The macroeconomic impact of artificial intelligence

February 2018

Contents

1. Executive summary 1
2. What is AI and how could it impact the economy? 6
3. Overview of our approach 14
4. Econometric analysis 18
5. Job automation study 26
6. AI Impact Index 31
7. S-CGE model analysis 35
8. Sensitivity analysis 56
9. Conclusion 62
10. Existing and forthcoming insights on AI insights from PwC 64

Appendices 65

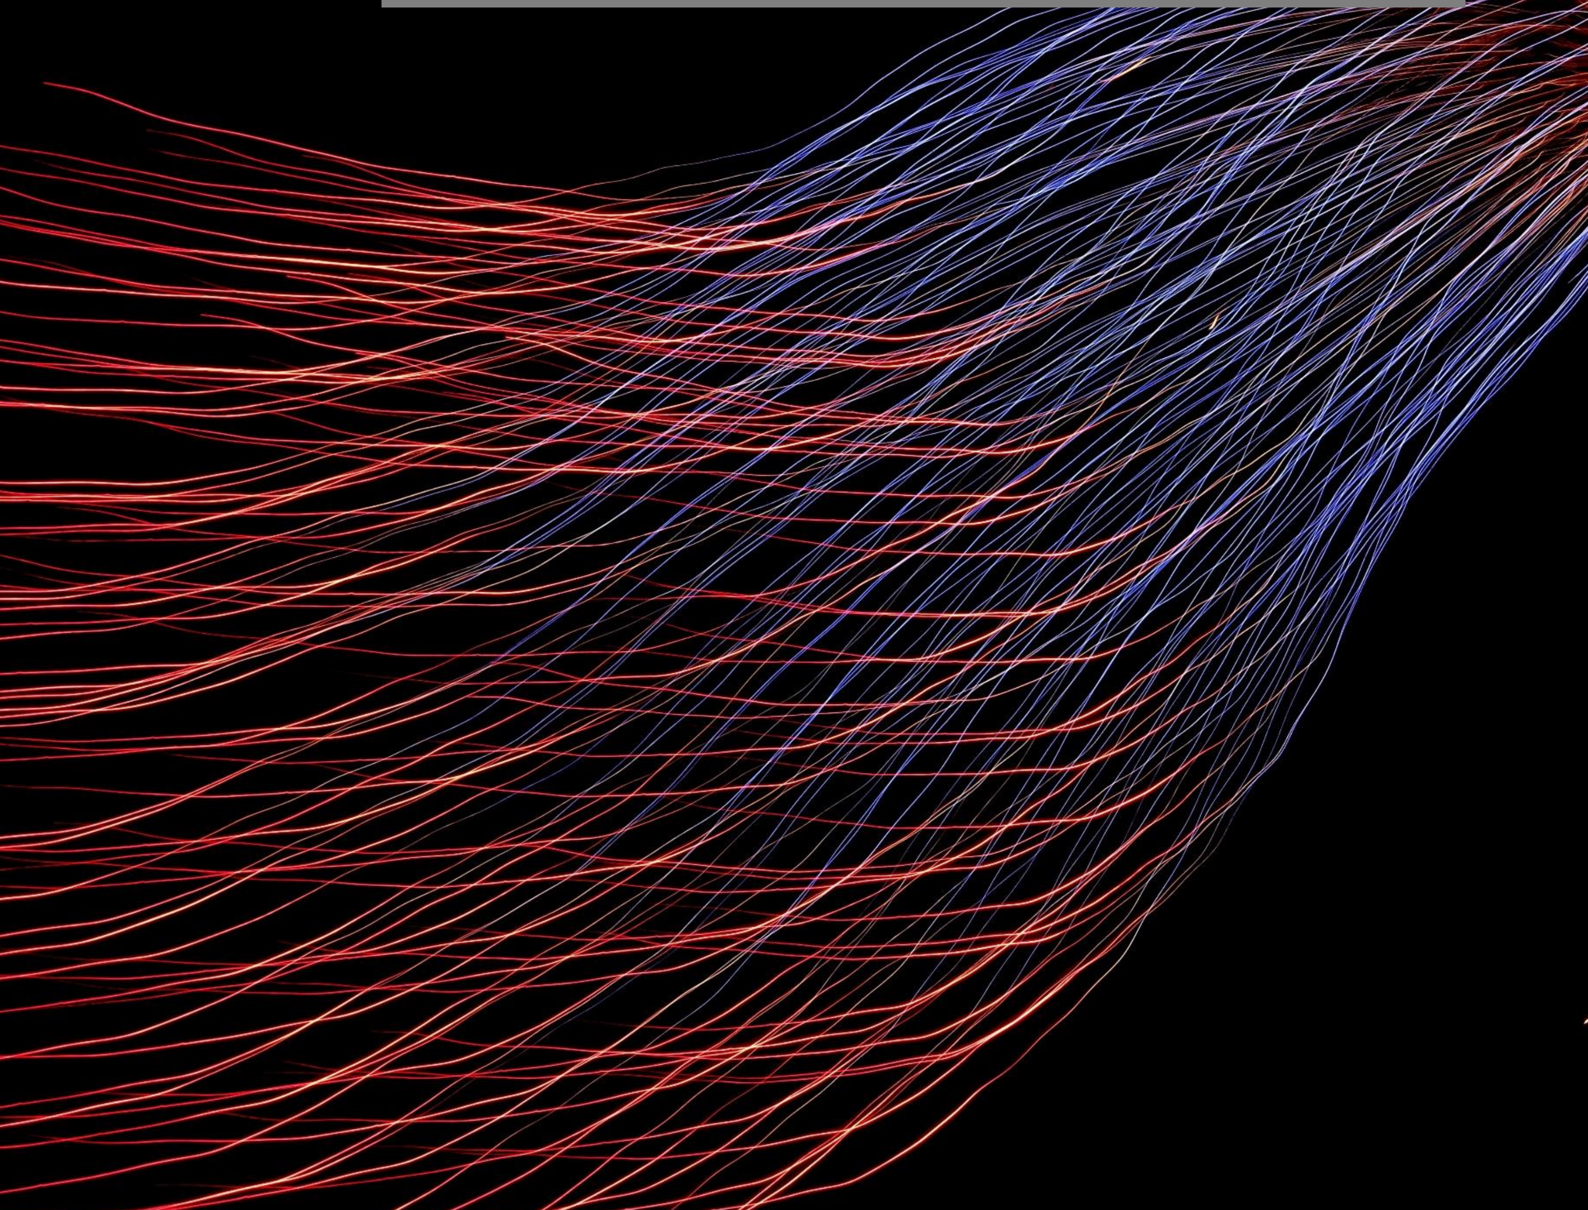

1. Executive summary

1.1. Purpose of this report

With artificial intelligence (AI) set to transform the way that we live and work, it raises the inevitable question of how much the technologies will impact businesses, consumers and the economy more generally. Employees want to know what AI means for their job and income, while businesses are asking how they can capitalise on the opportunities that AI presents and where investment should be targeted. Cutting across all these considerations is how to build AI in the responsible and transparent way needed to maintain the confidence of customers and wider stakeholders.

Traditionally, the research into the impact of AI, such as Frey and Osborne (2013) and Autor (2003), has focused on the effects on employment, as some jobs and tasks become automated and firms seek to make their business run more efficiently. More recently, some authors have focused on the benefits that could come from productivity gains associated with this automation. However, the possible benefits and opportunities of AI go much further. The ability to collect, store and analyse data at the scale, speed and in the ways facilitated by AI technologies will allow firms to improve the quality of products and tailor them to consumers, increasing their value. AI can also reduce the amount of time that consumers spend doing low-value tasks or reduce frictions in the consumption process, all leading to increased consumer demand.

We seek to provide a clearer picture of the full economic potential of AI globally, extending the exploration of AI's potential beyond the simple replacement of workers, to AI that augments the workforce and productivity. We also explore the impact of AI-driven consumption-side product enhancements on the economy, which, to our knowledge, has not yet been explored in any great detail within the AI literature.

These are the topics that we discuss and address in a series of PwC reports. In June 2017 we published our report, Sizing the prize: *What's the real value of AI for your business and how can you capitalise?*[1], which highlighted how AI can enhance and augment what enterprises can do and provides a clear and compelling case for AI investment and development. In this report, we provide detailed insight into the approach that we took to complete the analysis of the global economic impact, as well as a more in-depth look at the results of our analysis and an exploration of the robustness of those results. Our research has already created numerous insights into AI's possibilities and potential impacts in different sectors and regions. We are using these insights to help our clients leverage them in an effective, efficient and intelligent way – helping distinguish them from their competition and ensure they are ready for the age of AI.

Our study not only captures AI's impact through more channels than previously covered in a single study, but also presents detailed findings on both the geographic and sectoral distribution of these results. We look at which regions are set to gain more or less, and also examine in detail how the different AI-driven impacts on the economy will unfold over time in practice. Whilst we place a lot of confidence in the robustness of our main scenario figures, we have also run a number of scenario-based sensitivity tests to test not only whether our least robust assumptions had a material impact on the results, but also to provide insight into a world where AI impacts the economy differently due to heterogeneous possible business and consumer responses to AI's introduction. The following subsections discuss the research methodology, key findings and sensitivity tests we conducted as part of this study.

---

[1] https://www.pwc.com/gx/en/issues/data-and-analytics/publications/artificial-intelligence-study.html. This study formed part of a series of PwC publications related to artificial intelligence which address the key questions discussed in this section.

1. What AI means for jobs and income – https://www.pwc.co.uk/economic-services/ukeo/pwcukeo-section-4-automation-march-2017-v2.pdf.
2. How to build AI in a responsible and transparent way – https://www.pwc.co.uk/services/audit-assurance/insights/responsible-ai-how-to-build-trust-and-confidence.html.
3. How businesses can capitalise on the opportunities – https://www.strategy-business.com/article/A-Strategists-Guide-to-Artificial-Intelligence?gko=0abb5.

1.2. Research methodology

Estimating the global economic impact of AI is a pressing and challenging task. As a result, we have used a complex, multi-stage modelling approach that focussed on leveraging the global PwC network's resources and research capabilities to provide sufficiently accurate and convincing answers to such a compelling topic.

Our approach follows three stages. Figure 1.1 summarises each of these stages and the types of analysis that underpin each.

- The first two stages involved undertaking primary research aimed at estimating the relationship between AI and both consumer products and firm productivity, before combining these with estimates of the pace and profile of AI adoption to evaluate the direct impacts of AI on each of these two elements.
  - We used existing PwC research into the potential for AI-driven job automation as well as new econometric analysis assessing the relationship between AI and labour productivity to identify the key drivers of productivity growth, understanding where AI fits into this picture and specifying models capable of picking up the true causal effect of AI on productivity.
  - PwC's AI Impact Index, developed by our AI experts in partnership with Fraunhofer, scored different product lines according to five key criteria especially developed to evaluate AI's impact on products. In particular, the research captured the potential that AI has to improve the quality of products, the potential for products in an industry to be more personalised and the amount of time that consumers could save from using AI.
- The final stage focussed on bringing the analysis together and converting these results into AI-driven 'inputs' into our Spatial Computable General Equilibrium (S-CGE) model[2] – a dynamic model of the global economy that we have used to estimate the net global impact of AI on the economy up until 2030. Beyond the initial impacts on productivity, job displacement and consumer choice, these net effects account for secondary impacts such as the creation of new jobs, increased consumer demand (from attractive goods first and more affordable goods second), the increased supply of labour to the market, and trade flow patterns.

Figure 1.1: Our multi-stage approach to assessing the total economic impact of AI

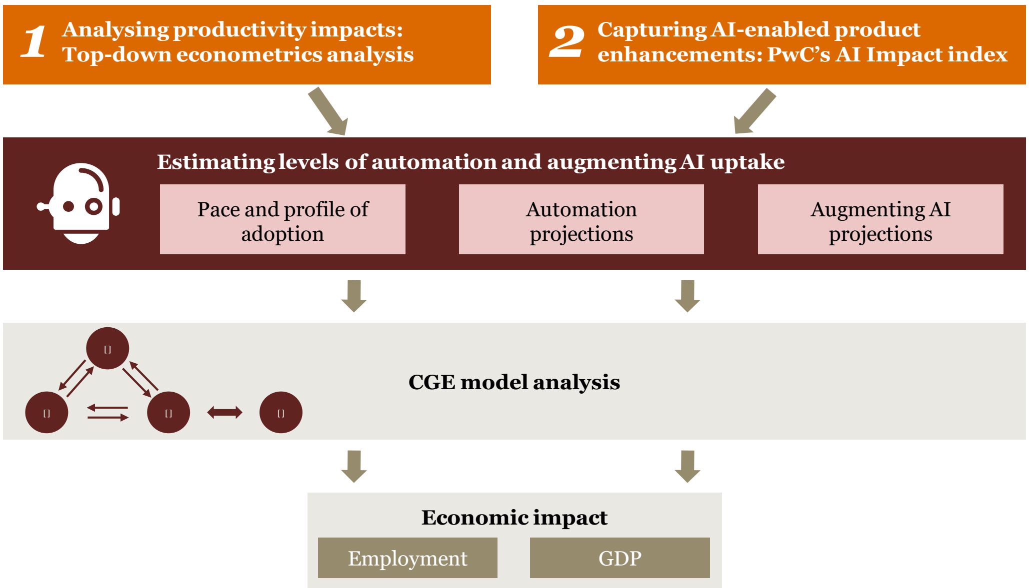

Source: PwC Analysis

[2] The S-CGE model is a dynamic, computable general equilibrium model, which models economic interactions between different players in the economy, namely: firms, households, and the government. The 'general equilibrium' nature of the model means that it represents a closed system which tracks flows of resources from one area or player to another (i.e. there is natural accounting within the model). The model captures a number of complexities of the real world economy including, but not limited to: household expectations about the economy and its development, passive government policy, household utility optimisation, trade flows between sectors within and across countries (based on historic data), and investment patterns within and between countries.

1.3. Key findings

- **Global economic impact:** Global GDP is estimated to have been approximately $75 trillion in 2016[3]. Our baseline projections suggest that that figure is estimated to be approximately $114 trillion by 2030. Our S-CGE model analysis suggests that global GDP could be up to 14% higher than this figure in 2030 as a result of AI – the equivalent of up to $15.7 trillion.
  - The economic impact of AI will be driven by (a) productivity gains from businesses automating processes as well as augmenting their existing labour force with AI technologies (assisted, autonomous and augmented intelligence) and (b) increased consumer demand resulting from the availability of personalised and/or higher-quality AI-enhanced products and services.
  - We estimate that approximately 58% of the 2030 GDP impact will come from consumption side impacts, or $9.1tn of additional GDP. However, over the entire period 2017-2030, approximately 55% of the GDP impact will be due to productivity increases. This is reflective of the faster (total) transmission mechanism on the production side of the economy, as the consumption-side GDP effects rely more heavily on the more delayed, indirect effect of dynamic firm entry which increases the supply of personalised, high quality AI-augmented products and makes these goods more affordable.

Figure 1.2 – Global GDP impact by effect of AI in main scenario

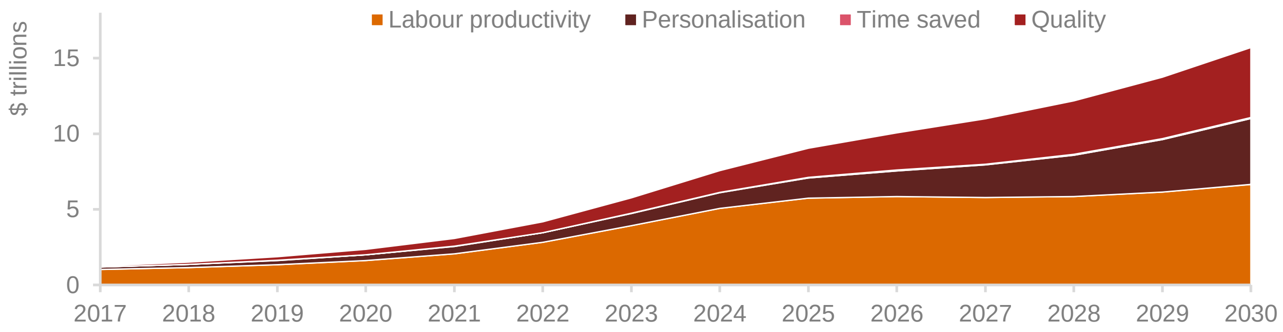

Source: PwC Analysis

- **Geographical impacts:** North America and China stand to see the biggest economic gains in percentage terms from AI.
  - In North America, readiness for adoption and a high fractional increase in replacement AI are the largest drivers of their economic impact, reflecting the region's leading stance on AI and its implementation, as well as a high automation potential that is expected to occur at a regional level between now and 2030.
  - In China, a high estimated marginal impact of AI-uptake on productivity is the largest driver, likely reflecting the region's low starting labour productivity base relative to other countries and the recent push to implement both replacement and augmenting AI technologies on a large scale. The altogether marginally less competitive landscape in China also enhances the consumption-side effects more than in other regions, as new firm entry supplying the market with AI-enhanced products has a resultantly larger downward impact on prices.
  - As a result, AI will specifically enhance GDP by 26.1% (China) and 14.5% (North America) in 2030, equivalent to a total of $10.7 trillion and accounting for almost 70% of the global impact.
- **Sectoral impacts:** Economic gains from AI will be experienced by all sectors of the economy, with each industry expected to see a gain in GDP of at least approximately 10% by 2030.

[3] According to data from the Global Trade Analysis Project (GTAP) - https://www.gtap.agecon.purdue.edu/.

- The services industry that encompasses health, education, public services and recreation stands to gain the most (21%), with retail and wholesale trade as well as accommodation and food services also expected to see a large boost (15%).
- Transport and logistics as well as financial and professional services will also see significant but smaller GDP gains by 2030 as a result of AI (10%), but with financial services benefiting particularly quickly in the short term.

- **Impact on labour:** We estimate that 326m jobs will be impacted by AI in 2030 and also find supporting evidence of skills-biased technological change. As with our analysis of GDP impacts, the impact on labour doesn't necessarily signify the new jobs that will be created as a direct result of AI, but the number of jobs that will come to depend on and be heavily impacted by AI. Interestingly, most of these jobs will be unskilled, however proportionally, skilled jobs will be more positively impacted, supporting a bias towards skilled labour. Specifically,67% of the jobs in 2030 that will depend on AI will be unskilled jobs, however unskilled labour accounts for 69% of jobs in the baseline scenario.

1.4. Sensitivity testing

We have embarked on a number of scenario-based sensitivity tests to provide further confidence around our results. In particular, we focus on the more assumptions-driven part of our approach and also vary our assumptions over the uptake time-scale and profile to reflect practical barriers to AI-uptake and a different event sequence over the next 10 years.

- **Sensitivity 1 – Pace of adoption:** We calculated an alternative scenario with a slower pace of AI adoption whereby the level of AI adoption originally expected by 2030 in each geographical region was actually to only be reached by 2040. We conceive of this scenario as representing a 'slow-uptake' world scenario, where inertia and other prohibitive factors slow down rate of AI-uptake.
- **Sensitivity 2 – Shape of the S-curve of AI adoption:** We have developed an alternative scenario for AI adoption during the period to 2030 which yields the same amount of AI adoption in aggregate but alters the profile of how this technology is introduced over time. Specifically, the profile in this scenario remains an 'S-shape' but with a later and steeper upturn in the middle of the period. This reflects a scenario whereby somewhat less AI technology is adopted in the short-term but in the medium-term significant advances take place rapidly in a 'delayed revolution.'
- **Sensitivity 3 – Quantification of our product quality model input:** We have used insights from the academic literature, economic theory and the insights of subject matter experts from our economics practice to quantify the outputs of the AI Impact Index into the consumption-side model inputs for the S-CGE analysis. Translating the results with respect to product quality proved the most challenging due to the lack of quantitative literature applicable. In this alternative scenario we supposed that the impact of AI on utility was 35% smaller than in our main scenario. This scenario is best interpreted as examining the global AI impact in a world where consumers value the benefits to AI-enhanced product quality much less than we anticipated.

Sensitivity results

In all our scenarios, the minimum global economic impact of AI in 2030 is expected to be $11.2tn – 9.8% of global GDP in 2030. This $11.2tn is calculated from our 'slow uptake' 2040 scenario, where the world takes 10 years longer to reach the level of AI uptake (and impact) predicted in 2030 within our main scenario.

In our other sensitivity scenarios, the global economic impact of AI in 2030 is predicted to be $14.2tn (12.5% of GDP) in our 'delayed revolution' scenario and $15.2tn (13.3% of GDP) in our scenario where product quality impacts are considerably smaller.

The impact of these sensitivity scenarios on both the geographic and sectoral distribution of impacts is also negligible, ensuring our main scenario results at a more granular sector and region specific level are valid under a number of alternate outcomes.

Positioning our results

Our economic model results are compared to a baseline of long-term steady state economic growth. The baseline is constructed from three key elements: population growth, growth in the capital stock and technological change. The assumed baseline rate of technological change is based on average historical trends.

Therefore, since AI has already been introduced prior to the starting evaluation period of this study, the component of these forecasts driven by technological change will already have factored in past trends in AI's GDP impact. As a result, it is difficult to quantify the exact fraction of AI's GDP impact that will be additional to historical average growth rates (i.e. additional to the baseline forecast).

However, our study is specifically focussed on the AI technologies that are yet to be implemented and are conceived to be implemented between 2017 and 2030. As a result, an underlying but reasonable assumption we make here is that the scale and impact of these AI technologies will be above the current trend in AI's uptake and impact. Under this premise, our study is centred on estimating the total marginal economic impact of yet-to-be-implemented AI specifically between 2017 and 2030 – not including the AI that has already been implemented prior to this study (which is implicitly included in the baseline growth assumption). This also means that, whilst our study results imply that average economic growth rates between 2017 and 2030 will be raised due to AI's impact, we do not make claims outside of this time interval. As a result, we do not interpret that AI will impact the fundamental long run growth rate of the global economy.

Finally, our results estimate the upwards pressure on GDP as a result of AI only, under the *ceteris paribus* assumption.[4] Our results may not be directly reflected in future economic growth figures, as there will be many positive or negative forces that either amplify or cancel out the potential effects of AI (e.g. shifts in global trade policy, financial booms and busts, major commodity price changes, geopolitical shocks etc.).

4 This implies that all other factors of the economy remain 'as expected' and do not suffer shocks which could deviate the economy from its predicted outcome in this study.

2. What is AI and how could it impact the economy?

Artificial intelligence (AI) is a rapidly growing market, with revenues expected to reach almost $50 billion by 2020[5]. It is poised to have a transformative effect on consumer, enterprise, and government markets around the world.

PwC released Bot. Me[6], a report that was constructed from insights gathered at a 2017 AI Expert Salon, and a survey of 2,500 U.S. consumers and business decision makers on attitudes towards AI and its current and future implications on society. Key findings include that AI is no longer viewed as primarily a corporate tool to increase automation, but instead as an emerging technology to be leveraged to handle global challenges – 63% of consumers agree that AI will help solve complex problems that plague modern society, such as closing the education gap, establishing cures for cancer and other diseases, and even gender inequality challenges.

	Human in the loop	No human in the loop
Hardwired/ specific systems	***Assisted intelligence*** AI systems that assist humans in making decisions or taking actions. Hard-wired systems that do not learn from their interactions.	***Automation*** Automation of manual and cognitive tasks, both routine and non-routine. This does not involve new ways of doing things – automates existing tasks.
Adaptive systems	***Augmented intelligence*** AI systems that augment human decision making and continuously learn from their interactions with humans and the environment.	***Autonomous intelligence*** AI systems that can adapt to different situations and can act autonomously without human assistance.

Source: PwC Analysis

2.1. What is AI?

In our broad definition, AI is a collective term for computer systems that can sense their environment, think, and in some cases learn, and take action in response to what they're sensing and their objectives. Forms of AI in use today include: digital assistants, deep question and answering, machine vision and many others. AI works in four ways as illustrated above. As humans and machines collaborate more closely, and AI innovations come out of the research lab and into the mainstream, the transformational possibilities are staggering.

5 https://www.idc.com/getdoc.jsp?containerId=prUS42439617

6 https://www.pwc.com/us/en/press-releases/assets/img/bot-me.pdf.

Figure 2.1 – AI areas by layer

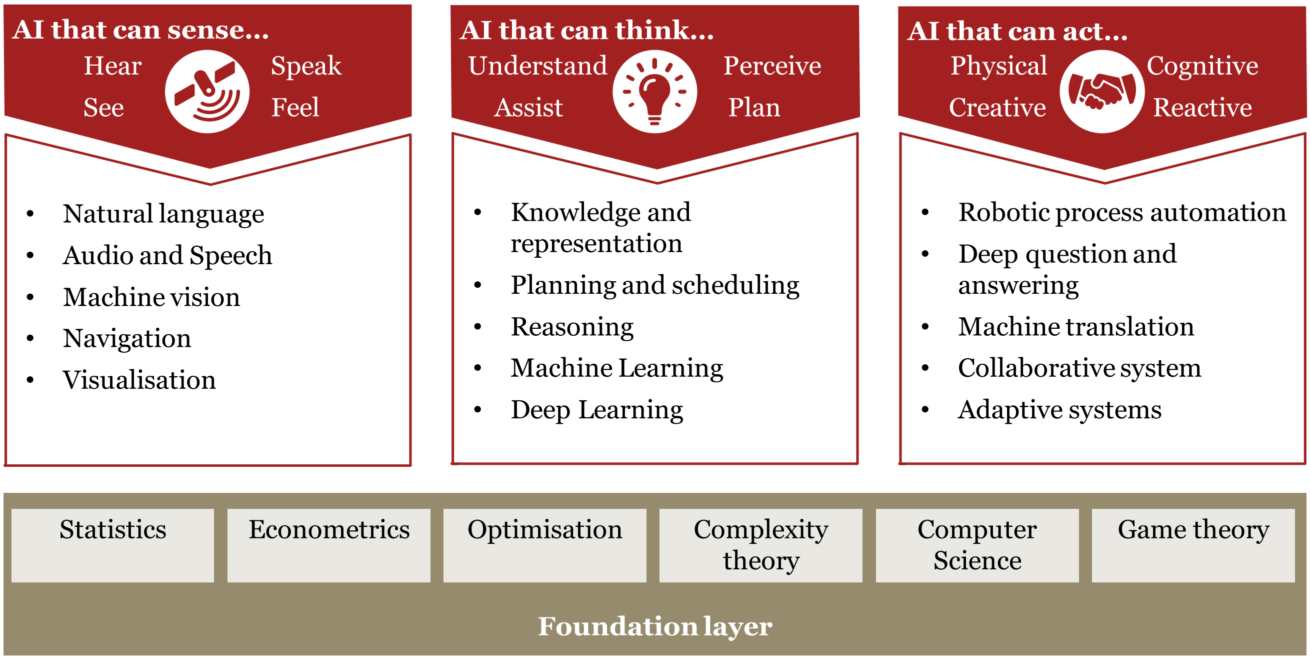

Source: PwC Analysis

2.2. Overview of AI and the economy

With AI set to transform the way that we live and work, it raises the inevitable question of how much the technologies will impact businesses, consumers and the economy more generally. Traditionally, research that seeks to quantify these impacts has been focussed on the effects on jobs and productivity as firms seek to make their businesses run more efficiently at lower costs. In this section we summarise this evidence base and report a range of estimates which vary significantly according to the scope of the different studies (e.g. timescale under consideration, definition of AI, consideration of consumer impacts or not).

2.2.1. Existing public studies on the economic impact of AI

A study published by Analysis Group in 2016[7] and funded by Facebook, considers the effects on jobs and productivity as two separate streams of impact: direct effects on GDP growth and indirect effects. They consider that the direct impacts of AI on GDP come from increased revenues and employment in firms and sectors that develop or manufacture AI technologies. Secondary 'indirect' impacts will come from other sectors employing some AI technologies that could make their processes and decisions more efficient as well as increase access to information. They conclude that over the next 10 years, a reasonable range of economic impact would be $1.49tn to $2.95tn, an average of $149bn to $295bn a year. While these may seem like significant amounts, in a global economy of about $74 trillion currently, this translates to much less than a 1% increase.

The Executive Office of the President in the United States also published a report in 2016[8] focussing on the economic impacts of automation driven by AI. The report recognises that technology traditionally increases productivity through the reduction of the amount of manpower needed to create the same amount of output and that increases in productivity typically lead to wage increases that benefit consumers' living standards.

The authors expect that the impact of AI on the economy will follow a similar channel but that these effects will not be felt uniformly. AI will automate some tasks and jobs, but the authors believe that it is not yet clear which jobs will be most affected. However, research tends to find that lower-skilled jobs that include routine manual and cognitive tasks tend to be the most likely to be automated. This research does not provide an estimate of the

[7] Chen, N. Christensen, L. Gallagher, K. Mate, R. & Rafert, G. (2016), 'Global economic impacts associated with artificial intelligence', Analysis Group.

[8] https://obamawhitehouse.archives.gov/sites/whitehouse.gov/files/documents/Artificial-Intelligence-Automation-Economy.PDF

costs or benefits but points out that despite the potential threat to some jobs, technology and innovation has improved quality of life and created jobs. It also points out the US has adapted a number of times to changes in the workplace and job market, which in the long-run has yielded large benefits. It recognises that these could be the same for new AI technologies.

More recently, as examined in Section 1, both Accenture and McKinsey have aimed to estimate the impact of AI on the global economy going forward (to 2035 and 2050+ respectively). Both studies focus on productivity-led effects of AI, but take different philosophical approaches towards their positioning on AI's economic impact. Whilst McKinsey infer that AI will reinforce global long run growth trends as the next key source of economic growth potential, Accenture believe AI's effect will take long run economic growth rates firmly 'above and beyond' our current forecasts. Details on the methodologies, key messages and assumptions of each of these studies is presented in Section 8 of this study, alongside our own for comparison

In light of the research discussed above, it is clear that AI is likely to have some significant impacts on both jobs and productivity. Our study aims to take further steps towards capturing the full economic potential of AI and the opportunities that it presents. In addition to the more traditionally examined productivity channel, we seek to provide a clearer picture of the full economic potential of AI globally, extending the exploration of the potential of AI beyond the simple replacement of workers to automation that augments the workforce and productivity. We also include a new area of focus on the consumption side of the economy, by evaluating the GDP impact of enhancements to products resulting from AI, how consumers respond to these and how they proliferate through the economy.

Table 2.1 – Comparison of other studies assessing the impact of artificial intelligence

	PwC: Sizing the prize	McKinsey Global Institute: A future that works	Accenture: Why AI is the future of growth
Key messages	• Global GDP will be 14% higher in 2030 as a result of AI. Our analysis focusses on the period to 2030 and we do not conclude that the long-run growth rate of the global economy will fundamentally shift. • North America GDP will be 14.5% ($3.7tn) higher in 2030. • Productivity impact accounts for 6.7%. • Consumption impact accounts for 7.9%. • Direct labour productivity impacts of 4-57% depending on country (North America 57%). These are only the direct impact of AI and further impact could occur in the general equilibrium. • Total GDP impact by channel of impact for North America (productivity only): replacement (5.7%) and augmentation (0.8%).	• In the not-so-distant future, without an acceleration in productivity growth, there will not be enough workers for countries to meet their aspirations for growth in GDP per capita. • Automation could help serve as a new productivity engine for the global economy, bridging that economic growth gap. • In the next 50 years, automation will increase economic growth by 0.8 to 1.4%. • They do not make a claim as to whether this sits above or as part of the 1.8% productivity forecast based on historical productivity growth.	• AI will redefine 'the new normal' as a period of high and long-lasting economic growth. AI has the potential to be not just another driver of TFP, but an entirely new factor of production. • Growth rates will be doubled by 2035. US growth rate will be 4.6% with AI, instead of 2.6%. • US GDP will be $8.3tn/35% higher in 2035. • AI-induced productivity impact: 3.8% of US GDP • Additional AI-induced growth impact: 31% of US GDP. • Boosts in labour productivity of 11-37% in 2035 depending on country (US 35%). • Total GDP impact by channel of impact for the US: TFP impact (3.8%) Intelligent automation (16.9%) Augmentation (14.1%).

