th July 2019.

The workshop was conducted under Chatham House Rules so both the data responses and the views of experts are anonymous.

The quantitative data collected in the workshop is used as training data for the machine learning model described in Annex 3 to assess automation probabilities.

Workshop participants were asked to consider the key points related to two questions when filling in the questionnaire:

1. *"Can all or most of the tasks of this job be automated by digital technologies available within the next 20 years?"*
2. *"Will adoption of these digital technologies be on a scale sufficient for the majority of the people currently employed in the occupation not to be required to deliver the same level of production/service in 5, 10, 20 or more years?"*

How did the questions asked in our workshop differ to the questions Frey & Osborne put to their experts?

In their workshop, Frey and Osborne asked participants 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"*

This question is ambiguous in a number of ways:

- If we interpret "the tasks of this job" as referring to every task in a given occupation, then an expert may struggle to identify occupations for which every task is automatable. To avoid this problem, we amended the question to refer to "most of the tasks of this job".
- "Computer-controlled equipment" may be interpreted narrowly as excluding some forms of AI and related automating technologies. We used the more general term "digital technologies" instead.
- It does not specify a time horizon.

As such, we instead asked our experts:

> *"Can most of the tasks of this job be automated by digital technologies available within the next 20 years?"*

We note that the question above (as with the original Frey and Osborne question) pertains to the technological feasibility of automation. But just because a job could potentially be automated does not imply that it will actually be automated, as this will also depend on the economic, regulatory, political and legal barriers to rolling out these technologies across the economy, as well as the willingness of consumers to accept goods and services delivered by AI. This important distinction is recognised in the text of most previous studies, including those by Frey and Osborne, the OECD and PwC, but adjustments to the final estimates to reflect actual rather than potential probabilities of automation tend to be rather ad hoc or not attempted at all.

To try to estimate when technologies may actually automate jobs in a more systematic way, we asked our experts a follow-up question about the practical feasibility of automation (if they answer yes to Q1):

> *"Will adoption of these digital technologies be on a scale sufficient for the majority of the people currently employed in the occupation not to be required to deliver the same level of production/service in 5 years, 10 years, 20 years or over 20 years?"*

We used these responses in order to form a view on the likelihood of automation in the short term (c.5 years), medium term (c.10 years) and long term (c.20 years).

We asked 13 experts these questions with respect to the exact same 70 occupations that Frey and Osborne asked their experts about, so as to allow for direct comparisons between our results and those of previous studies.[89]

Quantitative results from the workshop

We used a simple majority rule (i.e. at least 7 out of 13 experts should hold a certain view) to generate a single binary label for each occupation to serve as training data for the job displacement model. We also ran a sensitivity with labels generated from a super-majority rule (i.e. requiring more than two-thirds of experts to hold a certain view - i.e. at least 9 out of 13). In the main report we use the label generated from the majority rule, but in Annex 5 we also present some results using labels generated from the supermajority rule as a sensitivity test.

[89] Experts were not shown the original Frey and Osborne labels when casting their votes, although some will have been familiar with that paper in general terms.

Figure A2-1: Proportion of occupations that the majority of workshop participants expect to be automated in the given time horizons

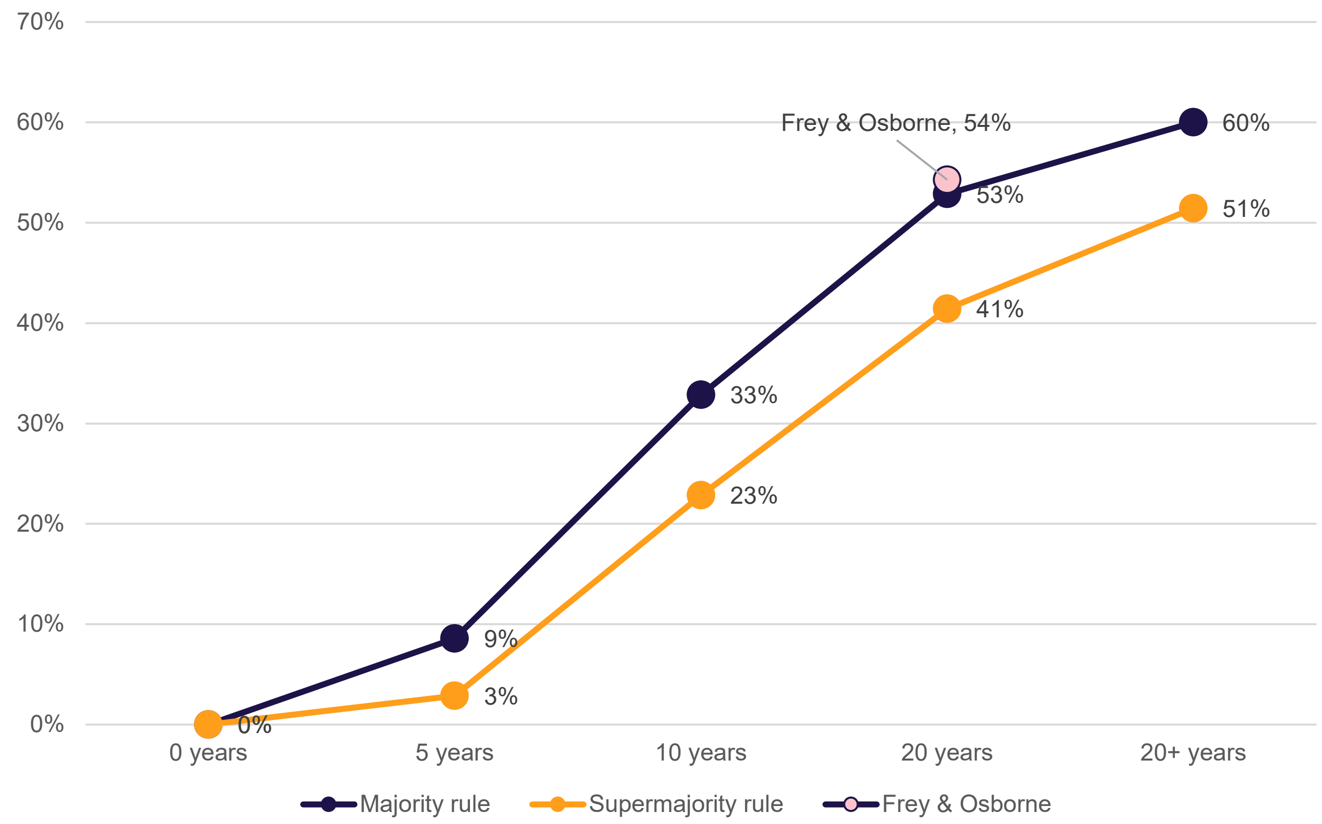

Figure A2-1 shows the quantitative results in terms of overall vote shares.

- In Frey and Osborne's 2013 study, which also drew on the results of an expert workshop, 38 out of 70 occupations (54%) were labelled as being potentially automatable from the perspective of technological feasibility. For our corresponding question (Q1) participants in this workshop labelled 42 of the same 70 occupations (60%) as automatable, based on majority voting.
- The timescale over which participants projected that AI-related automation would be rolled out across the economy traces out an S-shaped curve. More precisely, the majority of participants projected that 9% of the existing occupations considered (6/70) would be automated at scale in five years from now; 33% (23/70) in 10 years time; 53% (37/70) in 20 years time; and 60% (42/70) at some point beyond the 20 year time horizon for this study.
- The majority vote from our workshop produced a label that was different to Frey and Osborne's labels in the case of 8 occupations. In particular, the majority of participants labelled as automatable 6 occupations that Frey and Osborne labelled as non-automatable: 1) Transportation, storage and distribution managers; 2) Meeting, convention and event planners; 3) Education managers/Administrators; 4) Concierges; 5) Flight Attendants and 6) Waiters and Waitresses. Conversely, the majority of participants in our workshop labelled as non-automatable two occupations that Frey and Osborne labelled as potentially automatable: Surveyors and Farm Labour Contractors.
- Taken together, this may suggest that confidence about the ability of AI and related technologies to automate some managerial or customer service roles has increased

somewhat since 2013 as these technologies have advanced further. But, in the great majority of cases (62/70), the assessment from this workshop and that of Frey and Osborne have remained the same. This is consistent with the fact that the underlying characteristics of most of these occupations have not changed materially since 2013, while technology has developed along broadly similar lines to what was expected in 2013 (though perhaps faster in areas like deep learning and less rapidly in areas like robotics).

- The results of the workshop are based on majority voting. There was no decisive majority on the automatability of some occupations, so we caution over-interpreting the results in these cases. There was only one occupation for which all participants agreed that it was not technically feasible to automate: athletes and sports competitors. There were 14 occupations for which there was consensus that the occupation was automatable, but views of when that would happen in scale (10,15, 20 years) varied.

Figure A2-2: Vote share for each possible response for every occupation in the questionnaire

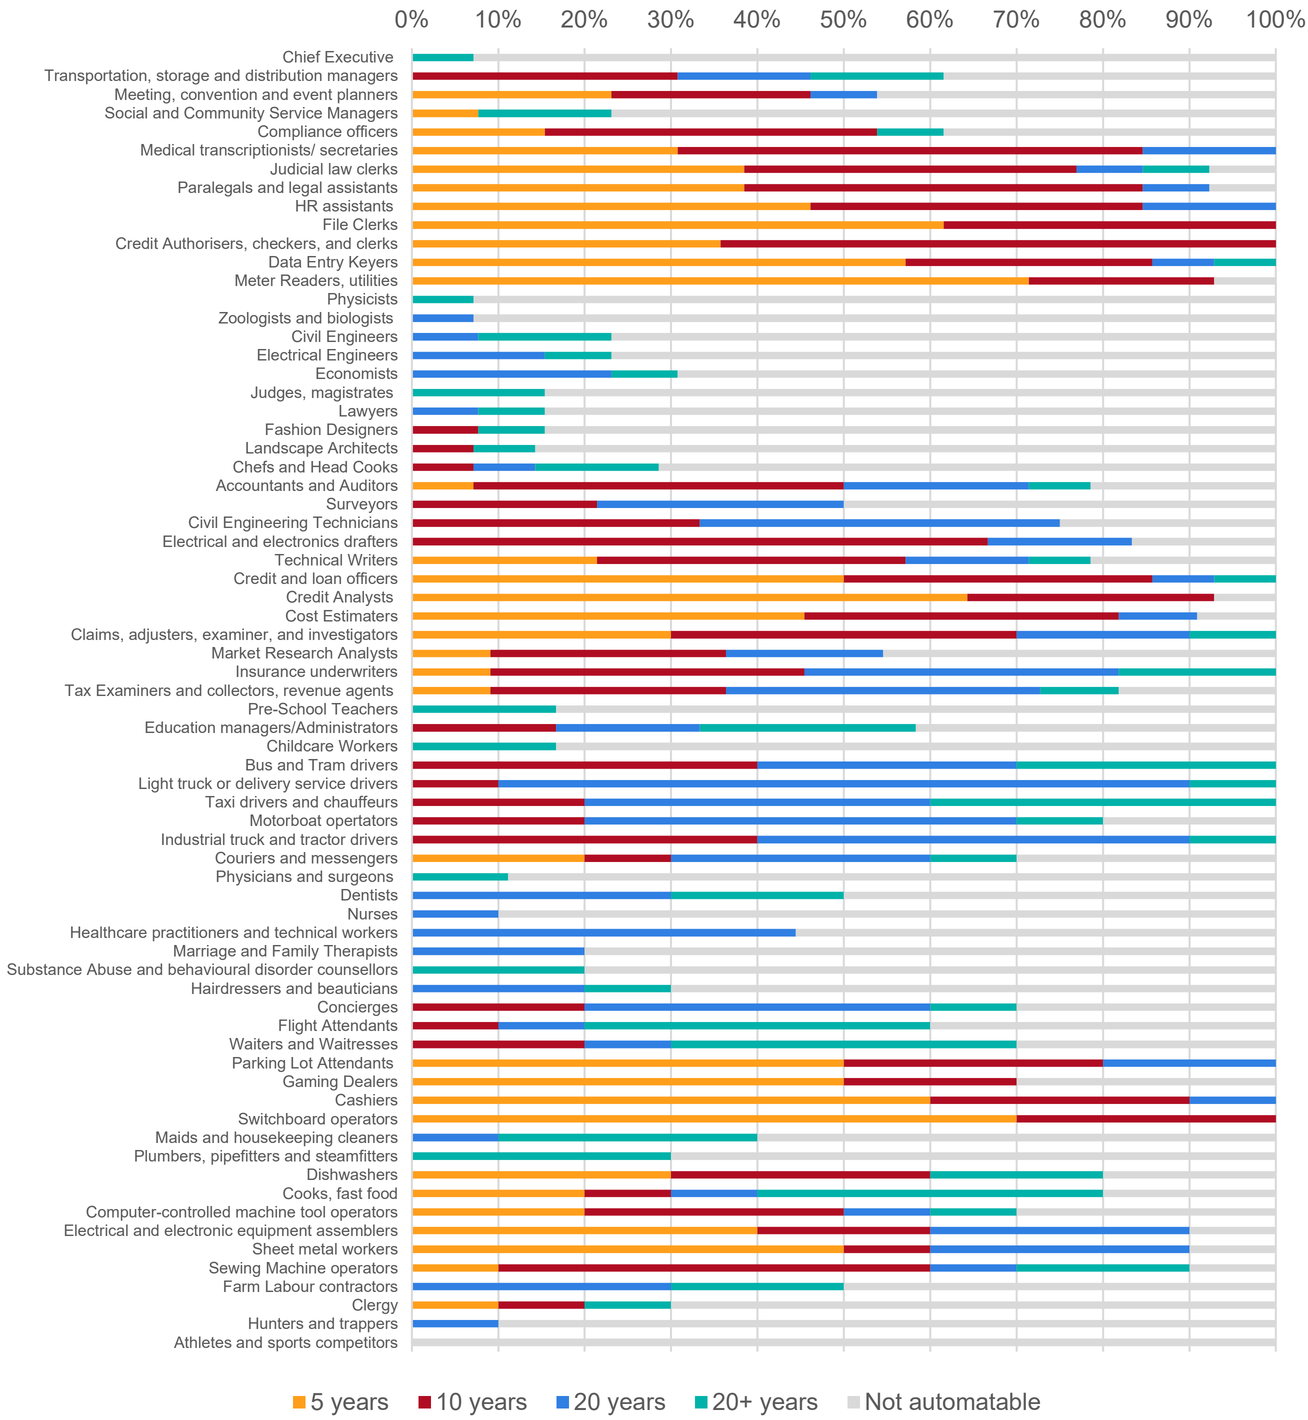

A full breakdown of the responses for each occupation are shown in Figure A2-2.

