The Economist Intelligence Unit

# RISKS AND REWARDS

## Scenarios around the economic impact of machine learning

A report from The Economist Intelligence Unit

Commissioned by Google

# CONTENTS





© The Economist Intelligence Unit Limited 2017

© The Economist Intelligence Unit Limited 2017

# ABOUT THE REPORT

Risks and rewards: Scenarios around the economic impact of machine learning is a report from The Economist Intelligence Unit sponsored by Google. It is based on the results of econometric modeling of three scenarios covering five countries and Developing Asia as a region. In addition, the report also presents qualitative scenarios for four industries: manufacturing, healthcare, energy, and transportation.

The econometric modelling and qualitative conducted by a team led by Trisha Suresh, with contributions from Simon Baptist, Patricia Morton, John Ferguson, Michael Frank, Vaibhav Sahgal, Ana Nicholls, and Matthew Kendall. In constructing the model, the team was advised three external experts (see below). The paper was written by Paul Kielstra and edited by Chris Clague. Findings from the modelling were supplemented with wide-ranging research and in-depth interviews with experts. Our thanks are due to the following people for their time and insights (listed alphabetically by affiliation):

## Expert advisory panel

**Jim Minifie**, director, Productivity Growth Program, Grattan Institute

**Beth Webster**, professor, Swinburne University of Technology

**Seth Benzell**, postdoctoral associate, MIT Initiative on the Digital Economy

## Interviewees

**R. Kim**, chief marketing officer and chief retina services, Aravind Eye Hospital

**Jim-Hyung Kim**, Artificial Intelligence Research Institute

**Manuela Veloso**, professor, head of machine learning department, Carnegie Mellon University

**Amy Abernathy**, chief medical officer, Flatiron Health

**Scott Otterson**, probabilistic wind forecasting, Frauenhofer Institute for Wind Energy and Energy System Technology


© The Economist Intelligence Unit Limited 2017

**Colin Parris**, vice president, GE Software Research

**Satyam Priyadarshy**, chief data scientist, Halliburton

**Andy Chun**, senior IT consultant, MTR, Hong Kong

**Masauyuki Morikawa**, vice president, Research Institute of Economy, Trade, and Industry (REITI)

**Blake Moret**, chief executive officer, Rockwell Automation

**Sujeet Chand**, chief technology officer, Rockwell Automation

**Sung-ho Lee**, research fellow, Science and Technology Policy Institute

**Jerry Kaplan**, CodeX fellow, The Stanford Center for Legal Informatics

**Ajay Agrawal**, professor, University of Toronto

© The Economist Intelligence Unit Limited 2017

# EXECUTIVE SUMMARY

There is more uncertainty around advances in artificial intelligence (AI) and one of its major sub-sets, machine learning, than the current debate suggests, particularly with regard to the technology's impact on society and the economy. No doubt the advances have indeed been incredible and advocates are right to highlight them. A decade ago few believed that a car could drive on its own, even in a controlled environment, or that an algorithm could learn how to label and organise photographs. Yet both of those are now possible and various forms of AI are performing new tasks it seems on a weekly basis.

Not everyone views this as an unalloyed good, however. In fact, there is great concern that AI poses a threat to jobs, privacy, and, eventually, even humanity. These concerns are not without merit, although the degrees to which they comport with reality vary, at least in the near-term. AI does indeed have as much potential to roil society as it does to improve it. Every new task that a machine learning algorithm—let alone an entire occupation—masters could mean a job lost. That doesn't mean a new job or entire occupation won't be created elsewhere because of it, but it raises important questions about the future of the labour market.

The problem is that both sides—the supporters of AI and its opponents, the champions and the naysayers—often deploy hyperbole to further their view. As a result, much of the current debate AI has become an either/or proposition. Either it will lead inexorably towards a utopian future or it will be the cause of our demise.

The truth probably lies somewhere in between and The Economist Intelligence Unit, with sponsorship from Google, has conducted research to identify the middle ground by developing quantitative and qualitative scenarios on the impact of machine learning for a select number of countries and industries. The findings are based on econometric modelling, desk research and interviews with academic and industry experts.