	PwC: Sizing the prize	McKinsey Global Institute: A future that works	Accenture: Why AI is the future of growth
Channels of impact	• Productivity (replacement and augmentation). • Consumption.	• Only labour substitution gains (other gains in the form of improved quality, fewer breakdowns etc. would come on top).	• Intelligent automation • Labour and capital augmentation. • Innovation diffusion.
Approach and assumptions	• Computable general equilibrium model developed by Adam Blake. • Initial, direct employment effects of AI will fall in line with PwC job automation predictions before general equilibrium effects affect job creation and net job effects. • Realised rate of adoption and realisation of potential using Global Innovation Index. • Labour productivity growth determined through econometric panel-data models per region of factors of production, AI technologies and human capital supply.	• Based country-level GDP projections on their Global Growth Model. • Complete their own analysis of technical automation potential using World Bank and US BoLS O*Net database. • Bass diffusion model to determine adoption. • Assume that labour displaced would re-join the workforce and be as productive as in 2014.	• Modified growth model developed by Robin Hanson, George Mason University. • Assume that employment will be constant in the long-term. • AI substitution is assumed to achieve 50% of its technological potential. • Determine county-by-country adoption using a measure of 'national absorptive capacity'.

Source: PwC Analysis, MGI[9] and Accenture[10] reports

2.3. Productivity channel and automation

As the research above highlights, a firm's productivity – how much they can produce using a given level of input – is expected to be heavily impacted by AI (like most emerging technologies of the past). There are applications for AI across the value chain, using various types of AI discussed in our definition above. Although there are many different AI-types as outlined in Section 1, there are two main strategies for introducing and applying these technologies:

1. 'Human-in-the-loop' technologies: investing in software, systems and machines that 'assist' or 'augment' the workforce, helping them to perform their tasks better and more efficiently and freeing up their time to focus on more stimulating and higher value-adding activities; or
2. 'No-human-in-the-loop' technologies: automating processes with robotics or other technologies, or creating autonomous agents, removing the labour input altogether.

In reality, it is likely that many businesses will implement a combination of autonomous intelligence and technologies that include humans 'in-the-loop' and yield benefits across the entire value chain from generating insights in R&D to generating better quality outputs with greater accuracy to enhancing consumer engagement. Table 2.2 below details the impact that AI can have at each stage of a firm's value chain and illustrates specific examples across various industry sectors.

[9] https://www.mckinsey.com/global-themes/digital-disruption/harnessing-automation-for-a-future-that-works

[10] https://www.accenture.com/gb-en/insight-artificial-intelligence-future-growth

Table 2.2 – Applications and the impact of AI on productivity along the value chain

Value Chain Element	*Impact of AI*	*Examples*
Strategy, business model, products and services The 'brains' of a company's operations, decision making about offerings, pricing and go-to-market strategy.	Reducing the risk, time and capital expended in the process of moving from strategy to execution.	• Simulating market conditions for production forecasts and pricing strategy. • Creating digital mock-ups of product features based on historically successful features/user preferences.
R&D and innovation Discovery of new information and trends.	Reducing the runway required before insights are generated.	• Drug repositioning – scanning scientific and clinical research data to identify other uses for drugs already approved.
Purchasing and production Sourcing raw materials and manufacturing.	More output or better quality output using fewer resources.	• Robotics automating assembly lines. • On-demand manufacturing: adjusting to produce goods based on order specifics or turning on/off autonomously.
Supply chain and logistics Getting production resources from A to B and getting the final product to the customer.	Reducing the time and resources required in these processes.	• Auto-ordering raw materials based on sales patterns and known lead/production times. • Routing emergency vehicles to hospitals based on case criticality, staffing, expertise, traffic and patient load.
Marketing, sales and customer service Increasing customer engagement and conversion of customers.	Reducing the information asymmetry between producer and consumer and tailoring messaging accordingly.	• Personalised recommendations of products and services. • AI chatbot customer service agents. • Call centre emotion detection and sales practice monitoring.
Enabling functions (finance, IT, risk) Back-office supporting activities.	Reducing costs and reducing risks with better planning and forecasting.	• Adverse event monitoring in pharmaceuticals (trends in doctor visits, social media reporting etc.).

Source: PwC Analysis.

Once firms adopt AI, the transmission mechanism from the firm to consumer is complex and is outlined below in Figure 2.2. Generally, the first stages of the effect of increased productivity result in increases in factor payments following increases in factor efficiency, which stimulates income to the owners of these factors. However, it is important to note that in each part of the economy there are winners and losers of these initial stages. For firms that supply intermediate inputs, the winners will be those firms that supply inputs to sectors which see the largest AI impact on productivity, and where productivity enhancements do not allow substitution away from their inputs towards a cheaper alternative.

For households, the picture is slightly less straightforward. Households which own capital stand to benefit relative to non-capital owners, and those capital owners in high AI-productivity impact sectors will see a particularly large income rise as the rent price of capital increases in line with its productivity. Complicating this picture is the other key source of income: labour income. As AI is uptaken, part of the labour force in each sector becomes automated. This leads to a divergence in labour income between those remaining part of the labour force in increasingly productive jobs, and those made redundant. However, those unemployed will likely find work in new areas of the economy as new AI-augmented jobs are created, so this does not necessarily translate through to increased inequality over time.

It is important to note that once consumption rises (in aggregate) as a result of this collective income and firm profits increase, this stimulates dynamic firm entry, which, in turn, increases consumption further through income effects on consumers. These effects continue to reiterate themselves until the economy reaches its new long-run equilibrium. This secondary, indirect channel is outlined in more detail within the consumption channel transmission mechanism in Section 2.4.

Figure 2.2 – Transmission mechanism of firm-side AI-driven productivity impacts through to consumption and GDP

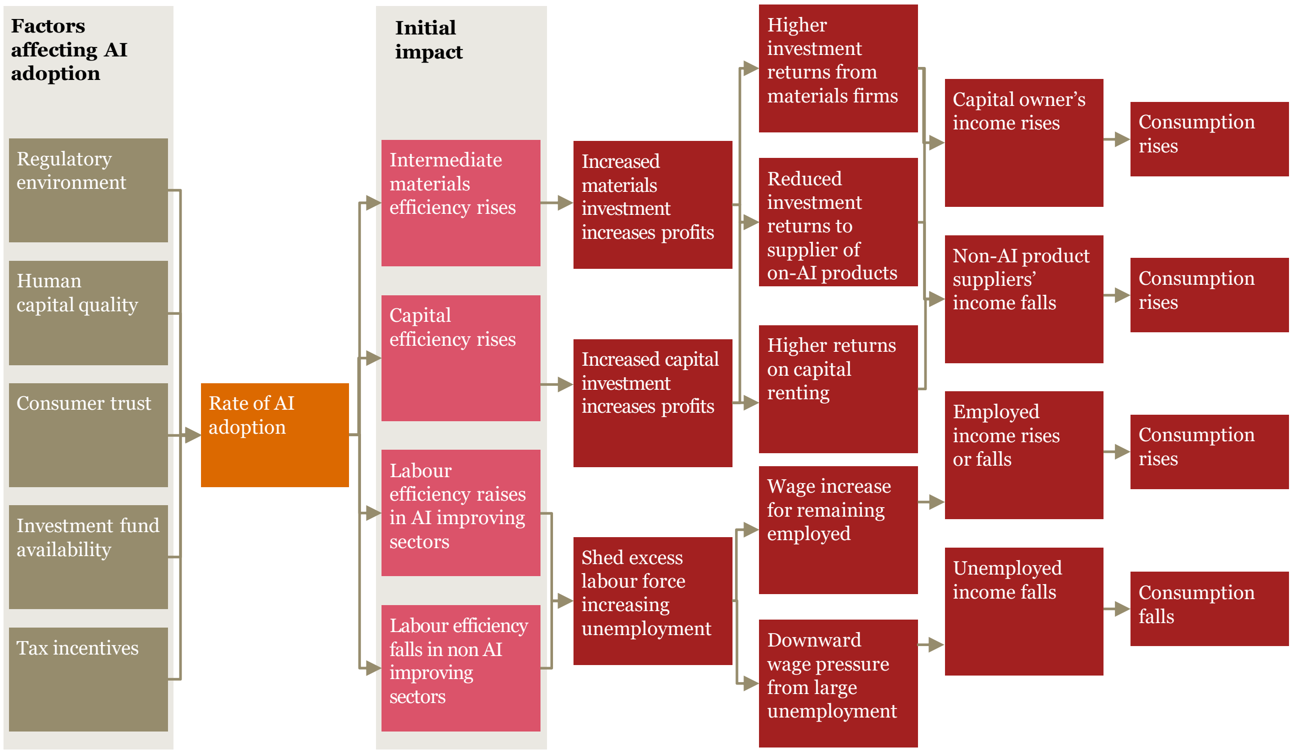

Source: PwC Analysis

2.4. Consumption-side channels

Although the most commonly discussed and well understood channel for AI to impact on the economy is through firms using AI technologies to make things more efficiently (increasing their productivity and stimulating a multiplier in the economy), less well understood but equally important is the potential for firms to use AI to enhance consumer products and services. The ability to collect, store and analyse data at the scale, speed and in the ways facilitated by AI technologies will allow firms to improve the quality of products and tailor them to consumers, increasing their inherent value. AI can also reduce the amount of time that consumers spend doing low-value tasks or reduce frictions in the consumption process, all leading to increased consumer demand.

As our study reveals, these product enhancements are expected to have a large impact on GDP. This is partially because AI could increase consumer spending on more attractive – i.e. more personalised and higher quality –) products (direct, substitution effects), but is most importantly the result of additional firms entering the market following the stimulation in consumer demand supplying new and enhanced AI products, leading to higher supplies of production and more affordable goods (indirect, income effects). We outline the key three product channels of focus below, before exploring their respective transmission mechanisms.

Product personalisation

The world has seen the introduction of some personalised products from book and film recommendations to custom homes and haute couture fashion but these tend to be expensive and challenging, or even risky for firms to produce. The level of personalisation offered by producers depends on many factors, but one critical determinant is the degree of visibility companies have into the true preferences of their consumers. Guessing incorrectly on these preferences can have huge ramifications for producers (the cost to design, ramp up production, and sell these goods for more narrowly defined populations is no small feat), which is one reason why we see very low levels of personalisation in some product categories.

AI enables better and more efficient discovery of consumer preferences by gathering more data points across a customer journey and drawing conclusions about drivers of consumption. In financial services, for example, with the help of AI institutions can access a wider range of more non-traditional data sources and determine which are most important with respect to creditworthiness. A loan package can then be developed with terms that will be appealing to the specific applicant.

Other recent technology developments, particularly in data storage, have facilitated this enhanced data collection, but AI, as put by SAP, is the '*missing link between massive data stores, mass personalisation and the ability to win in today's business environment*.'[11] AI is the key to being able to tailor products and services to customer at scale. Demis Hassabis, CEO of Deep Mind said that '*personalisation cannot be scaled without machine learning algorithms and AI embedded within all aspects of the customer experience*.' AI technologies enable lower risk test-and-learn methods and use the insights to recommend or even design a product best suited to the marketplace, with a lower likelihood of failure and cost to prototype. If firms become better equipped to know customers' preferences and produce affordable offerings made-to-measure in cost-effective ways, consumers will be more likely to consume.

Increased product personalisation can be conceived to impact consumers in two ways. First, more personalised products can increase the marginal utility of consumption for a given product. Second, more personalised products can increase the 'real' variety available to consumers, since previously homogenous goods become heterogeneous in nature due to differential personalisation features. In this study we focus on the second conception but ensure the scale of impact captures the improvements in marginal utility as well.

Product quality

AI, as well as facilitating the personalisation of products, can also simply increase the inherent value of product or service to consumers. This can manifest either through expanded scope of functionality of a given product or through the scope remaining the same but the quality, or utility yielded from the given scope increasing. For example, AI is already facilitating the provision of better TV and music recommendations through popular on-demand services. Looking forwards, AI technologies could help in film production and in identifying the optimal combination of performers or storyline[12]. AI technologies could analyse existing preferences, sift through scripts and assist in resource planning and creative direction to produce films that are most likely to draw audiences. Beyond increment value increases, at the very extreme, AI could mean the difference between life and death through enhanced accuracy or safety of products or services in healthcare, transport etc. As products' scope of functionality is expanded, or as the level of criticality of a product increases, we can expect consumption to increase.

Time

AI and many AI-enabled products could save consumers time and lead to increased consumption. Gartner predicts that, in 2018, half a billion users will save two hours a day as a result of AI-powered tools[13]. This could manifest itself in a variety of ways. Firstly, AI technologies could reduce the search costs, or general effort involved in identifying the ideal product or service, therefore reducing friction in the purchasing process leading to more consumption or greater utility derived from existing consumption. For example, the more a beauty brand knows about a shopper's skin tone, complexion, age, and lifestyle through a combination of image analysis and data mining, the better it can effectively offer products that will please the customer.

From another perspective, AI assists consumers to avoid having to perform non-value adding tasks, which also frees up their time, some of which is likely to be spent consuming additional products or services. For example, autonomous vehicles which do not require the driver to actively engage in operating the vehicle yield extra time for the driver to undertake an alternative activity as they move from A to B. This could include working, consuming media or entertainment content, shopping online or using a telecommunications platform to communicate with family or friends.

---

[11] Distilling Data – Machine Learning and the Promise of AI in Consumer Products.

[12] As one example, AI has been used to storyboard in production companies: https://www.dama.org/content/production-company-launches-first-film-use-ai-storyboarding-kickstarter

[13] http://www.gartner.com/smarterwithgartner/ai-gives-customers-a-valuable-resource-time.

As highlighted by the previous example, not only could saving time cause consumers to demand more products and services to consume in their increased spare time, but some of that time could be used for additional work. As with all of our time not spent sleeping, we can choose to offer our time to the labour market in return for a wage, or we can use the time for leisure. Though the final result will depend on decisions on the part of firms too, it is likely that time saved as a result of AI would increase, to some extent, both working time and leisure time.

A full transmission mechanism from these product enhancements to consumer impacts is outlined below including the indirect mechanism driven by the dynamic entry of firms in response to initial demand increases. We have included this secondary channel here due to its importance in contributing to the product enhancement driven GDP impact in our study.

Figure 2.3 – Transmission mechanism of consumer-side AI-driven product enhancements through to consumption and GDP

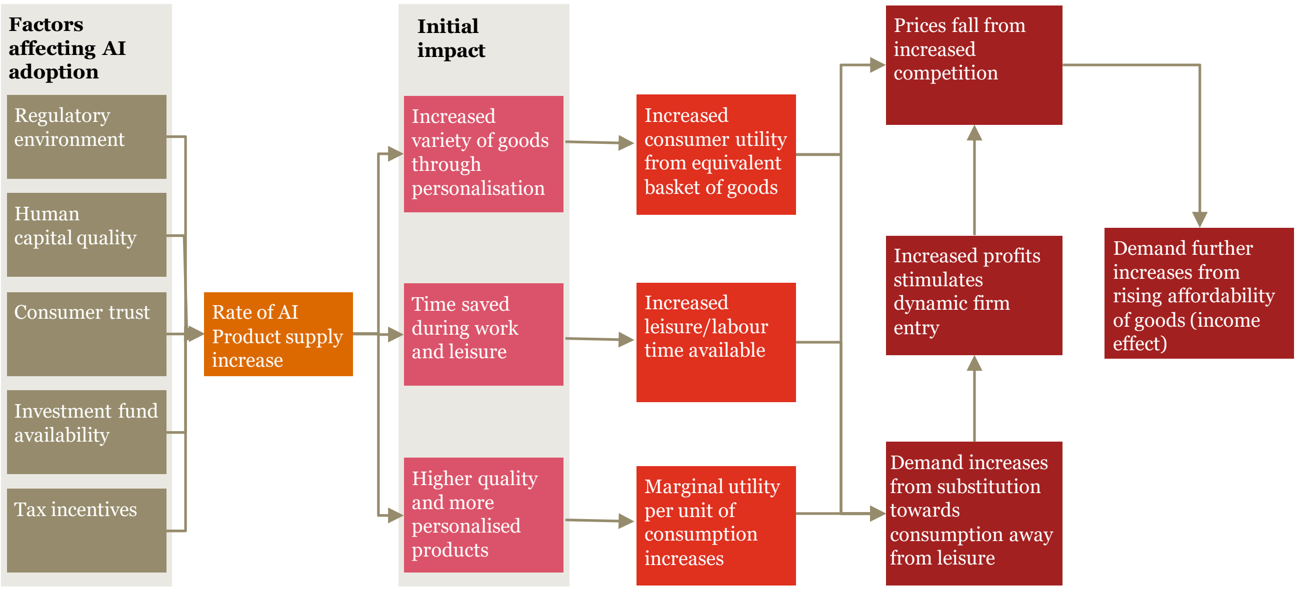

Source: PwC Analysis

As is depicted, increases in both the quality and the variety of goods first stimulate demand by making goods relatively more attractive than leisure (the substitution effect). However, without real incomes rising substantially these consumption impacts can only have a limited impact on GDP. The bulk of the GDP impact therefore instead comes from the dynamic entry of firms in response to rising profits. This increased competition and stimulation in supply leads to downwards pressure on goods prices, which in turn stimulates demand considerably further as consumers have more disposable income over time to spend on these more attractive goods (the income effect).

3. Overview of our approach

3.1. Section overview

In this section we outline the approach that we have used to derive our estimates of the economic impact of artificial intelligence (AI) on the global economy. First, we provide an overview of the approach that we have taken. We then provide more detail on each stage of the analysis including data used, specific methodology and intermediary results that led to our final estimates.

3.2. Our modelling approach

Our approach seeks to go beyond the previous literature focused on solely on productivity and job losses through automation as highlighted in Section 2. Our analysis quantifies the total economic impact (as measured by GDP) of AI on the global economy via both productivity gains and consumption-side product enhancements over the period 2017-2030. We have therefore used a dynamic economic model (known technically Spatial Computable General Equilibrium Model[14]) of the global economy which facilitates the evaluation of the net impact of each channel of AI's impact on GDP and the economy as a whole. As mentioned in Section 2, our definition of AI is broad and includes a number of emerging robotics and smart technologies that will be labour augmenting or replacing. Using a general equilibrium model to capture the net impact provides a comprehensive picture of the economic impact capturing further rounds of impacts as well as the effect of displacement within the economy. This analysis and approach, to our knowledge at the time of publication has not yet been applied to assessing the total economic impact of AI.

In order to determine what to input into the economic model, our team conducted an ambitious, dual-phased top-down and bottom-up analysis of the way in which AI could improve worker productivity and enhance consumer experiences. In addition to drawing on input from our extensive network of clients, and sector and functional advisors within PwC, we've been working with our partners at Forbes and Fraunhofer[15], a global leader in emerging technology research and development. The remainder of this Section describes the different stages in this approach.

3.2.1. Step 1: Capturing product enhancements

PwC's experts in AI, together with Fraunhofer, set out to identify the most compelling examples of potential AI applications across the value chain of 8 different industry sectors, and designed a framework to assess the degree and pace of impact of each. In total, we identified and rated nearly 300 use cases, which are captured in our AI Impact Index.

Using these use cases, the AI Impact Index is able to score different sectors, subsectors and product lines according to five key criteria especially developed to evaluate AI's impact on products. In particular, the research captured the potential that AI has to improve the quality of products, the potential for products in an industry to be more personalised, the amount of time that consumers could save from using AI, the consistency of products as AI technologies 'level the playing field,' and the improvements to data availability enabling producers to react to new consumer preferences. These criteria and their scores provided together the basis for us to quantify the consumption-side product enhancement impact of AI.

---

[14] The S-CGE model is a dynamic, computable general equilibrium model, which models economic interactions between different players in the economy – namely firms, households, and the government. The 'general equilibrium' nature of the model means that it represents a closed system which tracks flows of resources from one area or player to another (i.e. there is natural accounting within the model). The model captures a number of complexities of the real world economy including, but not limited to: household expectations about the economy and its development, passive government policy, general consumer optimisation, trade flows between sectors within and across countries (based on historic data), and investment patterns within and between countries.

[15] Fraunhofer are a European based research organisation and have assisted PwC in developing a series of internal workshops to better understand how AI can be applied in different industries.

3.2.2. *Step 2: Analysing productivity impacts*

For the second stage, analysing the impact of AI on productivity required combining two significant pieces of research on AI:

- **Existing PwC research into the potential for AI-driven job automation** – a study which used machine learning algorithms to predict probabilities and levels of workforce automation for select OECD countries by 2030;
- **New econometric analysis assessing the relationship between AI and labour productivity** – a robust exercise conducted for the purpose of this study which involved identifying the key drivers of productivity growth, understanding where AI fits into this picture and specifying panel-data models capable of picking up the true causal effect of AI on productivity.

The result of our econometric research was not only a clear relationship between AI uptake and productivity, but also a framework we used to calculate the total cumulative impact of AI on productivity using a number of factors such as estimates of: the scale of augmenting AI uptake, potential job losses from replacement automation, the scale of replacement AI uptake, labour cost savings, trends in AI expenditure and the marginal impact of AI on productivity. The existing PwC job automation study played a key role in determining the estimated job losses from replacement automation by 2030, and so was extended for this present study. These two pieces of research were then combined to estimate the cumulative productivity impact of AI, as outlined in more detail below.

3.2.3. *Step 3: Combining our research to estimate the global GDP impact of AI*

We next developed our primary impacts of AI on both the consumption and productivity side to feed into our dynamic, global S-CGE model, to assess the net impact of AI on the economy worldwide. This required our econometrics and modelling team to convert our calculated labour productivity impacts and product enhancements from AI into a percentage increase above a baseline world without these effects. On the product enhancement side, this specifically involved converting AI impacts on **product quality, product personalisation** and **time saved** using the literature on willingness to pay for product attributes, and welfare effects. On the productivity side, we used our econometric framework and predicted automation levels to directly work out the cumulative productivity impact per region and sector, before specifying a profile of AI's impact using existing research on technology diffusion.

The S-CGE model combines economic data and a complex system of equations in order to capture the interactions of the three main elements of the economy – households, businesses and the government. Each element is defined and linked through the labour market, capital market flows, household consumption, intermediate product demand, taxes, government transfers and international trade. Beyond the initial impacts on productivity, job displacement and consumer choice, these net effects account for secondary impacts such as the creation of new jobs, increased consumer demand (from attractive goods first and more affordable goods second), the increased supply of labour to the market and trade flow patterns.

All of this work was informed by extensive review of the academic literature and similar studies and used data from reputable global datasets including the EU and World KLEMS database, Global Trade Analysis Project (GTAP), Programme for the International Assessment of Adult Competencies (PIACC) and the Global Innovation Index – which is co-published by Cornell University, INSEAD and the World Intellectual Property Organisation, an agency of the United Nations.

Figure 3.1 – Our multi-stage approach to assessing the total economic impact of AI

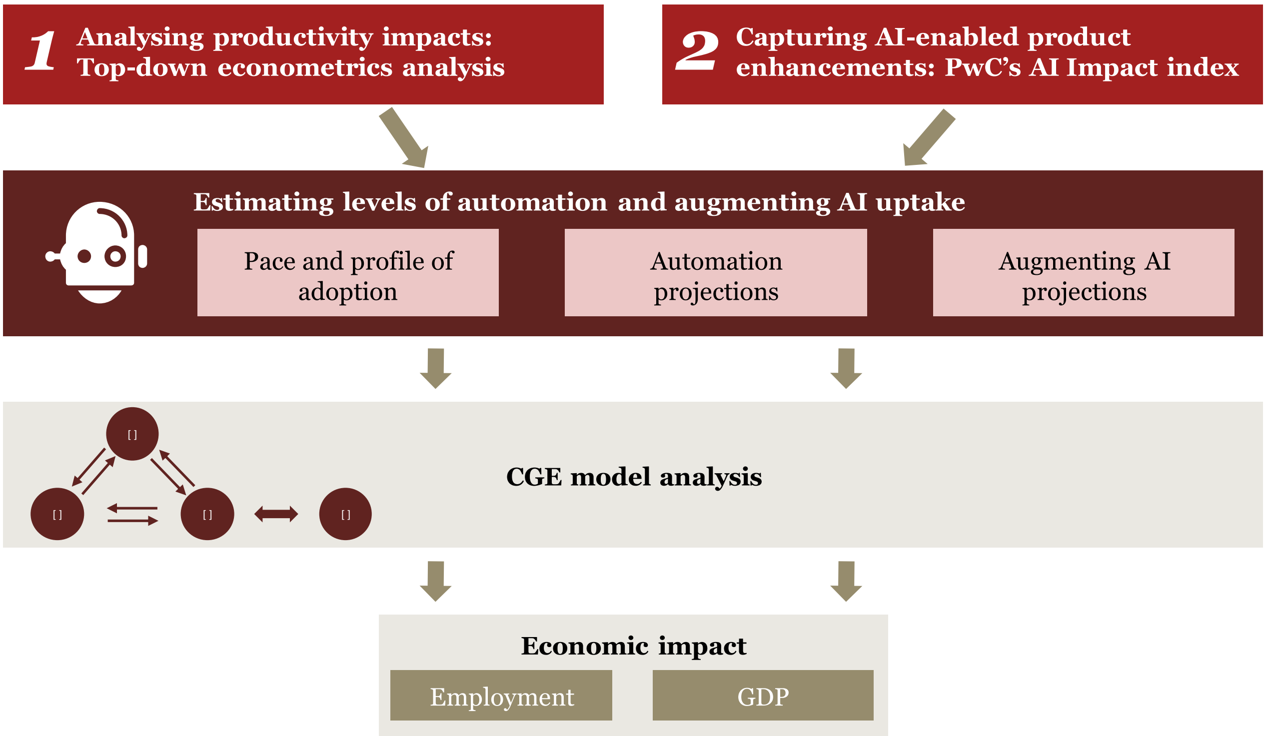

Source: PwC Analysis

3.3. *Timeline for modelling economic impacts*

Our approach assesses the economic impact of AI technologies over the period from 2017 to 2030[16]. The scope includes all AI technologies that fit under the definition outlined in Section 2.1 of this report that either (a) have been adopted already, (b) are in the process of development for future adoption or (c) have been conceived of and are likely to be adopted to some extent before 2030. The analysis does not account for any technologies that have not yet been conceived of. Though we understand that there are further waves of AI technology that are being discussed, and that may come to fruition before the end of 2030, the potential costs and benefits of these technologies for the economy cannot yet be fully understood and these later waves of AI technologies are therefore out of scope of this study.

We assume an 'S-shaped' adoption curve of AI technologies. In practice this means that adoption is relatively slow in the near-term then picks up more rapidly when the major obstacles have been overcome before slowing down again as the limit approaches.

We have chosen to use an S-shaped adoption curve to reflect much of the academic research on the life cycle of technology adoption and innovation diffusion. Indeed, it is almost a stylised fact that new technology-use typically follows an S-shaped adoption curve over time, as is reviewed in Geroski (2000).[17] There are two main competing models that explain this: the epidemic model and the probit model. Whilst the epidemic model is founded on the belief that the S-curve is a result of information asymmetries on new technology, how to use it and its purpose, the probit model focuses on how firm heterogeneity in goals, quality and ability result in different usage requirements at different points in time.

However, to reflect uncertainty about the extremity of the adoption-path tails, our S-curve is actually a hybrid curve which takes a flat average between a 'pure' S-curve which we create using a logistic cumulative distribution function (scaled between the interval [0,X], where X = the percentage impact on a given channel

---

16 The S-CGE model used to calculate our economic impact of AI runs out to a much farther date (2060) in order to prevent the rational expectations assumption used within the model to distort consumer behaviour and encourage capital shedding as 2030 approaches in the finite horizon setting.

17 Geroski, P.A. (2000) 'Models of Technology Diffusion', Research Policy, Vol. 29, No. 4-5, pp.603-625.

from AI created from our model inputs), and a linear adoption curve. In Section 8.4 of this report we explore the impact that assuming this particular shape of S-curve has on our findings by considering an alternative S-curve specification. However, we always use an S-curve to reflect the literature and empirical findings on technology diffusion and innovation.

Notably, we have also ensured that 2030 is not the limiting period of the impact for all AI-based technological enhancements. Although the enhancements considered are those that are implementable by 2030, it would be unrealistic to assume that all conceived enhancements to productivity or products would be completed by 2030 exactly. Therefore, we have considered completion by 2030 as a 'perfect' benchmark for every sector and region, and measured prohibitive factors that would prevent some fraction of the adoption cycle from being completed by 2030 using the Global Innovation Index (GII 2016), generating estimates of the AI adoption that is expected to occur in each geographic region over the study period.

Throughout our analysis we focus on the economic impact that will be realised by the year 2030 and the increase in GDP that could result from the adoption of AI technologies on an annual basis. We have explored how these effects will be accrued over time and how these trajectories may differ for the different channels of impact and across the different geographical regions of our model.

3.4. Geographic focus of economic impacts

Our global, dynamic Spatial Computable General Equilibrium (S-CGE) model splits the world into seven regions and analyses the interactions within each of those geographical regions as well as between each of them, through international trade. The geographic regions are: North America (excluding Mexico), China, Developed Asia (which includes Japan, South Korea, Taiwan and Singapore), Northern Europe, Southern Europe and Latin America. A full description of country allocations to each region can be found in Appendix A. Through our econometric analysis and AI Impact Index, we have evaluated the potential uptake of AI in each of these regions and the direct impact that this will have on productivity, jobs and product enhancements. The S-CGE model then accounts for net impact on each of these geographical regions once all interactions have taken place.

The S-CGE model also includes a 'rest of world' region. We have not modelled any adoption of AI technologies in this region due to paucity of data, but there are benefits accrued in these countries which come about through trade with AI adopting countries. We recognise this as one limitation of our analysis and is one reason why our figure could be a conservative estimate of the full global potential. This is particularly the case due to India being included in the rest of world region, who are likely in reality to see tangible and significant GDP gains as a result of AI over the next 10 years and beyond. However, unfortunately neither automation data, data for econometrics analysis nor the AI impact index were available or could be applied to India to provide an accurate estimate of AI's impact.

4. Econometric analysis

4.1. Overview

In this section, we outline the econometric approach we have used to assess the relationship between artificial intelligence (AI) and productivity, one of our key channels of impact on the economy. We discuss in detail the methodology that we followed, the data used in the modelling process, how our study fits in with the existing literature on the topic of productivity, and the key findings coming from this analysis.

We have used econometric analysis of historical data on jobs, emerging technology adoption and productivity to estimate the direct impact that uptake of AI technologies is likely to have on productivity. We have used comparable data across a number of industries and countries to assess the elasticity of productivity to usage of these technologies (i.e. the percentage increase in productivity associated with a 1% increase in the stock of AI technologies) for each industry sector in each our geographical regions.

These elasticities are later combined with our estimates of expected levels of automation and projections of investment in AI technologies to augment the workforce in order to provide an estimate of the direct impacts of AI on productivity over the period to 2030. Section 5 of this report discusses how we determine the expected levels of automation, Section 7 elaborates on how we arrived at our estimates of replacement and augmenting AI uptake and Section 7.5 provides further technical details on how these elements are combined.

4.2. Literature

In economics, the production function of a firm indicates that output increases if any of the economic inputs (capital, labour and materials), their individual efficiency, or total factor productivity increases. Typically this relationship is expressed in 'Cobb-Douglas[18]' form, which allows for interactive effects between the different inputs (and their efficiency) on each input's productivity, and we have followed this specification in our study given its commonality in the academic literature. Due to the well-established and agreed nature of these relationships, we do not discuss any of the related academic literature here in detail, but key contributors to, and users of, this conceptualisation of firm productive technology include Solow and Swan (1956)[19] and Ramsey (1928)[20] as well as the original creators: Paul Douglas and Charles Cobb (1928)[21].

Instead, we turn to the academic literature to inform our econometric methodology to assessing the relationship between emerging technologies and labour productivity, since this is the closest literature to the assessment of the potential relationship between AI and labour productivity.

Choudhry (2009)[22] is one of the most recognised and well-cited papers in this sphere. This research provides evidence of the key determinants of productivity that we should include in our analysis as control variables to prevent their exclusion from potentially confounding our results. It also indicates the validity of using a fixed-effects panel data approach to empirical estimation of productivity relationships. The paper concludes that education, investment in ICT and other forms of investment such as foreign direct investment (FDI), are the key factors that determine the growth of productivity across 45 different countries. This is somewhat reflective of

---

[18] The Cobb-Douglas production function used to express the relationship between output (or productivity), and the inputs used along with total factor productivity, usually takes the following functional form: $Y = AK^{1-\alpha}L^{\alpha}$, where A is total factor productivity, Y is output, K is the capital stock, L is labour, and $\alpha \in [0,1]$ is the labour share of income. The Cobb-Douglas production function is designed to capture the importance of inputs, their respective shares of income, unobservable total factor productivity (uncorrelated with capital and labour) and their interaction effects as determining output in the economy. Note that Cobb-Douglas production functions can be extended to include other inputs, such as land and intermediate materials.

[19] Solow, r. & Swan, T.W. (1956) 'A Contribution to the Theory of Economic Growth', *The Quarterly Journal of Economics*, Vol. 70, No. 1, pp. 64-94.