- Some general trends can be seen, such as the high degree of projected automation over 20 year time horizons for clerical workers, factory workers, drivers and some technical roles involving data analysis. Participants remarked that the first two groups of workers have already seen considerable automation, but the latter two groups may see more change, though again not all these roles will disappear - many will evolve in a way complementary to new technologies. In general, most such changes are only expected to be seen at 10-20 year time horizons, rather than in the shorter term.
- Other occupational groupings, such as senior managers, professionals, and health and education workers, are projected to see lower probabilities of automation, though here too there were divergent views. Some roles may disappear, but most may just evolve in line with technological advances, as in past such episodes.

Qualitative insights from the workshop

In addition to the quantitative assessment of automatability of different occupations, the workshop generated qualitative insights from the discussion - related to the technological feasibility of automation generated by AI but also to the non-technological factors that might affect the pace of automation in practice.

Below are the key messages from the discussion related to Question 1:

- In order to conceive what is technically feasible to automate, it is sometimes necessary to radically reimagine the existing workplace.
- It is also necessary to reimagine how humans will work alongside machines in order to conceive of which tasks are technically feasible to automate.
- Some technologies have advanced more rapidly than many experts would have expected over the last few years (e.g. Natural Language Processing and Computer Vision).
- Some tasks are not amenable to automation.

Below are the key messages from the discussion related to Question 2:

- Attitudes and preferences for human vs machines may speed up or slow down automation.
- Commercial viability will determine investment levels and uptake of labour-replacing technology and this will also depend on the cost/availability of labour.
- Automation could have profound ethical implications, but it remains to be seen how seriously these will be taken.
- Regulation in the form of labour protections may be needed to ensure the quality of jobs created.
- Politics will play a large role in determining the uptake of AI and related technologies.

Workshop participants

Christina Colclough

Director of Platform and Agency Workers, Digitalisation and Trade at UNI Global Union

Fabian Wallace-Stephens

Researcher in the RSA's Economy, Enterprise and Manufacturing Team and part of the RSA Future Work Centre research team

Jacob Beswick

Assistant Director, Adoption and Diffusion at Office for Artificial Intelligence

Jonathan Boys

Labour Market Economist at Chartered Institute of Personnel and Development

Dr Jonnie Penn

AI researcher at Harvard and Cambridge/Project Development Lead at The Young Workers Lab at Future World of Work division of UNI Global Union

Dr Julio Amador Diaz Lopez

Imperial College Research Fellow

Jyldyz Djumalieva

Data Science Research Fellow at NESTA

Olly Buston

CEO of Future Advocacy

Professor Terence Tse

Co-Founder and Executive Director at Nexus Frontier Tech and Associate Professor of Finance at the London campus of ESCP Europe Business School

Dr Xuxin Mao

Principal Economist at National Institute of Economic and Social Research (NIESR), Fellow (LSE)

PwC senior experts

John Hawksworth

Chief Economist at PwC

Euan Cameron

UK Artificial Intelligence Leader, PwC

Jonathan Gillham

Director of Econometrics and Economic Modelling, PwC

Sheetal Vyas

Lead Director for the Disruption & Innovation Team at PwC

BEIS representative (non-voting)

Riccardo Zecchinelli

Lead on Investment Analysis in innovation and technology at BEIS

Technical Annex 3. Detailed methodology for estimating job displacement and job creation

This research is based on a holistic approach to estimating the employment effect of AI (and related technologies). The approach is built upon two methodological pillars - each involving a self-contained methodology:

- **Job displacement effect:** we estimate the potential magnitude of job displacement by occupation and industry.
- **Job creation effect:** we estimate the potential magnitude of job creation by occupation and industry.

The **net employment effect** by occupation and industry is obtained by subtracting job displacement from job creation. Then we use the estimated net effect by occupation to derive estimates for net employment impacts by region, socioeconomic group and demographic group, as well as the implications for future skills demand.

Figure A3-1 gives an overview of the methodological steps. The rest of this annex explains these steps in more detail (except for the potential impact on future skills demand, which we defer to Technical Annex 4 where the general topic of skills is discussed in more detail).

Figure A3-1: Summary of our approach

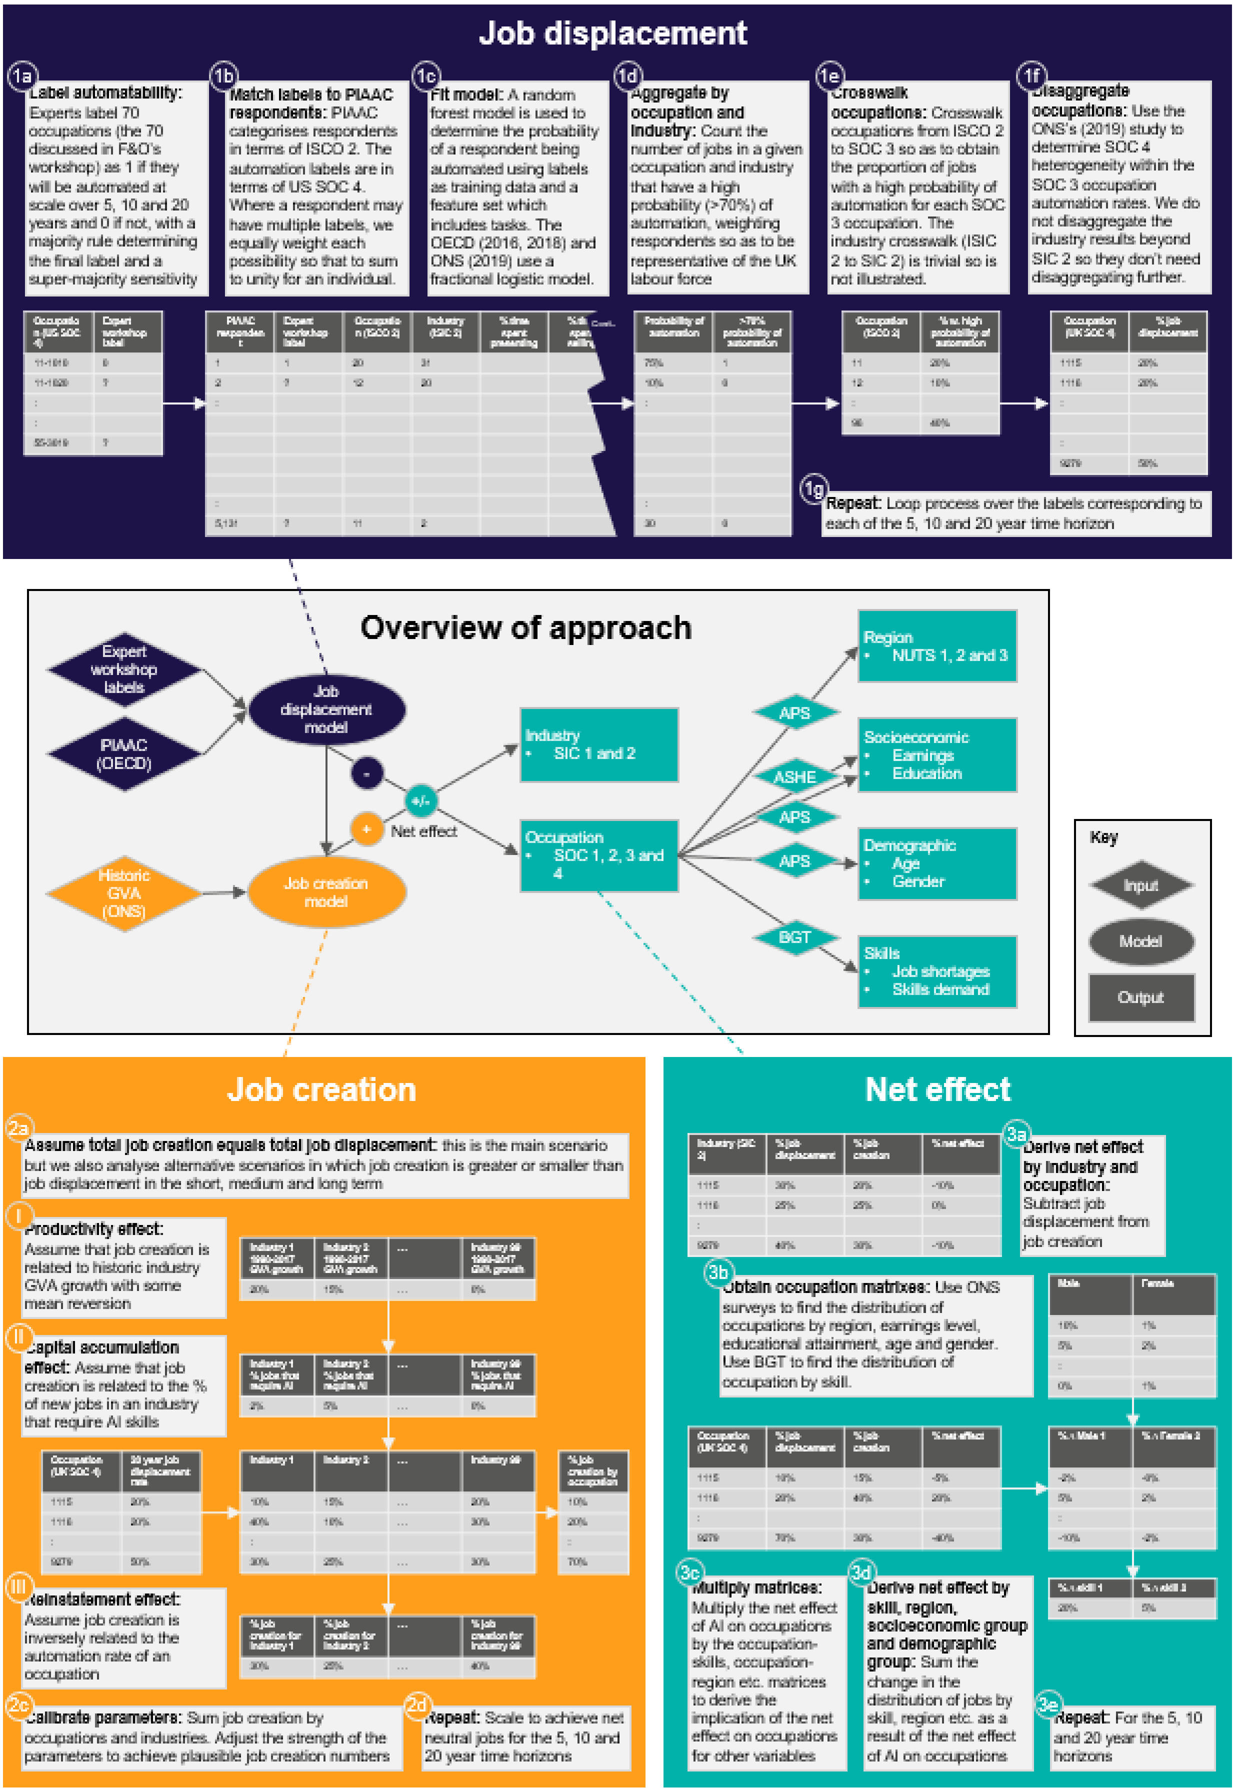

Approach to estimating job displacement

Our approach to estimating job displacement builds on a strand of literature initiated by Frey and Osborne (2013). The general approach is to:

(i) hand-label a subset of occupations as either automatable or not based on expert judgment;

(ii) fit a regression or machine learning model to explain the labels in terms of, among other features, the task composition of each occupation using an appropriate database (e.g. O*Net for the US or PIAAC for OECD countries); and then

(iii) use the model to estimate the probability for automation for all occupations (also including the subset of hand-labelled occupations, where the probabilistic model estimates overwrite these earlier binary expert judgements).

The diagram in Figure A3-2 illustrates how this report builds on selected previous studies that have taken this general approach.