## Impact on GDP and productivity

The Economist Intelligence Unit ran three econometric scenarios, using our current forecast to 2030 as the baseline. We covered five countries—the US, UK, Japan, South Korea and Australia—and the countries of Developing Asia as a group. The scenarios are:

- Scenario #1: Greater human productivity through upskilling


© The Economist Intelligence Unit Limited 2017

- Scenario #2: Greater investment in technology and access to open source data

- Scenario #3: Insufficient policy support for structural changes in the economy

**Scenario #1** assumes a higher degree of complementarity between human skills and AI than does the baseline and that governments will invest more in upskilling than current trends suggest. In the results, every country or grouping covered benefits, but some more than others. Australia, where growth in services is becoming more important for incremental economic growth than commodity exports, would see the greatest gains. The gains elsewhere would be more modest by comparison; although in this scenario, the UK's productivity rises to slightly positive from our baseline forecast, which is for a slight decline.

**Baseline v Scenario #1, GDP changes by country**

| COUNTRY | BASELINE | SCENARIO |
|---|---|---|
| US | 1.84% | 2.04% |
| UK | 0.63% | 1.29% |
| Australia | 1.03% | 3.11% |
| Japan | 1.57% | 1.96% |
| South Korea | 1.78% | 2.07% |
| Developing Asia | 4.34% | 5.04% |

**Scenario #2** assumes investment in access to open source data, tax credits to spur private sector adoption of machine learning, and advances in computing efficiency drive hardware costs down. This scenario yields the most encouraging results insofar as economic growth is concerned. Each of the five countries, as well as Developing Asia as a group, experience higher levels of growth relative to our baseline forecast. Australia, again, along with Developing Asia, reap the greatest rewards from promoting investment in this scenario, but all of the countries covered see GDP rise by at least 1% above the baseline between now and 2030.

**Baseline v Scenario #2, GDP changes by country**

| COUNTRY | BASELINE | SCENARIO |
|---|---|---|
| US | 1.84% | 3.00% |
| UK | 0.63% | 1.94% |
| Australia | 1.03% | 3.74% |
| Japan | 1.57% | 2.43% |
| South Korea | 1.78% | 3.00% |
| Developing Asia | 4.34% | 6.47% |

**Scenario #3,** which is the one negative scenario among the three, assumes the substitution effect for labour dominates due to inaction in workforce development—or more simply, skills—and a lack of national data sharing schemes. The losses are substantial compared to the baseline. The UK and Australian economies actually shrink in US dollar terms versus today, as a result, with the UK's economy becoming US$420bn smaller in absolute terms and the Australian economy US$50bn. The US, Japan and Developing Asia still grow in this scenario, but their economies are all significantly below

© The Economist Intelligence Unit Limited 2017

the baseline, with the US and Developing Asia both off by around US$3trn.

**Baseline v Scenario #3, GDP changes by country**

| COUNTRY | BASELINE | SCENARIO |
|---|---|---|
| US | 1.84% | 0.84% |
| UK | 0.63% | -1.20% |
| Australia | 1.03% | -0.24% |
| Japan | 1.57% | 0.53% |
| South Korea | 1.78% | 0.02% |
| Developing Asia | 4.34% | 3.20% |

## The industries

The qualitative scenarios look at four industries: manufacturing, healthcare, energy and transportation.

### Manufacturing

Employment in manufacturing has become a headline issue with the rise of populism in certain developed countries. When discussing AI, it's important to differentiate automation in hardware, such as robotics, from automation in software, AI and its sub-sets. The former has already had a significant effect on labour demand in the sector and while the latter may contribute to this trend, its impact has been less direct, at least to date.

When manufacturing firms talk about AI, they talk about creating greater efficiencies in their supply chains, reducing maintenance costs and moving towards batch production. Each of these may or may not result in the elimination of low wage jobs, but they will almost certainly create high wage ones, albeit not at a one-for-one ratio. The speed at which firms turn towards automation in both hardware and software depends on the "payback period", a measure which weighs the cost of investment in automation versus that of the cost of local labour.

### Healthcare

As a knowledge industry, healthcare is ripe for AI and there are a variety of applications already in place. It's being used in the discovery process for new drugs, to save costs in both prevention and treatment, and to augment the abilities of practicing physicians and clinicians.