[20] Ramsey, F.P. (1928) 'A Mathematical Theory of Saving', *Economic Journal*, Vol. 38, No. 152, pp. 543-559.

[21] Douglas, P. H. & Cobb, C. W. (1928) 'A Theory of Production', *American Economic Review*, Vol. 18, pp. 139-165.

[22] Choudhry, M.T. (2009) Determinants of Labor Productivity: An Empirical Investigation of Productivity and Divergence.

the fact that investment and supply of different factors of production play a key role in productivity growth, which is consistent with the standard conceptualisation of the firm's production function.

The other key paper used to help determine drivers of productivity to be included in our econometric models was Stiroh's (2001) Federal Reserve Bank of New York Economic Policy Review: 'What drives productivity growth?'[23]. In the review, Stiroh reviews what different key economic theories and empirical studies suggest about the long run drivers of productivity growth, in particular contrasting predictions from the neoclassical (exogenous growth) literature, with that of the endogenous growth literature. The review finds that investment in (or equivalently innovations in the stock of) physical capital, research and development, and the quality of human capital are all important aspects of productivity growth. Moreover, the evidence presented supports the use of measuring AI uptake through the quantity of emerging technologies (i.e. AI-related types of capital). In particular, the review finds that the empirical and theoretical impact of technological progress on productivity is consistent with the importance of capital deepening driving this effect. In other words, the quantity –,not just quality –,of capital is important in explaining the effect of technology uptake on productivity.

In addition to the academic literature, in 2007 the UK Office of National Statistics (ONS) have also published on the theory and drivers of productivity growth in the ONS Productivity Handbook[24]. In the handbook, the government's productivity framework identifies five key drivers of productivity growth: investment, innovation, skills, enterprise and competition. Our model is consistent with these drivers, having explicitly included controls for investment and skills whilst modelling innovation. However, data on enterprise opportunities and competitive landscape was not available at a Standard Industrial Classification (SIC) level over our defined sample period.

4.3. Methodology

Our strategy was to build heterogeneous coefficients panel-data models for each geographical region to estimate the impact of the uptake of AI technologies on labour productivity for each industrial sector group defined in our study.

We proxied the relationship for each geographical region using a representative, key country from each region. This approach was necessitated by the inconsistency of data across different countries and general lack of high-quality data globally. Notwithstanding data constraints, we chose countries for each region that were both large and representative – a larger country representing a larger fraction of regional activity and a representative in the sense that it is reflective of the relationships at play in the geographic region. Therefore, we used the following countries as a proxy for our geographic regions: US (North America excl. Mexico), UK (Northern Europe), Spain (Southern Europe), China (China), South Korea (Developed Asia), Argentina (Latin America).

Using existing data on emerging and AI-related technologies to capture uptake of AI

Since AI is a relatively new phenomenon, there is no exact data-series available that directly measures the uptake of AI by country and/or industry over time. As a result, one challenge for our analysis was to define a variable that could proxy for AI uptake, or at least help us estimate the size of the impact of AI uptake on labour productivity. The EU and World KLEMS databases contain more aggregated data-series of capital stock groupings that contain AI technologies and therefore can capture the potential effect of AI on productivity (assuming the impact of AI is similar to other emerging technologies in the grouping). Specifically, the KLEMS database contains data for each country and by SIC code, on the stock of capital categorised as software, databases, computer hardware and machinery, which together cover the different types of AI discussed in our definition. We have used these to create a variable to proxy for the stock of AI technologies and therefore the adoption and usage of AI technologies. As long as the relationship between all emerging technologies and productivity remains similar within the sub-groups, our variable should give us an unbiased and consistent estimate of AI's impact on productivity.

Econometrics specification in detail and controls

Each of our panel-data models (for each geographical region) follows an identical functional form with the same control variables to reflect economic theory and the empirical literature on the drivers of labour productivity. More specifically, we consider the following general functional form outlined below. Labour productivity in a given industry (i), sector group (j) and time period (t) is regressed on its lag, a time-trend which is specific to

---

[23] Stiroh, K.J. (2001) 'What Drives Productivity Growth?' Federal Reserve Bank of New York, Economic Policy Review.

[24] The ONS Productivity Handbook: A Statistical Overview and Guide (2007), Palgrave Macmillan.

each industry, our defined measure for AI uptake, the list of control variables and a fixed effect for each industry.

Equation 4.1 - Specification of regression equation used to assess the relationship between artificial intelligence and labour productivity

$$\Delta \ln(y_{ijt}) = \mathrm{a}_{ij} + \delta_{ijt} + \beta_j \Delta \ln\left(\frac{a_{ijt}}{l_{ijt}}\right) + \gamma' \Delta ln(z_{ijt}) + \varepsilon_{ijt}$$

Productivity/ Labour share

Industry specific time-trend and fixed effect

Elasticity between dependent variable and automation

Automation variable

Control variables

Idiosyncratic error term with mean zero

Source: PwC Analysis

Note that this specification not only reflects the empirical literature on labour productivity drivers, but also the general economic theory of productivity, which states that factor productivity is a function of inputs, their efficiency and total factor productivity. In our model, total factor productivity and input efficiency can be thought of as being approximately captured through the residual and linear time-trend (effectively decomposing these aspects into a deterministic and stochastic component), whilst our model also includes factors of production directly. Importantly, this does not confound any effects estimated by our model since input efficiency and total factor productivity are conceived to be orthogonal to factors of production.

Although emerging technologies are often conceptualised to impact total factor productivity, the Cobb-Douglas production function specification, amongst others, shows that this impact can still be captured through the effect on labour productivity. Moreover, total factor productivity is notoriously harder to predict. For these reasons, labour productivity is our preferred choice of productivity measure to capture the effect of AI-uptake.

Our AI-uptake variable is specified as the ratio between the stock of AI technologies (as defined above) and workforce jobs for three reasons. First, by using the AI-technology per worker, we can prevent periods of industry expansion (possibly correlated both with increasing output per hour and AI-uptake) from positively biasing the effect of AI-uptake on productivity. Second, we recognise that the processes of workforce replacing and augmenting AI-uptake take place simultaneously. Rather than trying to separate these channels, we are interested in capturing the aggregate effect of these channels together through our estimated coefficient. Although in practice this restricts the coefficient to capture both the impact of replacement automation and augmenting AI uptake on productivity, if the share of activity from each of these channels remains relatively constant over time then our estimated labour productivity elasticities should not be biased. Moreover, the alternative is to specify each variable separately in the regression, which leads to problematic interpretations, since each coefficient would be defined as the marginal effect holding all other variables constant. This makes for difficult interpretation since jobs would not be held constant when automated and autonomous intelligence are used to replace some workers. Finally, it is also conceptually consistent to look at the effect of the ratio of AI-uptake to jobs (effectively AI per worker) on labour productivity (output per hour), since the two terms are both rate variables.

Our controls can be considered as extending the standard Cobb-Douglas production function (in its reduced form approximation) to account for different factors of production and the indirect impact of some variables on productivity. Specifically, we use the following control variables: 1) residual capital stock (to capture the effect of other non-AI related capital on productivity); 2) lagged and present R&D expenditure (to capture the impact of R&D on total factor productivity and the potential correlation with periods of AI-technologies uptake); and 3) the fraction of the workforce with tertiary education (to capture the effect of human capital on output and productivity).

There are a number of other key aspects of our model specification that are important for its ability to capture the effect of AI on labour productivity accurately:

- The model has been specified in first differences[25] which ensures our data is stationary and limits the possibility of reporting spurious results – this helps us prevent a situation where we incorrectly find statistically significant effects.
- The variables are expressed in logarithms in order to prevent the presence of any non-linear, interactive relationships from biasing our results, and to give us an approximate constant elasticity – this helps the model capture the true relationships in the data whilst providing a simple percentage relationship between AI-uptake and productivity.
- The model uses fixed effects to capture the impact of industry specific attributes which do not change over time on productivity growth – this helps the model to capture the true effect of AI on productivity growth, and avoids picking up effects that are explained by other (industry specific) factors.
- The model uses an industry specific time-trend to pick up long-run deterministic trends in productivity growth – this prevents the model capturing long-run correlations between growth in AI uptake and productivity growth, which may confound the causal impact of AI on productivity.
- All variables are specified in 'real' not nominal terms, to eradicate any noise or correlations induced from price and wage inflation.

Estimator

All of our models were estimated using a standard Least-Squares Dummy Variable (LSDV) estimator. Despite considering implementing a dynamic panel-data approach and including a lagged dependent variable to model stochastic autoregressive dynamics in productivity growth, we did not end up taking this approach. This is because specifying a dynamic panel data model using our estimator would have resulted in bias (the Nickell (1981)[26]) bias that comes from the correlation between the residual and the lagged dependent variable when using the classic fixed-effects estimator. However, the estimators proposed to overcome this problem (Arellano & Bover; 1995) (Blundell & Bond; 1998) require 'Large-N, Small-T' panels to ensure that the observations-to-instrument ratio is large enough for 1) tests of instrument validity to be meaningful and 2) to prevent overfitting the lagged dependent variable, which would fail to expunge the endogenous effects as a result. Our panel-data dimensions are instead best described as 'Moderate-N, Moderate-T' and using these estimators would therefore mean a unity ratio between observations and instruments (in the limit) and subsequently would heavily over-fit the endogenous variable.

As a result, we have used the LSDV estimator (mathematically equivalent to the classic fixed effects, or within estimator) as the 'lesser of two evils' and have excluded the lagged dependent variable whenever not significant (i.e. in all cases) to prevent any source of bias from entering the regression. Whilst we recognise this estimator cannot reject the null hypothesis of the lagged dependant variable with 100% coverage rate due to its inconsistency in estimating the lagged dependent variable coefficient, our finding that in every country model the lagged dependent variable was not even significant at the 20% level is suggestive enough that stochastic dynamics are not important in our model, irrespective of whether the coverage rate is 100%. Instead, productivity growth is best captured through a deterministic trend and its predictors, as our model results in Section 4.5 show.

4.4. Data

The majority of the data used in this analysis comes from the EU and World KLEMS datasets. These data aim to facilitate the analysis of growth and productivity over time – typically 1997 to 2014 – in a comparable way across several countries of the world including all EU member states, the US, Japan, Canada, Russia, China, Korea, India and Argentina. They include data on output and productivity as well as detailed capital and labour inputs by industry sector according to the ISIC Rev. 4 industry classification. The capital input data includes the

[25] The combined presence of fixed-effects and industry specific time-trends in first-differences is an interesting and unusual aspect of the model. This implies that productivity levels follow a non-linear deterministic trend over time and that productivity growth follows a linear trend over time.

[26] Nickell, S. (1981) 'Biases in Dynamic Models with Fixed Effects', Econometrica, Vol. 49, No. 6, pp. 1417-1426.

value of the real capital stock for various different categories of capital, e.g. software and IT equipment etc. which have formed the basis for the measure of AI in our econometrics as explained in section 3.

In particular, we used the KLEMS dataset to obtain the real capital stock (AI related and residual), real R&D investment[27] and jobs for the following countries. For Argentina, only national data was available and so we therefore used a time-series regression model in place of a panel-data model.

For our human capital supply control, we used OECD data on fraction of the workforce between 21-65 years old with tertiary education for the UK, US, Italy, South Korea and Argentina (the OECD countries in our study). Meanwhile for China and Argentina, the best available similar variables were the ratio of higher education graduates to the working population using data from the China Statistical Yearbooks and data from the World Bank on the enrolment ratio in tertiary education for Argentina.

For each country, we set the sample limit as 1997-2015. This is not only to ensure consistency between different countries due to data availability constraints, but also to ensure our AI-technologies variable only considers the impact of the most recent types of high-technology capital on productivity. These are most likely to be directly relatable and similar in nature to AI of the future yet to be implemented. The results table in Appendix B includes the specific sample size for each proxy country.

4.5. Results

Below we present our econometric results for our AI variable (AI-uptake per worker) in our regression, by industry and geographical region. These results are the elasticity of productivity to AI – for example if AI per worker were to increase by 1% in the transport and logistics sector in China, productivity in that sector would be expected to increase by 1.14%. If the same were to happen in that sector in Southern Europe, productivity would increase by 0.22% by comparison. A full set of results are included in Appendix B.

Table 4.1 – Coefficients representing elasticity of productivity with respect to AI, by country and sector (0.1 means a 1% increase in AI-uptake (per worker) increases labour productivity by 0.1%)

	Northern Europe	North America [28]	Developed Asia	Southern Europe	China	Latin America
National estimate	**0.21** ***	**0.51** **	**0.50** ***	**0.20** ***	**0.94** ***	**0.10**
Industry estimates						
Energy, Utilities, mining	**0.23** ***	**0.88** ***	**0.89** ***	**-0.06**	**0.86** ***	**NA**
Manufacturing and construction	**0.37** **	**0.60** **	**0.53** **	**0.51** **	**0.44**	**NA**
Consumer Goods, Accommodation and food services	-0.05	0.38	0.11	0.08	1.53 ***	NA
Transport and logistics	**0.42** ***	**0.53** *	**0.99** ***	**0.22** ***	**1.14** ***	**NA**

[27] For the US and South Korea we used real gross fixed capital formation as an (aggregate) proxy for real investment in R&D since R&D specific data was not available.

[28] For North America, the World KLEMS database only ran until 2007. This means the initial econometric model estimated for the US (proxying North America) did not capture the post-financial crisis developments and uptake of AI, therefore possibly underestimating the marginal impact of AI. To counter this, we used BEA data on similar (AI-related) components of the national US capital stock, to build another econometric model running until 2015. We then additively scaled the initial industry-specific elasticities by the difference between the 2015 model and 2007 model national elasticity, to arrive at our final estimates in table 4.1.

	Northern Europe	North America [28]	Developed Asia	Southern Europe	China	Latin America
Technology, Media and communications	**0.42 *** **	**0.69 *** **	**0.75 *** **	**0.50 *** **	**1.00 *** **	**NA**
Financial and professional services	0.19 *	0.38	0.21 ***	0.21 ***	0.94 ***	NA
Health, Education and other public and personal services	0.07	0.73	0.64 ***	0.33 ***	1.15 ***	NA

*Standard errors clustered by industry for each regional regression, * = 90% statistical significance, ** = 95% statistical significance, *** = 99% statistical significance. US significance levels are from original 1997-2007 model.*

Source: PwC Analysis

One of the standout findings from these results is the larger national elasticities present in China and developed countries in Asia and North America than in Europe and Latin America. There are a number of aspects which we believe have impacted the coefficient sizes in each region. The first key source of the discrepancies is likely to derive from the (unobservable) quality of AI-technologies introduced by these regions over time. In regions where the AI-technology being up-taken is of higher quality, one would expect that the impact on labour productivity would be greater. This helps explain the large national elasticities for our country proxies in North America and Developed Asia, but does not alone explain why China's elasticity is so much larger than all other regions.

The other key aspect that is most likely to contribute to China's elasticity being so large, is starting productivity levels. In countries uptaking AI technologies with lower levels of initial productivity levels, the impact is likely to be much greater since there is effectively a 'leapfrog' effect as these countries catch up to (and potentially overtake) other regional productivity levels – holding equal and constant uptake of AI technologies. This is consistent with the well-documented literature on growth convergence (see Nakaoka; 1987[29], and Sokoloff & Engerman; 2000[30]) between economies.

It is tempting to assume that the national elasticities may simply reflect periods of simultaneous AI uptake and productivity growth over the sample period, and therefore that they simply represent pure (spurious) correlation, not causation. However, this overlooks the fact that the model is not only specified in first differences (removing the possibility of a spurious result from non-stationarity of the residuals), but also includes linear, heterogeneous time-trends per specified industry group. This means that even if there was considerable persistent and simultaneous growth in AI-uptake and productivity over time and per region in the sample, the AI-uptake coefficient would only relate to the residual movement of productivity around its deterministic growth trend.

Turning to the sector specific coefficients estimated by our models, we find those industries which are more capital intensive are generally estimated to see greatest productivity gains from AI uptake (in bold). Our explanation for this result is that capital intensive industries are typically able to benefit from AI uptake since the productivity of the capital-driven processes can be enhanced considerably through the use of smart AI-technologies. Additionally, we find that the Technology, Media and Telecommunications (TMT) sector sees substantial marginal productivity gains from uptaking replacement and augmenting AI, with the largest minimum elasticity across regions at 0.42. Whilst TMT is not necessarily as capital intensive as the other high performing industries, the products and services offered naturally benefit from AI greatly due to their use of high-technology and so it is not surprising that AI-uptake is estimated to impact productivity powerfully in this sector.

29 Nakaoka, T. (1987) 'On technological leaps of Japan as a developing country', *Osaka City University Economic Review*, Vol. 22, pp. 1-25.

30 Sokoloff, K. L. & Engerman, S. L. (2000) "History Lessons: Institutions, Factor Endowments, and Paths of Development in the New World', *The Journal of Economic Perspectives*, Vol. 14, No.3, pp. 217-232.

It is important to note that for the S-CGE model inputs, we used the point estimate of the marginal impact of automation on productivity, irrespective of statistical significance. Although this risks potentially including estimated marginal impacts which may in fact be zero in reality, it is theoretically difficult to justify why the impact would be exactly zero in one sector if there is automation that is certain to take place. Moreover, statistical hypothesis tests only tell us the probability of observing a given dataset for a given hypothesised coefficient size. Even if the test suggests that the evidence cannot reject with much confidence that the true coefficient size is not zero, the test also tells us that our estimated parameter is that which maximises the probability of observing our data. As a result, they are the best estimate we have of the effect. Equally, one can consider our estimated elasticities as the best non-parametric point estimate for the marginal impact of automation on productivity – which, given the data and estimation, they are. We also constrain the coefficients to be non-negative when using them to construct the S-CGE inputs, since it is nonsensical to conceive of negative automation impacts on productivity (which are most likely the result of the model capturing insignificant 'noise' rather than 'signal')

4.6. Robustness checks and post-estimation tests

We present below a number of post-estimation tests as a robustness check on our model specification and standard errors:

Durbin-Wu-Hausman test[31] [32] [33] – The Hausman specification test – as it is commonly known – is used to test between two models to see whether a more robust model is required. This chi-square test is often used to compare the consistency of a more efficient estimator with a less efficient but potentially more robust one. The null hypothesis states that both models are consistent, whilst the alternative hypothesis states that the more efficient model is inconsistent. The test is frequently used to consider whether fixed or random effects models should be used to estimate causal effects. Whilst this is a very important test in practice, since we used fixed-effects estimators for each region in this present study, failing to reject the null hypothesis of this test would only mean our estimator choice was inefficient, but would not affect our AI impact. This is because it would imply that the use of the less efficient (fixed effects) estimator over its more efficient (random effects) alternative does not come with the benefit of increased robustness.

Ramsey RESET test[34] – The RESET test is used to check whether the model equation is of the right explanatory format. More specifically, this specification test helps evaluate whether non-linear functions of the regressors and their coefficients help explain the dependent variable. Intuitively, if this is true, then the model is mis-specified since a non-linear model specification can be used to better explain the dependent variable. The null hypothesis of this F-test states that all additional non-linear terms have a coefficient of zero. Whilst mis-specification tests are also important in research designed to evaluate causal effects, rejecting the null hypothesis of this test would only imply a potentially (further) non-linear form of model is more accurate. Since our models are at best an approximation of the true relationship between AI and productivity, this would not dramatically change our results.

Wooldridge test[35] – This test is principally used to make sure correct inference can be drawn from model results regarding statistical significance. Designed to test for serial correlation in panel data, the test is the closest alternative to a test for the importance of clustered standard errors (if clustering is by panel indicator). The test looks for autocorrelations between the residuals from a first-differenced panel-data estimator. Under the null hypothesis the residuals are serially uncorrelated. Note that clustered standard errors are also robust to heteroskedasticity, meaning a separate test of heteroskedasticity is not required since we use clustered standard errors in all cases.

---

31 Durbin, James (1954). 'Errors in variables', *Review of the International Statistical Institute*, Vol. 22, No. 1, pp 23-32.

32 Wu, De-Min (1973) 'Alternative Tests of Independence between Stochastic Regressors and Disturbances', *Econometrica*, Vol. 41, No. 4, pp. 733–750.

33 Hausman, J. A. (1978) 'Specification Tests in Econometrics', *Econometrica*. Vol. 46, No.6, pp. 1251–1271.

34 Ramsey, J. B. (1969) 'Tests for Specification Errors in Classical Linear Least Squares Regression Analysis', *Journal of the Royal Statistical Society Series B*, Vol. 31, No. 2, pp. 350–371.

35 Wooldridge, J. M. (2002) Econometric Analysis of Cross Section and Panel Data. Cambridge, MA: MIT Press.

As can be seen below, for every (proxied) region in the model the hausman test rejects the null hypothesis of a random effects specification being identical to that of fixed effects at the 99% significance level or higher (pointing towards fixed effects being the preferred estimator of the two), reaffirming our choice of using fixed effects. The Ramsey RESET test results show that the null hypothesis of correct model specification is not rejected at 90% significance or higher in all cases except for Southern Europe and China. As explained above, this is not overly problematic as our models are intended as an approximation regardless.

The Wooldridge test of serial correlation of the residuals in panel-data shows heterogeneous results. In North America and China, the null hypothesis of no serial correlation is rejected at the 92% and 98% confidence level respectively, whilst in Northern Europe it is rejected at the 86% confidence level. However, in Developed Asia and Southern Europe, the null hypothesis can only be rejected at the 53% and 34% confidence levels respectively. However, this only tells us that there is weak evidence of residual autocorrelation in the data, and since we use the point estimate of the estimator in our AI impact calculations, the test result does not affect the validity of our coefficients.

Table 4.2 – Results of robustness and post-estimation tests on econometrics specification

	North America	Northern Europe	Southern Europe	Developed Asia	China	Latin America
Hausman test	-21.12	-10.37	27.61 [0.01]	45.01 [0.00]	43.04 [0.00]	NA
Ramsey RESET F-test	1.22 [0.31]	0.21 [0.89]	2.72 [0.05]	0.96 [0.42]	3.72 [0.01]	0.46 [0.73]
Wooldridge test	3.81 [0.08]	2.40 [0.14]	0.20 [0.66]	0.54 [0.47]	7.51 [0.02]	NA

P-values in brackets. Note the Hausman specification test statistics for North America and Northern Europe are negative, implying the estimated parameter variance-covariance matrix is not positive semi-definite. This has been documented to occur in small samples or in large samples even asymptotically under the alternative hypothesis. This means the variance of the typically more efficient random effects estimator is larger than the typically less efficient and more robust fixed effects alternative we are testing. We therefore conclude that fixed effects is at least as consistent as random effects and more efficient, supporting our choice of fixed effects.

Source: PwC Analysis

We also ran a two-way fixed effects estimator as an additional robustness check, and found that including time-fixed effects to account for potential cross-sectional shocks (and prevent any sources of spatial dependence in the error terms) does not significantly change results for each region. Notably, we opt to use the one-way fixed effects estimator as our main input into the S-CGE model. This is because the theory and empirics suggest that our econometric model has captured the key covariates important in determining productivity, and since some of these are not sector specific (i.e. fraction of the workforce with tertiary education), a two-way fixed effects estimator would prevent the inclusion of these variables, potentially biasing our results.

5. Job automation study

5.1. Overview

PwC's 2017 study on job automation in the March UK Economic Outlook – 'Will robots steal our jobs? The potential impact of automation on the UK and other major economies[36]' – also played a crucial role in helping us measure the direct impact of artificial intelligence (AI) on productivity. The PwC job automation study used machine learning and expert judgement from AI specialists to predict the fraction of the current workforce for a number of countries which are expected to be at high-risk of automation by 2030, and in using a new approach were arguably able to provide more accurate estimates across a multitude of countries.

As our AI-impact study aims to capture the cumulative impact of both replacement and augmenting AI for different global regions, part of this exercise requires an estimate of both the fraction of the current workforce expected to be working in by 2030 and the number of jobs that will be lost in the process (which is used to calculate initial firm labour cost savings in our study). As a result, the PwC job automation study results were extended to other countries for the AI global impact study. This allowed us to create regional (and sector specific) estimates of workforce automation and total job losses using weighted averages of country-specific estimates.

Note, the PwC job automation study does not take into account job creation as a result of AI, so we do not consider these jobs estimates as net of AI's effect. Rather, they are designed to capture the scale of replacement automation that takes place as a result of AI, which has a direct impact on labour productivity through more than one channel (as discussed in Section 7.4). At the end of this section of the report, however, we consider the literature relating to job creation enabled by AI and therefore the potential outcomes related to net job creation in 2030.

5.2. Literature

The body of literature related to AI and its potential to automate jobs is dominated by two key studies, which reach starkly different conclusions. In 2013, Oxford University researchers Frey and Osborne (2013)[37] estimated that 47% of jobs in the United States were at a 'high risk of computerisation' by around 2030. More recently though, Arntz, Gregory and Zierahn (AGZ) (2015)[38] of the OECD completed similar research concluding that only 9% of US jobs were high risk.

In the original study, Frey and Osborne used a sample of occupations taken from the US Department of Labour and had these labelled by AI experts as strictly automatable or not automatable. They then developed a machine learning algorithm to learn which occupation features determine automation potential and generated the probability of the remaining occupations being automated.

The OECD researchers purported that it would not be entire occupations that would be automated but, rather, that computers and algorithms may take over particular tasks within certain occupations. To account for this, they used an alternative OECD-compiled dataset from the Programme for the International Assessment of Adult Competencies (PIAAC).

In order to obtain an accurate estimate of the number of jobs that could be automated over our study period in our countries of focus, we turned to PwC's study of job automation in the March 2017 UKEO outlined above and in more detail below, which had sought to critically assess these approaches and determine the correct figures.

---

[36] https://www.pwc.co.uk/economic-services/ukeo/pwcukeo-section-4-automation-march-2017-v2.pdf

[37] Frey, C.B. & Osborne, M.A. (2013) 'The future of employment: how susceptible are jobs to computerisation?', United Kingdom Department of Engineering Science, University of Oxford, Oxford OX1 3PJ.

[38] Arntz, M. Gregory, T. & Zierahn, M. (2016) 'The Risk for Automation of Jobs in OECD Countries: A Comparative Analysis', OECD Social, Employment and Migration Working Papers, No. 189.

5.3. Methodology

The PwC job automation study sought to replicate the OECD study, preferring the task-based approach. The study followed a multi-stage process to estimating automation potential per country and sector. The first stage involved replicating the AGZ dataset and simulating their distribution of automation probabilities for the US for both the occupation-based and task-based approaches. Once this had been achieved, the next step was to show that the predictive model used by AGZ for their task-based approach for predicting automation probabilities was heavily dependent on the feature set used for prediction. In particular, it turned out that the more accurate the predictive model was (based on the feature set), the higher the automation probabilities.

The final stage was to change the data reweighting scheme used in prediction, and also build an enhanced classifier for optimised prediction, using a random forest approach. As a result, the study has arguably more accurately predicted probabilities of automation for a given standard occupational classification (SOC) by 2030. As before, the algorithm and classifier used could then be applied to any country – data permitting – and the SOC level probability predictions were then aggregated into SIC level predictions for 2030 automation probabilities. Finally, by making the assumption that those occupations with a predicted probability of automation by 2030 above 0.7 will be automated (as we do in this study), the study was able to estimate the fraction of the current workforce and total number of jobs by sector and country expected to be automated by 2030 (notwithstanding new job creation as a result of AI and other emerging technologies)

For the present study, as mentioned, the PwC job automation study was extended to capture multiple countries' data, so that we could approximate regional estimates of job automation. Unfortunately the study was not able to provide estimates for each country per region, but two key countries' data were used when possible to ensure the automation estimates for the region were representative – where one of these countries were typically the proxy country used for the econometric modelling. To create a regional automation estimate where more than one country was used, an industry-specific weighted average (by industry employment) was taken over the two countries' estimates. The countries chosen were typically large and representative to ensure that as much of the region could be accounted for, and also to ensure we capture the employment share and skill distribution that best represents the region. A full breakdown of the countries used and the GDP share of the region they capture is presented below in Table 5.1

Table 5.1 – Countries used to proxy regions and regional GDP share captured

Region	Country proxies	GDP share captured (%)
North America	US	92
Northern Europe	UK, Germany	46
Southern Europe	Spain, Italy	73
Developed Asia	Japan, South Korea	88
China	South Korea	N/A
Latin America	Chile	5
Rest of World[39]	N/A	N/A

Source: PwC Analysis and Global Trade Analysis Programme

One limitation of our approach is that since the PwC job automation study uses a data source that is generally only available for OECD countries, we were unable to simulate automation estimates in non-OECD countries – notably impacting our ability to capture the GDP share of China and Latin America. In China's case, we did not have automation results on any country in the region (China or Hong Kong), and so had to use data on another country that would in theory best describe automation potential in China. After careful evaluation, we decided South Korea would be the best country to represent China's automation potential. This is largely because the automation potential in South Korea estimated by the PwC job automation study was low and we expect this to

[39] As the (direct) AI impact on rest-of-world was not explicitly modelled as part of this study, we did not use any automation data or more generally compute any AI impacts due to this region's uptake.

equally be the case for China, albeit for different reasons. In China, the level of AI and robotics uptake is reported to have accelerated dramatically in recent years, and a lack of worker bargaining power has enabled firms to replace workers swiftly and with little red tape getting in the way. However, due to the huge scale of the Chinese workforce (and general population size), the current low robot density (robots per 10,000 workers) and wide geographic distribution of labour, the ability for China to automate a large fraction of its workforce by 2030 is likely to be very limited.

5.4. Data

As mentioned, the main data source used in the PwC job automation study was the OECD dataset compiled by the Programme for the International Assessment of Adult Competencies (PIAAC) database on ISCO-8 codes (standard classification of occupations). This dataset includes task structures of individuals and more detailed data on the characteristics of jobs and the individuals doing them. However, we also used the original 702 O*Net dataset used by Frey and Osborne (2013) which uses SOC codes to classify different occupations. Cross-walks from the US census bureau were also used to map the respective codes used from each of these datasets, to create a new extended dataset with many-to-one relationships from the Frey and Osborne data to PIAAC data. In each many-to-one case, weights were given to the Frey and Osborne data which added up to unity.

5.5. Results: Estimating job automation

We find that the percentage of jobs at a high risk of automation by 2030 varies significantly across geographic regions and industrial sectors.

North America and Europe have the highest potential for automation with the number of jobs at high risk varying between 23% and 76% (in varying industries) while economies in Asia, including China and other developed countries have a much lower potential varying between only 11% and 29%. The study found that differences across geographical regions are typically driven by a) the industry composition and b) differences in the automatability of jobs within sectors related to variances in the task composition in different countries. Countries with a greater focus on manufacturing typically have a higher risk of automation. However, for example, despite having an industry composition susceptible to automation, workers in Japan spend a lower proportion of their time conducting manual tasks compared with management-type tasks which therefore reduces the likelihood of automation.

From a sectoral perspective, the key determinants of the potential for automation are a) the proportion of time spent on manual tasks, routine tasks or simple computations compared to tasks involving management, social and literacy skills as well as b) the level of education typically required for the role. Sectors with a large proportion of manual and routine tasks, such as transport and logistics, or simple computation, such as manufacturing, are the most likely to be automated. Whereas in many service sectors, where there is an increased focus on social and literacy skills and where those employed are typically more highly-educated – ,such as healthcare, technology itself –,financial services and education, the likelihood of automation is lower.