Figure A3-2: How our approach to estimating job displacement relates to previous studies

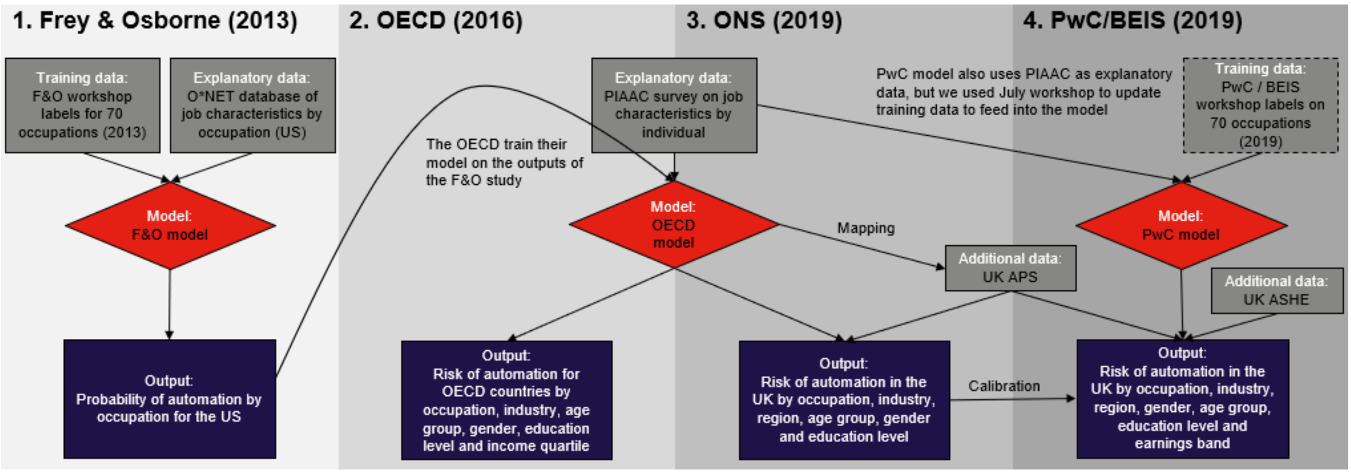

Source: PwC

1. For their 2013 study, **Frey and Osborne** (i) held an expert workshop to label 70 of 702 4-digit US SOC occupations; (ii) estimated a model based on O*NET data on the task composition of occupations; and then (iii) used this model to estimate the probability of automation for all 702 occupations in the US.

2. In 2016, the **OECD** published a study by Arntz et al. (2016) which (i) used the occupational automation probability estimates[90] for the US from Frey and Osborne as inputs to (ii) build another model based on data from the OECD's Programme for the International Assessment of Adult Competencies (PIAAC) survey, which allowed them to (iii) estimate the probability of automation for occupations and countries across the OECD.

[90] A later OECD study (Nedelkoska and Quintini, 2018) instead used the Frey and Osborne occupational label directly and estimated their model based on the Canadian PIAAC data only due to its relatively large sample size. Our approach is somewhat closer to their study in using occupational labels directly as training data, albeit based on labels from our own, more up-to-date expert workshop in July 2019 and estimating our model using PIAAC data for all countries rather than just for Canada (which has pros and cons, but using a Canadian data set for a UK-specific study did not seem appropriate).

3. The **ONS** recently (i) used Arntz et al's model to analyse the UK PIAAC data, but then (ii) regressed estimated automation rates from the OECD PIAAC analysis on demographic variables from the UK's Annual Population Survey (APS), which allowed them to (iii) estimate the probability of automation for 4-digit UK SOC occupations as well as granular industries, regions and demographic groups across England. Note that, since they used the model developed by Arntz et al, which was derived from the Frey and Osborne estimates, which in turn were based on the expert workshop labels from the Frey and Osborne study, the ONS were also indirectly using these same labels in their estimates, albeit in an indirect way due to the many intermediate steps in the process.

4. For **this study** we (i) held another expert workshop in order to refresh the input labels from Frey and Osborne, but this time specifying different labels for specific short, medium and long time horizons; (ii) estimated our model using the OECD PIAAC data but using what we consider to be a superior model to Arntz et al; and (iii) used additional data from the ONS to produce more disaggregated UK estimates of job displacement probabilities.

In summary, our approach to estimating job displacement (i) updates the training data for the probability of automation model; (ii) uses a superior modelling approach; and (iii) increases the granularity and comprehensiveness of estimates relative to past studies that we and others have done. We elaborate on these points below.

1a. Labelling the automatability of occupations through an expert workshop

The job displacement approach is based on a supervised machine learning (ML) model that, as with all such models, needs to be 'trained' on an initial set of data labels. In this case, the labels refer to whether an occupation is judged by human experts to be automatable or non-automatable. In previous studies, we and others (including the OECD, the Bank of England and the ONS) have relied directly or indirectly for initial training data on the labels for 70 selected occupations (out of a total of 702 US 4-digit SOC codes) used by Frey and Osborne based on a workshop of machine learning experts held at Oxford University in 2013.

There have, however, been many technological advances in AI and related technologies such as robotics and autonomous vehicles since 2013. There has also been considerable further practical experience on what might determine the speed of adoption of these technologies across the economy - regulations, social preferences etc. In light of these developments we held another expert workshop to revisit the original Frey and Osborne labels, resulting in revised automatability labels for the same 70 occupations, as described in detail in Annex 2 above.

1b. Mapping labels to PIAAC respondents

The expert workshop labels are for a selection of 70 4-digit US SOC occupations. Respondents in the PIAAC survey are classified in terms of 2-digit ISCO 08 occupations. It was therefore necessary to crosswalk the labels from 4-digit US SOC occupations to 2-digit ISCO 08 occupations, but doing so is a 'many-to-one' mapping process and therefore assigns some PIAAC respondents to multiple labels. For example it is ambiguous whether a PIAAC respondent who identifies as a "nursing professional" in terms of 2-digit ISCO codes is in fact a "registered nurse", "nurse anaesthetist" or "nurse midwife" according to the 4-digit US SOC codes which we have labelled. When fitting the model we assign ⅓ weight to each possibility. In general we assign 1/n weight to each possible occupation, where n is the number of possible occupations, so that the weight sums to unity for each individual.

Arntz et al. (2016) adjusted the weights using an Expectations Maximisation (EM) algorithm. This involved an iterative process in which Arntz et al. predict the probability of automation of each occupation using a fractional logit model and then adjust the initial weights based on these predictions, which they then use to re-predict the probabilities and so forth. However, PwC 2017a found that adjusting the weights did not improve model performance. In other words the model did not perform any better in the last iteration compared with the first iteration. Consequently, we resolved to use the non-adjusted weights. By doing so PwC were able to separate this calculation from the logit regression model that estimates the probability of automation.

1c. Fitting model on labels

To estimate probabilities of automation by occupation, Frey and Osborne (2013) used the O*NET database, which details the task composition of each occupation listed in the US Labor Department's Standard Occupational Classification (SOC). In their paper, Arntz et al. point out that by using this database Frey and Osborne are forced to assume that all jobs within a given occupation have the same task composition and are therefore equally likely to be automated. Arntz et al. avoid this assumption by using the OECD's Programme for the International Assessment of Adult Competencies (PIAAC) survey, which is for individuals not jobs. The task composition of jobs in PIAAC are self-reported by individual respondents, which means that individuals can be given unique automation rates based on the unique task compositions that they report doing for their particular job.

We agree that this is a methodological improvement. However, as argued in PwC 2017a, Arntz et al.'s logistic regression model lacked predictive power. PwC 2017a illustrates this graphically by showing that that estimated probabilities of automation are very contingent choice of feature set, and when a more predictive feature is used, Arntz et al.'s headline result - that only 9% of jobs in the US have a high probability of automation - rose[91] to around 30% across the OECD or around 38% for the US, closer to - but still below - the original Frey and Osborne estimate of 47% for the US.[92] Indeed when the distribution is calibrated with the labels from Frey & Osborne's expert workshop the logistic model is able to recover Frey & Osborne's probabilities of automation, which are used as an input to the logistic model.[93] PwC 2017a therefore concludes that Arntz et al.'s result "was more an artefact of their methodology than a true representation of the data".

PwC (2018b) seeks to correct a more fundamental issue with using the logistic regression model approach. The logistic model uses as predictors the probabilities of automation for occupations estimated by Frey & Osborne. In doing so, Arntz et al. (2016) build a 'model on a model' i.e. they use Frey & Osborne's outputs as their inputs.[94] This means that any biases in Frey & Osborne's model are carried over to their model. For example, if the task composition in the O*NET survey are not representative of an occupation, this would cause Frey & Osborne

[91] The reason a weakly predictive model caused Arntz et al. to underestimate the proportion of jobs with a high probability of automation is because "high" probability of automation is, consistent with Frey & Osborne, defined as a probability above 70%. A weakly predictive model places most individual jobs in the medium probability range (30-70%), therefore reducing the proportion of individuals above 70% probability of automation at the tail of the distribution.

[92] The smaller difference that remains could be explained by the fact that we are using the PIAAC data for individual jobs, not the O*Net data for occupations, but this difference is of a more plausible magnitude than that implied by the Arntz et al. estimates, or indeed the later OECD study by Nedelkoska and Quintini (2018), which suffered from similar methodological limitations.

[93] PwC re-calibrated the distribution of probabilities according to the ratio of 0s to 1s over the relevant occupations from Frey & Osborne's expert workshop.

[94] Indeed another way to see that Arntz et al.'s model suffers a lack of predictive power is to observe that their inputs (Frey & Osborne's probabilities of automation) do not resemble their outputs.

to under- or over-estimate the automatability of that occupation, and this error is carried through to their model. The process of using as inputs the probability of automation by occupation to derive estimates on the probability of automation by occupation (albeit for countries outside the US) is also quite a convoluted process. To simplify the end-to-end modelling process and to avoid carrying over errors PwC 2018b train their model on Frey & Osborne's input labels rather than their output probabilities. For this present study we updated these input labels, as discussed above.

In PwC 2018b we developed a random forest model to estimate the probability of an individual in the PIAAC survey being automated. The random forest model is fitted on the 70 occupations labelled by our experts in the workshop. Whilst each tree in the forest estimates an individual in PIAAC as 1 or 0, we obtain the probability of automation from the proportion of trees in the forest that assign an individual to 1 or 0.

Feature set

Our model uses the following key features from PIAAC. This list is ranked based on the importance of the variable in the random forest model. The below ranking relates to the majority labels over for the 20 year time horizon.[95] We have selected the top twelve most important variables.

1. Educational job requirements
2. % time reading books
3. % time planning activities of others
4. % time presenting
5. % time influencing people
6. Responsible for other staff
7. % time using internet for work related activities
8. % time working physically for long
9. Hugh numeracy skills
10. % time calculating shares or percentages
11. % time organising own schedule
12. % time reading professional publications
13. % time writing articles

The most important driver of automation probability was an individual job's task composition.

We implemented the random forest model using the randomForest package in R.

[95] The importance ranking is similar for the other models.

As the model uses binary inputs as training data it does not risk carrying over or compounding estimation errors from Frey and Osborne's model. Moreover as training data has been updated in June 2019 it reflects an up-to-date view of the state of AI and related automating technologies.

In addition, we believe that the random forest model used in this present study is superior in its explanatory power to the logistic model developed by Arntz et al. which has been adopted by ONS (2019) and was refined by Nedelkoska and Quintini (2018). The random forest model is also capable of capturing complex interactions between features in the data, improving the fit of the model significantly relative to Arntz et al's model. The latter tends to push more occupations into the medium automation probability category (30-70%) due to its lack of explanatory power, while our model results in more occupations still being estimated to face a high (> 70%) or low (< 30%) probability of automation, which is more in line with the original Frey and Osborne study. Since the initial labels from our expert workshop are not all that different from those used by Frey and Osborne, this makes more sense intuitively than the Artnz et al. study finding very different results, despite also using the Frey and Osborne estimates of automation probabilities as their starting point.

1d. Aggregating estimates by occupation and industry

As explained above, our model estimates the probability of automation for each individual respondent in the PIAAC survey sample. As we want to estimate the automation rates by occupation and industry, it is necessary to aggregate individuals by occupation and industry. To do so we first calculate whether an individual has a high probability of their existing job being automated and, consistent with the literature as noted above, we define 'high' as greater than 70%.[96] We then calculate the proportion of individuals in a given occupation (or industry) that have a high estimated probability of being automated.[97] Our model can do this for all 29 countries with adequate PIAAC data, but for the purposes of this study we focus on the UK results only.

1e. Crosswalking estimates to UK SOC and SIC codes

As the PIAAC survey classifies occupations in terms of 2-digit ISCO occupations and ISIC industries, so too the output of our job displacement model is in terms of 2-digit ISCO and ISIC codes. To make the results consistent with the UK's standard occupation and industry classification systems it is necessary to crosswalk the results so that they are in terms of the UK's standard occupational classification (SOC) and standard industrial classification (SIC).

The ISIC to UK SIC crosswalk is trivial as the codes map to each other on a one-to-one basis. No further data manipulation is therefore needed to obtain job displacement rates for the 87 2-digit SIC industries and the 19 1-digit SIC industries (note that these are alphanumeric - A, B, C etc). However, we caveat that the results for some of the industries, especially some smaller 2-digit SIC industries but also the 1-digit industries 'T' and 'U', are based on small sample sizes in PIAAC and are therefore unreliable, which we flag later in the results section.

The ISCO to UK SOC crosswalk is more complex as the codes do not have one-to-one correspondences. As there are around 40 2-digit ISIC codes and only 25 2-digit SOC codes, we crosswalk to the 90 3-digit SOC occupations, so as to create a one-to-many mapping which

[96] It is important to consider that AI will likely still automate many of the tasks involved in the jobs that are not displaced, just not so much that that workers currently in the job couldn't plausibly retain the job with their existing core skill set.

[97] When aggregating individuals by occupation we weight respondents so that the aggregation is representative of the UK economy (i.e. if a respondent is under-represented in the survey then we will give them a greater weight).

will ensure that each 3-digit SOC occupation is assigned a single value, albeit non-unique where there are multiple 3-digit SOC occupations that map to the same 2-digit ISCO code.

1f. Disaggregating estimates to 4-digit SOC

We have disaggregated the automation rates by occupation further from the 3-digit to the 4-digit SOC level using the results of the 2019 ONS automation probability study.[98] This study used a pooled dataset of seven years (2011-2017) of APS data to ensure the sample sizes were large enough to calculate a probability of automation for each 4-digit SOC code. We used the ONS's probability of automation results at the 4-digit SOC level relative to their estimates at the SOC 3 level to estimate the relative probabilities of automation of 4-digit SOC codes within each of the 3-digit SOC codes that we estimated as described above. We did this in a way that constrained the weighted average automation rate at the 3-digit SOC occupational level to be consistent with the predictions of our job displacement model, with employment data from the APS as the basis for the weights in this calculation.