Yet there are constraints. The healthcare sector has traditionally been slower than most sectors to adopt innovations. That may be changing, however slowly, but there are other hurdles that need to be overcome. One is the issue of privacy. Patients are understandably sensitive about their personal data being shared and unless they can be assured their data will only be used for specific and agreed purposes, they may not agree to sharing it at all. That would hinder the use of AI considerably, dependent as it is on data for developing solutions.

### Energy

AI is expected to have the most significant impact on the energy sector in transforming generation, transmission and distribution
into a more coherent system. This means, among other things, creating pricing systems


© The Economist Intelligence Unit Limited 2017

based on probabilistic models and developing smart grids that can better deal with the issue of intermittency, or the fact that the wind doesn't always blow and the sun doesn't always shine. Solving for intermittency will allow electricity providers to maximize their use of green energy.

This does not come without risks. Smart grids, while more efficient than current analogue grids, are exposed to new and greater risks from cyber-attacks, which, in turn, creates national security concerns. That has knock-on effects on the willingness of local governments to integrate their grids and share their data.

**Transportation**

While autonomous vehicles have captured much of the public's attention, even though they may still be far off, AI is already making major contributions to the speed and safety of public transport. In many cities, AI being used to balance the flow of passengers across different modes of transport and data received from sensors around cities, combined with AI, are helping to make traffic flow more smoothly.

The advent of autonomous vehicles nevertheless remains in the fore of people's minds in this area. Besides the obvious issue of its impact on employment, there are regulatory and privacy concerns, as well as the question of liability when, inevitability, a driverless vehicle becomes involved in an accident, fatal or otherwise.

## Conclusion

The debate around AI is sure to intensify in the coming years. At present, sound analysis and information of the issue appears to exist on the outside of a broad core of misunderstanding and misinformation, a situation that ultimately benefits no one. Based on the results of our research, we have identified five approaches to grounding the discussion in reality.

**Managing expectations.** In the near-term, AI will be neither utopian nor dystopian. It will provide new benefits and it will create new problems. Exaggerating its upsides is as detrimental to the debate as is exaggerating its downsides.

**Better communication.** There are many understanding gaps when it comes to AI, but one of the most important to bridge is that between developers and businesses and government institutions. The former are often only dimly aware of what the latter two really need, and the latter, in turn, are often only dimly aware of the potential solutions the former could provide. A more robust and frequent exchange of information, capabilities and needs would help to remedy this.

**Acknowledging the risks.** It is important to acknowledge that AI presents risks to employment, as well as privacy, and to start finding solutions to these and other issues rather than encourage complacency or resignation through unshakeable confidence.

© The Economist Intelligence Unit Limited 2017

**Improving trust and transparency.** "Trust us" or trust the algorithm is not a viable strategy for gaining widespread acceptance of AI and its various subfields. Developers and users alike need to make known what they are doing and how they are doing it, in a way that is both meaningful given the usage context and practical given technological constraints.

**Educating the public.** Gaps in knowledge and understanding are filled more quickly than ever with misinformation and distortion. The public needs an explanation of what AI is and does, and as simply as possible.

Policymakers, for their part, face a number of choices as regards AI and its impact.

**Investing in skills and training**. That there is going to be churn in labour markets as a result of AI is widely accepted. Vocational education, now lacking is most countries, will need to become more prevalent. The growing focus on STEM education is important, as well, but the expected rise in demand for "soft skills" such as team building, cooperation, and critical thinking means that liberal arts should not be neglected. The right mix of these three, and others, will require constant monitoring and close cooperation between inudstry, educators, and policymakers.

**Dealing with data**. The uses and misuses of data are going to be among the defining issues of the 21st century. More needs to be done to assuage citizens' concerns about privacy and security through measures such as regulations that enable and support the use of anonymised data sets. These measures also need to be inter-operable across borders so that data can flow around the world.

**Investing in R&D and technology**. Public sector investment in R&D has decreased in many countries, creating a gap that has been filled, in part, by the private sector. That could prove unsustainable and if countries are to have the intellectual capacity to capitalise on new technologies, the public sector will need to be more involved.