Figure 5.1 – Percentage of jobs at high risk of automation by 2030, by geographic region and industry sector (adjusted using the Global Innovation Index – see Section 7.7 for more details)

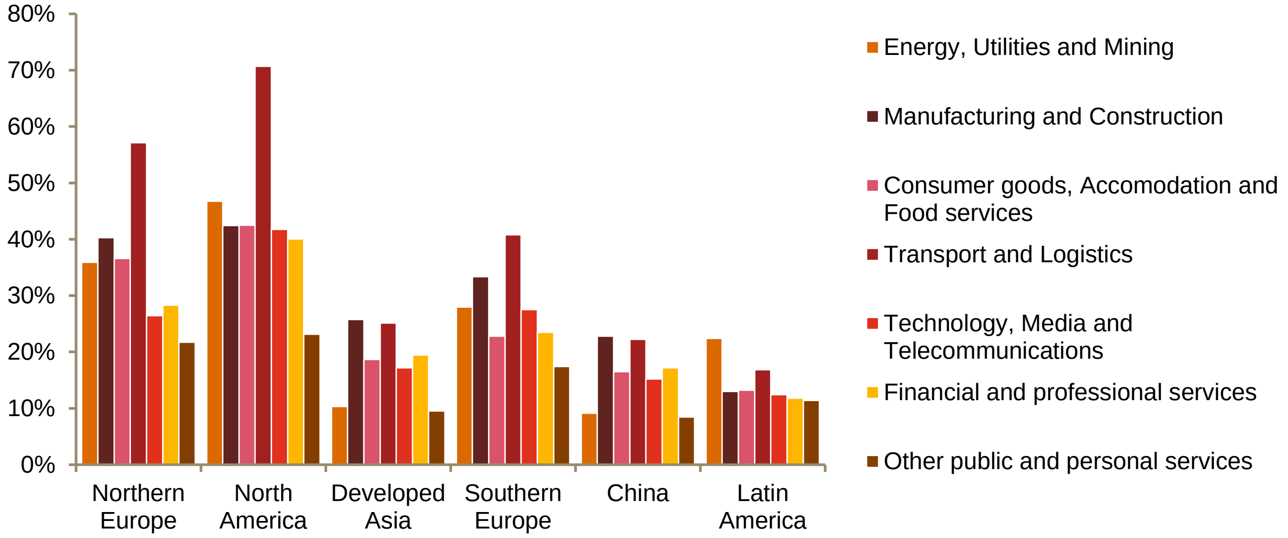

Source: PwC Analysis

5.6. Potential for job creation

Although there is widespread agreement that a significant amount of jobs could be automated by AI technologies –,and we have found similar results –,concluding at that point would focus too narrowly on one element of the impact of AI on jobs and labour. More recently, researchers in the space have turned their attention to look beyond the threat that AI poses to jobs to explore the opportunity that AI presents for new job creation. Even in the OECD's Arntz, Gregory and Zierahn (2015) paper the authors recognise that 'new technologies may also exert positive effects on labour demand if they raise product demand due to improved competitiveness and a positive effect on worker's incomes.' Autor (2015)[40] purports that previous research ignores what he describes as the 'strong complementarities between automation and labour that increase productivity, raise earnings, and augment demand for labour. This impact on labour demand is also amplified as it is filtered along the value chain of the businesses that experience the initial productivity and consumer demand increases.

Even without the uplift in labour demand resulting from economic factors, AI will in and of itself require new jobs and new roles. Along with jobs in the development and application of AI, the technologies will need to be built, maintained, operated and regulated. Recently, Wilson, Daugherty and Bianzino[41] sought to identify and classify the key jobs that AI will create, indicating that many of these jobs will 'look nothing like those that exist today.' They propose that the jobs will fall into three main categories:

- *Trainers:* People that will teach AI technologies how to perform and where possible to mimic human behaviours including how to show compassion, detect sarcasm and use humour in appropriate situations.
- *Explainers:* Technical professionals who can explain how algorithms and AI technologies work and understand why the response and output is a certain conclusion or action. Essentially these jobs will act to 'bridge the gap between technologies and business leaders.' Specifically jobs in this category could include Context Designers and AI Strategists.
- *Sustainers:* Individuals who ensure that AI systems are operating effectively and appropriately, such as economists and ethicists.

So, artificial intelligence may facilitate the creation of jobs that would not have existed in a world without AI both through the economic boost that it provides to productivity and consumer demand, and through itself and the requirements of training, explaining and sustaining the technologies. As such, despite the fact that we

40 https://economics.mit.edu/files/11563.

41 http://sloanreview.mit.edu/article/will-ai-create-as-many-jobs-as-it-eliminates/.

estimate that a large proportion of jobs could be automated by 2030, there are three potential scenarios for the impact on net jobs:

1. Artificial intelligence creates some new jobs but not enough to keep pace with the number of jobs automated and the net impact on labour in negative;
2. Artificial intelligence creates the same amount of jobs as are automated and the net jobs impact is neutral.
3. Artificial intelligence creates more jobs than are automated and yields a positive net impact on jobs and labour.

The final layer within this complex issue is the distributional aspects of the impacts, specifically, to what extent the upwards and downwards pressures on jobs affect the different groups of skilled and unskilled labour. This question has been at the centre of the debate. Both Frey and Osborne[42] and the OECD researchers[43] find that in contrast to the typical theory and recent trend of polarisation of the labour market, automation is likely to impact low-skilled occupations the most.

As part of this study and within our S-CGE model analysis, we analyse the net impact of jobs resulting from the automation that AI is likely to bring, in conjunction with the positive impact on jobs resulting from AI's productivity and consumer demand increases. Section 7.12 outlines our findings related to net job creation, including the specific impacts on geographic regions and how AI could affect high and low-skilled labour.

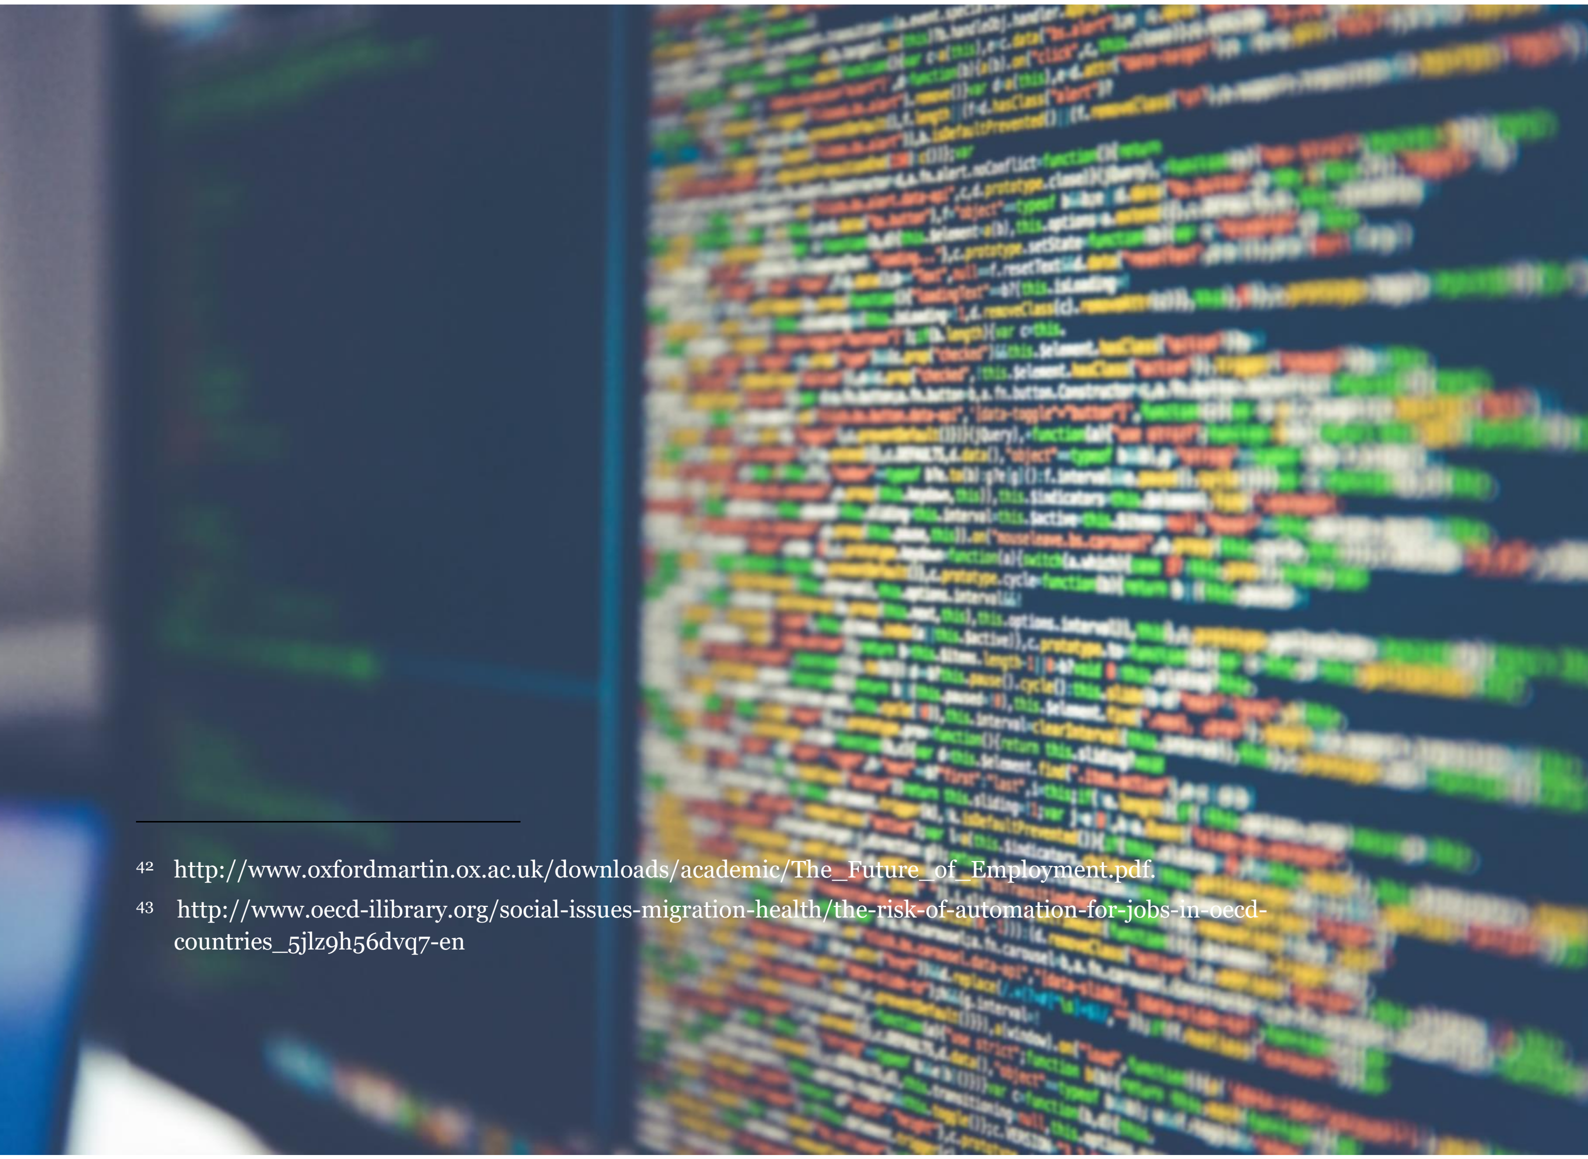

42 http://www.oxfordmartin.ox.ac.uk/downloads/academic/The_Future_of_Employment.pdf.

43 http://www.oecd-ilibrary.org/social-issues-migration-health/the-risk-of-automation-for-jobs-in-oecd-countries_5jlz9h56dvq7-en

6. AI Impact Index

6.1. Overview

The AI Impact Index's aim is to identify and evaluate the impact of different 'use cases' of artificial intelligence (AI) across firm value chains for products included in the industry sectors in scope of the study. The Index captures the potential that AI has to transform specific products and therefore provided a score for the degree to which the product offering in an industry could be transformed. These scores are later quantified using techniques justified by academic literature to calculate the initial impact that AI will have on some elements of consumer choice. These initial impacts are then fed into the S-CGE model to determine the net economic impact.

6.2. Methodology

Together with our AI specialists, industry specialists and our partners at Fraunhofer – a global leader in emerging technology research and development – we set out to gather a collection of AI use cases across firm value chains for the sectors in scope. First, we identified major product lines across the seven industry sectors that we have focused on in this study. For the product lines agreed upon, our research teams investigated how companies and institutions were currently using AI, and what emerging techniques and applications were on the horizon. Armed with this knowledge, we catalogued firms' current and emerging areas of opportunity, but we also worked to envision the art of the possible –AI use cases that might be conceivable for a given product, but that had not yet been seen in practice. In the end, we compiled around 300 of the most attractive applications of AI across product categories.

After gathering these use cases, we ventured to form an Impact Index, which would classify applications of AI across a few dimensions: value chain element, sector, sub-sector, product line, type of impact (revenue driver vs. cost reducer), potential consumption impact, and time to maturity. The remainder of this section outlines the different dimensions we use to determine potential consumption impact and as well as the time to maturity.

The AI Impact Index is meant to be both a stand-alone AI impact assessment tool, and also a source of input for the economic impact analysis using the S-CGE model. Two of the AI Impact Index's parameters: consistency and data availability, were represented by other means in the modelling. Three of the Impact Index's parameters personalisation, utility, and time saved were incorporated as direct inputs into the S-CGE model analysis and we discuss how we have quantified these as model inputs in Section 7.3 of this report.

6.3. AI Impact Index Parameters

For all AI Impact Index ratings, we considered the state of the average product produced in developed economies and considered how the introduction of AI by firms in these industry sectors could personalise the products, increase their quality and therefore consumer's utility, save consumers time, increase consistency in production and how much data was available to facilitate this. Each use case was given a score for each element on a 1 to 5 basis. We discuss the specific methodology and scoring for each element of the Index here.

Potential for product personalisation

We measure the extent to which there is scope for AI to personalise products within each industry sector, given how much they are currently personalised through existing technologies or through the general consumer offering in that industry. We looked at existing products and their average level of personalisation in mature economies. If products were already very personalised, they received a low potential uplift score, and if a product was capable of being personalised by AI and was currently produced at a commodity level, the highest score was given. Further details of this scoring framework are below.

Scoring

5 = Product is currently a commodity (e.g. uniform call centre customer service irrespective of caller),

4 = Product is currently a slightly tailored commodity (e.g. mutual funds based on target retirement date).

3 = Product is currently available and tailored at a demographic level (e.g. Super Bowl ads targeting white women in the USA).

2 = Product is currently available and tailored at a segment level (e.g. direct marketing ads targeted to people who live in a certain zip code).

5 = Product is currently available and tailored at the individual level (e.g. home inspection).

Some products are not capable of being personalised, and are given a score of 'NA'.

Product quality and utility

Every product and service has an inherent value, or level of utility associated with it that drives consumer demand for that product or service. Included in this utility is a set of functionalities that defines what the product can and cannot offer the consumer. We considered the future state of the AI-enabled product and compare its utility with the current state. When we examined AI's impact on products, we considered how each use case could have a vastly different degree of impact on utility, and commensurately, on consumer demand and followed the scoring framework as set out below.

Scoring

5 = This product's utility impacts life and death (e.g. a car that avoids accidents, a drug that prevents disease).

4 = The product's scope of functionality is enhanced (e.g. mining soln. that can operate AND relocate itself).

3 = The product's scope of functionality is the same, but its utility is greatly enhanced (e.g. it can do the same thing but much better).

2 = AI has some observable effect on utility, but it is minimal.

1= AI has virtually no impact on the end result for the consumer – they still end up with a comparable version of the non-AI enabled product produced.

Time saved

We looked at each use case and evaluated how much time the next generation of a product would save the customer. We had to differentiate over the frequency of the time saving as well as the scale of the time saving per given use. As a result, the following scoring mechanism was used. The scoring is as follows:

Scoring

5 = saves them a significant amount of time every day (hours).

4= saves them some time every day (minutes).

3= saves them a lot of time on an infrequent basis (hours).

2= saves them minimal time on an infrequent basis (minutes).

1= will not save them much time at all.

Consistency

The AI Impact Index also included a measure of consistency of production for goods and services. Since the aim of this indicator overlaps to some extent with that of the econometric analysis for the productivity relationship, we do not include this as a specific input for the economic impact S-CGE model analysis..

Data availability

Data availability at scale and specificity is a fundamental requirement for full AI potential to be realised. More available data will enable better AI, thus higher total impact potential. Though an important element of the Index, this is a measure of an enabler of AI rather than an output so we do not include this as a specific input for the economic impact S-CGE model analysis..

6.4. Time to maturity in the AI Impact Index

In addition to the specific parameters to capture product enhancements and other aspects discussed above, the AI Impact Index analysis also captures the estimated time to maturity, when a use case might be adopted by the average company in developed economies. In looking at time to maturity, our team of technology and sector experts considered both technological feasibility of the use case, and unique sector adoption drivers and inhibitors.

Technological feasibility gave us a sense for whether the use case was in fact something that could currently be achieved with existing capabilities and infrastructure. We made sure to consider how some use cases, such as autonomous driving, are currently technologically viable in limited applications, but are not fully technologically feasible. In looking at time to maturity, the categorisations of time frames were as outlined below.

Short term: Currently in use or adopted within 3 years.

Medium term: Adoption between 3-7 years.

Long term: Adoption beyond 7 years.

In arriving at these determinations, we considered many factors, such the regulatory environment surrounding the product, intra-sector adoption dependencies, and the risk that AI failure would pose on human lives and the environment. As each use case is highly specific to the product and sector in which it would be deployed, extensive consideration was given to these timelines; rationale for these times scales is available in our proprietary AI Impact Index.

We used the results with respect to time to maturity to help inform our assumptions about the shape of the S-curve of AI technology adoption and therefore the profile of the productivity and product enhancement impacts captured in our modelling. See Section 7.5 for further details.

6.5. Results

The AI Impact Index results for personalisation, time saved, and utility/quality, all of which have been used to assess the economic impact of AI through the channel of product enhancements, are summarised below. Sector scores represent the average of the use cases included in that sector.

Table 6.1 – Summary of AI Impact Index scores for industry sectors

Sector	Personalisation	Time Saved	Utility/ Quality
Energy, Utilities and Mining	1.0	2.0	3.1
Manufacturing and Construction	1.9	1.7	3.7
Consumer Goods, Accommodation and Food services	2.9	2.6	3.1
Transport and Logistics	3.4	2.9	3.0
Technology, Media and Communications	2.2	2.6	3.1
Financial and Professional Services	2.8	2.4	3.5
Health, Education and Other Public and Personal Services	4.3	3.0	3.7

Source: PwC Analysis

Health, education and other public and personal services, as well as transportation and logistics scored highest in personalisation and time saved, while manufacturing, healthcare, and financial services scored highest in utility.

Personalisation

Many of the products and services included in healthcare and automotive tend not to be personalised based on the individual, but show great personalisation potential. In healthcare, much of care delivery and medicines available are designed for the masses, and personalised goods and services (e.g. custom nutrition supplements, health coaching) are rendered *after* the commoditised experience has been exhausted, and at a great cost. While we are starting to see personalised medicine come into maturity with the advent of things like genetic screening, there is still significant improvement to be expected. Similarly, transportation and logistics is lacking in personalisation for the specific mode of transport a person or good needs, depending on the use case. Currently, most people rely on one or two modes of transportation for daily life (perhaps a car for leisure and a subway train for commuting). In the future, the mode of transport can be matched to the specific use case, optimising the asset utility and the consumer experience.

Time Saved

Transport and Health sectors scores particularly highly in the time saved category. Autonomous transportation has extensively been covered as a major time saver for consumers. AI in healthcare has a wider range of potential time savings. There are minor time savings, in automating check in or appointment scheduling, and more significant opportunities. For example, experiencing a major health incident can reduce an individual's quality of life, and cost them significant unexpected time and resource investment. In some cases, hospital admissions could be avoided if patients and their caretakers had better access to real-time health data that would enable them to take corrective action. Take, for example, acute renal failure, with US hospitalisations jumping to 404,000 in 2010, up a staggering 265% from 1997[44][1]. Given that lengths of stays for such a condition can be close to a week, this is a significant amount of time that is lost for the individual and society[45].

Utility/Quality

Manufacturing and healthcare scored the highest in the utility category. The manufacturing sector had many use cases that would expand the scope of current aspects of production. For example, technologies that can both identify a possible increase in demand for a product, and also (either by augmenting a human's decision or autonomously) adjust and run production of that product to meet demand. Enhanced production efficiencies will reduce costs to produce by manufacturers, and lead to fewer delays in getting products in the hands of customers. Healthcare use cases would in many cases extend the lives of patients, which led to high scores.

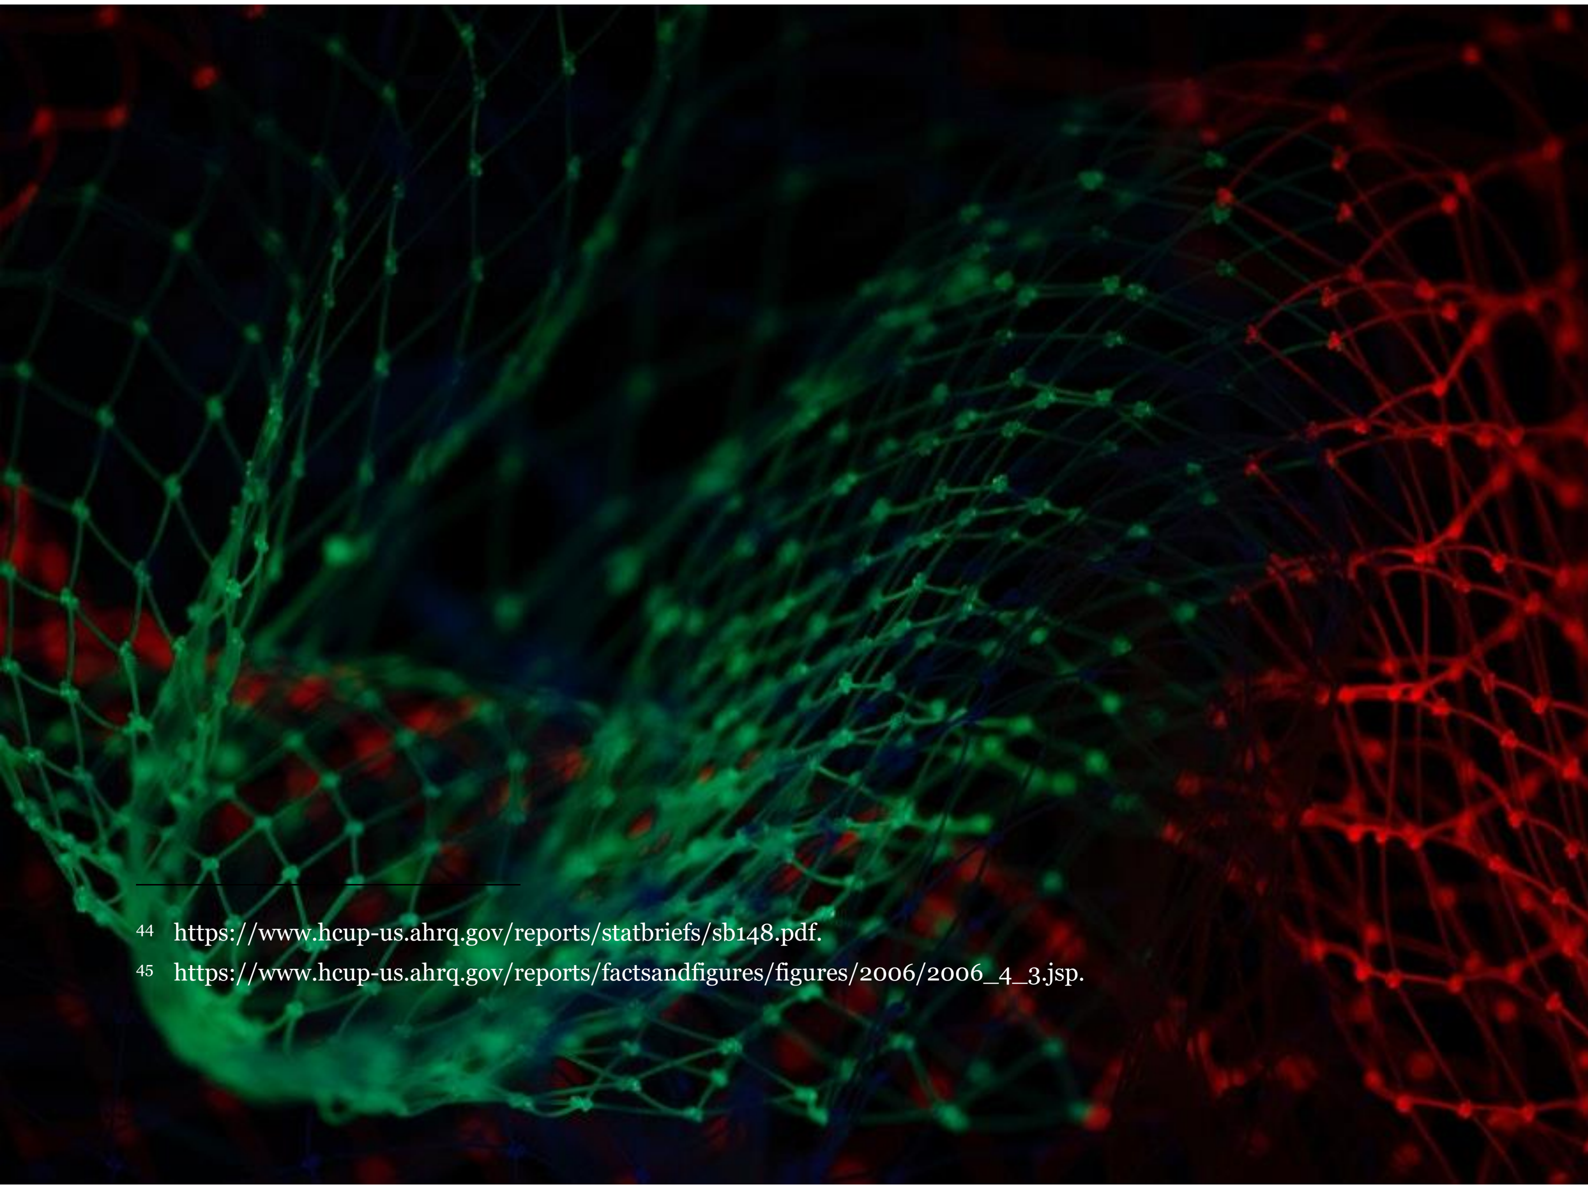

44 https://www.hcup-us.ahrq.gov/reports/statbriefs/sb148.pdf.

45 https://www.hcup-us.ahrq.gov/reports/factsandfigures/figures/2006/2006_4_3.jsp.

7. S-CGE model analysis

7.1. Overview

To our knowledge, S-CGE models have not yet been used to assess the economic impact of AI. However, S-CGE models have become a standard tool for certain types of empirical economic analysis. Their primary domestic use is to investigate the effects of different economic scenarios (e.g. a change in real exchange rates, trade police, productivity, or the level of consumer demand) or to assess the impacts of different government or institutional policies (e.g. changes in tax policy, government spending and the economic effects of CO_2 emissions). S-CGE models are used widely by academics, institutions and government departments. Notable examples include: the World Bank, OECD, European Commission, the UK Ministry of Finance and the Congressional Budget Office in the US.

In this Section we describe the structure and core assumptions underpinning our S-CGE model.

7.2. The model

A S-CGE model combines economic data and a complex system of equations in order to capture the interactions of the three main agents of the economy – households, businesses and the government. Each agent is defined and linked through the labour market or capital market flows, household consumption, intermediate product demand, taxes and government transfers and international trade.

The economic systems that S-CGE models proxy are complex. The multiple households and businesses that are defined in each model engage in repeated local microeconomic interactions that in turn give rise to macroeconomic relationships affecting variables such as employment, investment and GDP growth. These macro relationships also feed back into the determination of local micro interactions. Because of this relationship, S-CGE models are often referred to as micro-macro models (Sue-Wing and Balistreri, 2012)[46]. Figure 7.1 illustrates the interactions in the model.

Our global S-CGE model allows us to assess the net dynamic impact of AI, via the primary impacts on labour productivity and product enhancements, on GDP at a global level, in specific geographical regions as well as within specific industry sectors.

Figure 7.1 – Economic interactions in the S-CGE model

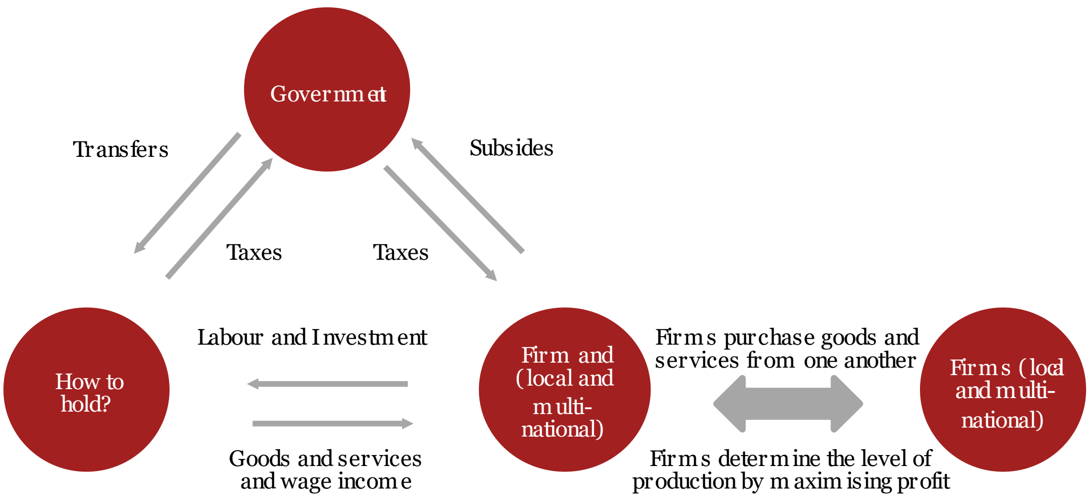

Source: PwC Analysis

[46] Sue-Wing, I. & Balistreri, E.J. (2012). Computable General Equilibrium Models for Economic Policy Evaluation and Impact Analysis. Working Paper, *Department of Earth and Environment, University of Boston.*

The dataset at the heart of our model is the GTAP (Global Trade Analysis Project) database. GTAP is an international modelling project based at the Purdue University, Indiana. Since its inception in 1993, GTAP has rapidly become a common ‘language’ for many of those conducting global economic analysis. Initially GTAP was used to support analysis of international trade agreements (starting with the Uruguay round in the early 1990s), but has more recently been used to analyse a wide range of policy scenarios e.g. the effects of Carbon taxes.[47]

The GTAP database primarily consists of a group of input-output (IO) tables for 140 different countries. Input output tables (and their hybrid formats – Social Accounting Matrices and Supply Use Tables) form the core of any S-CGE model. All data presented in an IO table are given in cash terms (US dollars) for each of the 140 countries for 57 commodities and for the year 2011.[48] A stylised illustration of a standard input output table is presented in Figure 7.2 below.

Figure 7.2 – A standard input-output table

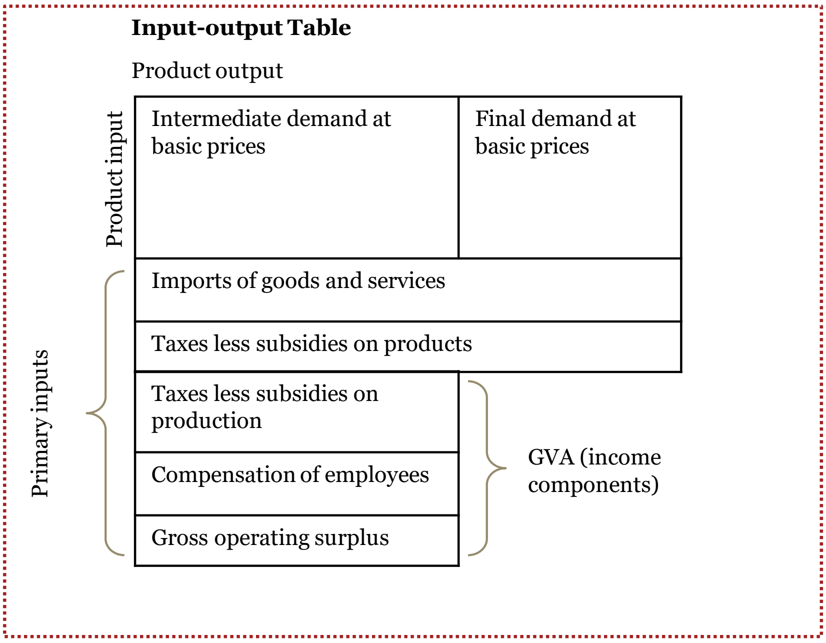

Source: PwC Analysis

A useful way to view an IO table is as a large asymmetric matrix of industry, consumer and government interactions on a commodity and industry basis. IO tables are constructed by combining and transforming two important data sources – the Use table and the Supply tables (together Supply Use Tables, or SUTs); the Use table provides data on the inputs consumed by each sector of the economy, while the Supply table provides data on the outputs produced by each sector of the economy. The IO table formed using the SUTs provides data on the goods and services produced and consumed by each sector of the economy.