The resulting estimates at 4-digit SOC code level are subject to greater margins of uncertainty than the 3-digit SOC estimates and should be interpreted with appropriate caution. However, it seemed worthwhile to go to this extra level of detail since the alternative to using the relative 4-digit vs 3-digit estimates from the ONS study would have been to assume no difference in 4-digit SOC estimates within a given 3-digit SOC category, which seems unlikely to be the case.

1g. Repeat for 5, 10 and 20 year time horizons

The above process is repeated for the expert workshop labels associated with the 5, 10 and 20 years time horizons, so as to generate a comprehensive set of estimates of the proportion of jobs with a high probability (70%+) of being automated by 1-digit and 2-digit SIC and 1-digit, 2-digit, 3-digit and 4-digit SOC. Consistent with the literature, we interpret these as estimates of the proportion of existing UK jobs that we estimate may be displaced by AI and related technologies at these different timescales.

Note that the results of this exercise for a 20 year time horizon are consistent for the results in PwC (2018b) at a 20 year time horizon if we instead use the original Frey and Osborne workshop labels as inputs. The change in inputs is due to updated inputs, not a different model.

Approach to estimating job creation

Even though there is no consensus in the literature on the proportion of existing jobs that will be automated over the next few decades, there is at least a recognised approach that we have followed. The same cannot be said for estimating the number of jobs that will be created by AI and related technologies. For example, McKinsey estimate the global economic value of AI using their "micro-to-macro" model. PwC (2017b) estimates the economic value of AI using a UK focused dynamic computable general equilibrium (CGE) model. Unsurprisingly, the different approaches that have been taken have generated very different estimates. Indeed, an article in the MIT Technology Review lists 18 such attempts to estimate (either or both) the number of jobs that will be created and destroyed by AI (Winick, 2018).

[98] ONS (2019), The probability of automation in England: 2011 and 2017.

2a. Job creation scenarios

Since the primary focus of this study is on the distributional impact of AI on jobs, we do not attempt to quantify the total number of jobs that will be created by AI. Instead, we consider a range of scenarios. In our main 'base case' scenario we assume that the total number of jobs created by AI and related technologies is equal to the total number of jobs displaced by AI. We consider this a plausible scenario because it is consistent with our previous attempt to quantify the net effect of AI on jobs in the UK, it places us in the middle of other estimates (some of which suggest a positive net effect and others a negative net effect of AI), and it is consistent with the fact that previous technological revolutions have not led to mass technological unemployment. We can see this from estimates by the Bank of England, extending official ONS estimates back to the 1860s (see Figure A3-3), that there has been no very clear upward and downward trend in the UK employment rate over this 150 year period. Instead there have been some significant cyclical dips in employment rates due to severe recessions or depressions (most notably in the 1920s and 1930s, and to a lesser degree in the downturns seen since the 1970s) but no clear evidence of technological unemployment despite the many advances seen over this long time period.

Figure A3-3: UK employment rate since 1861

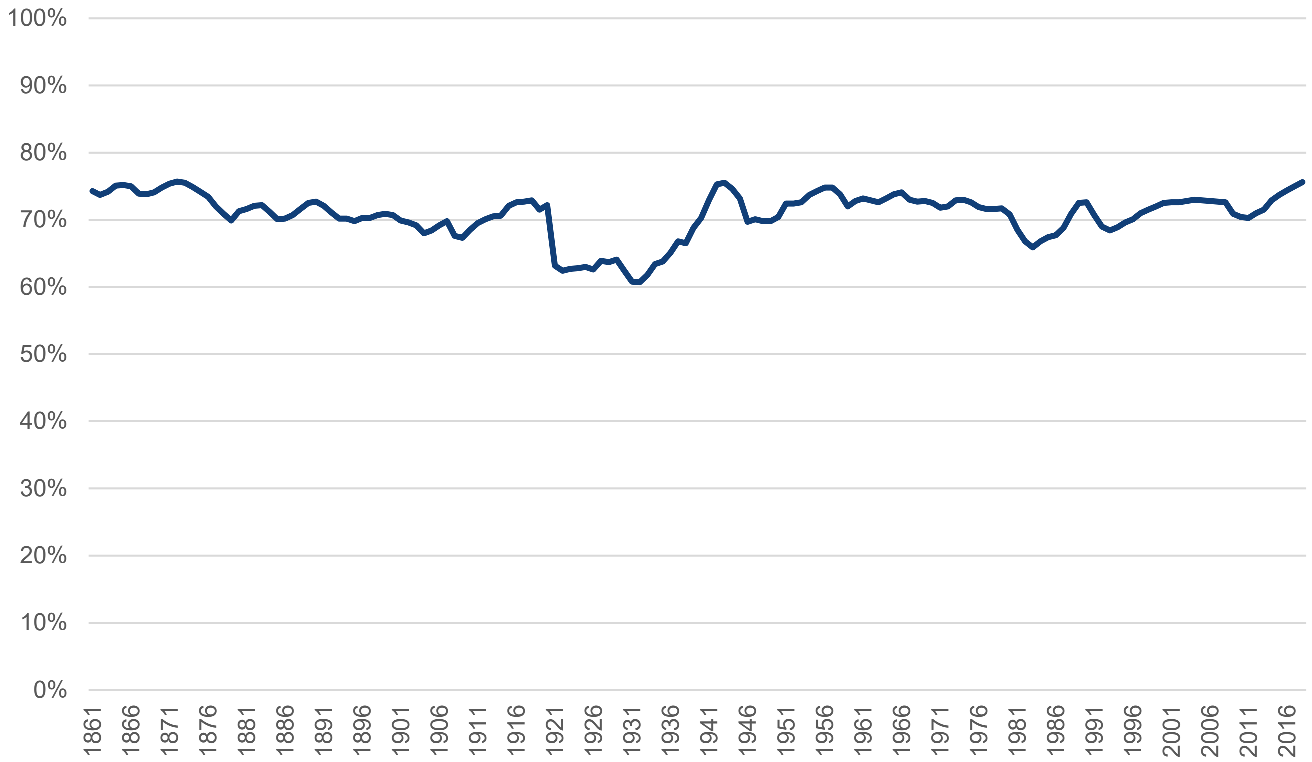

Source: PwC analysis of ONS data

Instead, what we have seen over long periods of time have been major shifts first from agriculture to manufacturing and services in the 19th century, and then from manufacturing to services since around 1960. In future, the major shifts will be from services sectors that are more automatable to those that are less automatable, particularly where the latter are producing goods and services that are in high demand (e.g. health and social care due to an ageing population and a relatively high income elasticity of demand; the latter may also be true for personal services more generally that still require the 'human touch').

We therefore believe that both past research and long-term historical trends in the UK labour market do not point to AI causing mass technological unemployment.[99] A broadly neutral net employment effect seems the most plausible assumption to make, as in our base case estimates.

However, we recognise that this is inevitably subject to significant (and irreducible) levels of uncertainty, so we also examine alternative scenarios in Annex 5 below: one in which job creation is 25%[100] less than job displacement at all time horizons; one in which job creation is 25% greater than job displacement at all time horizons; and another in which job creation lags behind automation initially, but eventually 'catches up' over 20 years.

These scenarios provide a test of whether the sign of the estimated net employment effect (i.e. whether it is positive or negative for different occupations, industries and demographic groups) is robust to different macroeconomic scenarios. In general, we find that this to be true for broad groupings, though it will not apply in every case for more disaggregated categories. However, our sensitivity analysis offers some comfort that the broad qualitative conclusions we would draw from our analysis (e.g. the impact of education levels on net employment estimates) remain valid in alternative scenarios - see Annex 5 for details.

2b. Parameterising effects

Given our base case assumption that job creation will balance job displacement from AI at the aggregate macroeconomic level for the UK, the key modelling challenge is therefore to apportion the jobs created to different occupations and industries. To do so we have adopted a framework that, while not based on any detailed theoretical model, can be related in general qualitative terms to a framework set out by Acemoglu and Restrepo (2018), who argue that AI. creates jobs through three effects: the productivity effect (which we refer to in past studies as the 'income effect'), the capital accumulation effect and the replacement effect.

Our job creation model uses three datasets to proxy the size of these effects (either by industry or occupation), and we then adjust the relative strength of each effect to produce plausible job creation estimates over 5, 10 and 20 year time horizons. In Table A3-1 we describe each effect and how it has been modelled.

Table A3-1: Job creation modelling approach

Effect	Modelling approach
Productivity (or income) effect: "By reducing the cost of producing a subset of tasks, automation raises the demand for labor in non-automated tasks" (Autor, 2015; Acemoglu and Restrepo, 2018). Productivity gains occur in sectors undergoing automation and this leads to higher real incomes and thus to greater demand for all goods and services, including those not experiencing (much) automation.	The productivity effect increases demand for all products, and we assume that the demand for all products will increase broadly in line with recent trends. We therefore assume that the productivity effects will be felt the most in the industries that have seen the highest growth (as measured by average GVA growth by 2-digit SIC industry for 1998-2007). To avoid unrealistic divergence we assume some degree of mean reversion, the strength of which determines the strength of the productivity effect.

[99] This is also the conclusion of a recent book by Bootle (2019).
[100] Given the high level of uncertainty involved, we adopt a relatively large variation here for the purposes of sensitivity testing, as in PwC (2018c).

Capital accumulation effect: "The adoption of new technologies implies rising demand for new machines and intangible capital, which increases demand for knowledge-based tasks and for labour tasks that involve producing, implementing, maintaining and upgrading the new technologies in use" (EU JRC, 2019).	We assume that industries with the most jobs created as a result of capital accumulation are the industries in which the demand for AI and related skills is currently highest, as measured using data on all UK job advertisements scraped by Burning Glass Technologies in 2018. This assumption is applied to 2-digit SIC industries.
Replacement effect: "New technologies induce the creation of new tasks for workers as ... the displacement of workers from old tasks could make more workers available to take over new, more productive tasks." (EU JRC, 2019).	In general workers will be reinstated from occupations where there tasks are highly automatable into occupations which are less easily automated. We therefore assume that more jobs will be created in occupations that are relatively less automatable, as measured by our estimates of job displacement by 4-digit SOC occupation.

In summary, our job creation model assumes that the productivity/income effect, proxied by historic GVA growth, modifies job creation by 2-digit SIC industry; the capital accumulation effect, proxied by the current demand for AI and related skills, modifies job creation by 2-digit SIC industry; and the reinstatement effect, assumed to be inversely related to automatability, modifies job creation by 4-digit SOC occupation. This implies that the highest job creation rate will be in occupations that are not automatable (high reinstatement effect) that reside in industries that are high growth (productivity/income effect) and AI-intensive (capital accumulation effect).

Estimating job creation: Illustrative example

To illustrate how this works in practice, let us assume for simplicity that the economy has only two industries (health and agriculture) and two occupations (technicians and administrators). In this case the data required to estimate the job creation effect would take the form of a two-by-two matrix with job creation rates for technicians and administrators in the health industry as one row and technicians and administrators in agriculture as another row. The aggregate job creation rate is determined by the aggregate job displacement rate - and in the base case scenario is equal to it. Suppose that the aggregate job displacement rate over 20 years is 20%, and therefore that the aggregate job creation rate over is also 20% over the same period.

Let us provisionally assume that all occupations across all industries have the same job creation rate i.e. every cell in the matrix is 20%. The productivity (or income) effect has the effect of adjusting the job creation rates by industry. If, for example, the historic GVA growth in the health industry was higher than in agriculture, then this serves to increase the job creation rate for all occupations in the health sector relative to agriculture (e.g. the job creation rate for technicians and administrators in health is, say, 25% and 25%, offset by lower job creation rates of 15% and 15% in agriculture).

The capital accumulation effect also serves to modify the job creation rates by industry. If, for example, there was proportionally more AI investment - and therefore more capital accumulation - in health than agriculture, then the capital accumulation effect would serve to further increase the divergence in job creation rates in the health and agricultural sectors to, say, 30% and 10%, respectively.

Finally, the replacement effect influences job creation rates by occupation. If, for example, administrators are estimated to be more automatable than technicians, then the replacement effect will increase the job creation rate of technicians in health and agriculture to, say, 35% and 11%, respectively, and will decrease the job creation rate of administrators to, say, 25% and 9%.

In the full-scale model, these effects are much more complex because they are applied to a table of employment of 369 4-digit SOC occupations by 87 2-digit SIC industries, obtained from the ONS. All effects are implemented such that the employment matrix balances, which in the base case scenario means that the total number of jobs created must equal the total number of jobs displaced. One implication of this is that, if (in the illustrative example above) there are currently more administrators than technicians, then a 1% decrease in the job creation rate of administrators (e.g. 10% to 9%) must be offset by a greater than 1% increase in the job creation rate for technicians. Another feature of the model is that if there is an unequal distribution of occupations by industry, then effects that alter industry job creation rates also alter occupation job creation rates, and vice versa.

2c. Calibrating parameters

The job creation model consists of three effects which ‘compete’ with one another to collectively determine estimated relative job creation by industry and occupation, subject to the constraint. Clearly, it is necessary to calibrate the model by deciding the strength of these competing effects, and we parameterise each effect between 0 and 1. In our base case, we set the productivity (or income) effect to be strongest i.e. have the biggest impact in terms of which industries jobs are estimated to be created in. Nevertheless, the parameter that determines the strength of this effect is set below 1, implying some degree of mean reversion. The replacement effect is assumed to be the next strongest effect, and finally the capital accumulation effect is set to be the weakest of the three effects since the data underlying the assumptions on this latter effect are less strong than for the other parameters. We therefore prefer not to give this effect too much weight in our job creation estimates.