GTAP provides data for 57 different sectoral groupings covering the whole economy and ranging from agriculture, through to extractive industries, manufacturing, construction and services. It is technically possible to build S-CGE models that can be solved for all 140 countries and 57 sectors in the GTAP database, but these models are very basic and also very difficult to interpret given the huge volume of results that they generate. Moreover, these more basic models lack realism – they are static (i.e. do not present results over time) and lack key functionality such as any treatment of market power or household minimum consumption requirements. These functions add substantial levels of complexity to the model’s solution and mean that the number of regions and sectors in the model must be heavily restricted. To allow for the introduction of these key functions our model aggregates the GTAP data set to 7 regions and 7 sectors.

---

47 Scott, M. Robinson, S. & Thierfelder, S. (2008) ‘Leveling the Global Playing Field: Taxing Energy Use and Carbon Emissions’, 2008 Conference paper at the 11th Annual Conference on Global Economic Analysis.

48 Many of the countries captures in the GTAP database will produce data on a more frequent basis, but IO tables are complex to produce and many countries only update their IO tables on 5 year cycles.

These are detailed in Table 7.1 below:

Table 7.1 – Regions and sectors in the PwC Global S-CGE model

Regions in the PwC Global S-CGE model	Sectors in the PwC Global S-CGE model
North America	Energy, Utilities and Mining
Latin America	Manufacturing and Construction
Northern Europe	Consumer Goods, Accommodation and Food Services
Southern Europe	Transport and Logistics
Developed Asia	Technology, Media and Communications
China	Financial and Professional Services
Rest of World – Africa, Oceania and other Asian markets	Health, Education and Other Public and Professional Services

Source: PwC Analysis

We now turn to describing the key elements of the IO table presented in Figure 7.2 and how they drive the outcomes from our AI scenarios:

Intermediate demand/intermediate consumption

The intermediate demand/intermediate consumption block is a matrix of the different products used as 'raw materials' in the production process. For instance, it details the commodities used in the production of a sector's output, for example the amount of manufacturing and financial services products used as inputs to produce output in the construction sector (or vice versa). This data comes from the GTAP database.

This matrix is central to guiding how sectors expand and contract in the S-CGE model – for instance, if the construction sector expands as a result of AI driven demand, then it will need to purchase more inputs from its suppliers in the manufacturing and financial services sectors. This demand boost leads to secondary (or indirect) economic reactions – for instance, as direct demand increases in the construction sector, this will in turn increase the manufacturing sector's demand from its suppliers to meet the increased construction demand. These interactions create a ripple effect from AI adoption across the whole economy. They capture the fact that the economic stimulus associated with AI will not simply impact on businesses that adopt AI, but also on businesses where AI uptake will not be the central focus.

Final demand

This section of the IO table captures the core elements of business, government and consumer expenditure.

- From a consumer perspective it represents household demand for the 7 different aggregated industrial products produced in the economy. In our model scenarios we capture how AI will enhance the available varieties of these different products and also their quality as well.
- From a government perspective it represents government purchases of the output of different sectors. The government could purchase goods and services that have been enhanced by AI.
- From a business perspective the IO table represents how businesses purchase different goods and services for investment – for instance if a business wanted to expand its premises, it might purchase construction services or manufacturing services.

In addition, final demand also captures demand from overseas consumers, or exports as well as imported goods. These data are provided at the sectoral level and include both goods traded in their final form (e.g. cars or a service supplied over international borders) and some intermediate goods elements (e.g. raw oil exports).

The sum of intermediate demand and final demand provides the total demand for each sector's output. We now turn to look at total supply to the market.

Total output (or total supply to the market)

The IO tables are driven by a fundamental relationship: expenditure equals income (or supply equals demand). GDP by expenditure can be calculated from the IO table as the total demand less imports of goods and services and intermediate product demand.

Correspondingly GDP by income source can be calculated from the bottom three blocks within the IO table (taxes less subsidies on production, compensation to employees and gross operating surplus). This relationship is also critical to the S-CGE model and GDP (expenditure) always equals GDP (income).

In practice this means that all output produced in the model is in turn consumed. The production of output generates income for owners of land and capital goods (gross operating surplus, which often includes returns to land, or else land is treated separately as is the case in GTAP), as well as wages for workers (known technically as compensation of employees as the data also includes additional worker benefits such as payments in kind) – these elements form the factors of production and they are paid in line with their respective productivity levels.

The relationships that govern the supply of different goods and services are detailed in Figure 7.3 below. Starting at the very bottom of the Figure, skilled and unskilled labour combines to provide a total supply of labour to the market. Firms' earnings from capital goods (buildings, machinery technological equipment – including AI tech) combine with labour (earnings from labour are paid in the form of employee compensation) and land to produce Gross Value Added (there are some small additional taxation elements that are not shown in the figure). These earnings are combined with intermediate goods (i.e. the raw materials described above) to produce total output. Intermediate goods can be sourced from domestic and overseas sources.

Depending on the relative price of factors or production firms can substitute between them. For instance, people can be replaced with machines etc. However, factors cannot act autonomously; they are complementary: if a firm wants to expand then it will need people, capital equipment and land. However, if capital goods become cheaper firms may use more capital and less people per extra unit of output produced. Generally it is difficult for firms to use less raw materials, but it is possible within the model. There are strict controls in place in the model around the substitution of raw materials e.g. car tyres must be made from rubber, they cannot be produced using cement even though it may be a cheaper alternative. However, the rubber could be sourced from any supplier across the world in the event of a price change.

Figure 7.3 – Value added nesting structure in the S-CGE model

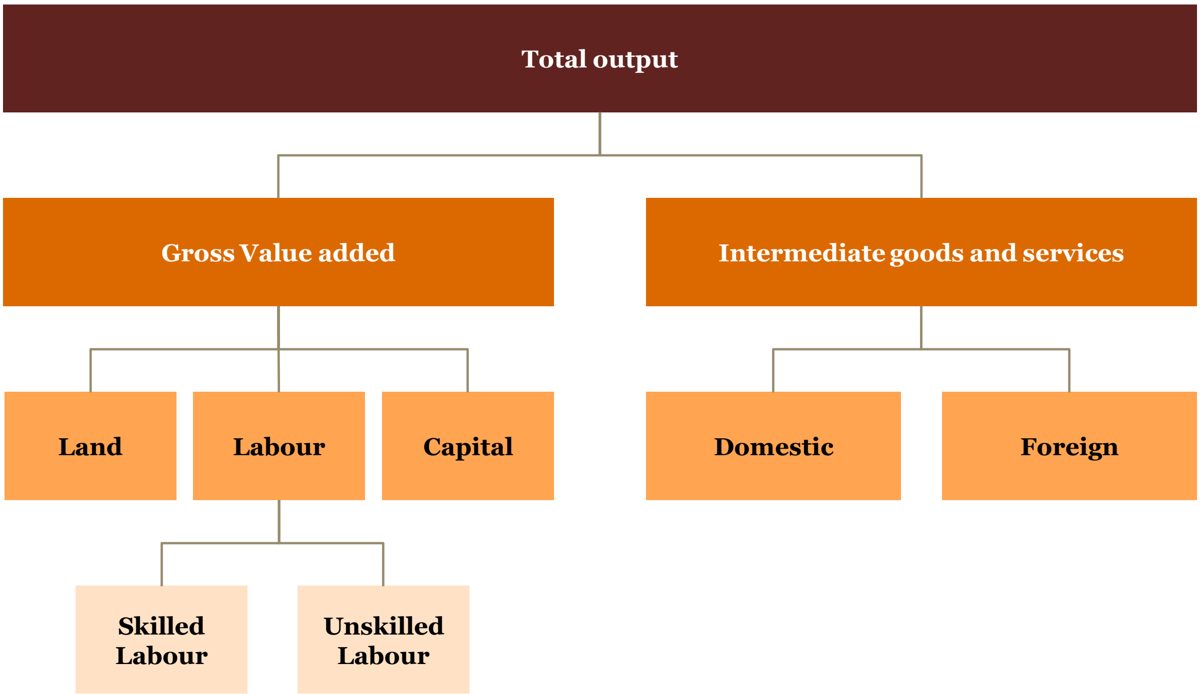

Source: PwC Analysis

7.3. Limitations of S-CGE models

Ultimately the S-CGE model used in this project should be taken at face value. Economic models are built on a range of complex assumptions that interact and need to be monitored carefully. When conducting global modelling exercises this problem can be made more acute due to the large volume of inter-sectoral and multi-geography data, which can be difficult to track and verify. There is a significant amount of quality assurance of this data within the GTAP community – this gives us confidence to use the model.

Further, the main assumptions in the model (e.g. elasticity values or functional forms) are chosen to track the expected global evolution of the economy and are not tailored for individual countries or sectors. These assumptions are made because there is simply not the availability of suitable country studies to provide a suitable range of parameter information. Such problems are often solved at individual country level, but more work needs to be done to produce, collate and share this information on a global level.

Both of the above listed problems are typical for global macroeconomic modelling in general – they are not exclusive to S-CGE models. S-CGE models in themselves can help mitigate these problems as they facilitate better options for sensitivity testing than other macroeconomic models (e.g. alternative functional forms can be tested more easily). They are also useful in that the data used is 'balanced' i.e. income equals expenditure across all sectors, households and government accounts which forces a much stronger degree of internal consistency than is observed in other macroeconomic models.

Finally, widely noted limitation of S-CGE models is that very few of them are validated against historical experience-although GTAP is one such model where external validation has been conducted, useful examples include: Valenzuela *et al.* (2007)[49] or Liu e*t al.* (2004)[50].

7.4. Model inputs

This subsection explains how we converted each separate piece discussed above in sections 4 to 6 above into percentage impacts, or 'shocks,' to productivity and other model parameters related to consumer demand and consumption as a result of anticipated AI uptake over time. We explain 1) how we quantified each of these cumulative input shocks, 2) how we assumed the uptake path to look over time, and 3) how we adjusted the uptake path to account for mitigating region-specific factors not captured by each work stream. First though, we discuss the exact components of the S-CGE model that were shocked using our model inputs.

We have conceptualised shocks to the economy from AI-uptake to either take effect at the production-side of the economy (through productivity and jobs impacts from replacement automation and augmenting AI), or the consumption-side of the economy (through product enhancements from inbuilt AI). More specifically, on the production side of the economy we shock labour productivity (output per hour worked), and on the consumption side of the economy we shock the variety (personalisation) of goods, quality of goods (marginal utility per consumption) and time saved from consumption of AI-enhanced products (available labour supply per worker). The focus on these specific channels is justified by the discussions in sections 2.3 and 2.4 respectively.

---

49 Valenzuela, E. Hertel, T. Keeney, R. & Reimer, J. (2005). *Assessing Global S-CGE Model Validity Using Agricultural Price Volatility* (GTAP Working Paper No. 32; American Journal of Agricultural Economics 2007, 89(2):383-397). Purdue University, West Lafayette, IN: Global Trade Analysis Project (GTAP). Retrieved from https://www.gtap.agecon.purdue.edu/resources/res_display.asp?RecordID=1875.

50 Liu, J. C. Arndt, and T.W. Hertel. 2004. 'Parameter Estimation and Measures of Fit in a Global, General Equilibrium Model.' Journal of Economic Integration 19(3):626- 649.

Figure 7.4 – Inputs into the S-CGE model to capture the impact of AI through productivity and product enhancements

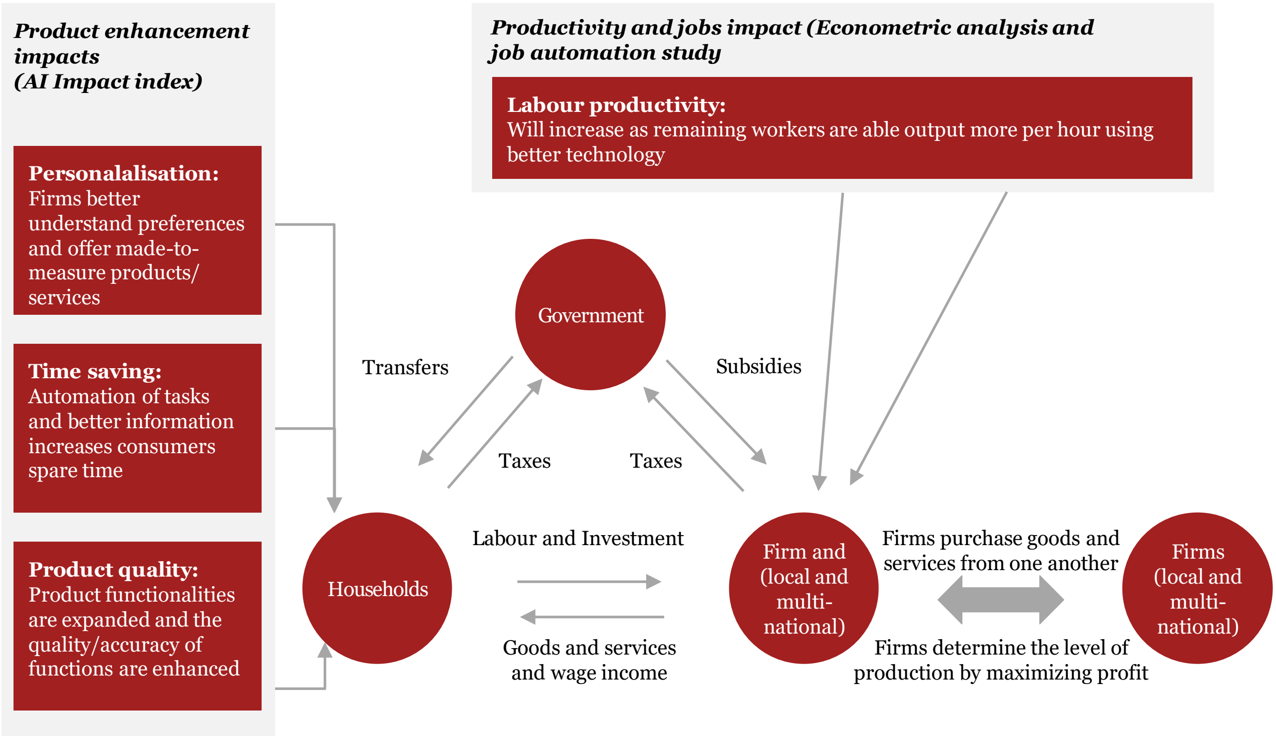

Source: PwC Analysis

7.5. Quantification of model inputs

Turning to how we quantified the model inputs, there were a number of parameters and moving pieces that had to be estimated to turn our AI index into a consumption product enhancement on the one hand, and our job automation study and econometric analysis into a labour productivity shock on the other. We discuss each of these in turn.

Quantifying the productivity inputs

To calculate cumulative labour productivity shocks, we multiplied the impact of each unit of AI (the elasticity evaluated in our econometric analysis) with an estimate of the scale of AI uptake per worker by 2030 for each region and industry considered in our study. This provided us with a cumulative labour productivity shock per region and industry between the current period and 2030. However, whilst the elasticity is directly available from our econometric models, estimating the expected scale of AI-uptake per worker (and automation) required a number of separate models and assumptions.

More specifically, according to our panel-data model specification, the cumulative shock to labour productivity in a given industry and region is a function of the following elements: the fraction of workforce remaining by 2030, the elasticity between productivity and AI-uptake (marginal impact of AI), the amount of real expenditure on replacement AI between now and 2030, and the amount of real expenditure on augmenting AI between now and 2030. The functional form of these inputs used to estimate productivity changes is laid out in Equation 7.1. This equation is derived from the specification of our econometric models, and the drivers of changes in AI-per-worker.

Equation 7.1 – Breakdown of cumulative labour productivity impact by region and sector as determined from econometric specification

Labour cost savings/ replacement effect

Augmentation effect

Cumulative labour productivity shock in period t from o

$$\%\Delta LPj_{t,o} = \hat{\beta}_j ln\left(\frac{\left[\frac{\Delta Ljtw_{jo}\left(\frac{A_{jt}}{K_{jt}}\right)+\Delta fj(t)}{A_{jo}}\right]+1}{\lambda_{jt}}\right)$$

Elasticity of productivity to automation

Fraction of remaining workforce

Source: PwC Analysis

In the diagram above for each sector group (j) and time period (t): L is the stock of labour, w is the real wage, A represents the stock of AI technologies, K is the capital stock, f(t) is a time-trend predicting augmenting AI uptake (see Figure 7.6), 'λ' represents the fraction of workforce remaining and 'β' is the estimated marginal impact of AI-uptake (per worker) on labour productivity. Each of the components listed above were derived from a different model or piece of work. As mentioned, the elasticity is estimated directly from our econometric models. Meanwhile, our estimates of each industry group workforce fraction remaining by 2030 were calculated using the machine-learning methodology as used in PwC's previous job automation study, as discussed in Section 5 of this report.

To estimate real expenditure on replacement AI, we used the real wage per industry and region in the base year used (in our proxy countries) to measure the real value of the stock of AI technologies per worker in our econometric models, and combined this with the estimated number of employees lost per industry indicated by the automation predictions of PwC's job automation study. We then achieved our final estimate of labour cost savings spent on replacement AI by multiplying this number by the current fraction of capital expenditure on AI technologies. This assumption represents 'business as usual', where industries are expected to spend on AI at least in line with their current pattern. This might be considered a conservative assumption since one might expect relative AI expenditure to increase over time as the technology spreads, but due to our inability to forecast their expenditure change accurately, we abstract from these changes where no quantification is feasible.

Finally, to estimate real expenditure on augmenting AI, we built separate, simple econometric models aimed at forecasting expenditure trends on augmenting AI over time. We specified a simple model of AI-uptake, as a function only of the labour force and a linear time trend as shown in Figure 7.6 (all variable definitions as before with alpha as an industry specific constant):

Equation 7.2 – Econometric model used to predict augmenting AI uptake over time

AI uptake

$$a_{jt} = \alpha_j + \beta_j l_{jt} + f_j(t) + \epsilon_{jt}$$

Industry group effect on level of AI uptake

AI uptake related to labour expansion or replacement

Time-trend in augmenting AI uptake unrelated to labour input

Source: PwC Analysis

The idea behind these models is to capture long-run (deterministic) trends in AI expenditure that are unrelated to labour force decisions. Under the assumption that a) AI uptake is limited to the augmenting and replacement types specified and b) augmenting AI is the component of AI uptake that is unrelated to labour force decisions, then the linear trend should purely capture the long-run trend in augmenting AI uptake, since it is the trend in AI uptake holding labour force constant. Additionally, we imposed the constriction on the regression that the trend must be positive or zero to remain consistent with the study's assumptions. Whilst we recognise this modelling is not as robust as we prefer, we have imposed a restrictive and simple model (notably with no non-linearities) to minimise overfitting and to improve external validity for long-run forecasting. Again, we do not pay attention to significance of the time trend due to the small sample used in estimation (T = 10-15 depending on the region), so the procedure can, like our econometric models for productivity, be thought of as providing the best, non-parametric point-estimate of the long-run trend in augmenting AI uptake.

One immediate caveat to note is that we apply our labour productivity shock to both types of labour in the model equally. In practice it is possible that that automation will be a substitute for low skilled labour and a complement to high skilled labour, which would mean the productivity enhancement from replacement and augmenting automation would only affect high skilled labour, thereby affecting the income distribution. However, if we believe some automation has already taken place, our econometric results should have captured the net productivity impact when spread across all types of labour still in employment, since we only had access to aggregate productivity across all skill types. Therefore the total impacts should be approximately correct in our S-CGE model, despite the income distribution varying differently in reality.

Quantifying the personalisation product enhancement inputs

We used a combination of academic literature and qualitative judgement to convert the AI Impact Index into relative percentage impacts of S-CGE model parameters to effectively capture the product enhancements specific to product quality, saving consumers time and offering more tailored and personalised products and services.[51]

To quantify increased personalisation, we adopted a two-stage approach which used the academic literature on willingness-to-pay for personalised goods in conjunction with known estimates of the (equivalent) effect on marginal utility of increasing goods varieties. The first stage involved developing a conversion mechanism between the AI Impact Index personalisation score to an equivalent percentage impact in marginal utility from personalisation. To do this, we first assume that a maximum AI Impact Index score of 5 corresponds to a fully personalised good, whilst a score of 1 corresponds to no personalisation at all. Doing this enables us to interpolate between estimates of marginal willingness-to-pay for personalised products, since we can consider willingness-to-pay as a stated preference approximation for marginal utility.

To calculate marginal utility estimates we used willingness-to-pay impacts from a previous study[52] on the impact of personalisation across a number of different product lines to create an average willingness-to-pay impact from personalisation for each of the industry groups considered in our study. This provided us with a personalisation impact per industry at the AI Impact Index 'bounds' (the impact being zero at the lower bound). Notably, we find that on average consumers are willing to pay an extra 25% for a personalised good over its non-personalised equivalent. We then interpolated between these bounds using a convex polynomial of order two[53].

---

[51] We consider changes in the personalisation of goods to equally reflect changes in goods' variety. This is because personalisation can either be conceptualised as the enhancement of existing products, or the replacement of the existing products with more personalised ones as new innovative firms enter the market. The former effect is less well documented in the literature, and does not capture the pro-competitive effects of dynamic firm entry over time as new (personalised) goods varieties enter the market. For this reason we conceptualise personalisation impacts as variety of goods impacts within the S-CGE model to capture this channel's prominence and to ensure competitive effects are captured.

[52] We used Deloitte's Consumer Review study on personalisation: *'Made-to-order: The rise of mass personalisation'* which, amongst other things, looked at the percentage increase in willingness-to-pay by consumer groups and for 15 different product lines: clothing, footwear, furniture and homeware, jewellery and accessories, holidays, beauty products, health and wellness, electrical products, food and groceries, flights, hotels, alcohol, restaurants, entertainment, and soft drinks.

[53] This is to capture the idea that consumers tend to view personalisation in discrete terms – with a product being 'not personalised' or 'personalised'. Therefore, our non-linear function used to interpolate between the bounds and attribute

The second stage was to convert these quantified marginal utility impacts into an equivalent shock to goods variety (which has a similar welfare impact – holding marginal utility constant), which would then be fed into the S-CGE model. Whilst in theory shocking marginal utility in the S-CGE model directly would provide a cleaner estimate of the personalisation impact, for reasons mentioned in footnote 51, we prefer to proxy this effect through the goods variety channel.

We were able to use existing estimates[54] of the impact of new product varieties on demand to calculate the percentage increase in expenditure at a given marginal utility of consumption due to a percentage increase in goods variety. This can proxy as a constant elasticity between goods variety and marginal utility. Since the study does not show exactly by what percentage goods varieties increased, under standard competitive market structure assumptions, we were able to calculate the elasticity to be roughly 1 between the two measures. This elasticity was then combined with our results from the first stage to finally map the AI impact index scores on personalisation to a (cumulative) percentage impact on goods variety over time.

Quantifying time saving S-CGE model inputs

To create a percentage impact of AI-uptake on time saved per region and sector, we used data from the mobile application 'Sleep Alarm' on average sleeping times per region (using our proxy countries as in the econometric modelling). This provided us with a proxy for average available time (in hours per day) per region. Since we consider time saved in this context as increasing the labour supply constraint for consumers, we then used a conversion mechanism put together by the US Data Analytics firm that equates each AI impact index score (in units) to a fraction of daily hours saved. Again, we interpolated between units using a convex polynomial of order two.

Quantifying the product quality enhancement S-CGE model inputs

Finally, for utility/quality, we had to use qualitative, reasoned judgement to create a conversion system between AI-impact index scores and percentage impacts to marginal utility to be input into the S-CGE model; this was created by canvassing expert opinions amongst PwC's UK Economics team about the **lowest** possible marginal utility impact that could reasonably be associated with a given AI impact index score. This was largely shaped by the scoring criteria itself, which can be found in Section 6.3. Since this impact was derived using a less robust methodology than other shocks in the S-CGE model, we have performed sensitivity analysis on this where we further reduced the shock by 35%.

7.6. Impact path over time

Whilst our inputs were used to calculate a cumulative percentage impact on labour productivity and product enhancements by 2030 (above baseline), they could not tell us how the global economy would arrive at its predicted state in 2030. Therefore, as mentioned previously, we have assumed that the impact path of each shock over time follows an S-curve in line with the literature and evidence on technology diffusion and innovation. To create the S-curve in practice, we simulated a scaled version of a Logistic cumulative distribution function (CDF), where the CDF was scaled to cover the interval [0,X] where X = the cumulative percentage impact for a given shock (in a given region and sector).

Although this gives us an S-shape uptake profile, in our view the tails of the profile are excessively thin and are unlikely to reflect the smoother reality of AI-uptake and its impact over time. As a result, we took a flat)linear combination of this uptake profile, with a linear uptake profile, to arrive at what we have called a 'hybrid' uptake profile. This final hybrid profile has much fatter tails and is smoother than the original simulative Logistic CDF and has been used in the main scenario of our analysis. A direct comparison of these different profiles is available below in Figure 7.7.

a marginal utility increase to each AI impact index score, captures this by ensuring that most personalisation impacts to marginal utility are only captured if the product is sufficiently personalised.

54 In particular, we use a paper Arnade, Gopinath and Pick (2011) which looks at the impact of new brand introductions on demand and consumer welfare in the US using data on 11 major cities. They are able to separate out the compensating variation measure of utility into a 'price' effect and a 'variety' effect.

Figure 7.7 – Logistic, linear and hybrid impact paths for S-CGE model

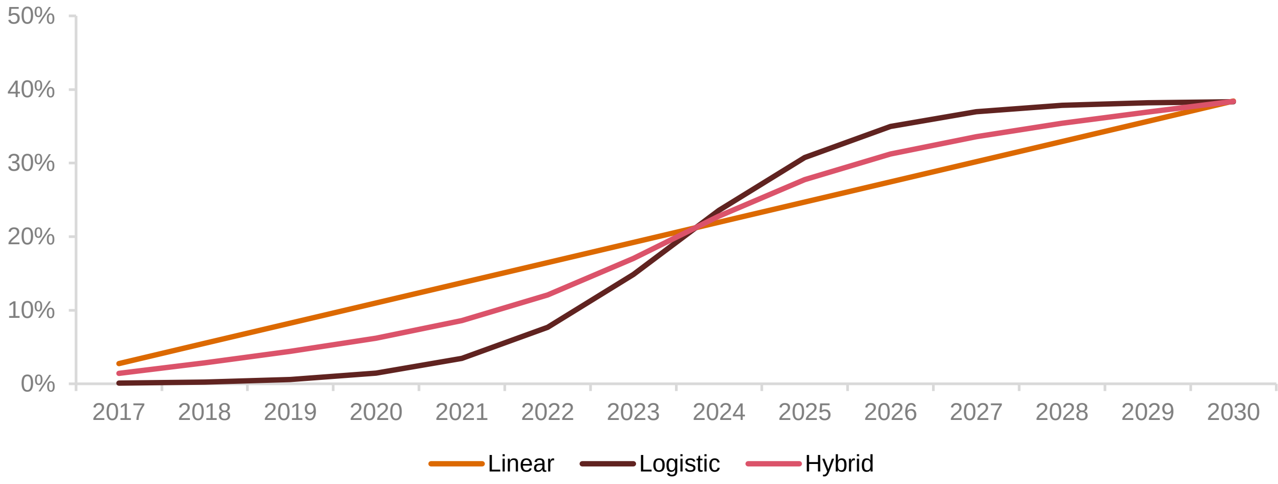

Source: PwC Analysis

7.7. Mitigating factors used for path adjustment

We have also taken steps to evaluate what fraction of predicted impacts from AI will actually take place in each region, and have corrected our model inputs to reflect this. This is because the fundamental drivers of our shocks do not account for potential mitigating factors that could limit the uptake of AI in practice – such as regulatory environment, corruption, and business 'AI-readiness'. As a result, we have used the Global Innovation Index (GII)[55] as a way to proxy the relative difference between our initial predicted rate of AI-uptake over time and the actual expected rate in the different geographical regions. The index combines 81 indicators to 'provide a broad vision of innovation,' from the perspective of both innovation inputs and outputs, including the political environment, infrastructure, human capital and business sophistication.

In our scaling methodology, we assume that the top GII score is effectively a 'perfect world' score, where there are no barriers that would cause less AI-uptake to occur between now and 2030 than predicted. As a rough scaling proxy, we then use the inverse of the fraction of each region's relative GII score to the top score as a proxy for the gross percentage of time it takes this region to complete all AI-uptake it was initially predicted to uptake by 2030 (i.e. if this inverse fraction was 130% for some region A, then region A would take 30% longer than predicted to complete the AI-uptake that was predicted to occur by 2030). Note that since the GII only provides country-specific scores, we created regional average scores by taking weighted averages over countries per region (weighting by GVA in 2016).

The results are relatively unsurprising and lead us to model North America, Northern Europe and Developed Asia as the fastest adopters of AI with Latin America lagging behind and China and Southern Europe somewhere in the middle.

7.8. Model results

In our main scenario, Global GDP could be up to 14% higher in 2030 as a result of AI – the equivalent of up to $15.7 trillion, more than the current output of China and India combined.

Labour productivity improvements are expected to account for over 50% of all GDP gains related to AI over the period 2017-2030. However as new technologies are gradually adopted and consumers respond to improved products with increased demand, the share of impact from product innovation and enhancements increases over time. By 2030, although GDP could increase by $6.6tn (5.8%) as a result of productivity gains, consumption-side inputs are expected to account for the bulk of the impact at $9.1tn (8.0%).

Of product enhancement impacts, nearly all the GDP impact is expected to derive from increases in the varieties of goods as well as increases in goods quality, with only a negligible impact from increases in time savings for

55 Cornell SC Johnson College of Business, INSEAD and World Intellectual Property Organisation, 'The Global Innovation Index 2017: Innovation Feeding the World,' Tenth Edition.

consumers using AI-enhanced products. This not only reflects the fact that the base (%) increase in time savings is relatively small on average over the year (per consumer), but also that this labour supply (availability) increase is not significant enough to actually incentivise an increase in the supply of labour meaningfully (i.e. this reflects a reasonably low – but empirically driven – Frisch elasticity of labour supply in the model).

Figure 7.8 – The global GDP impact by effect of AI 2017-2030

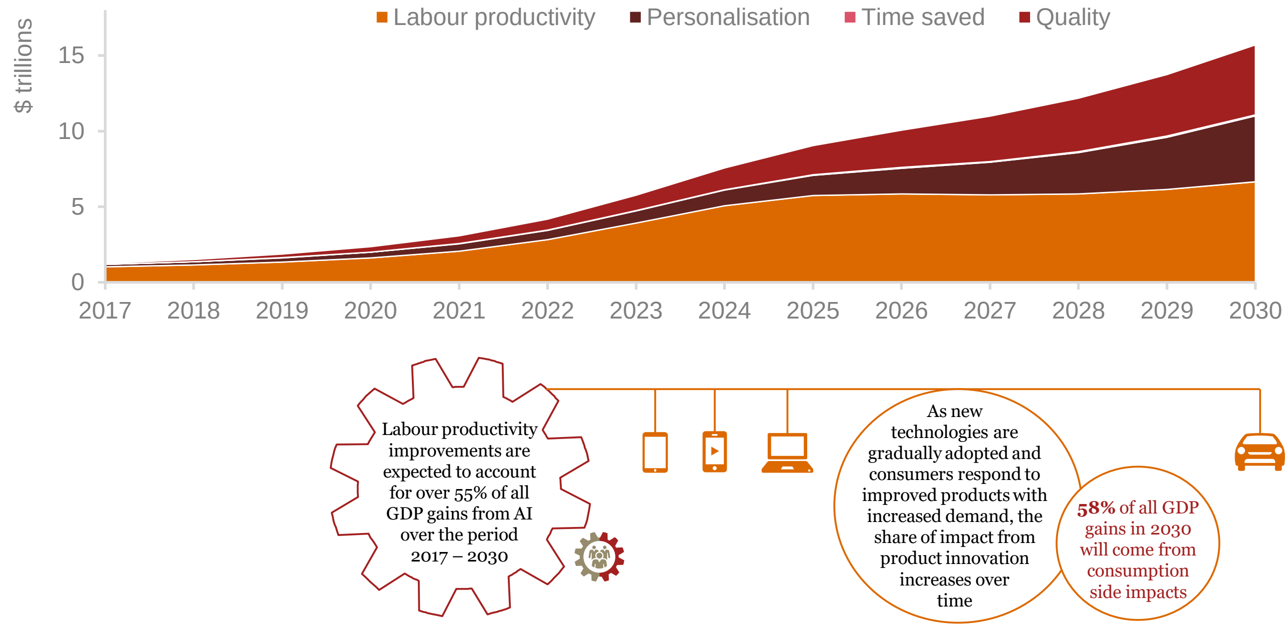

Source: PwC Analysis

Notably, the consumption-side impacts are more delayed but are surprisingly large in nature, overtaking the labour productivity contribution to GDP gains in the late 2020s. Both the delayed and large nature of these impacts can be explained by the complex (initial) transmission mechanism from these product enhancements to consumption – outlined in detail in Section 2.4.

Since the consumption-side product enhancements transmission mechanism – and particularly firm entry – takes considerable time, the effect of product enhancements on GDP takes considerably longer to appear than the productivity driven effects. But, the impact on GDP is substantial once the transmission has taken place, predominantly reflecting the size of the goods affordability increase from new firm entry.

However, there are two important caveats which must be considered when interpreting our analysis.

1. Our results estimate the upwards pressure on GDP as a result of AI only, under the *ceteris paribus* assumption.[56] Our results may not be directly reflected in future economic growth figures, as there will be many positive or negative forces that either amplify or cancel out the potential effects of AI (e.g. shifts in global trade policy, financial booms and busts, major commodity price changes, geopolitical shocks etc.).
2. As mentioned previously, our economic model results are compared to a baseline of long-term steady state economic growth. The baseline is constructed from three key elements: population growth, growth in the capital stock and technological change. The assumed baseline rate of technological change is based on average historical trends. Therefore, since AI has already been introduced prior to the starting evaluation period of this study, the component of these forecasts driven by technological change will already have factored in past trends in AI's GDP impact. As a result, it is difficult to quantify the exact fraction of AI's GDP impact that will be additional to historical average growth rates (i.e. additional to the baseline forecast).

However, our study is specifically focussed on the AI technologies that are yet to be implemented and are conceived to be implemented between 2017 and 2030. As a result, an underlying but reasonable assumption we

---

[56] This implies that all other factors of the economy remain 'as expected' and do not suffer shocks which could deviate the economy from its predicted outcome in this study.

make here is that the scale and impact of these AI technologies will be above the current trend in AI's uptake and impact. Under this premise, our study is centred on estimating the total marginal economic impact of yet-to-be-implemented AI specifically between 2017 and 2030 – not including the AI that has already been implemented prior to this study (which is implicitly included in the baseline growth assumption). This also means that, whilst our study results imply that average economic growth rates between 2017 and 2030 will be raised due to AI's impact, we do not make claims outside of this time interval. As a result, we do not interpret that AI will impact the fundamental long run growth rate of the global economy.

These two factors mean that our results should be interpreted as the potential 'size of the economic prize' associated with AI over the period of our study, as opposed to direct estimates of the impact of AI on long-run economic growth.

One other economic caveat which must be addressed is that it is assumed that firms can enter and compete in the future under identical circumstances as today. This assumes that data ownership – as a mechanism for doing business – will be regulated under the same anti-trust considerations that the pre-AI economy enforced. However, if firms maintain almost exclusive ownership of data, or are able to build a sufficiently large moat around them, increased dynamic entry and competition may be less feasible.

7.9. Geographical impacts

All geographic regions of the global economy will experience economic benefits from AI. North America and China stand to see the biggest economic gains with AI enhancing GDP by 26.1% and 14.5% in 2030 respectively, equivalent to a total of $10.7 trillion and accounting for almost 70% of the global economic impact.

Beyond North America and China, other countries such as those across Europe and the more developed countries in Asia are also likely to experience significant GDP gains of around 9.5-11.5% of GDP by 2030. Although the adoption of AI is projected to be slower in these countries than in the North American region, the potential for automation is high in Europe, while the marginal impact of AI technologies on productivity is particularly high in developed Asian economies, as is projected investment in workforce augmenting technologies.

On the consumption side, fast adopters of AI are likely to see the greatest gains (North America and Northern Europe), although China will see a disproportionately large share of the benefits due to the altogether slightly lower level of competition in their landscape, which increases the marginal impact of firm entry on prices, discussed in more detail below. Latin America and other less developed markets are expected to lag behind somewhat, though despite lower uptake of AI themselves are still expected to see GDP gains of approximately 5% of GDP in 2030.

Figure 7.9 – Economic impact of AI by geographical region

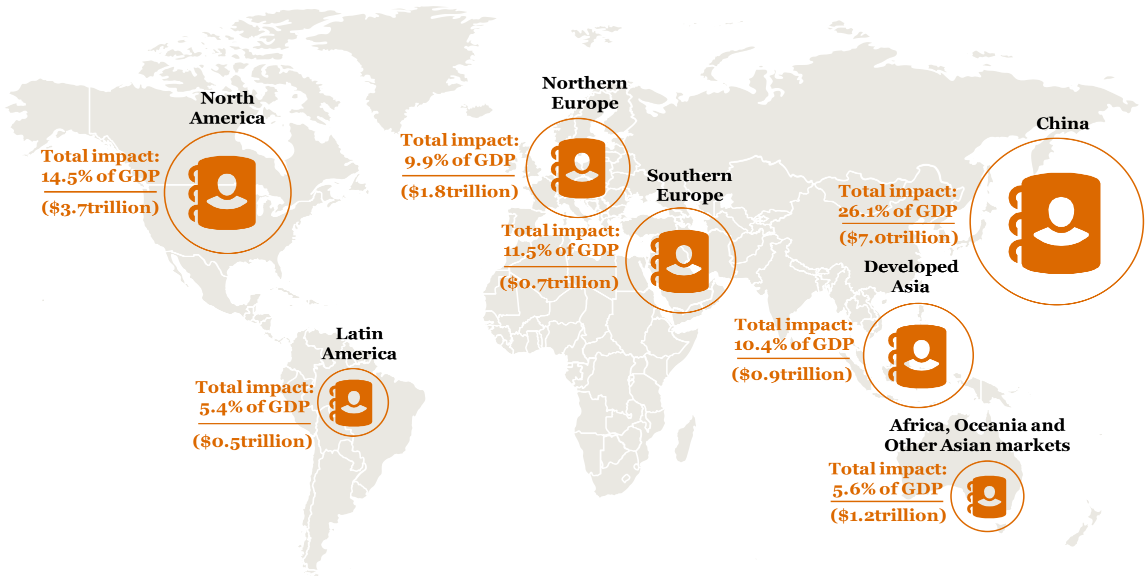

All GDP figures are reported in market exchange rate terms

All GDP figures are reported in real 2016 prices, GDP baseline based on market exchange rate basis

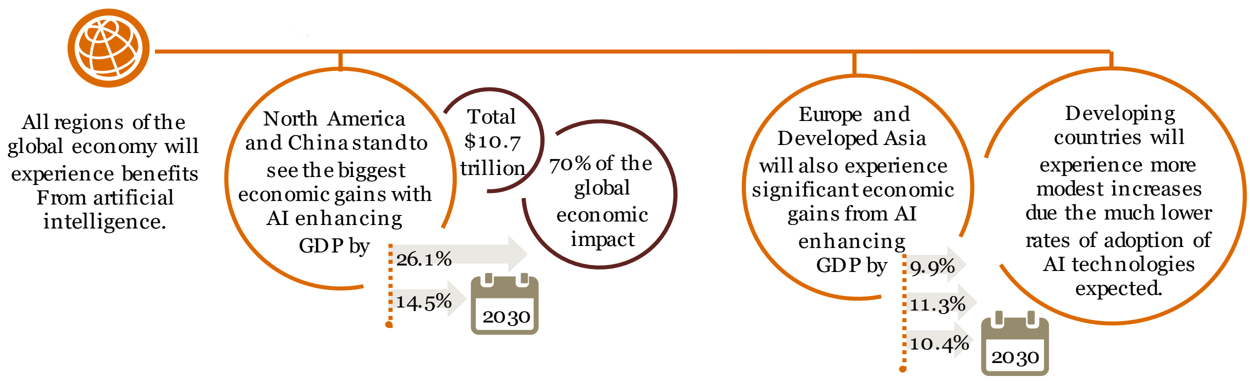

Source: PwC Analysis

The distribution of the impacts of AI by channel of impact, through productivity or through product enhancements on the consumption side provide further explanation for the geographical impacts. Globally, the impact resulting from productivity gains accounts for approximately 42% of the impact of GDP by 2030. However, the figure varies widely across geographic regions as shown in Table 7.2 below.

China and North America follow a similar pattern to the global average, with the split between the two channels close to half and half, as expected given that together they account for 70% of the global impact. GDP impacts resulting from productivity gains are expected to account for 46% of the impact in 2030 in North America and 51% in China.

However, in all of the other geographic regions, GDP impacts associated with product enhancements are expected to far outpace the GDP impacts associated with productivity by 2030. For example, productivity gains are only expected to account for 23% of the impact on GDP in Northern Europe and only 31% in Latin America.

The main reason for this fluctuation is that the scale and GDP impact of product enhancements vary less by region than the productivity based impact – as measured by the regional coefficient of variation in Figure 7.10. This is mostly driven by the high coefficient of variation in productivity inputs per region, since the productivity-driven GDP impacts of AI are mostly driven by the size of the inputs. This is shown by the high cross-sectional correlation across regions between productivity input sizes and productivity-driven GDP impacts, as well as their respective coefficients of variation in Figure 7.10 (by region).

In Northern Europe and Latin America, automation potential is somewhat lower, as is the marginal impact of AI on productivity, limiting the AI impact from productivity increases. Furthermore, a large share of the GDP impacts in these regions associated with product enhancements come from gains through trade in AI-enabled products with other AI leading countries.

Figure 7.10 – Coefficient of variation of productivity GDP impacts, consumption-side GDP impacts and their respective inputs across regions, and correlation between productivity and consumption-side GDP impacts with their inputs across regions.

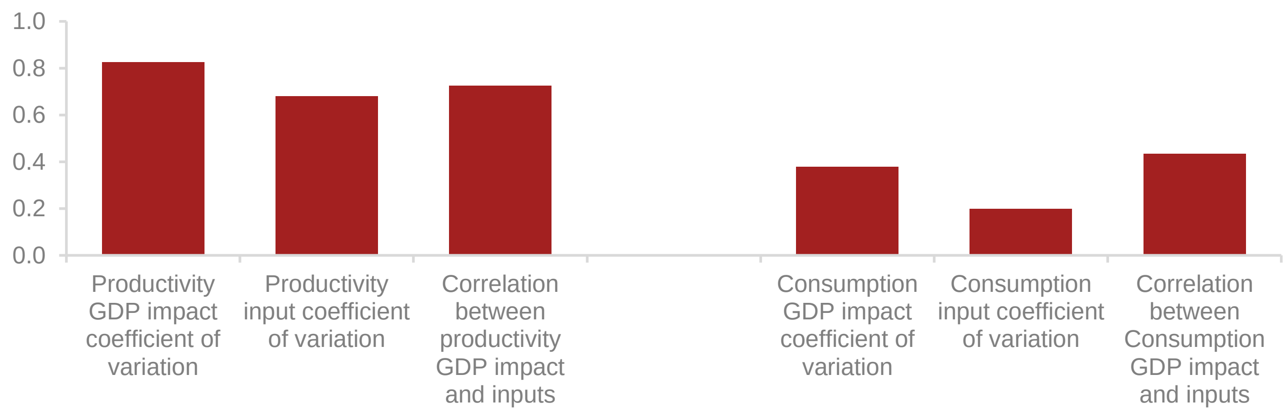

Source: PwC Analysis

Table 7.2 – GDP impact of AI by geographical region and channel of impact

(%)	GDP impact associated with productivity	GDP impact associated with product enhancements	Total GDP impact
North America	6.7	7.9	14.5
China	13.3	12.8	26.1
Developed Asia	3.9	6.5	10.4
Northern Europe	2.3	7.6	9.9
Southern Europe	4.1	7.5	11.5
Latin America	1.7	3.7	5.4
Africa, Oceania and other Asian markets	1.1	4.5	5.6

Source: PwC Analysis

7.9.1. Regional focus: why is the impact in North America and China so large?

The impact in North America and China is much larger in total than other regions and so worth examining in more detail. These impacts can come from input-related reasons (outside the model) and structural economic reasons (inside the model).

North America

In North America, the potential uplift to GDP from AI is mainly amplified by the larger-than-average GDP gains the region is expected to see as a result of AI-driven labour productivity enhancements (outside the S-CGE model). Figure 7.11 breaks down the sources of AI's impact on productivity itself (the productivity input) into five components, and shows which of these components are particularly high for North America (and China) relative to the other regions modelled.

Looking at Figure 7.11, it is clear that in North America, advanced technological and consumer readiness for AI is particularly high, enabling a faster effect of AI on productivity and overall larger effect by 2030. Meanwhile, high real wages and the largest automation potential by 2030 mean that firms will be able to automate a huge number of existing jobs and replace these with top-quality AI-technologies using the generous labour cost savings as a result. These factors are key in driving the expected labour productivity gains in North America and by extension the GDP impacts.

China

Whilst in North America it is purely a large productivity-driven GDP impact that drives their sizeable gains, in China both productivity and product enhancement (consumption-side) driven GDP effects are considerably larger than in all countries. This is driven by two factors: 1) labour productivity driven factors in China that mean the AI impact on labour productivity is particularly large in China (outside the model), and 2) indirect, underlying economic factors that enhance both the consumption-side and labour productivity driven GDP impact of a given AI input (inside the model).

Turning first to aspects driving the labour productivity input, these gains are driven by somewhat different factors than in North America. In particular, as discussed in Section 4, the marginal impact of AI-uptake on productivity in China is estimated to be the largest out of all regions by some distance. This – as mentioned – is most plausibly driven from implementing high-quality AI against a low labour productivity base, which is likely to have a larger productivity impact than in regions where labour productivity is already much higher.

Moreover, labour cost savings are expected to be large in China and, although they are equally so in the US, this is not driven by high real wages but rather a large workforce. Although one would expect this effect to be balanced out since workers need to be proportionally replaced by AI, it is the fractional increase in automation and AI-uptake that matters most to labour productivity. Since China's current stock of AI-technologies is low relative to their workforce size, these labour cost savings are expected to be spent increasing this stock by a larger fractional amount than in other regions, therefore stimulating the labour productivity impact over and above the other regions.

Turning to the structural economic factors in China (inside the model), the main benefit over other regions comes in the form of lower starting levels of competition. As mentioned in Section 2, since both labour productivity and product enhancements drive profit rises, this leads to dynamic firm entry over time. This entry of new firms – bringing new, better, more personalised products to market with efficient production processes – places downwards pressure on goods prices, making them more affordable to consumers, which in turn stimulates demand further as more consumption becomes possible at any given level of income. The relatively less competitive landscape in China means that the marginal impact of firm entry on prices is much larger than in other regions. This also explains why the impact in China is, overall, more delayed than in North America – even after controlling for the slower uptake expected.

China's position as a key global exporter also enhances both product enhancement and productivity effects on GDP, since they can export their technologies to other regions – both leaders and laggards. This is examined in more detail below in Section 7.10.

Figure 7.11 – Relative performance (vs. global average) of each source of AI-driven labour productivity impact for North America and China (greater relative scores depicted by corners)

North America | **China**

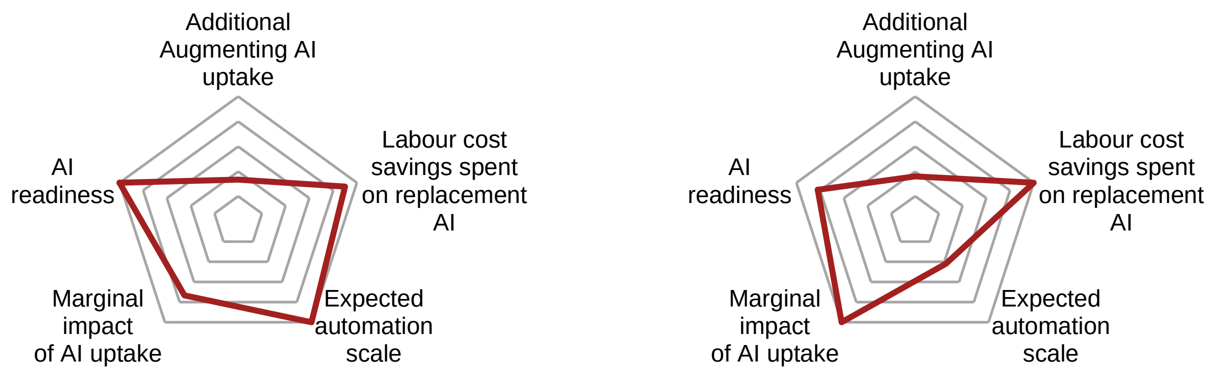

Source: PwC Analysis

7.10. Impacts over time

The impacts resulting from product enhancements on the consumption side follow a simple and steady pattern with the GDP impact in each geographic region beginning low and growing modestly over the first few years of the period. Once AI technologies begin to be taken up much more quickly, around 2021/2022 – the upturn of the S-shaped adoption curve – the GDP impact gathers pace (also following from the slower transmission mechanism, which relies on dynamic firm entry). In most geographic regions the GDP impact from product enhancements slows again slightly towards the end of the period – as the top of the S-curve is reached – except in China, where the impact is expected to grow in an exponential fashion after the upturn, largely because of China's status as a key net exporter and due to the larger impact on competition and prices in Chinese markets due to the market structure.

On the other hand, the impacts resulting from productivity gains follow considerably different patterns across geographic regions. As explained earlier, North America is expected to see large GDP gains related to productivity almost immediately and see this impact grow quickly within the next 8-10 years, peaking around 2025. Similarly, China will also see productivity impacts from AI grow quickly over the next years, but not at the same pace. However, after around 2025, China's percentage impact from productivity begins to outpace that of North America as productivity itself in China begins to catch up, stimulating exports of AI-enabled products from China to North America. To a lesser extent, these exports are also expected to slightly reduce the GDP impact resulting from productivity in Northern Europe towards the end of the period.

In general, all of the other geographic regions follow a steady pattern similar to the pattern of consumption side impacts with the impact growing over time with a slight increase in pace at the time of the upturn in the adoption S-curve.

Figure 7.12 – GDP impact by geographical region over time (% of GDP) associated with productivity and product enhancements (2017-2030)

GDP impact associated with productivity

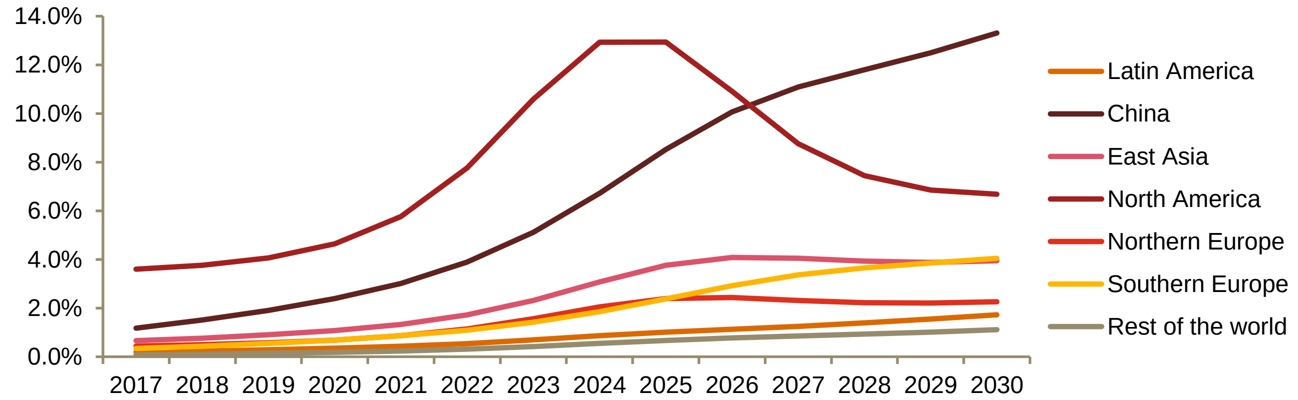

GDP impact associated with product enhancements

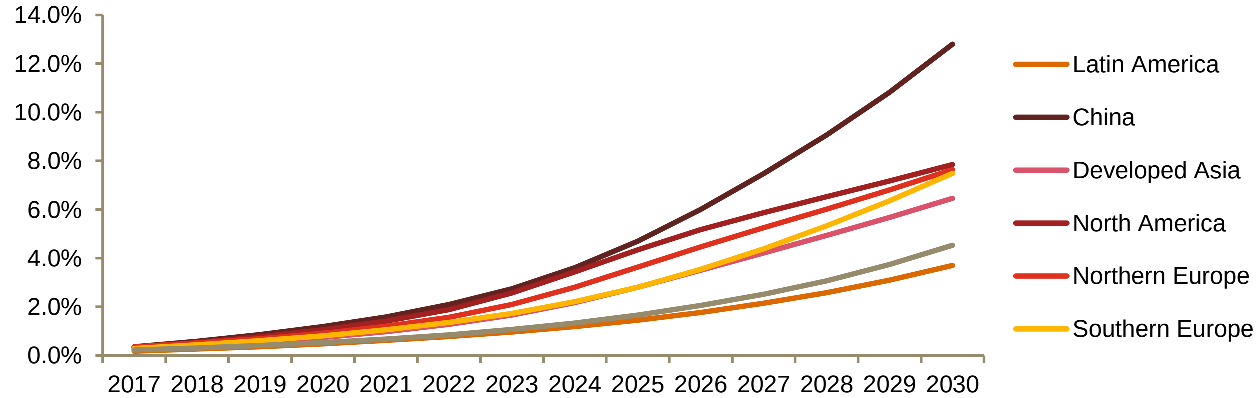

Source: PwC Analysis

Is AI Growth enhancing?

A key question regarding our results is whether AI is expected to have permanent growth enhancing effects on GDP or whether it just a factor that leads to a temporary change in economic growth. Whilst it is true that the shocks from AI to labour productivity and product aspects (quality, time saved, personalisation) are by definition permanent, this does not lead to permanent growth rate changes. There will be permanent effects on the GDP level over and above baseline by definition, but since there are no more permanent shocks considered in our study beyond 2030, the impact on growth rates over and above baseline will die out once the shocks have run through the system entirely. Therefore, any impact on growth rates implied from our study only hold over the period we consider (2017-2030), and should not be interpreted as impacts on the long-run growth rate of the global economy.

7.11. Sectoral impacts

Economic gains from AI will be experienced by all sectors of the economy, with each industry expected to see a gain in GDP of at least 10% by 2030. Figure 7.13 below illustrates how the impacts will be distributed. The services industry, which encompasses health, education, public services and recreation, stands to gain the most (21%), quite significantly more than any other sector. This largely reflects the scale of the AI impact on productivity and particularly product enhancements for this industry group. This is in itself influenced heavily by the presence of health in the sector group, which is expected to see large personalisation and quality improvements through the introduction of the likes of camera-based healthcare apps for analysing images and diagnosing condition, digital customer service assisted by virtual assistants as well as clinical workflow optimisation.

Meanwhile, retail and wholesale trade and accommodation and food services, as well as other labour-intensive sectors, are also expected to see a large boost (15%). Transport and logistics and financial and professional services will gain slightly less, but are still expected to see 10% higher GDP in 2030 as a result of AI. However, financial services are set to see fast, tangible gains in the short to medium term. This comes as a result of key expected AI innovations such as Deep Learning and speech analytics enabling asset managers to better understand sales and customer behaviours, ensuring they deliver higher customer satisfaction and a more generally superior customer experience.

The contribution to the total economic impact from productivity versus product enhancements is relatively constant across industry sectors. Transport and logistics as well as financial and professional services are expected to gain the largest impact share from product enhancements, with 60% of the impact coming through that channel. The other side of the same coin means that those sectors are expected to see the smallest share of their GDP impact derive from productivity gains. Since the automation potential and marginal AI impact on productivity is not low in these sectors, we believe the limited productivity impact comes from the trade flows within and between countries (between sectors in the form of intermediate input demand) which, as outlined in Section 2, can result in winners and losers both between and within sectors depending on changes in the patterns of input demand.

At the other end, capital-intensive sectors, such as energy, utilities and mining as well as manufacturing and construction are set to see the largest AI impact share from productivity gains, with this channel of impact accounting for approximately 44% of the total GDP impact in 2030. This not only reflects the fact that these sectors are expected to reap greater productivity rewards from AI due to their capital intensive nature, but also that as a result of these sectors having little front-end consumer contact, their products are generally less enhanceable as a result of AI and therefore contribute to a smaller share of the overall impact.

It should be noted that all of these sectoral results are in percentages. Because of the relative sizes of sectors, some sectors may experience larger absolute gains. For example, manufacturing and construction in China is expected to experience a GDP impact of $2.5tn, the largest absolute GDP gain as a result of AI by far. This is largely the result of the size of the industry, which accounts for 60% of China's economy, though the high level of relative automation potential also plays a role.

Figure 7.13 – GDP gains in 2030 resulting from AI by industry sector (% of GDP)

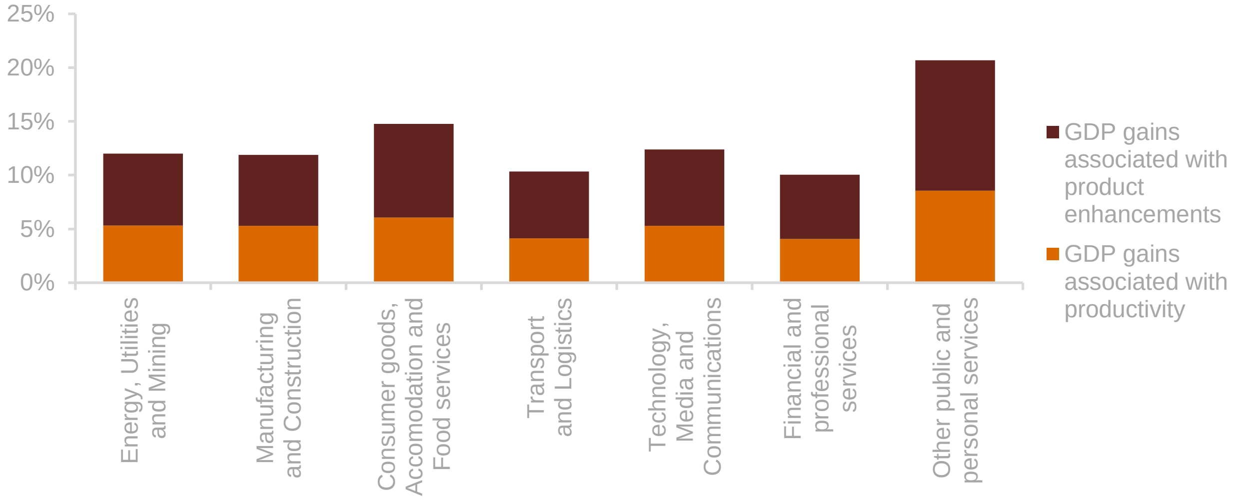

Source: PwC Analysis

Interestingly, the distribution of impact across industry sectors remains relatively consistent across the different geographical regions. Figure 7.14 below indicates that the health, education and other services sector is expected to see the largest GDP percentage impact in all regions except for Latin America and Southern Europe where the consumers goods/accommodation and food services sector is expected to see slightly larger impacts. In general, the pattern of sectoral distribution by country tends to be driven predominantly by the relative size of the productivity versus product enhancement impacts by country (not least since product enhancement impacts vary much more minimally by country). For example, the manufacturing and construction sectors in North America and China, where productivity is expected to account for more of the GDP impact, are expected to see the largest boost relative to the average across all sectors in those countries. In comparison, in geographic regions where product enhancements account for relatively more of the GDP gains, such as Southern Europe and Oceania, Africa and in other Asian markets, manufacturing and construction will gain the least of any sector for reasons outlined in Section 8.5 and as shown in Figure 7.14.

Figure 7.14 – GDP impact in 2030 resulting from AI by industry sector, geographical region and channel of impact (% of GDP)

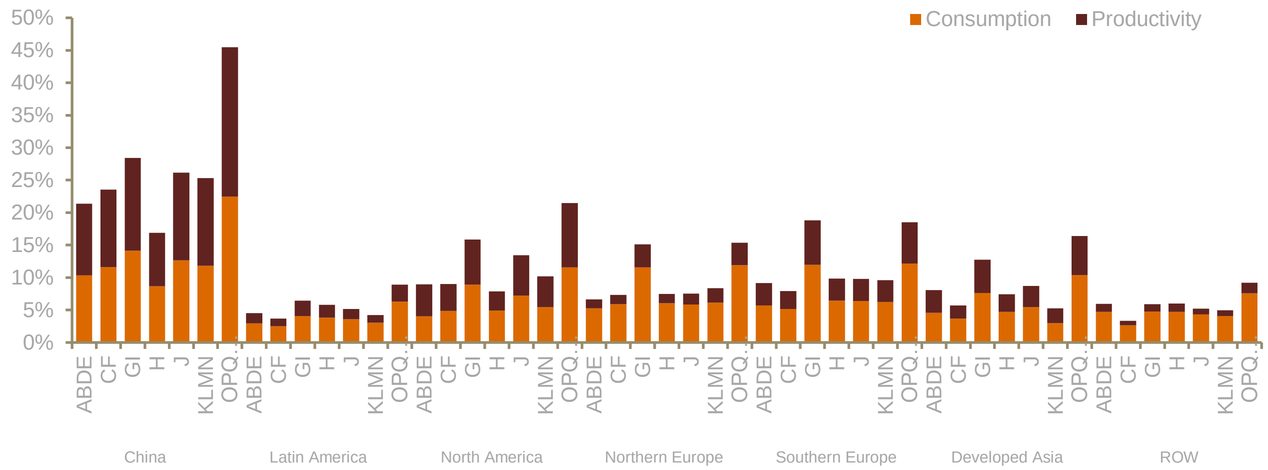

Key:

ABDE = Energy, utilities and mining **CF** = Manufacturing and Construction	**GI** = Consumer Goods, Accommodation and Food Services **H** = Transport and Logistics	**J** = Technology, Media and Communications **KLMN** = Financial and Professional Services	**OPQRS** = Health, Education and Other Public and Personal Services

Source: PwC Analysis

However, market results also reflect the prevailing market structures and trade patterns. More competitive markets will see less of a GDP impact from new (personalised) product entry and higher quality products, whilst any relatively fast-adopting countries with trade surpluses are expected to be able to export their technology and enhanced products in return for a sizeable extra GDP gain. As an example, China (40%) are expected to see a particularly large impact in the health, education and other services sector, which reflects the less competitive market structure in Chinese markets, as well as the size of expected product enhancements and productivity in this sector. For comparison, the impact on this sector in North America is only 26%, but the size of the AI impact on product enhancements and productivity is comparable to China, as Figures 7.15 and 7.16 show.