2d. Repeat for 5, 10 and 20 year time horizons

We project job creation over the same time horizons used to project job displacement (i.e. 5, 10 and 20 years). For each time horizon in our base case scenario for job creation is scaled such that total jobs created equals total job displaced. For the alternative scenarios we scale job creation accordingly.

Approach to estimating net employment effects of AI

3a. Deriving net effect by industry and occupation

Since we have estimated job creation and job displacement by industry and occupation as described above, we can also immediately calculate the estimated net employment effect of AI by industry and occupation for 5, 10 and 20 year time horizons.

3b, c, d. Using occupational matrices to produce net employment effect estimates for regional, socio-economic and demographic groups

Our general approach to disaggregating the net effect of AI by region and demographic groups is to use the net effect of AI by occupation combined with knowledge of how occupations are distributed by region and demographic group in order to derive the net effect of AI across regions and demographic groups. Suppose, for example, that the net effect of AI is +5% for chefs and +15% for cleaners, and suppose further that 50% of 16-19 year olds are chefs and 50% are cleaners, then it follows that the net effect of AI for 16-19 year olds is 10%.

To do so it is necessary to obtain matrices which show the distribution of occupations by region, earnings level, educational attainment level, age and gender. We use matrices from the ONS for 4-digit occupations. The occupational distribution by region, educational attainment level, age and gender are sourced from the Annual Population Survey (APS) and we obtain earnings data from the Annual Survey of Hours and Earnings (ASHE). There are limitations to this approach, so the results should be taken as illustrative of broad patterns, rather than definitive estimates for particular sub-groups of workers.

3e. Repeat for all time horizons

Again this process is repeated for the 5, 10 and 20 year time horizons. Taken together this provides us with estimates on the net effect of AI by industry, occupation, region, earnings level, educational attainment level, age and gender over a 5, 10 and 20 year time horizon. Although we focus our analysis of the estimates on the net effect, we present charts which decompose the net effect into job creation and job displacement in order to understand what is driving the results.

Technical Annex 4: Approach to estimating the future demand for skills

The methodology we have used to estimate the future demand for skills is illustrated in Figure A4-1 below. As described in sections 3 and 4 of the main text, we have produced estimated net effects of job displacement and job creation for each occupation down to SOC 4 level (albeit with some caveats about the reliability of estimates below SOC 3 level). These estimates suggest strongly, for example, that the numbers of nurses needed in the UK will increase due to the impact of AI in boosting the size of the economy (and so tax revenues to fund the NHS) over the next 5, 10 or 20 years. By contrast, the estimates suggest that the number of customer service workers is likely to decline due to automation. The next step is to map the skills required by each SOC 4 occupation through an occupation/skills matrix.

By combining the net changes by occupation (in absolute numbers) with the occupation/skills matrix (as illustrated by the rectangle in Figure A4-1), we are able to estimate the potential increase (or decrease) in the demand for different skills over 5, 10 and 20 years.

Figure A4-1: Model used to estimate the future demand for skills

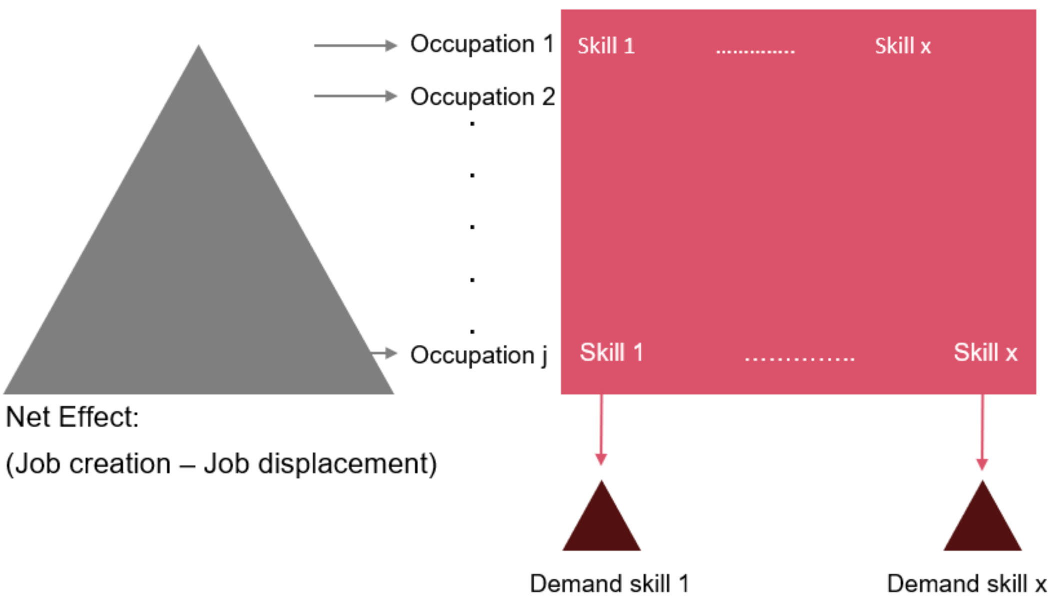

Source: PwC

This approach is in line with previous studies analysing the evolution of tasks and skills. Several studies rely on the O*NET database, which was developed under the sponsorship of the U.S. Department of Labor and contains hundreds of standardised and occupation-specific descriptors on almost 1,000 occupations covering the entire U.S. economy. The parameters of the database (mainly describing skills needed by occupations in the U.S.) have been applied in different country contexts to build occupation/skills matrices. But this data set, and others available from official sources in the UK and elsewhere, suffer from not being very up to date, which has led us to look instead at online data.

We have used data on online job ads from Burning Glass Technologies (BGT)[101] to obtain up-to-date information about the skills required by different occupations. This database has several advantages (although also some limitations as noted in Box A4.1). Firstly, the skills that are captured for each job advertisement are based on the needs of the market - and not informed by expert opinions. Secondly, online job postings are a great source of near real-time information on the labour market - in addition to vast volumes (our dataset contains 45,346,459 UK adverts for 2018), job adverts also offer more granular data than skill surveys as they allow employers to describe their skill needs more precisely. An additional advantage is that BGT database aggregates information by occupation using the SOC classification, which fits well with the occupational analysis carried out in this study.

Box A4-1: Limitations of BGT data

One limitation of the BGT data, or any similar data set, is that not all job openings are advertised online, and the proportion that are will vary systematically by occupation and industry and over time (as the online job market has become more mature quite rapidly since BGT began collecting its data in 2012). We believe that the online data market is now relatively mature and widespread across the economy, so the latest annual data for 2018 should not be subject to as much bias for these reasons as online job advert data for earlier years. This limitation does mean, however, that analysis of trends in the BGT data between 2012 and 2018 need to be treated with some caution. While we have looked at some such trend analysis, we prefer to base our forward skills demand estimates on the 2018 occupation/skills matrix rather than trying to project forward past trends in this matrix that may be misleading due to the biases noted above.

BGT has created a database tool that allows accessing in a structured way all the information extracted from the millions of job ads that it manages every year. Data mining techniques allow them to capture keywords from job ads. BGT then uses classifications of occupations and industries and different taxonomies of skills to aggregate the data. For occupations and industries, BGT uses standard ONS classifications - SOC and SIC respectively - which is helpful in linking to our net employment effect estimates defined on the same basis.[102]

The more disaggregated taxonomy of skills consists of what BGT calls ‘skills clusters’. There are some transversal skills captured such as communication or collaboration skills that the database calls baseline skills. Yet, the majority of skills captured by such skills clusters refer to what the database calls specialised skills. Examples of skills clusters include: ‘business intelligence’, ‘data analysis’, ‘graphic and visual design’ or even broader domain specialisations such as ‘cardiology’.

[101] BGT collects data on active job postings from thousands of web-pages on a daily basis. For each job posting, in addition to extracting job title, salary, education and experience requirements, Burning Glass identifies keywords from free text job descriptions. The full job descriptions are not available. Keywords include: skills, personal competences and knowledge required by employers

[102] There has been some interesting research, notably by NESTA (2018), on possible alternative occupational classifications derived from the BGT data. While this could be a useful avenue for future research, we could not follow it on this study given that our core employment effect estimates had to follow standard ONS SOC classifications for data availability reasons.

Figure A4-2: Overview of BGT job ads data

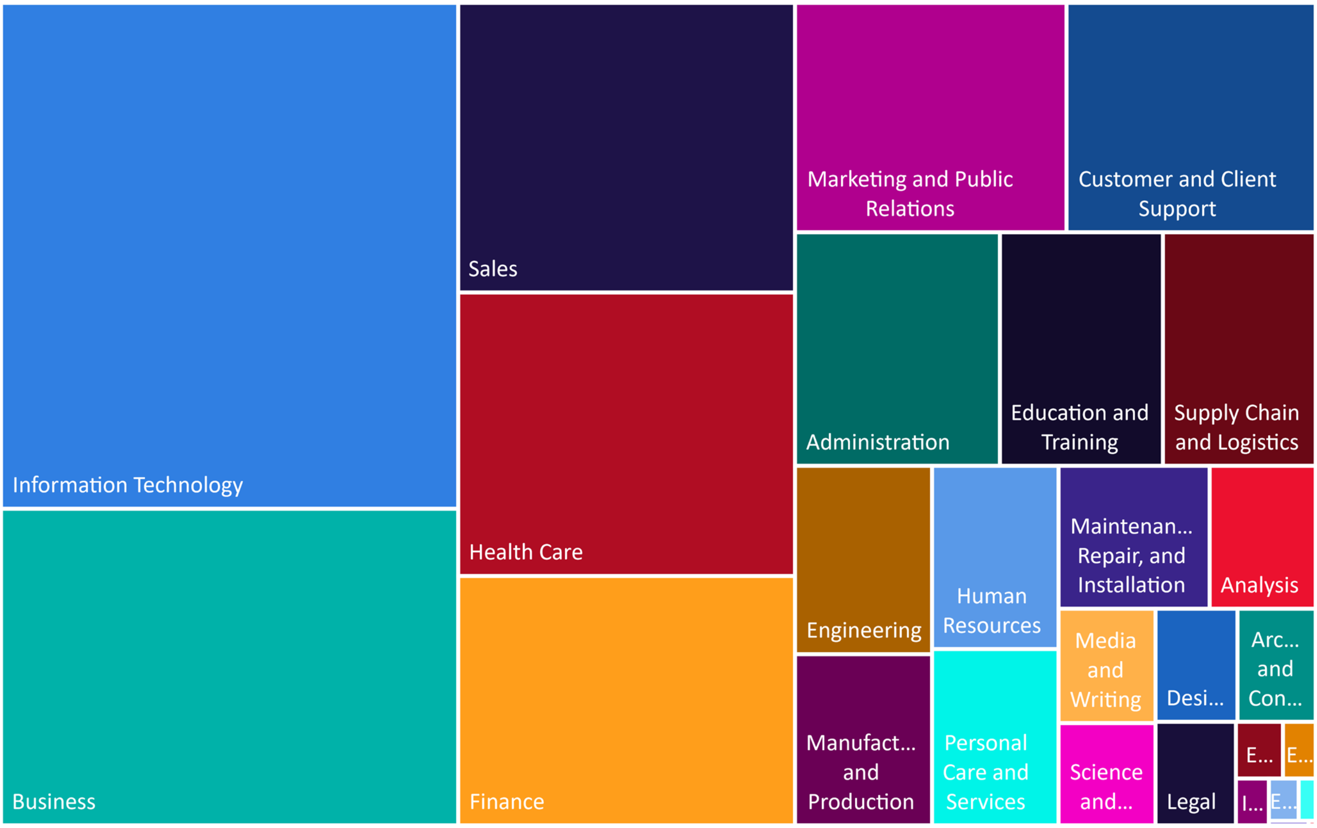

Source: PwC analysis of BGT data

BGT then has a more aggregated classification of skills. It groups the more than 500 skills clusters into 28 groups. The groups that are in most demand in the UK are 'Business', 'Information Technology', 'Health Care' and 'Sales' (see Figure A4-2). Then within each group there are some skills clusters that are in more demand than others: for instance, within the Business skills group, the skill cluster 'project management' is mentioned more times in job ads than 'customer and client support'.

These classifications of skills allow analysis of the frequency with which a skills cluster is mentioned in job ads (for specific occupations) for a given period of time. This information is the basis for creating an occupation/skills matrix, which can be presented at different levels of aggregation. For the purposes of estimating the demand for skills into the future (as presented in Section 4 of the main text), we use the most disaggregated matrix: at the SOC-4 level for occupations and at the level of the c.500 skills clusters (Figure A4-3 presents an extract of the whole matrix for illustrative purposes). Data on all job ads published in 2018 are used as the basis for the matrix.