Figure 7.15 – Cumulative labour productivity input by region and sector from AI uptake (by 2030)

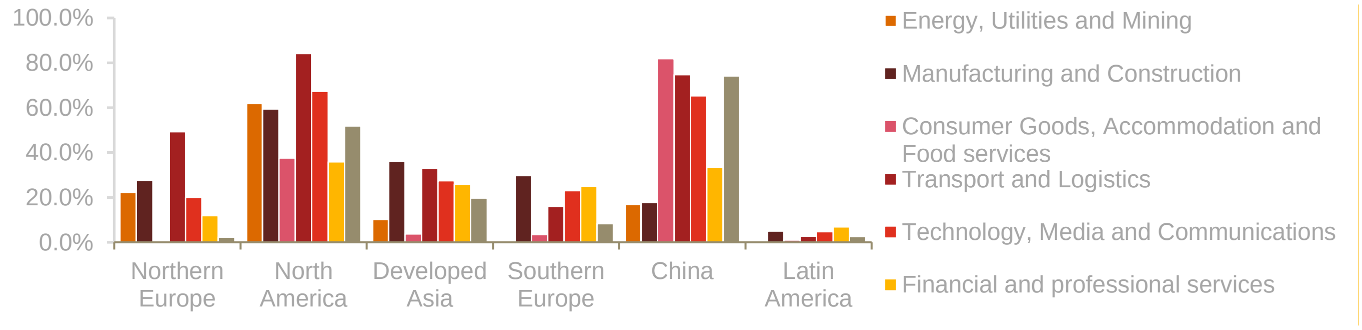

Source: PwC Analysis

Figure 7.16 – Cumulative product enhancement inputs by region and sector from AI uptake (by 2030)

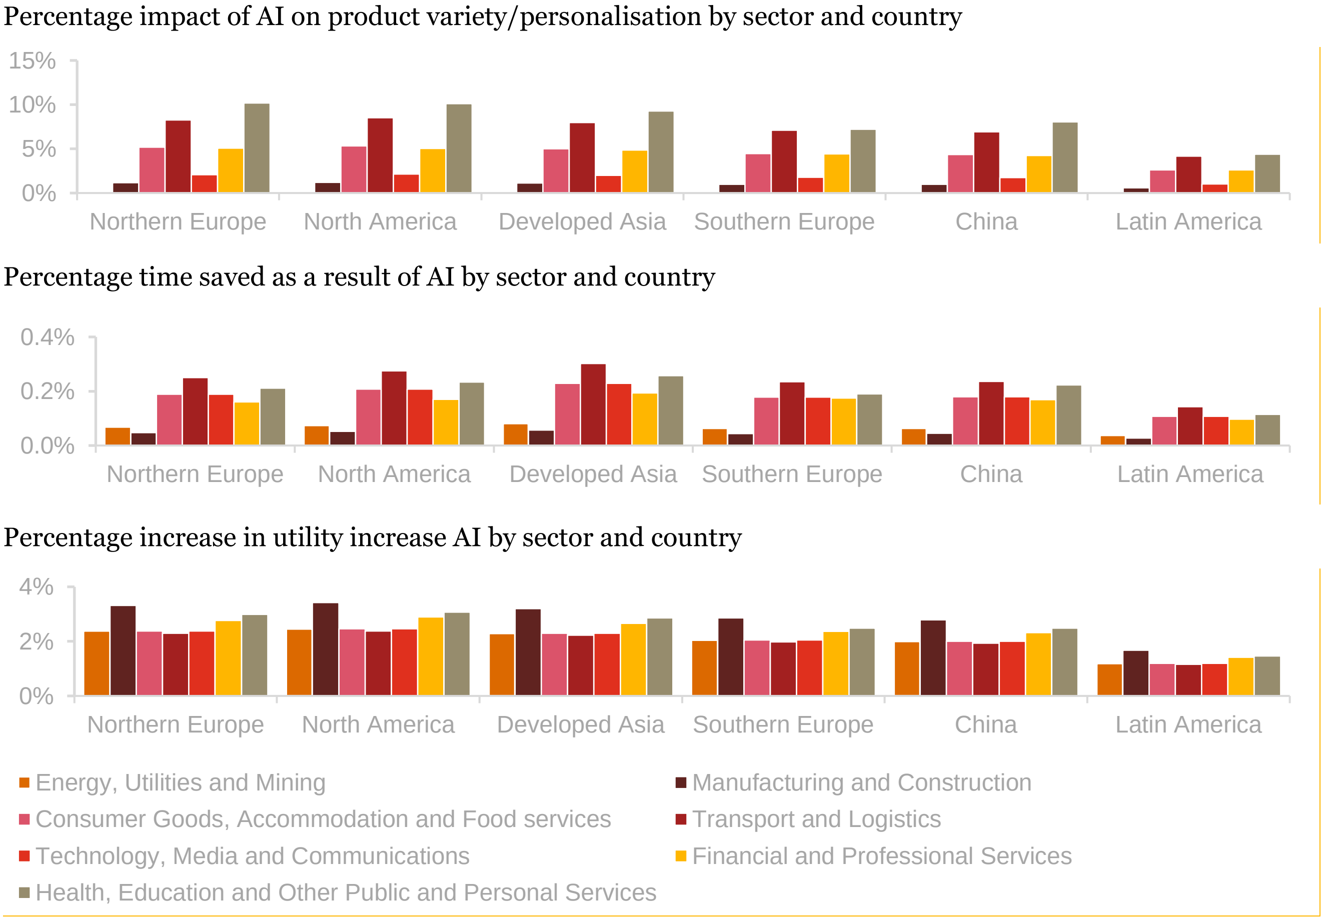

Source: PwC Analysis

7.12. Impact on labour

The impact of AI on jobs and the labour market is likely to be very significant and in general positive, but we cannot determine whether job creation alone will be net positive or negative. In Section 5.6 of this report we outlined the potential for AI to reduce the number of jobs through the creation of new roles and opportunities to related (a) the uplift in labour demand associated with AI's productivity and consumer demand boost and (b) the necessity of new roles for explain, train and sustain AI technologies. We estimate that 326m jobs will be impacted by AI in 2030. As with our analysis of GDP impacts, the impact on labour doesn't necessarily signify the new jobs that will be created as a direct result of AI, but the number of jobs that will come to depend on and be heavily impacted by AI. So our model suggests that by 2030 around 10% of jobs will to some degree be dependent on AI. By 'dependant' we mean that these roles are either created through AI or rely on AI to the extent that the role would no longer exist without AI, holding all other aspects of the economy constant.

Box 7.12

The total number of jobs impacted by AI is estimated to be over **326m** by 2030.

67% (218m) of these jobs will be unskilled whilst **33% (107m)** will be skilled.

Surprisingly, most of these jobs will be unskilled (though proportionally skilled jobs will be more positively impacted). 67% of the jobs in 2030 that will depend on AI will be the unskilled jobs, though this should be interpreted in the context of unskilled labour accounting for 69% of jobs in the baseline scenario. Our results therefore support the effect of skills-biased technological change, particularly since our labour productivity shock was not applied heterogeneously across different skills of labour. In practice, skilled labour would receive the bulk of the labour productivity gains, in which case the skills-bias in jobs would be much more significant. However, we are confident our aggregate numbers are robust to this nuance.

Impact on labour by geographical region

AI doesn't have a uniform impact on the labour markets across the geographic regions that we have focused on, as the impact depends on the degree of automation and capital levels. As seen in Figure 7.17, more than 50% of the jobs impacted by AI will be within China, with North America and Latin America making up the largest portions of the other regions. Southern Europe and the Developed Asia region will have relatively fewer jobs impacted and dependent on AI. Part of the reason for such a large impact on labour in China is the country's large base labour force of 768m in 2016. Another key reason is the large change to capital efficiency as technology levels 'catch up' with the US – boosting GDP and impacting on the labour market.

Figure 7.17 compares GDP impact of AI per job impacted. According to these results the impact of GDP per job impacted in the US ($113,332) could be significantly higher than in other geographic regions. This result is mainly due to the process of automation in the US predicted to result from AI and relates to the fact that the GDP figure here includes economic boosts through jobs that have been automated (GDP increase but job decrease) as well as the boost through new jobs (GDP increase and jobs increase). Because automation levels are expected to be high in North America, and the productivity increase resulting from this automation is also expected to be particularly high, the GDP relatively to impacted jobs is inflated.

GDP per job impacted are more evenly spread in the remaining regions, as the substantial $7 trillion impact on GDP in China comes with 181m jobs being impacted by AI, leaving the GDP per job impacted at around $38,735.

Figure 7.17 – Comparison of total jobs impacted by AI in 2030 across regions

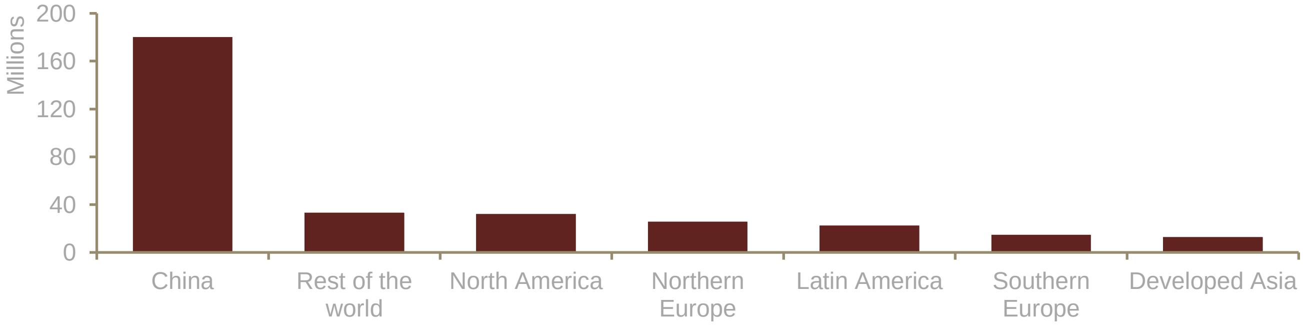

Source: PwC Analysis

Table 7.18 – Comparison of GDP impact in 2030 per job impacted by AI ($)

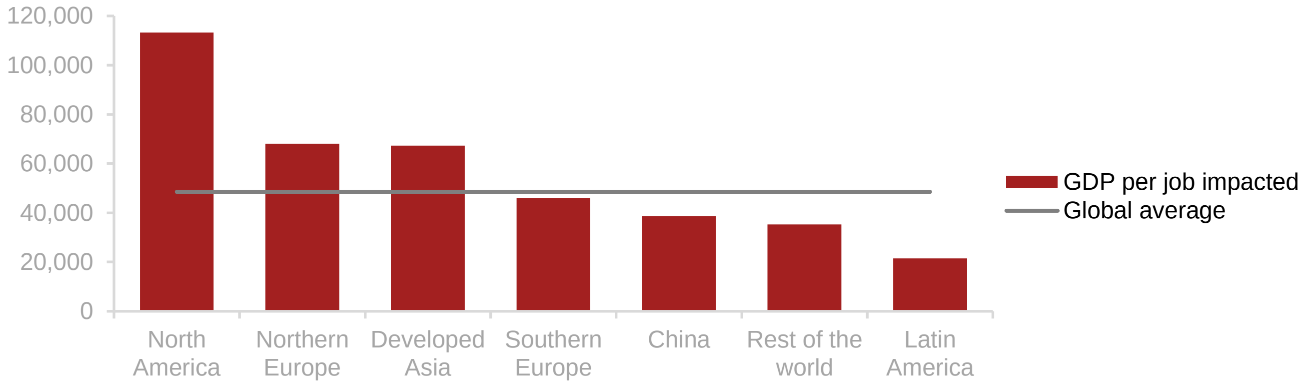

Source: PwC Analysis

Comparing the impact of AI on skilled and unskilled labour

The impact of GDP varies not only by geographic region, but also by type of labour. A key result of our model is that unskilled labour benefits more as a whole than skilled labour, as 67% of the jobs impacted by AI are unskilled, amounting to 187m jobs. This is to be expected due to the larger base of unskilled jobs, as in 2016 we estimate there to be 2.1bn unskilled jobs compared to just under 1bn of skilled jobs globally. So as a proportion of existing jobs, skilled jobs still benefit relatively more from AI. The model result, that all labour types will benefit overall, is a positive outlook given that some studies (as discussed in section 5.6) have suggested that automation would impact unskilled labour more and potentially perpetuate a polarisation of the labour force as seen in previous years.

However, as mentioned previously, it should be noted that our result is dependent on an average labour productivity shock we have fed into the S-CGE model. Despite our shock being engineered to favour capital efficiency over labour efficiency (discussed below), our average labour productivity shock favours unskilled labour, who in practice would likely see smaller productivity gains than skilled labour. Whilst our aggregate results are robust to this nuance, the AI-job impacts would be more balanced towards skilled labour than unskilled had we been able to quantify relative labour productivity impacts on differently skilled labour. Therefore, our results are best interpreted as heavily supporting skills-biased technological change, since we see evidence of this effect here regardless of applying a uniform labour productivity shock over the skill distribution.

Our result is also due to the benefits from technological advancement in the economy, which we project as increasing the efficiency of capital by a significant enough amount that more labour is used in the economy to facilitate the higher production levels. That is, the capitalisation effect of companies entering highly productive industries will dominate the destruction effect of machines making some jobs obsolete. Though this will inevitably cause a re-allocation of labour, if labour is sufficiently flexible then the overall impact on jobs could be positive.

Figure 7.19 – Impact on skilled and unskilled labour by geographic region

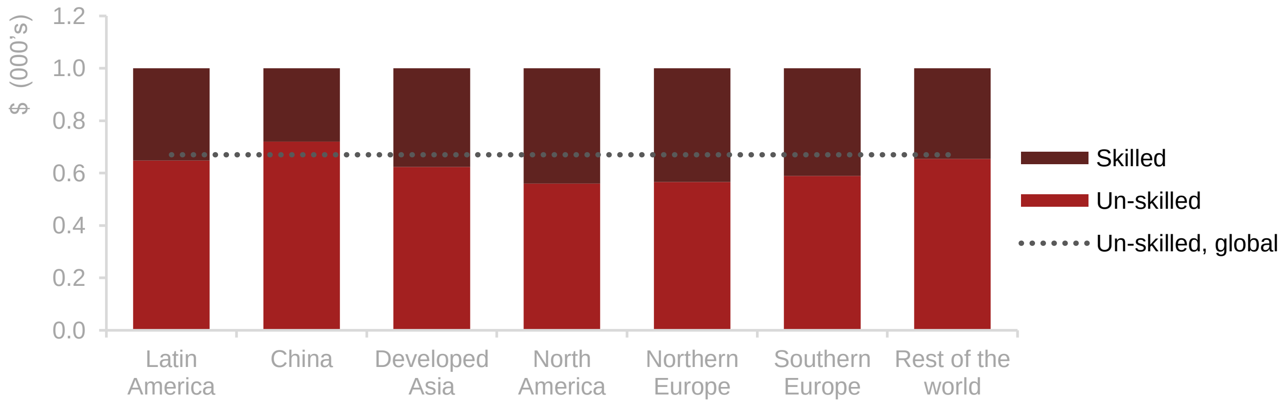

Source: PwC Analysis

8. Sensitivity analysis

8.1. Our approach to sensitivity analysis

In addition to the main economic impact results using all the analysis and inputs discussed in Sections 3 to 7 of this report, we have conducted a number of sensitivity tests on our modelling. The purpose of these sensitivity analyses is to understand how responsive our economic impact results are with respect to some of the assumptions that have been made in the process of the analysis and therefore how much the results would change if the assumptions were to alter slightly. To some extent, these analyses can also be used to determine confidence and reliability of the model outcomes.

In selecting the elements to focus on with our sensitivity tests, we sought to identify assumptions that were made in the process of estimating the economic impact. We chose to focus on these assumptions rather than the results of the econometric analysis, job automation study and the AI Impact Index. This is because the econometric analysis is based upon a well-regarded dataset of productivity, capital and labour data and uses a methodology developed in line with respected academic literature in the field, which we have further tested for robustness using a number of specification tests. Equally, the AI Impact Index collated insights from a number of experts in the field of artificial intelligence (AI) and uses a detailed bottom-up approach that yields a robust and defensible set of results, whilst the job automation study is also based off a robust machine learning approach.

Instead, we focus on the phase of the approach which turns the results from the AI Impact Index, the econometric analysis and the job automation study into S-CGE model inputs to capture the primary impacts of AI over time, since this process was arguably less robust than the creation of our model inputs themselves. We also vary our assumptions over the uptake time-scale and profile to reflect potential barriers to AI-uptake and a different event sequence over the next 10 years. In sum, we identify three key assumptions that were made in that process that we have flexed in our sensitivity analysis. The following outlines each of the areas, provides the rationale for assessing the sensitivity of the assumption and describes how we have approached assessing the sensitivity of that assumption.

Pace of adoption

Our model results focus on the economic impact of AI over the period from 2017 to 2030. The scale of AI uptake, though not featuring as a specific direct input, is captured through a number of elements of our analysis (as discussed in Section 7.3). As a reminder of these elements: the estimated productivity impact assumes that jobs with a 70% chance or higher of being automated by 2030 will be in reality (in line with the conclusion of the job automation study) and that the uptake of workforce augmenting technologies will increase in line with long run trends in AI uptake orthogonal to labour force choices. Moreover, the AI Impact Index accounts for all of the AI technologies that (a) have been adopted already, (b) are in the process of development for future adoption or (c) have been conceived of and are likely to be adopted to some extent before 2030. Resultantly, as discussed we consider the total AI impacts used in our main scenario to be the 'perfect' benchmark for every sector and region. To account for prohibitive factors that would prevent some fraction of the adoption cycle from being completed by 2030, we used the Global Innovation Index (GII 2016), to generate more realistic estimates of the scale of AI adoption that is expected to be completed in each geographic region by 2030.

As this demonstrates, we have thought carefully and made a deliberate and justifiable assumption about the pace and scale of AI adoption and its distribution globally. However, our assumption was made in light of the current global economic climate, current global sentiment towards technologies and the general consumer appetite. In practice, we recognise that the pace of AI adoption could be faster or slower, perhaps related to significant future events that affect any of the aforementioned factors.

Therefore we calculated an alternative scenario for the pace of AI adoption to assess the sensitivity of our model results with respect to this. In particular, we created a rule-of-thumb 'slow-uptake' scenario which supposed that the level of AI adoption originally expected by 2030 in each geographical region was actually to only be reached by 2040. We then recalculated the productivity and product enhancement impacts that would occur by 2030 under this scenario, before re-estimating the full economic impact again using the S-CGE model.

We conceive of this scenario as representing a 'slow-uptake' world scenario, where inertia-driven and prohibitive factors are envisaged to slow down rate of AI-uptake. In practice, key prohibitive factors are likely to be unionised resistance against replacement automation, public rejection of AI-enhanced products, and slow adaptation of the workforce to the uptake augmenting AI-technologies. Inertia-based factors are likely to relate to firms' willingness to uptake AI in practice due to the large-scale transformation nature of the changes involved.

Shape of the S-curve of AI adoption

We assume an 'S-shaped' adoption curve of AI technologies, in line with much of the research on the life cycle or technology adoption and innovation diffusion. In practice this means that adoption is slower in the near-term, before picking up more rapidly when the major obstacles have been overcome before slowing down again as the limit approaches. This shape was further justified by the insights provided in the time to maturity element of the AI Impact Index.

Again, whilst we have thought carefully and made a deliberate and justifiable assumption about the shape of AI adoption, we recognise that the nature of AI adoption could be heavily dependent on the success of leading companies and the timing or speed of significant technological breakthroughs. In light of this, we have developed an alternative scenario for AI adoption during the period to 2030 which yields the same amount of AI adoption in aggregate, but alters the profile of how this technology is introduced over time. Specifically, the profile in this scenario remains an 'S-shape' but with a later and steeper upturn in the middle of the period. This reflects a scenario whereby less AI technology is adopted in the short-term but in the medium-term many significant advances happen very quickly and adoption happens across many companies rapidly. This 'delayed revolution' scenario would be expected in the event that businesses and the regulatory infrastructure for AI globally is less AI-ready than anticipated, but that once sufficient resources have been invested in using and understanding AI and the regulatory environment is capable enough, there is a 'race to the top' as AI is rolled out globally.

Quantification of our product quality model input

As discussed in detail in Section 7.3 of this report, we used insights from the academic literature, economic theory and the insights of subject matter experts from our economics practice to quantify the outputs of the AI Impact Index into the consumption-side model inputs for the S-CGE analysis. Translating the results with respect to product quality proved the most challenging since less literature existed on the topic and the index results lent themselves less to a smooth transition into percentage impacts on one of our model parameters (by comparison to personalisation and time saved). As mentioned, we therefore used expert judgement to create a conversion system between AI-impact index scores and percentage impacts on marginal utility corroborated by several economists in our team.

Recognising the risks in this approach, we sought to take a conservative estimate for marginal utility impacts and are therefore confident that we have not overestimated the economic impact of AI from this perspective. However, since this is the least rigorous element of our quantification, we sought to understand the impact of this judgement on the final economic impact figures. We therefore supposed that the impact of on utility was 35% smaller than in our main scenario and ran the model based on that scenario. We have not considered the impact of a 35% larger shock, since we expect the impact to be symmetric. In other words, the (supposed) multiplier between the aggregate product quality shock and global GDP impact should be constant.

8.2. Results of our sensitivity analysis

Here we present the results of the economic impact analysis under each of the three alternative scenarios discussed above. We first look at the responsiveness of the aggregate results related to economic impact of AI including by channel of impact. We then turn to evaluate how the distribution of the impacts by geographic region and by sector could be different under the alternative scenarios.

8.3. Assessing the sensitivity of the pace of AI adoption

$11.2 trillion – 9.8% of GDP

In a world where AI adoption is slower than current projections and expectations suggest, global GDP could be up to 10% higher in 2030 as a result of AI, the equivalent of up to $11.2 trillion. This is in comparison to a global GDP increase of up to 13.8% in our main scenario.

Under this scenario, we expect that in 2030 GDP could increase by up to 4.4% as a result of productivity gains and up to 5.4% as a result of product innovations and product enhancements. The composition of GDP impact with respect to productivity and product enhancements is therefore almost unchanged relative to the main scenario. However, within the total impact as a result of product enhancements, the distribution of the impacts shifts somewhat. Almost 65% of those impacts come through from AI facilitating and enhancing product personalisation whereas the contribution is close to 50% in our main scenario.

Looking at the profile of the GDP impacts across the full period from 2017-2030, the profile of the impacts is significantly different in a world where the adoption of AI is slower than expected. In our main scenario, the impacts resulting from productivity gains plateau after approximately 2025 whereas in the slower scenario the productivity impact continues to increase throughout the period, as seen in Figure 8.1. This is likely because if AI were to be adopted more slowly and the adoption S-curve were to end instead around 2040 (rather than 2030), then the period 2017-2030 would not include the phase of AI adoption beginning to reach its limit and slow.

Figure 8.1 – Global economic impact of AI over time (by channel of impact), alternative pace of AI adoption scenario ($tn of GDP)

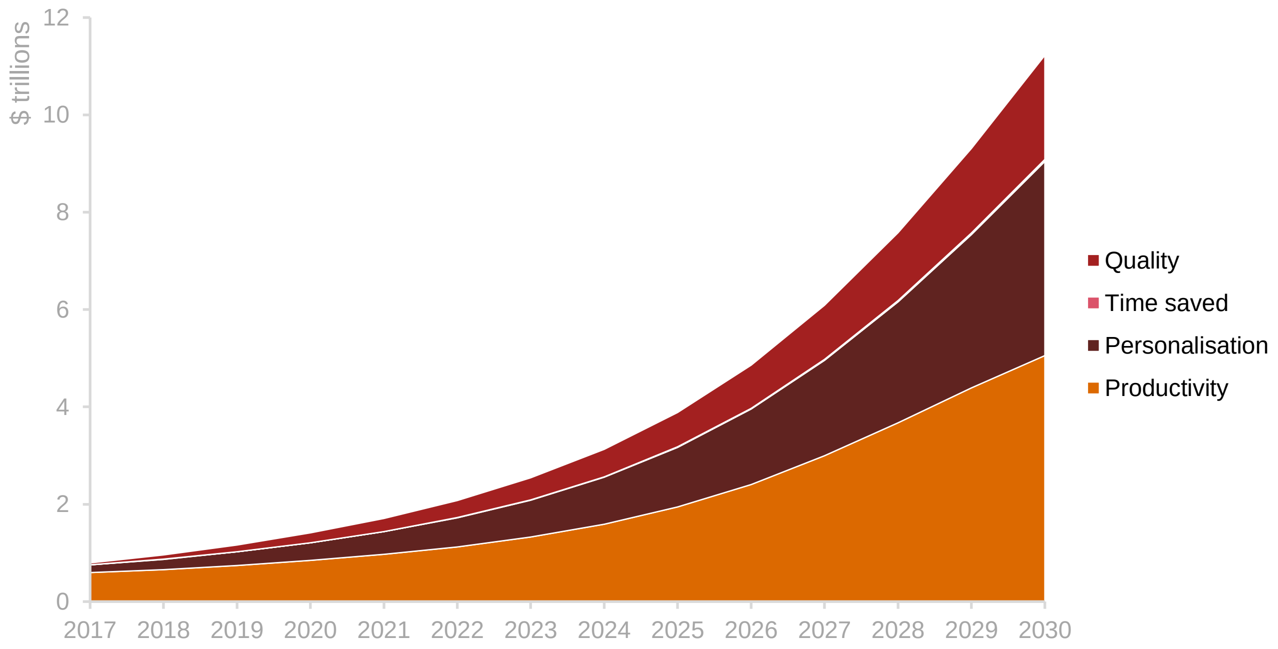

Source: PwC Analysis

8.4. Assessing the sensitivity of the shape of the AI adoption curve

$14.2 trillion – 12.5% of GDP

In a world where somewhat less AI technology is adopted in the short-term than currently expected but in the medium-term many significant advances happen very quickly and adoption happens across many companies rapidly, global GDP could be up to 12.5% higher in 2030 as a result of AI, the equivalent of up to $14.2 trillion. This is in similar when compared to the global GDP increase of up to 13.8% in our main scenario.

Under this scenario, we expect that in 2030 GDP could increase by up to 5.5% as a result of productivity gains and up to 6.9% as a result of product innovations and product enhancements. If the adoption of AI were to follow this alternative pattern, our analysis indicates that by 2030, impacts resulting from productivity would be similar to our main scenrio, whereas the impacts on the consumption-side could be 1.0% of GDP lower, and therefore the composition of impacts would be slightly different. The impact resulting from productivity gain is likely to remain similar because in our main scenario, the impacts resulting from productivity gains plateau after approximately 2025. These results suggest that despite the profile of the AI adoption, once a certain level of uptake is reached, the productivity impacts will also reach a similar 'limit.'

Looking at the profile of the GDP impacts across the full period from 2017-2030, the profile of the impacts reflects the steeper curve of AI adoption assumed under this scenario when compared to our main scenario.

Figure 8.2 – Global economic impact of AI over time (by channel of impact), alternative profile of AI adoption scenario ($tn of GDP)

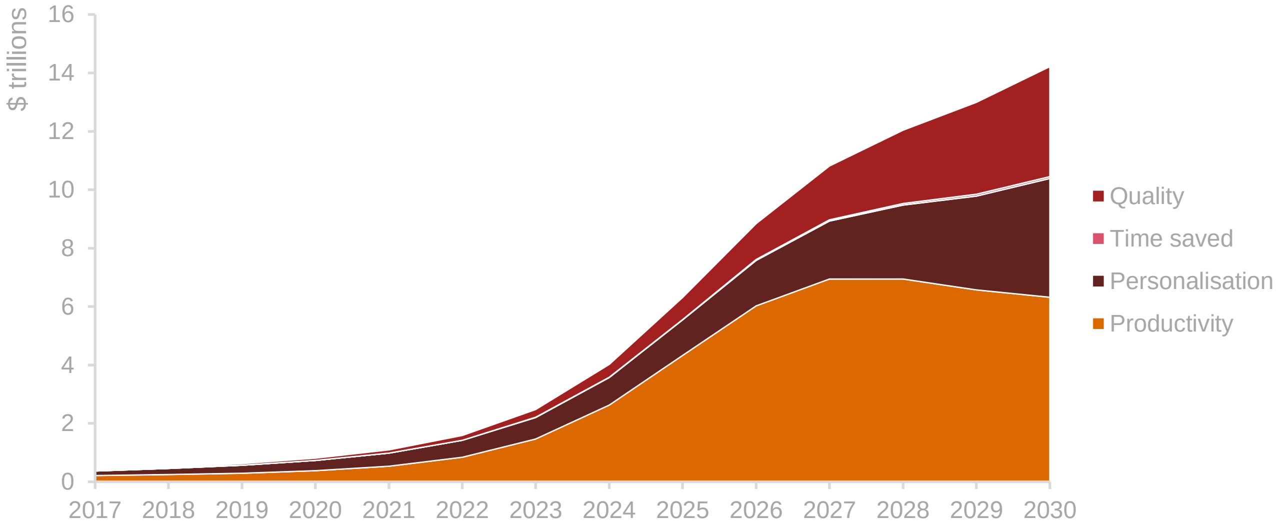

Source: PwC Analysis

8.5. Assessing the sensitivity of the quantification of our product quality model input

$15.2 trillion – 13.3% of GDP

If the impact of AI on marginal utility were to be 35% lower than in our main scenario, global GDP could be up to 13.8% higher in 2030 as a result of AI (against the baseline, not against the main scenario), the equivalent of up to $15.2 trillion, with no material change from our main scenario.

Under this scenario, we expect that in 2030 GDP could increase by up to 5.8% as a result of productivity gains and up to 7.5% as a result of product innovations and product enhancements. As expected, given the nature of the scenario altering the input quantification only on the product enhancements inputs, the productivity impacts are unchanged in comparison to our main scenario. However, the effect of the sensitivity test on the product enhancement-driven global GDP impact also looks surprisingly small. The reasons for this are two-fold. First, the multiplier between product enhancements and their GDP impact is actually of reasonable size, but this is disguised by the fact that product utility enhancements only make up around half of the 7.9% GDP impact from product enhancements in total. Therefore, the multiplier (between the cumulative product quality impact and 2030 GDP impact) is roughly 0.30, which, whilst small, is credible. Second, a smaller shock to product utility results in less product heterogeneity. This facilitates the trade of goods from regions less impacted by product enhancements and overall acts as a counterbalance to the lost GDP as a result of the smaller shock.

Figure 8.3 – Global economic impact of AI over time (by channel of impact), lower marginal utility of product quality scenario ($tn of GDP)

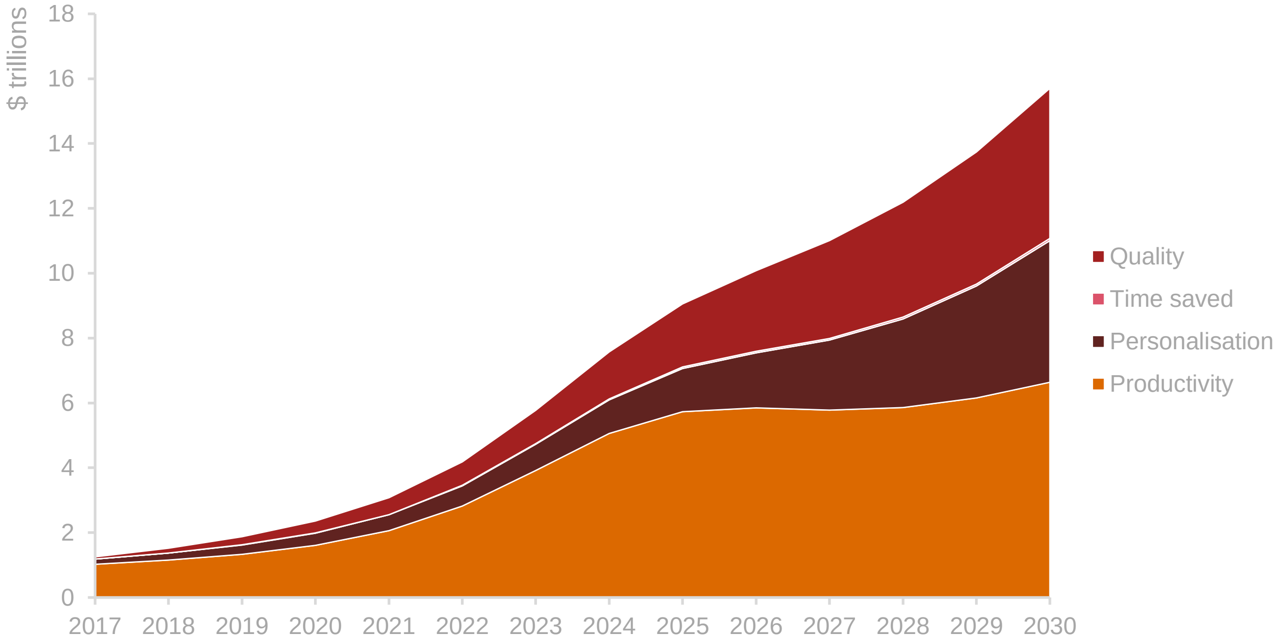

Source: PwC Analysis

8.6. Impact of sensitivities on sectoral and geographic distribution of results

Sections 8.3 to 8.5 indicate the range of estimates of total global GDP impact that could result from our alternative scenarios of AI uptake and impacts. As well as assessing the responsiveness of the aggregate level results, it is important to understand the sensitivity of the distribution of the economic impact of AI across various geographic regions and sectors and whether this could be different under these alternative scenarios.

Geographic distribution

First we turn to the distribution of the results across geographic regions, illustrated in Figure 8.4 below. Specifically, Figure 8.4 indicates the percentage of the total impact under each of our alternative scenario, as well as the main scenario, attributable to each geographic region, i.e. the percentages across geographic regions for each scenario add up to 100%.