Figure A4-3: Example of the occupation/skills matrix (base year: 2018)

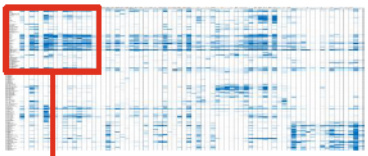

Occupation-skills matrix: each cell shows the % of job ads in each occupation (in columns) that cite the corresponding skill cluster

SOC / Skills Cluster Group: Skills Cluster	Chief executives and senior officials	Elected officers and representatives	Production managers and directors in manufacturing	Production managers and directors in construction	Production managers and directors in mining and energy
Business: Business Management	5%	1%	3%	2%	2%
Business: Business Process and Analysis	5%	5%	7%	7%	9%
Business: Business Solutions	1%	0%	1%	0%	2%
Business: Business Strategy	13%	2%	5%	3%	8%
Business: Category Management	0%	1%	0%	0%	1%
Business: Contract Management	2%	2%	2%	5%	6%
Business: Due Diligence	2%	0%	0%	0%	1%
Business: Event Planning and Management	0%	0%	1%	0%	0%
Business: Internal Controls	0%	0%	0%	0%	0%
Business: Key Performance Indicators	3%	1%	14%	3%	7%
Business: Knowledge Management	0%	0%	0%	0%	0%
Business: Operations Management	1%	1%	26%	2%	2%
Business: Optimisation	1%	0%	2%	0%	2%
Business: Order Management	0%	0%	0%	0%	2%
Business: People Management	7%	1%	12%	8%	9%
Business: Performance Management	2%	1%	4%	1%	5%
Business: Pricing Analysis	0%	0%	0%	0%	0%
Business: Process Improvement	1%	1%	3%	0%	1%
Business: Product Management	1%	0%	2%	0%	1%
Business: Project Management	12%	10%	8%	19%	20%
Business: Property Management	0%	0%	0%	0%	0%
Business: Quality Assurance and Control	2%	2%	6%	6%	3%
Business: Real Estate and Rental	0%	0%	0%	0%	0%
Business: Risk Management	10%	0%	1%	2%	4%
Business: Technical Assistance	0%	0%	0%	0%	0%
Customer and Client Support: Advanced Customer Service	1%	1%	1%	0%	1%
Customer and Client Support: Basic Customer Service	7%	5%	7%	8%	13%
Customer and Client Support: Cash Register Operation	0%	0%	0%	0%	0%
Customer and Client Support: Claims Processing	1%	0%	0%	0%	1%
Customer and Client Support: Payment Processing and Collection	0%	0%	0%	0%	0%
Design: Animation and Game Design	0%	0%	0%	0%	0%
Design: Art and Illustration	0%	0%	0%	0%	0%

Many of the skills that will be in demand in the future are already in short supply and skills gaps for many occupations are expected to increase: many of the occupations in the official list of shortages will see a net increase in demand according to our estimates of net employment effects in 10 and 20 years (Figure A4-4).

Figure A4-4: Top growing occupations and the skills they require

SOC 4	Description	Shortage list?	Net effect, (20 years, 000s)	Net effect (20 years, %)	Top demand skills (BGT)
2231	Nurses	Yes	338	49%	Medical Support; Advanced Patient Care; Basic Patient Care; Obstetrics and Gynaecology (OBGYN); Mental and Behavioural Health Specialties
7111	Sales and retail assistants	No	170	14%	General Sales; Basic Customer Service; Retail Sales; General Sales; Practices; Merchandising
6145	Care workers and home carers	No	142	18%	Basic Living Activities Support; Mental Health Diseases and Disorders; Medical Support; Food and

					Beverage Service; Mental and Behavioural Health Specialties
2211	Medical practitioners	Yes	132	46%	Surgery; Mental and Behavioural Health Specialties; Teaching; Emergency and Intensive Care; General Medicine
6141	Nursing auxiliaries and assistants	No	87	27%	Medical Support; Basic Patient Care; Advanced Patient Care; Basic Living Activities Support; Mental and Behavioural Health Specialties; Medical Support; Basic Patient Care
1190	Managers and directors in retail and wholesale	No	85	25%	Store Management; Knowledge Management; Cash Register Operation; General Sales; Performance Management; Merchandising
2136	Programmers and software development professionals	Yes	75	24%	Software Development Principles; SQL Databases and Programming; Microsoft Development Tools; JavaScript and JQuery; Java; Web Development
2139	Information technology and telecommunications professionals	Yes	58	29%	Cybersecurity; Software Development Principles; Real Estate and Rental; Test Automation; SQL Databases/ Programming; Software Quality Assurance
3562	Human resources and industrial relations officers	No	55	37%	Recruitment; General Sales; Microsoft Office and Productivity Tools; Employee Relations; Business Development; Human Resource Management and Planning
2314	Secondary education teaching professionals	Yes	55	11%	Test Administration; Physics; Chemistry; Biology; Urban Planning; Child Care
2413	Solicitors	No	54	41%	Litigation; Business Development; Customer Relationship Management (CRM); Cash Register Operation; Due Diligence; General Sales
2315	Primary and nursery education	No	49	10%	Test Administration; Child Care; Banking Services; Mental and Behavioural

	teaching professionals				Health Specialties; Medical Support; Childhood Education and Development
1131	Financial managers and directors	No	47	14%	General Lending; Cash Management; Financial Trading; Financial Management; Category Management; Property Management
1132	Marketing and sales directors	No	46	19%	General Sales; Business Development; General Marketing; Sales Management; Marketing Management; Cash Management
2133	IT specialist managers	Yes	46	22%	Property Management; Cash Register Operation; Cash Management; Performance Management; Business Solutions; Microsoft Office and Productivity Tools
1135	Human resource managers and directors	No	43	23%	Human Resource Management and Planning; Employee Relations; Property Management; Pricing Analysis; Talent Management; Performance Management; Recruitment
1136	Information technology and telecommunications directors	No	39	39%	IT Management; Cash Management; Technical Support; Business Solutions; Category Management; Performance Management
1115	Chief executives and senior officials	No	37	45%	Category Management; Property Management; Technical Assistance; Cash Management; General Accounting; Microsoft Office and Productivity Tools; Cash Register Operation
2423	Management consultants and business analysts	No	36	19%	Business Solutions; Property Management; SQL Databases and Programming; Microsoft Office and Productivity Tools; Enterprise Resource Planning (ERP); Data Mining; Cash Register Operation

1251	Property, housing and estate managers	No	32	20%	Cash Management; Quality Assurance and Control; Cash Register Operation; Facility Management and Maintenance; Performance Management; Property Management, Microsoft Office and Productivity Tools

Source: PwC analysis of BGT and ONS data

Note: Although the estimated net effect (e.g. 338,000 for nurses) should not be taken as a precise forecast given the uncertainties, it does give a clear indication of the tendency and magnitude of the growing demand for nurses as compared to other occupations.

Technical Annex 5: Scenario analysis

As we have noted throughout this report, any such estimates of future technological impacts are subject to significant margins of uncertainty. In terms of job displacement, one major source of uncertainty is the time frame over which automation will take place. Another major source of uncertainty is the extent to which job displacement will be offset by AI-induced job creation. As such, we have considered alternative job displacement and job creation scenarios.

These scenarios provide tests of whether the estimated relative net employment effects of AI for different occupations, industries, regions and demographic groups are at least directionally robust to different macroeconomic scenarios. In general, this is true for the scenarios we have considered, even if the absolute estimates of net employment effects are uncertain.

Job displacement scenarios

The alternative job displacement scenario is derived by using an alternative set of labels to train the job displacement model. The alternative set of labels are generated using a super-majority rather than a majority rule. For the supermajority rule we only label an occupation as ‘automatable’ (i.e. 1 rather than 0) if at least ⅔ of experts in our workshop agreed. For our base case scenario, by contrast, we assumed that a simple majority vote was sufficient to label an occupation as automatable. Using a supermajority rather than majority rule has the effect of slowing down the estimated speed at which automation takes hold. As such we refer to the scenarios formed from the majority and supermajority rules as the ‘base case’ and ‘slow adoption’ scenarios.

Under the slow adoption scenario almost no jobs are estimated to be displaced by AI and related automating technologies over the next 5 years.[103] This reflects the fact that only 2 out of 70 occupations were deemed by our experts as being automatable over 5 years under the supermajority rule: meter readers and switchboard operators. The total proportion of jobs estimated to be displaced by AI over 10 and 20 years in the slow adoption scenario is approximately two thirds of the proportion estimated to be displaced in the base case job displacement scenario (Figure A5-1).

[103] This is consistent with a world in which some tasks are automated, just not enough to displace human workers altogether.

Figure A5-1: Total job displacement in alternative job displacement scenarios

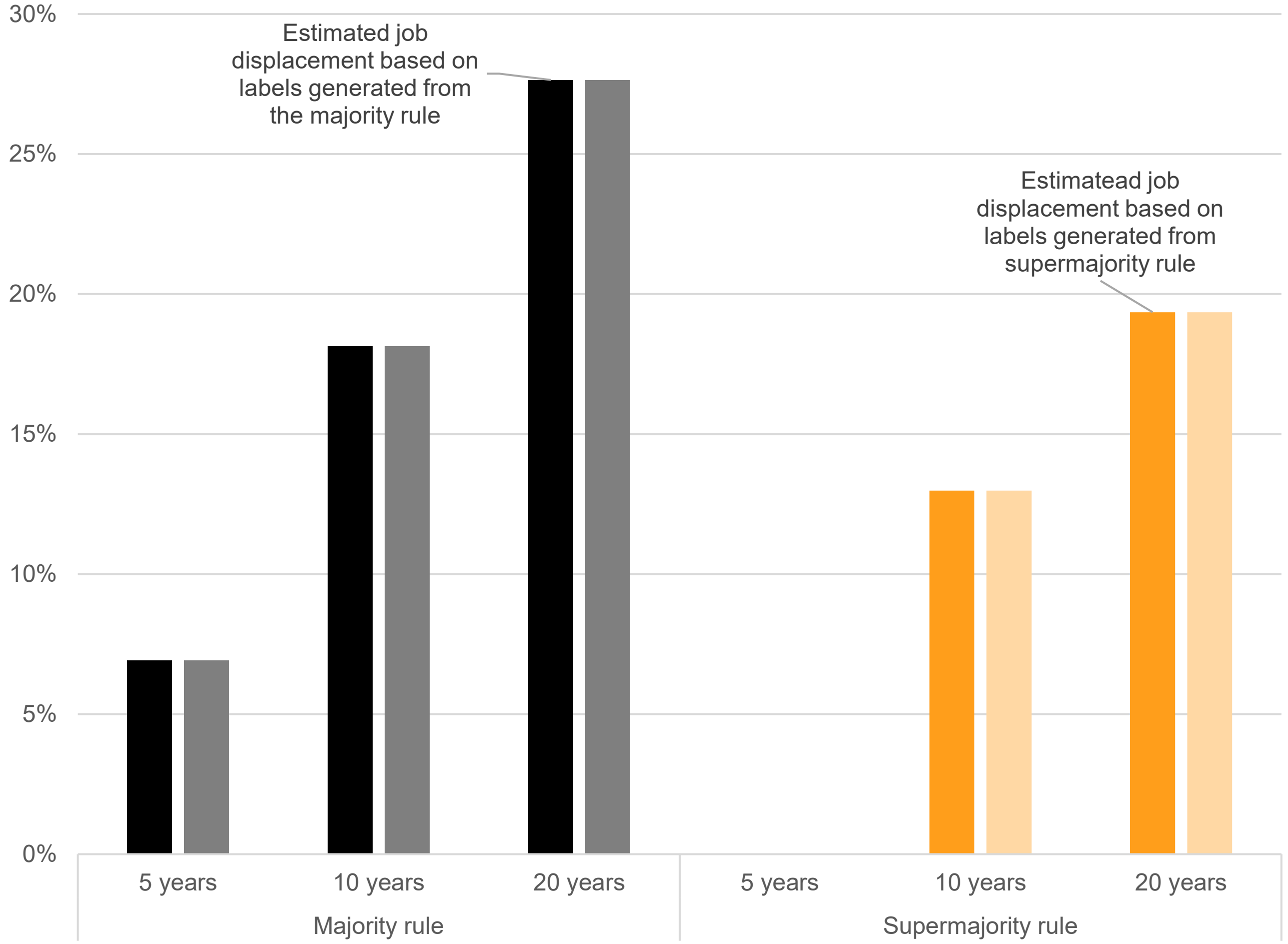

Source: PwC analysis of OECD PIAAC and ONS APS data

Below we set out a series of charts showing how the slow adoption scenario causes results to vary by occupation, region and education level. In general, as you would expect, job displacement is more muted in this scenario but so is job creation, so the net effect is not greatly changed in qualitative terms. The same kind of occupations see positive or negative net effects (Figure A5-2) and this also true for education levels (Figure A5-3) and regions (Figure A5-4). In the former case, higher education levels are still generally associated with more positive net employment effects from AI. In the latter case, net employment effects still tend to be somewhat more positive in London and the South East than in the Midlands or the North, although these regional variations are not that large in terms of net effects.

Figure A5-2: Job creation and job displacement by occupation (SOC 1) after 20 years assuming a majority vs supermajority rule

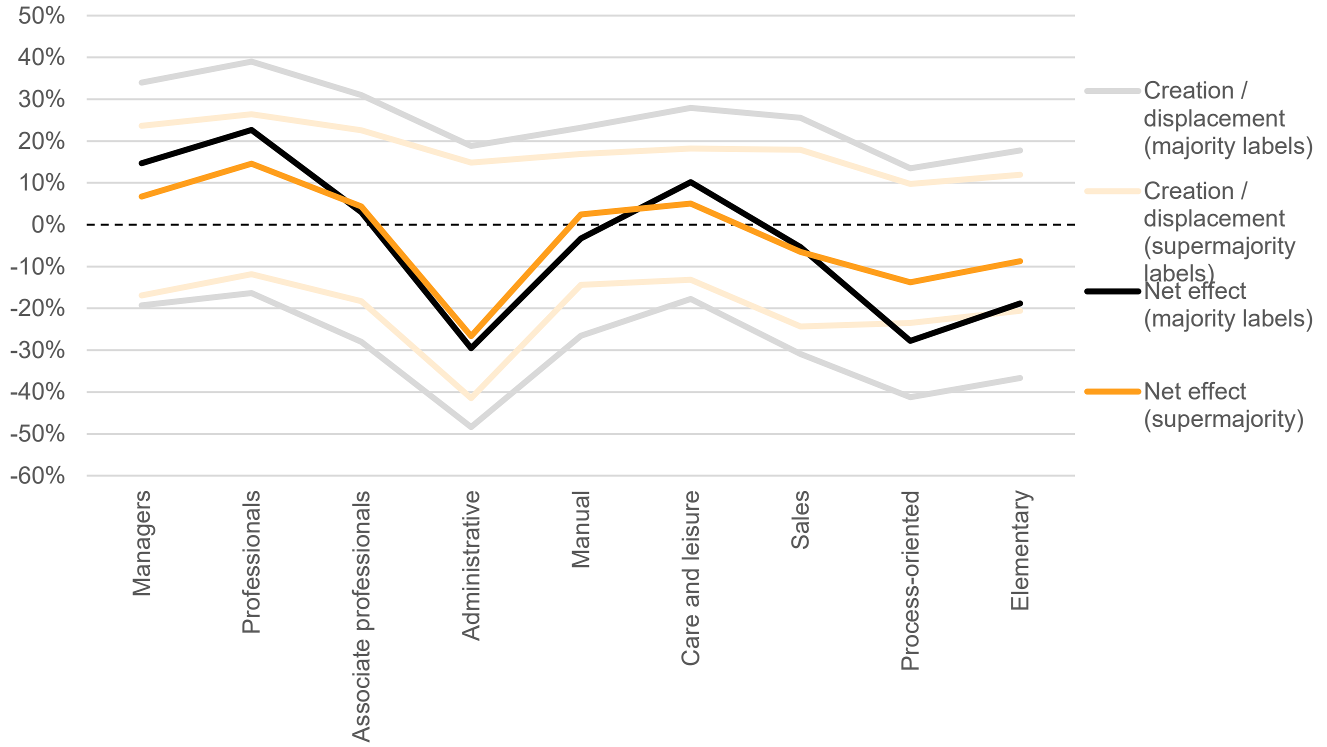

Figure A5-3: Job creation and job displacement by education level after 20 years assuming a majority vs supermajority rule

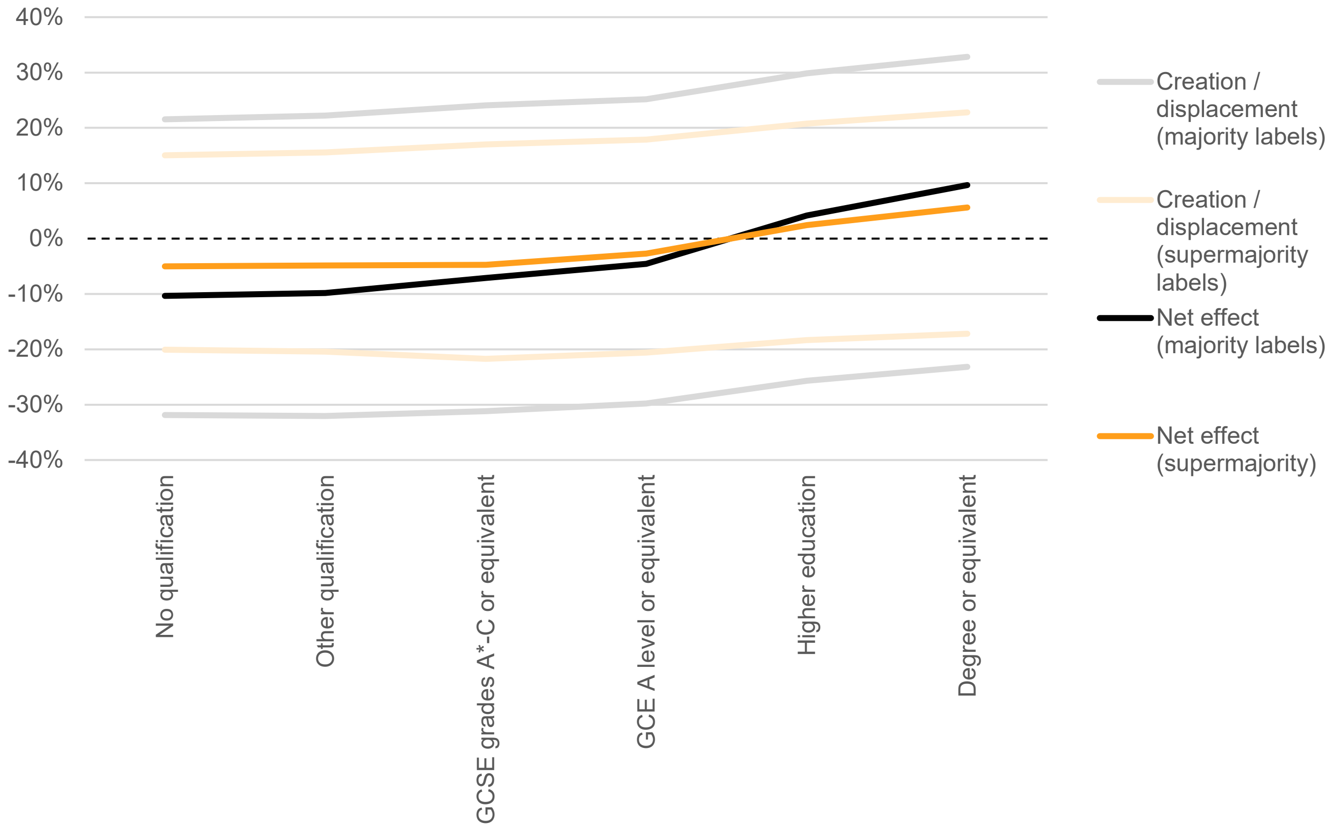

Figure A5-4: Job creation and job displacement by region (NUTS 1) after 20 years assuming a majority vs supermajority rule

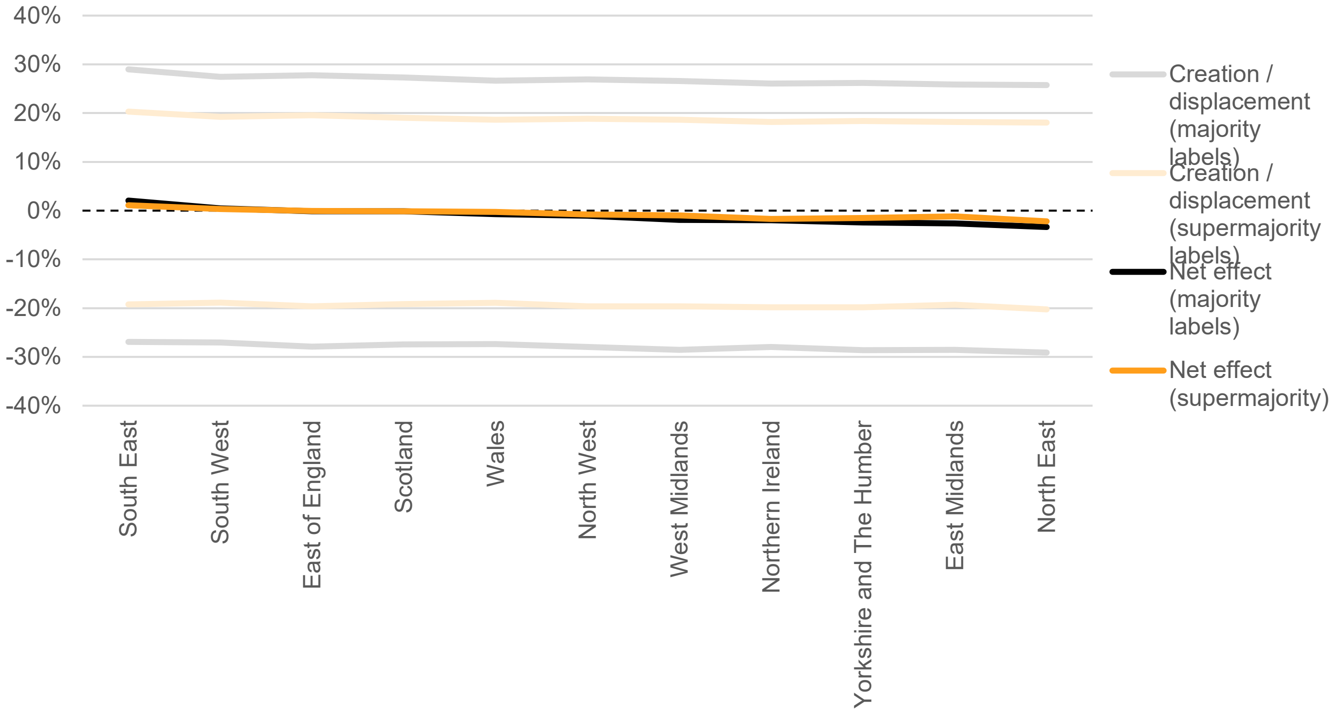

Source: PwC analysis of OECD PIAAC and ONS APS data

Overall, the alternative 'slow adoption' job displacement scenario, based on the super-majority voting rule for assigning labels to occupations based on our expert workshop results, suggests that the broad directional conclusions from our analysis remain reasonably robust.

Job Creation Scenarios

In our 'base case' scenario we assume, as discussed in Section 2 above, that the total number of jobs created by AI and related technologies is equal to the total number of jobs displaced by AI. Given the uncertainties around this assumption, we have also considered three alternative job creation scenarios alongside the base case (Scenario 1). In Scenario 2, job creation is 25% less than job displacement. In Scenario 3, job creation is 25% greater than job displacement. In Scenario 4, job creation is 40% below job displacement after 5 years, 30% below cumulative job displacement after 10 years, and catches up to the cumulative level of job displacement after 20 years. We vary our assumptions by relatively large magnitudes in these alternative scenarios to capture the potential large magnitude of the associated uncertainties.

Figure A5-5: Total job creation and job displacement in alternative job creation scenarios

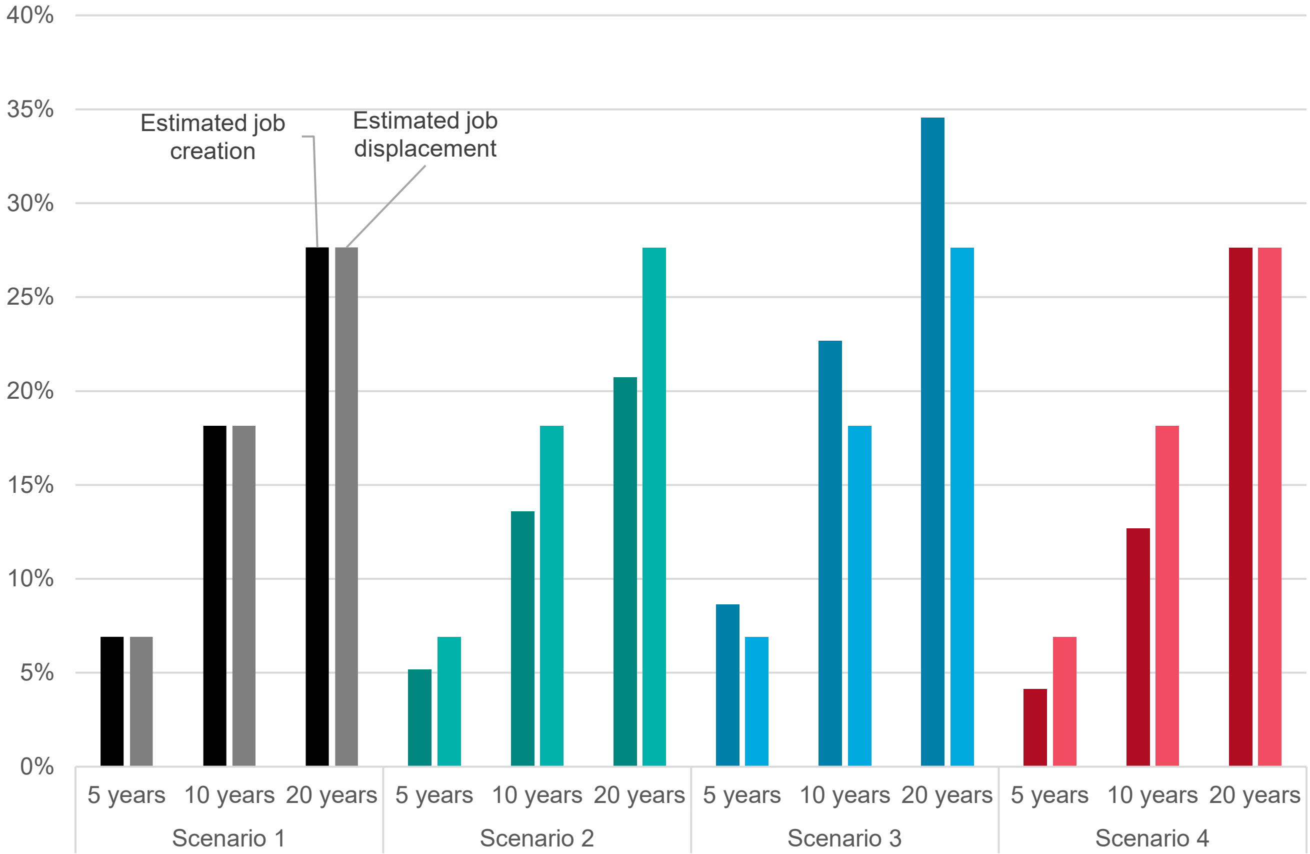

Source: PwC analysis of OECD PIAAC and ONS APS data

These scenarios account for the fact that automation will require displaced workers to retrain for new roles, existing workers to upgrade their digital skills, and real wages to adjust to incentivise the movement of capital and workers from lower to higher growth sectors and regions. These processes will take time to play out, so it is possible there could be a temporary rise in unemployment as new technologies such as AI are introduced, even if the long run employment impact is neutral.[104]

In the charts below (Figures A5-6-A5-9), we consider how the net impacts over 20 years vary across scenarios by occupation, industry, education level and region. The broad conclusion in each case is that the general pattern of results is similar in each case, even though the absolute estimates of the net effects vary in the different scenarios. This scenario analysis, as with that for job displacement above, therefore lends some support to the view that our qualitative conclusions on the relative distributional effect of AI on employment should be relatively robust, even if the precise estimates made are not.