Figure 8.4 – Percentage of total impact in 2030 attributed to each geographic region in main scenario and sensitivity scenarios

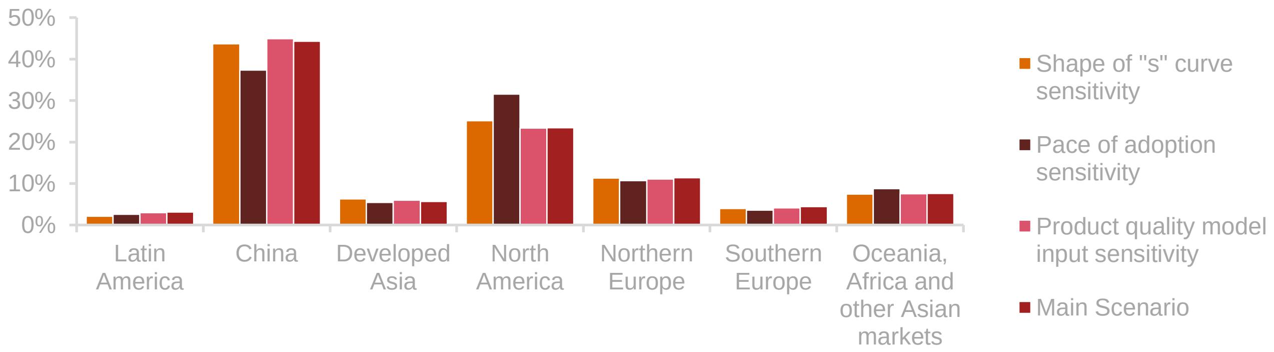

Source: PwC Analysis

Pace of AI adoption: In a world where AI adoption is slower than currently expected, the distribution of GDP impacts across most of the geographic regions remains the same, however, the share of impacts between China and North America could differ. Together, we would expect the two countries to still account for almost 70% of the global GDP impact, and for China to still gain more than North America, but to a lesser extent. In our main scenario, we think that China could account for 44% of the global GDP impact of AI whereas this could reduce to 37% in a slower AI adoption scenario. This relates to the profile of impacts over time as discussed in Section 7 of this report. In our main scenario, the GDP impacts in North America lessen as China begins to catch up and this stimulates exports of AI-enabled products from China to North America. In the slower scenario China 'catches up' to a lesser extent and therefore the North America impacts remain relatively high.

Profile of AI adoption: There is no material impact in the distribution of GDP impact across geographic regions under the scenario with a steeper S-curve of AI adoption.

Product quality model input: There is no material impact in the distribution of GDP impact across geographic regions under the scenario with a 35% reduction in the impact of AI on marginal utility.

Sectoral distribution

Next, we turn to the distribution of the results across industry sectors, illustrated in Figure 8.5 below. Here again we present the percentage of the total impact under each of our alternative scenario, as well as the main scenario, attributable to each geographic region i.e. the percentages across geographic regions for each scenario add up to 100%.

As the graph shows, the distribution of the GDP impacts across industry sectors changes only marginally under the different scenarios. The only difference of note is that in a world where the pace of adoption of AI is slower, the impact on the health and other services sector would account for slightly less of the GDP impact while the manufacturing and construction sector would account for slightly more, though this does not alter the distribution in a material way. This result follows from the time profiles North America and China productivity impacts shown in Figure 7.12. Put simply, this is because the 2040 'slow-uptake' sensitivity takes its AI GDP impact snapshot (in 2030) at a time when North America sees maximal benefits of AI, before China's catching up can erode North America's gain through exports to the North America. Since the m

Figure 8.5 – Percentage of total impact in 2030 attributed to each industry sector in main scenario and sensitivity scenarios

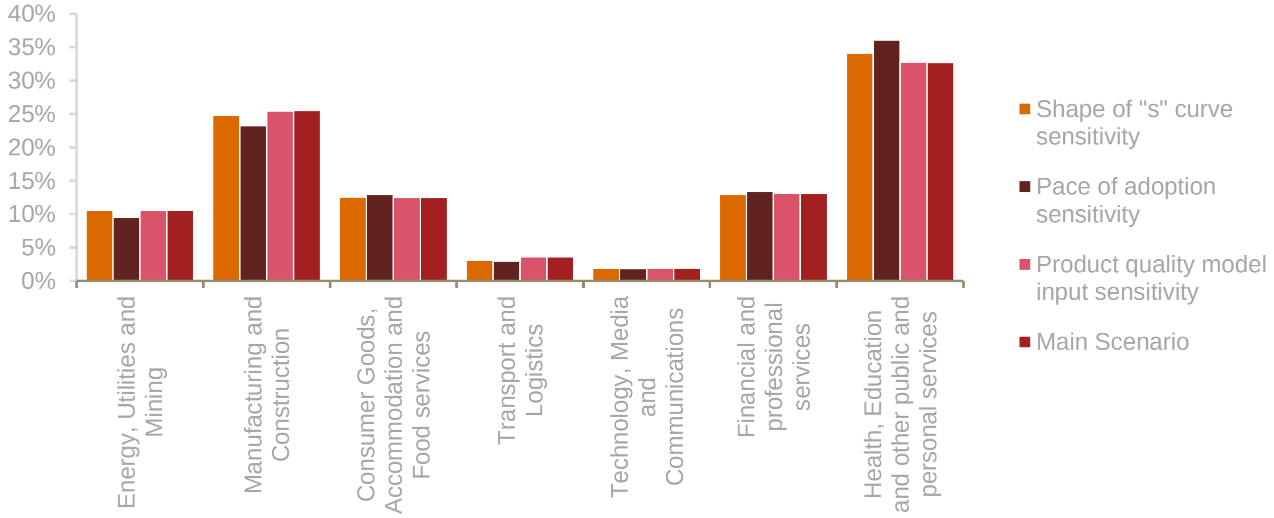

Source: PwC analysis

9. Conclusion

Summarising our results

Our study has demonstrated that the global macroeconomic impact of AI between now and 2030 is likely to be substantial. The results of our Spatial Computable General Equilibrium (S-CGE) model analysis indicate that global GDP could be up to 13.8% higher in 2030 as a result of AI's impact. Our study is the first, to our knowledge, that captures both the consumption-side and production-side effects of AI, therefore presenting a more holistic picture of the opportunities, threats, and potential economic value at stake. The importance of accounting for consumption-side product enhancements is reflected by our results, which indicate that 2030 productivity enhancements from AI account for less than half of this impact (40%), and only account for 55% of the total cumulative impact over the entire period (2017-2030).

Labour market results from our analysis indicate that although initial jobs displacement will occur through replacement automation, the creation of new roles in the economy could potentially offset the jobs lost through automation by 2030. Overall, we estimate that 326m jobs will be impacted by AI in 2030, however these should be interpreted as jobs both created and impacted by AI (i.e. jobs which have either been created by AI, are 'AI-dependent' or are heavily impacted by AI), rather than net jobs created by AI.

A geographic lens shows that the economic impact of AI is likely to be especially large in North America (14% of GDP in 2030) and China (26% of GDP in 2030). AI's impact in North America could grow relatively quickly in the coming years while, towards the end of the period, the percentage impact of AI on China's GDP could grow more quickly with Chinese productivity stimulating exports of AI-enabled products from China to North America.

In the sector-side of the story, the health, education, public services and recreation sector is set to see the largest GDP gains from AI (21%), which is particularly driven by the substantial AI-driven product enhancements expected in this sector. Encouragingly, retail and wholesale trade, accommodation and food services, as well as other labour intensive sectors are also expected to see a large boost (15%), while transport and logistics, and financial and professional services will also benefit a substantial 10%.

All of these results estimate the upwards pressure on GDP as a result of AI only, under the *ceteris paribus* assumption. Although our results should not be interpreted as absolute outcomes due to unforeseen shocks that may pull GDP growth away from our predicted levels, we do place strong confidence in our estimated size of AI's impact on the global economy, and this is supported by our sensitivity analysis in Section 8. In particular, we estimated three alternative scenarios for pace, profile and marginal utility impact of AI. We found that only the pace of AI adoption affected the GDP impact of AI materially, suggesting that the limit to AI's impact will be any lack of willingness to invest or other barriers to adoption.

The results of our analysis ultimately illustrate that AI should be interpreted as an exciting source of wealth and job creation for the future, provided businesses and workers invest to adopt and adapt in order to meet the challenges AI will present, and take advantage of the opportunity.

What this means for businesses, employees and governments

Our research underscores what many business executives, technology companies, and consumers intuit: that AI's potential is vast—and multidimensional. Yet our work is just the starting point for the real work to come: organisations taking the deliberate and strategic steps necessary to capitalise on AI's two-fold benefits: 1) productivity gains through automation and 2) greater consumer value through enhanced personalisation, time-savings, and quality.

We see four areas that require the focus of business leaders, the shepherds of AI in organisations:

1. Determine what AI means for your business

On the automation side, you'll want to consider what areas are ripe for improvement– pain points like rote data entry, time-consuming contract review, or error-prone processes like reconciliations across multiple systems. Identifying projects that yield short-term returns, such as automating identity verification via RPA, is a good way to demonstrate viability and get the organisation comfortable with automation technologies. In some cases, cost-savings from these quicker winners can be reinvested into longer-range AI efforts.

At the same time, you need to look outside the business and evaluate the technological developments and competitive pressures coming up within your sector, how quickly they will arrive, and how you might respond. For example, insurance companies who have used drones effectively to assess roof damage after storms will also be focusing on automation in other areas besides claims adjustment, such as claims processing or payment. Using AI and other technologies together can open up broad possibilities for improvement, but also raise the possibility of new implementation risks. A strategic plan that scans the overall technological and competitive landscape, considers scenarios and presents a range of options will be essential.

2. Prioritise your response — and your data

As you develop your AI strategy, it is critical to view and communicate initiatives in terms of business value. It also means demonstrating leadership commitment so that AI is not viewed solely as an array of disconnected technology projects. Every executive in the organisation, as well as board members, should be able to articulate what the organisation is setting out to accomplish with AI, how it connects to business strategy, and how its value will be measured. The goal is to develop a coherent response that brings to bear all of the organisation's resources.

As you set priorities, you'll need to focus on your organisation's data maturity. Consider what investments and changes would enable you to capture more data and use it more productively For example, are your activities capturing the right data to begin with? Do your data scientists have access to the data they'll need in a data lake, or is the data being held captive in data cartels? Is the data well described, provenanced and in a form that's appropriate for reuse? Is the appropriate level of data governance applied in a consistent and comprehensive manner so that personally identifiable information is adequately protected? What steps does your organisation need to take to comply with General Data Protection Regulation requirements by May 2018? Is the quality of the data reliable so that data scientists can be confident of the results?

3. Focus on people and culture

As data become even more central to the business, it's essential to instil a data-driven culture that blends intuition and analytical insights with a focus on practical and actionable decisions across all levels. Demand for employees with data science and analytic skills are in high demand, so many employers are looking to cultivate existing talent through long-term training and development. At the same time, as adoption of AI accelerates, skills like creativity, leadership, and emotional intelligence will continue to be at a premium. It's important to prepare for a hybrid workforce in which AI and human beings work side-by-side. The challenge isn't just ensuring you have the right systems in place, but judging what role your people will play in this new model. People will need to be responsible for determining the strategic application of AI and providing challenge and oversight to decisions.

4. Emphasise trust and transparency

Paramount to success is an unwavering focus on what we call responsible AI— that is, AI in which there is assurance and control. Some considerations here: Have you considered the societal and ethical implications of your AI initiatives? How can you build stakeholder trust in your solutions? How can you build AI that can explain its logic so that a layperson can understand? How can you build AI that is unbiased and transparent? It's important to put in place mechanisms to source, cleanse and control key data inputs and ensure data and AI management are integrated. Transparency is not only important in guarding against biases within the AI, but also helping to increase human understanding of what the AI can do and how to use it most effectively.

As you anticipate the tremendous changes that AI will bring, it's important to remember that we are only at the beginning of the AI revolution. We expect further disruption and transformation in all economies, geographies, and sectors, as described in this report. And while AI's economic impacts might not be fully realised for a decade or more, businesses must begin making the right strategic moves today. PwC is working with a wide range of businesses to address the immense opportunities and emerging challenges that AI will bring. We're helping our clients to reimagine what is possible, and ensure they exploit AI's full potential.

10. Existing and forthcoming insights on AI from PwC

In addition to this Macroeconomic Impact of AI study, PwC has released and will continue to release a series of reports based on the output of the Global AI Study and perspectives of our AI leaders.

Market view content

Sizing the Prize summarises the macroeconomic potential of AI across key sectors and geographies. In addition to discussing how productivity will be enhanced by AI, the report highlights for the first time an assessment of how consumption will be impacted by enabling consumers more personalised, useful, and time saving products and services. This market view makes clear that AI is the greatest commercial opportunity in our economy, and one that executives cannot ignore going forward.

Enterprise view content

After calling attention to AI as a key disruptor across all sectors, we offer executives in the driver's seat some navigational assistance on how to exploit the opportunity. *The Strategist's Guide to AI*, our first report in this series, clarifies our definition of AI and introduces our point of view on how executives should think about using AI based on business need. In *Responsible AI: how to build trust and confidence*, we've also shared a framework on how to think about trust when building deploying AI. As one of the most trusted advisory brands in the marketplace, PwC takes pride in our ongoing efforts to work with clients on making the black box of AI more transparent to all stakeholders.

In addition to these holistic examinations about the types of challenges and opportunities that company leaders face in adopting AI, we also offer some tactical content to inspire in two separate series. Through our collaboration with Forbes Insights, we've commissioned a collection of articles highlighting how innovators are using AI in their organisations. This 12-part series covers case studies of companies across sectors and the globe, and features both large MNCs and small start-ups. Separately, our global network of industry subject matter experts has met to produce a series of sector and functional white papers. These reports will build on the previous work done by our firm, and will serve as AI opportunity summaries emphasising the immediate, mid-term, and long-term use cases of AI. These practical AI reports will not only demonstrate how AI can be applied to your specific sector or function, but also emphasise the nuanced capabilities that are required to succeed.

Consumer view content

PwC released Bot.Me, a report that was constructed from insights gathered at a 2017 AI Expert Salon, and a survey of 2,500 U.S. consumers and business decision makers on attitudes towards AI and its current and future implications on society. The results reveal an enthusiasm about AI from both cohorts; 63% of consumers agree AI will help solve complex problems that plague modern societies, 72% of business leaders believe AI will be the business advantage of the future.

PwC's broad set of capabilities allows us to provide a diverse set of solutions for leaders looking to proactively exploit the promises of AI. Our global firm offers new-to-market AI strategy services for those seeking to identify the most attractive and feasible applications of AI, as well as myriad solutions to known challenges leveraging AI techniques. Our teams have experience working with drones, image, audio, and emotion recognition, and machine and deep learning across functions and sectors. In addition to helping companies harness the power of AI, PwC also stands alongside them in ensuring these initiatives keep trust and transparency top of mind. We look forward to conversations with current and prospective clients to discuss the emerging opportunities and threats for their businesses, and the chance to illustrate the many ways PwC is uniquely equipped to help.

Appendix A

Table A.1 Mapping of countries to geographic regions in our modelling

Region	Region Code	Countries
Latin America	AM	Argentina, Bolivia, Brazil, Chile, Colombia, Ecuador, Paraguay, Peru, Uruguay, Venezuela, Costa Rica, Guatemala, Honduras, Nicaragua, Panama, El Salvador, Mexico, Rest of South America, Rest of Central America
China	CHN	China, Hong Kong
Developed Asia	EAS	Mongolia, Japan, South Korea Taiwan, Singapore
North America	NA	Canada, United States of America
Northern Europe	NE	Austria, Belgium, Czech Republic, Denmark, Estonia, Finland, France, Germany, Ireland, Latvia, Lithuania, Luxembourg, Netherlands, Poland, Sweden, United Kingdom, Switzerland, Norway
Southern Europe	SE	Cyprus, Greece, Hungary, Italy, Malta, Portugal, Slovakia, Slovenia, Spain, Bulgaria, Croatia, Romania, Albania, Belarus, Ukraine, Rest of EFTA, Rest of Eastern Europe
Africa, Oceania and other Asian markets	ROW	Cambodia, Indonesia, Laos, Malaysia, Philippines, Thailand, Vietnam, Bangladesh, India, Nepal, Pakistan, Sri Lanka, Australia, New Zealand. Rest of Oceania, Rest of East Asia, Rest of Southeast Asia, Rest of South Asia, Russian Federation, Kazakhstan, Kyrgyzstan, Rest of Former Soviet Union, Armenia, Azerbaijan, Georgia, Bahrain, Iran, Israel, Kuwait, Oman, Qatar, Saudi Arabia, Turkey, United Arab Emirates, Rest of Western Asia, Egypt, Morocco, Tunisia, Rest of North Africa, Cameroon, Cote d'Ivoire, Ghana, Nigeria, Senegal, Rest of Western Africa, Central Africa, South Central Africa, Ethiopia, Kenya, Madagascar, Malawi, Mauritius, Mozambique, Tanzania, Uganda, Zambia, Zimbabwe, Rest of Eastern Africa, Botswana, Namibia, South Africa, Rest of South African Customs Union

Appendix B

Table B.2 Full econometric results, industry-level estimates of relationship between artificial intelligence and productivity

Variable	*UK*	*China*	*South Korea*	*Spain*	*United States*
Dependent Variable: Log Labour Productivity					
Year	0.00128***	0	0	-0.00043	0
	(0.001)	(.)	(.)	(0.695)	(.)
Fraction of Workforce with Tertiary Education (Log First Difference)[1]	0.579**	-0.00296	0.977*	0.284	0.698**
	(0.009)	(0.972)	(0.095)	(0.31)	(0.003)
Current Residual Stock (Log First Difference)	0.00289	-0.628	-0.0107*	-0.0261	1.328
	(0.837	(0.161)	(0.049)	(0.799)	(0.134)
R&D Expenditure (Log First Difference)	0.0126	-0.0623	-0.0151	-0.00327	
	(0.385)	(0.267)	(0.6)	(0.594)	
Lagged R&D Expenditure (Log First Difference)[4]	0.00928	-0.0479	0.011	-0.013	
	(0.414)	(0.154)	(0.55)	(0.274)	
Sector-specific estimates					
Technology, Media, and Communications	0.422***	0.535	0.734***	0.505***	0.645*
	0	(0.185)	0	0	(0.051)
Energy, Utilities, and Mining	-0.204***	0.305	0.239	-0.561***	-0.318
	0	(0.421)	(0.119)	0	(0.389)
Manufacturing & Construction	-0.0445	-0.169	-0.549***	0.0183	-0.286
	(0.789)	(0.68)	0	(0.907)	(0.237)
Consumer Goods, Accomm. and Food Services	-0.481***	1.130*	0.857***	-0.422***	-0.0963
	0	(0.031)	0	0	(0.574)
Transport and Logistics	0.00364	0.597*	0.488***	-0.280***	-0.899***
	(0.88	(0.061)	0	0	(0.001)
Financial and Professional Services	-0.242**	0.396	-0.201	-0.286***	-0.302
	(0.005)	(0.322)	(0.536)	0	(0.202)
Health, Education, and other Public and Professional Services	-0.339**	0.633	0.121	-0.175*	-0.669**
	(0.002)	(0.233)	(0.191)	(0.016)	(0.002)
Industry specific fixed effects					
Agriculture, Forestry, and Fishing	13.68***	-36.38*	-24.98***	-2.244	
	0	(0.011)	0	(0.237)	
Mining and Quarrying	13.37***	-15.81	-23.47***	-6.582***	-2.995**
	0	(0.174)		0	(0.005)
Manufacturing	5.754*	-37.90**	-24.48***	-4.263*	0.0382
	(0.022)	(0.007)	0	(0.012)	(0.88)
Electricity, Gas, Steam, and Air Conditioning Supply[2]	13.49***				
	0				

Water Supply, Sewerage, Waste Management, and Remediation Activities²	14.94*** 0				
Construction	4.539*** 0	-33.40* (0.046)	-35.29*** 0	-11.03*** 0	
Wholesale and Retail Trade	2.608** (0.001)	-14.94 (0.392)	-21.94*** 0	-2.774** (0.001)	0.383 (0.256)
Transport and Logistics	4.743*** 0	-2.253 (0.888)	-12.21*** 0	-4.187*** 0	1.047* (0.044)
Accommodation and Food Service Activities	9.878*** 0	-17.85 (0.353)	0 (.)	-6.571*** 0	-0.631*** 0
Technology, Media, and Communications	16.22*** 0	17.7 (0.319)	0.0574 (0.967)	-9.778*** 0	0.801 (0.327)
Financial and Insurance Activities	13.59*** 0	-72.27*** (0.001)	-24.52*** 0	12.73*** 0	2.075** (0.002)
Real Estate	0.778 (0.296)	-30.88* (0.047)	-51.15*** 0	-22.71*** 0	-0.702 (0.205)
Professional, Scientific, and Technical Activities³	1.751 (0.083)				
Administrative Activities³	1.816 (0.066)				
Public Administration	2.907*** 0	-13.57* (0.091)	-56.67***	-3.656* (0.028)	
Education	1.312 (0.164)	-30.25 (0.162)	-42.20***	2.643* (0.077)	
Human Health and Social Work	4.250*** 0	-9.292 (0.38)	-65.56***	-1.828 (0.154)	0 (.)
Arts, Entertainment, and Recreation	2.092*** 0	0 (.)	-37.14***		
Energy and Water		0.482 (0.965)	-10 (0.003)	5.540* (0.012)	6.049*** 0
Professional, Scientific, and Technical, Administrative Activities		-4.688 (0.667)	-41.27*** 0	-5.049* (0.042)	4.237 (0.204)
Industry Specific Time Trend					
Agriculture, Forestry, and Fishing	-.00681*** 0	0.0109** (0.002)	-0.00239** (0.001)	0.00113 (0.231)	
Mining and Quarrying	-.00669*** 0	0.000657 (0.93)	-0.00322** (0.009)	0.00328*** 0	0.000473 (0.134)
Manufacturing	-0.00286* (0.022)	0.0117* (0.053)	-.00266*** 0	0.00213* (0.012)	-.00104*** (0.001)
Electricity, Gas, and Steam and Air Conditioning Supply²	-.00673*** 0				
Water Supply, Sewerage, and Waste Manaagement²	-.00745*** 0				
Construction	-.00226*** 0	0.00947*** 0	0.00270** (0.004)	0	
Wholesale and Retail Trade	-0.00129** 0.0001	0.000179 (0.967)	-.00393*** 0	0.00138** (0.001)	-.00121*** 0
Transport and Logistics	-.00236*** 0	-0.00613* (0.04)	-.00874*** 0	0.00208*** 0	-0.00152*** 0
Accommodation and Food Services	-.00492*** 0	0.00161 (0.426)	-0.0148*** 0	0.00326*** 0	-.000710** (0.002)
	-.00806***	-0.0160***	-0.0148***	0.00487***	-0.00142*

Technology, Media, and Communication	0	0	0	0	(0.019)
Financial and Insurance	-.00677*** 0	0.0288*** 0	-0.00266 (0.231)	-.00633*** 0	-.00205*** 0
Real Estate	-0.000402 (0.281)	0.00817*** 0	0.0106*** 0	0.0113*** 0	-0.000672* (0.044)
Professional, Scientific, and Technical[3]	-0.000862 (0.086)				
Administrative and Support[3]	-0.000895 (0.069)				
Public Administrative and Defense	-.00145*** 0	-0.00048 (0.917)	0.0133*** 0	0.00183* (0.028)	
Education	-0.000666 (0.154)	0.00785** (0.006)	0.00612*** 0	-0.00132* (0.076)	
Health and Social Work	-.00211*** 0	-0.00261 (0.381)	0.0178*** 0	0.00091 (0.155)	-.00101*** (0.001)
Arts, Entertainment, and Recreation	-.00105*** 0	-0.00726 (0.37)	0.00362** (0.003)		
Energy and Water		-0.00746 (0.108)	-.00985*** 0	-0.00276* (0.011)	-.00404*** 0
Professional, Scientific, and Technical, Administrative Activities		-0.00486 (0.236)	0.00567** (0.0006)	0.00251* (0.042)	-0.00312* (0.078)
Constant	-2.578*** (0.001)	14.59 (0.366)	29.87*** 0	0.873 (0.694)	2.074*** 0
Number of Observations	304	192	222	272	244
R^2	0.379	0.704	0.684	0.432	0.369

All figures in parentheses are p-values. All variables are first-differences. Significance levels. ***p<0.01, **p<0.05, *p<0.1. Errors are clustered standard errors.

Notes:

1. For China and Argentina, data on fraction of workforce with tertiary education wasn't available from OECD. For China, ratio of graduates to working population from Chinese Statistical Yearbooks were used.
2. Except for the UK, data is available for the electricity, gas and air conditioning supply, and water supply sectors together.
3. Except for the UK, data is available for the professional, scientific and technical and the administrative services sector together.

Table B.2 Full econometric results, national-level estimates of relationship between artificial intelligence and productivity

Variable	UK	China	South Korea	Spain	United States	Argentina
Dependent Variable: Log Labour Productivity						
Artificial Intelligence Proxy (Log Automation/Labour)	0.207***	0.931***	0.832***	0.201**	0.349*	0.0958*
	0	0	0	(0.005)	(0.029)	(0.041)
Year	0.00153***	0	0	-0.00112	-.00161***	0.00617*
	(0.001)	(.)	(.)	(0.695)	(.)	
Fraction of Workforce with Tertiary Education (Log First Difference)[1]	0.574***	-0.0611	1.292*	0.136	0.733**	-0.225
	(0.01)	(0.366)	(0.015)	(0.664)	(0.001)	(0.263)
Current Residual Stock (Log First Difference)	-0.00462	-0.936***	-0.00647	-0.0315	1.388*	-0.169*
	(0.37)	0	(0.103)	(0.739)	(0.198)	(0.067)
R&D Expenditure (Log First Difference)	0.0135	-0.0295	-0.0285	-0.00231	-1.938*	-0.126
	(0.23)	(0.468)	(0.224)	(0.721)	(0.038)	(0.111)
Lagged R&D Expenditure (Log First Difference)[4]	0.00688	-0.0264	0.00218	-0.012	-0.0978	
	(0.32)	(0.001)	(0.881)	(0.305)	(0.07)	
Industry specific fixed effects						
Agriculture, Forestry, and Fishing	14.19***	32.29***	-31.89***	-2.446	-8.571*	
	0	0	0	(0.23)	(0.029)	
Mining and Quarrying	13.65***	43.34**	-31.74***	-2.276		
	0	(0.001)	0	(0.306)		
Manufacturing	8.455***	37.61***	-29.70***	-4.866**	-5.376*	
	0	0	0	(0.003)	(0.085)	
Electricity, Gas, Steam, and Air Conditioning Supply[2]	14.19***					
	0					
Water Supply, Sewerage, Waste Management, and Remediation Activities[2]	15.58***					
	0					
Construction	4.945***	29.33***	-39.08***	-16.01***		
	0	0	0	0		
Wholesale and Retail Trade	1.425***	32.45***	-30.68***	-3.846***	-3.12	
	0	0	0	0	(0.184)	
Transport and Logistics	5.865***	51.08***	-22.60***	-5.394***	1.35	
	0	0	0	0	(0.723)	
	8.817***	59.70***	-20.09***	-9.705***	-4.27	

Accommodation and Food Service Activities	0	0	0	0	(0.198)
Technology, Media, and Communications	18.72*** 0	87.72*** 0	0 (.)	-8.927*** 0	-10.40** (0.004)
Financial and Insurance Activities	14.23*** 0	0 (.)	-34.34*** 0	11.70*** 0	-1.133 (0.653)
Real Estate	1.546* (0.01)	37.73*** 0	-61.25*** 0	-24.27*** 0	-3.09 (0.276)
Professional, Scientific, and Technical Activities[3]	2.173*** 0				
Administrative Activities[3]	2.248*** 0				
Public Administration	3.636*** 0	47.89*** 0	-62.53*** 0	-3.905* (0.02)	
Education	-0.0352 (0.40)	44.86*** 0	-48.29*** 0	1.496 (0.294)	
Human Health and Social Work	4.482*** 0	55.36*** 0	-71.61*** 0	-2.883* (0.021)	-6.304* (0.051)
Arts, Entertainment, and Recreation	2.209*** 0	53.32*** 0	-43.18*** 0		
Energy and Water		63.88*** 0	-19.00*** 0	9.520** (0.003)	-6.112* (0.033)
Professional, Scientific, and Technical, Administrative Activities		60.29*** 0	-50.82*** 0	-6.384* (0.021)	
Industry Specific Time Trend					
Agriculture, Forestry, and Fishing	-0.00706*** 0	0.0116*** 0	-.00189*** 0	0.00123 (0.226)	
Mining and Quarrying	-0.00682*** 0	0.0061 (0.166)	-0.00204** (0.003)	0.00113 (0.307)	0.00113* (0.071)
Manufacturing	-0.00420*** 0	0.00898** (0.003)	-.00301*** 0	0.00243** (0.003)	-0.000466 (0.305)
Electricity, Gas, and Steam and Air Conditioning Supply[2]	-0.00708*** 0				
Water Supply, Sewerage, and Waste Manaagement[2]	-0.00777*** 0				
Construction	-0.00246*** 0	0.0131*** 0	0.00165** (0.007)	0	
Wholesale and Retail Trade	-.000707*** 0	0.0115*** 0	-.00251*** 0	.00192*** 0	-0.00158** (0.008)
Transport and Logistics	-0.00292*** 0	0.00226 (0.379)	-.00652*** 0	.00268*** 0	-.00379*** 0
	-0.00439***	-0.00204*	-.00778***	.00483***	-.00101***

Accommodation and Food Services	0	(0.018)	0	0	0	
Technology, Media, and Communication	-0.00931*** 0	-0.0160*** 0	-0.0178*** 0	.00445*** 0	0.00205*** 0	
Financial and Insurance	-0.00709*** 0	0.0277*** 0	-0.00071 (0.423)	-.00582*** 0	-.00257*** 0	
Real Estate	-.000785*** (0.01)	0.00892*** 0	0.0127*** 0	0.0121*** 0	-.00160*** 0	
Professional, Scientific, and Technical[3]	-0.00107*** 0					
Administrative and Support[3]	-0.00111*** 0					
Public Administrative and Defense	-0.00181*** 0	0.00386* (0.024)	0.0133*** 0	0.00195* (0.02)		
Education	0.0000017 (0.40)	0.00537*** 0	0.00622*** 0	-0.00075 (0.293)		
Health and Social Work	-0.00223*** 0	0.000141 (0.884)	0.0178*** 0	0.00144* (0.022)	0 (.)	
Arts, Entertainment, and Recreation	-0.00111*** 0	0.00114 (0.609)	0.00369** (0.006)			
Energy and Water		-0.00412 (0.105)	-.00831*** 0	-0.00475** (0.003)	-0.000101 (0.782)	
Professional, Scientific, and Technical, Administrative Activities		-0.00229 (0.19)	0.00749*** 0	0.00318* (0.021)	-0.00313* (0.051)	
Constant	-3.075*** 0	-55.47*** 0	35.77*** 0	2.257 (0.317)	9.557* (0.011)	-12.34* (0.012)
Number of Observations	304	192	222	272	185	12
R^2	0.364	0.687	0.662	0.393	0.405	0.796

All figures in parentheses are p-values. All variables are first-differences. Significance levels. ***p<0.01, **p<0.05, *p<0.1. Errors are clustered standard errors.

Notes: 1. For China and Argentina, data on fraction of workforce with tertiary education wasn't available from OECD. For China, ratio of graduates to working population from Chinese Statistical Yearbooks were used. For Argentina, World Bank data on enrolment ratio in tertiary education was used. 2. Except for the UK, data is available for the electricity, gas and air conditioning supply, and water supply sectors together. 3. Except for the UK, data is available for the professional, scientific and technical and the administrative services sector together.

Helping your business to make the most of AI

PwC is already working with companies across each of the different sectors highlighted in this report, to help them plan for and take advantage of AI to support their business strategy and improve performance.

If you would be interested in a consultation about the potential within your business, please feel free to get in touch.

Anand Rao
Global Leader of Artificial Intelligence, PwC

T: +1 (617) 530 4691
E: anand.s.rao@pwc.com
Twitter: @AnandSRao

Gerard Verweij
Global & US Data & Analytics Leader, PwC US

T: +1 (617) 530 7015
E: gerard.verweij@pwc.com
Twitter: @gverweij

Jonathan Gillham
Head of Econometrics and Economic Modelling, PwC UK

T: +44 (0)7714 567297
E: jonathan.gillham@pwc.com

Euan Cameron
Artificial Intelligence Leader, PwC UK

T: +44 (0)20 7804 3554
E: euan.cameron@pwc.com
Twitter: @euancameron55

Authors

Jonathan Gillham
Head of Econometrics and Economic Modelling, PwC UK

Lucy Rimmington
Economist, PwC UK

Hugh Dance
Economist, PwC UK

Gerard Verweij
Global and US Data & Analytics Leader, PwC US

Anand Rao
Global leader of Artificial Intelligence, PwC US

Kate Barnard Roberts
Management Consultant and AI specialist, PwC US

Mark Paich
SME in AI and Economic Analysis, PwC US

The authors would also like to acknowledge the contribution of:

- Alan Morrison
- Barbara Lix
- Cathryn Marsh
- Chengyao Gu
- Chris Curran
- Craig Scalise
- Cristina Ampil
- Edmond Lee
- Euan Cameron
- John Hawksworth
- John Sviokla
- Michael Schneider
- Pia Ramchandani
- Rob McCargow
- Shivangi Jain

We would also like to thank our academic and applied research partners at Fraunhofer, and especially Fraunhofer Big Data Alliance