[104] Empirical evidence for the latter is that there has been no clear upward or downward trend in the UK employment rate over the past 150 years, despite several waves of major technological change over this period.

Similar results have been derived for 5 and 10 year horizons, including for Scenario 4 (which is the same as the base case over 20 years). There are smaller absolute variations in net employment effects over these shorter time horizons, but the general pattern of relative effects is similar to the base case in the other three scenarios.

Figure A5-6: Estimated net effect of AI on occupations (SOC 1) in each job creation scenario over 20 years

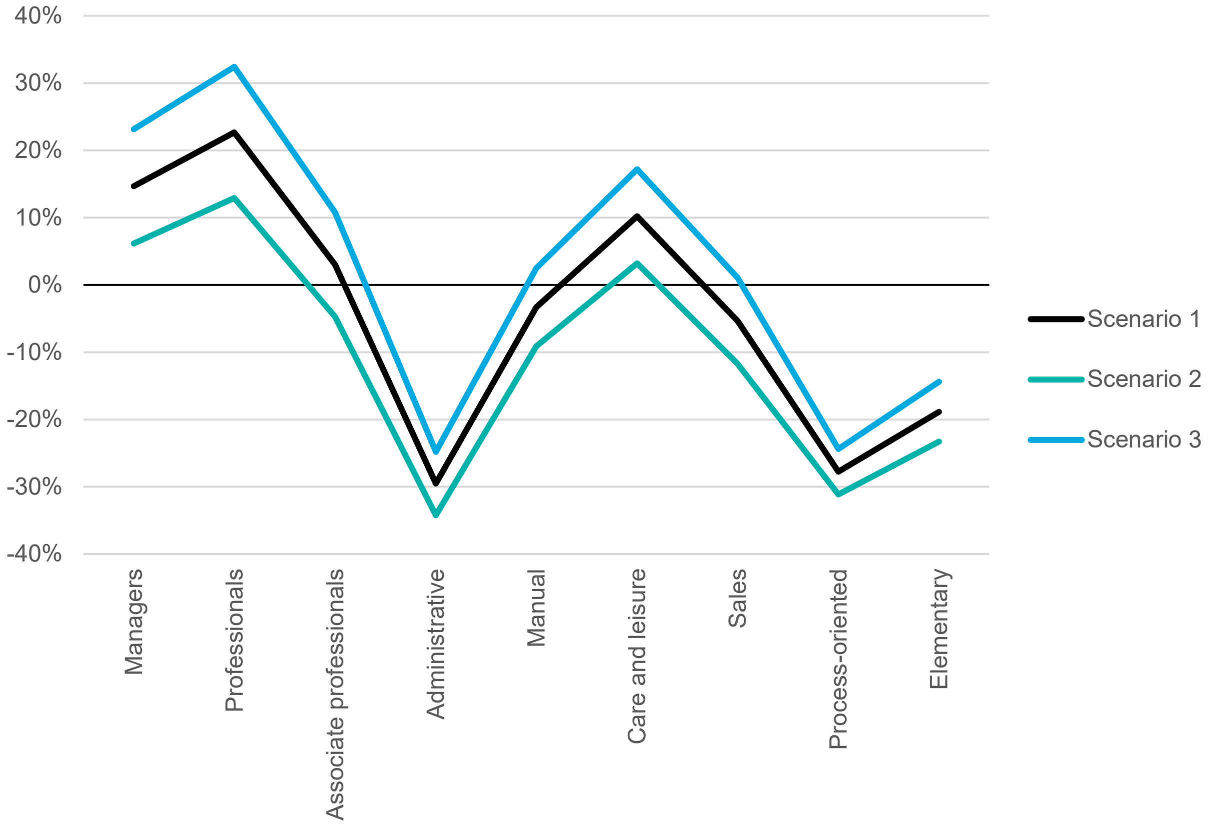

Figure A5-7: Estimated net effect of AI on industries (SOC 1) in each job creation scenario over 20 years

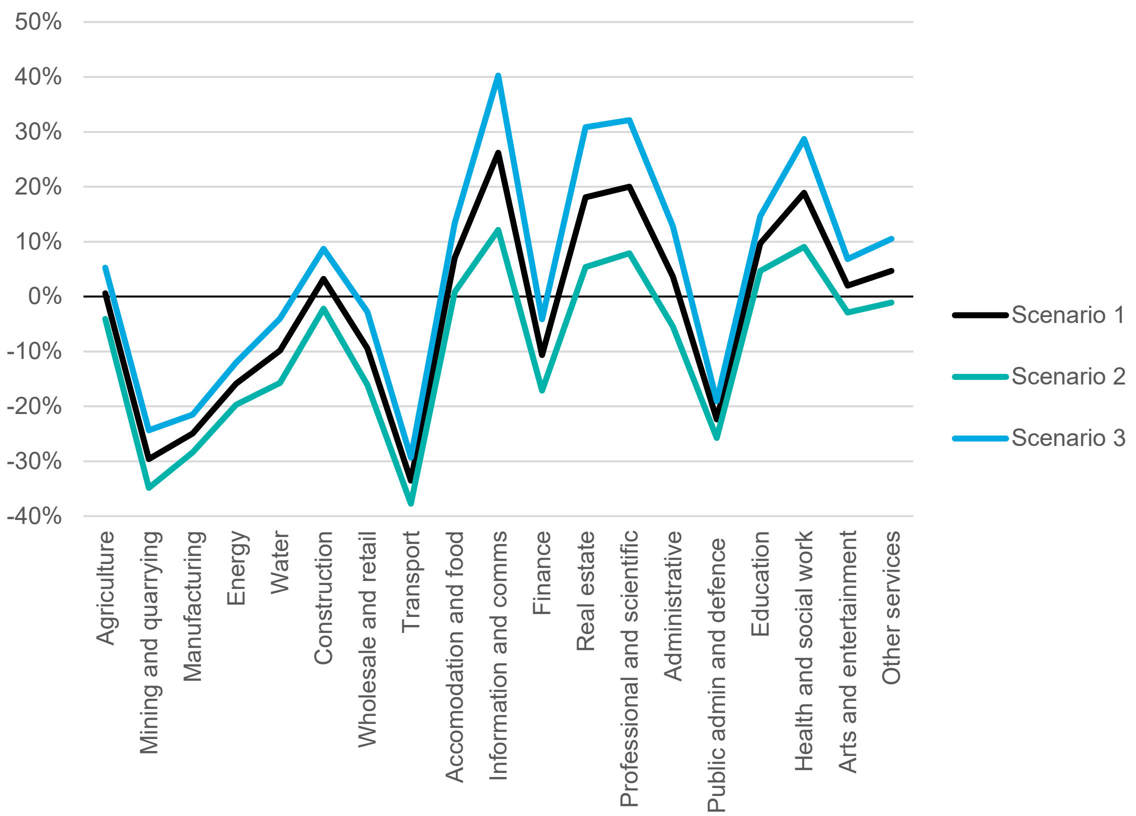

Figure A5-8: Estimated net effect of AI on different educational groups in each job creation scenario over 20 years

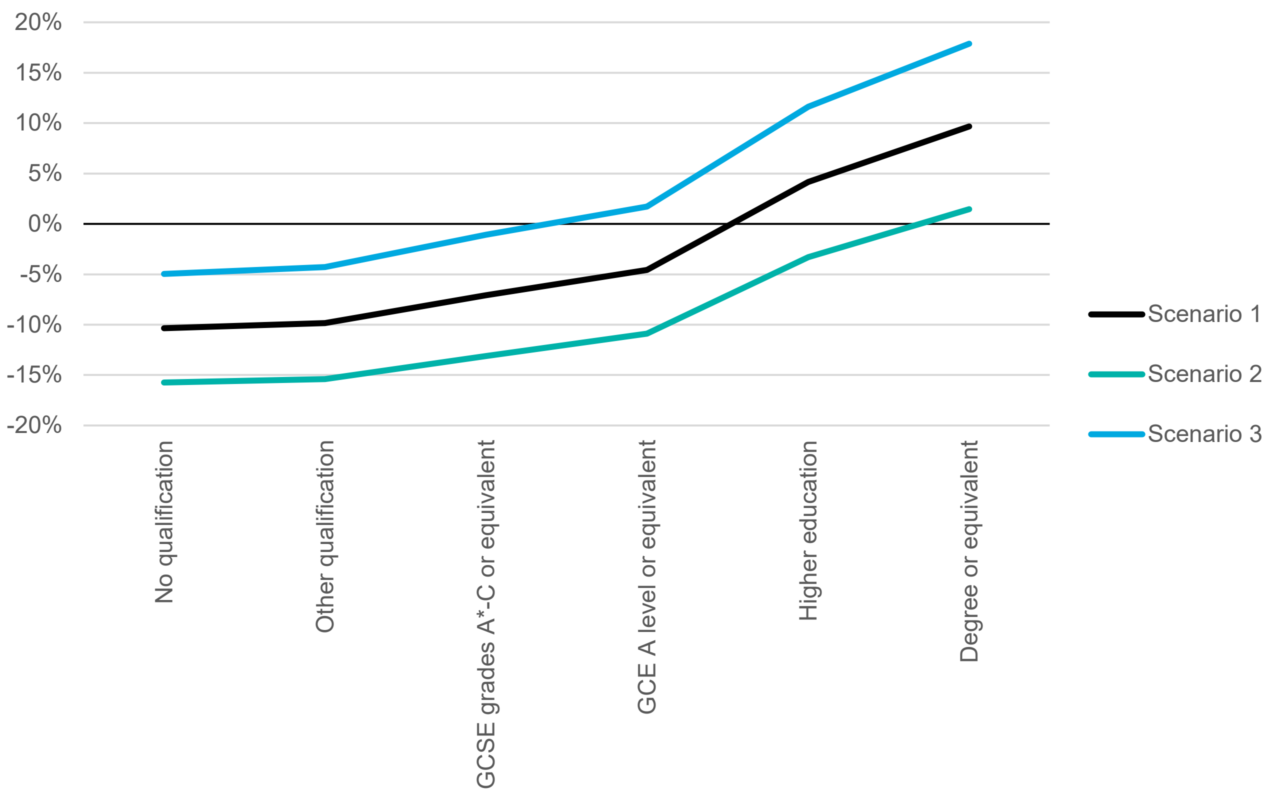

Figure A5-9: Estimated net effect of AI on regions (NUTS 1) in each job creation scenario over 20 years

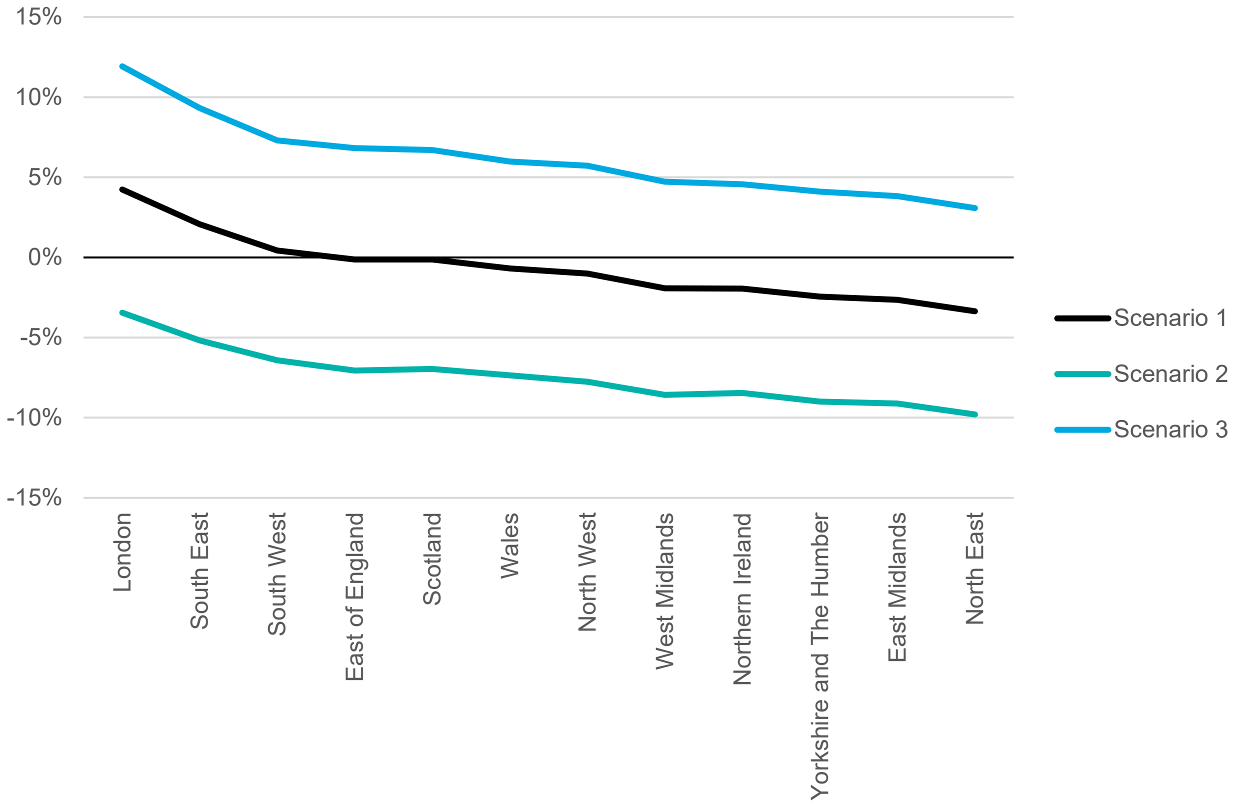

Source: PwC analysis of OECD PIAAC and ONS APS data

This publication is available from: www.gov.uk/beis

If you need a version of this document in a more accessible format, please email enquiries@beis.gov.uk. Please tell us what format you need. It will help us if you say what assistive technology you use.