© The Economist Intelligence Unit Limited 2017

© The Economist Intelligence Unit Limited 2017

# INTRODUCTION

In just the last few years, we have seen incredible advances in artificial intelligence (AI), especially in machine learning, one of its subsets. Besides beating humans in games like the US quiz show Jeopardy and the strategy board game Go, machine learning algorithms are now having a real impact on our daily lives, most sectors of the economy, and society at large.

To some, this is a welcome development, with wide-ranging and positive implications for humanity. For example, healthcare outcomes, already improving through the use of robotic procedures, will continue to do so on the back of new solutions developed through the use of machine learning and other forms of AI, which will also bring costs down and expand access. More broadly, proponents of AI contend that this and other applications will help to raise productivity, enable more efficient use of resources, change the way we live and work, and even provide a boost to creativity.

Not everyone shares this optimism. Those sceptical of AI, if not outright hostile to it, generally fall into one of two camps. There are critics who understand the technology and warn that in addition to being open to misuse, AI could render large swathes of the global workforce redundant and exacerbate the trend of rising inequality, which in turn could create massive social disruptions. Then there are the less reasoned, less informed opponents of AI, whose fears are founded more in science fiction than in fact and who think of only murderous computers like HAL 9000 from the movie 2001: *A Space Odyssey* or Skynet from *Terminator* movies when they think of AI.

To help keep the debate on the reasoned and informed side, The Economist Intelligence Unit has undertaken quantitative and qualitative research, sponsored by Google, to look at the impact of AI on the GDP and productivity of five individual economies—the US, UK, Japan, South Korea, and Australia—and developing countries in Asia as a group. The econometric modelling used our current forecasts as a baseline to develop three scenarios around AI between now and 2030:[1]

- **Scenario #1: Greater human productivity through upskilling**, which assumes that, through government investment, higher education and access to finance, there will be a higher degree of complementarity between labour and machine learning.

1 For a full explanation of the methodology, please see Appendix 2.


© The Economist Intelligence Unit Limited 2017

- **Scenario #2: Greater investment in technology and access to open source data**, which assumes investment in access to open source data, tax credits to spur private sector adoption of machine learning and advances in computing efficiency drive hardware costs down.

- **Scenario #3: Insufficient policy support for structural changes in the economy** is the one pessimistic scenario of the three and assumes that, due to government inaction on upskilling and no active role in national data sharing, the substitution effect dominates, creating large pools of involuntary unemployed.

The qualitative research focused on the impact of AI on four select industries: manufacturing, healthcare, energy and transportation. The results are based on an extensive review of the literature and a number of interviews with experts from industry and academia and are presented as a complement to the quantitative scenarios in the main body of the report, as well as in four case studies.

## Box out: Terminology

***AI and machine learning***

A major source of confusion in discussing AI and its impact is the terminology. AI, a category term, is often used interchangeably with its constituent subfields. In the scenario modelling, the focus was on the impact of machine learning specifically and not the whole of AI.

We define machine learning as a subfield of AI that leverages algorithms which learn and optimise from data without being explicitly programmed to accomplish a task with pre-defined rules. It incorporates a variety of techniques, including neural networks, decision tree learning and support vector machines, among others.

***Automation and robots***

These terms are not one and the same. In general, when the paper discusses automation, the distinction is made between automation in software, like machine learning, and automation in hardware, like robots in factories. When only the term 'robots' is used, it is referring to automation in hardware.

© The Economist Intelligence Unit Limited 2017

# SCENARIO 1: GREATER HUMAN PRODUCTIVITY THROUGH IMPROVED SKILLS

## Productivity ain't what it used to be, but why?

In much of the developed world, productivity growth has declined in recent decades and is now largely stagnant. According to total factor productivity (TFP) data from the Conference Board, between 1995 and the turn of the century, productivity in OECD countries rose on average by just under 1% per annum. Between then and 2007, it was already sliding lower, to 0.6% annually. Since 2008, however, TFP in these nations has actually declined, down 0.6% per year on average.[2] This is not merely an artefact of the downturn: the measurement of TFP adjusts for the amount of capital and labour used, so high unemployment will have a muted effect on the figures.

The OECD itself has data reaching back to 1985 on productivity for 20 of its member states. This paints a similar picture over a longer period. Although the numbers often vary widely by country and year, a clear underlying pattern exists: with a handful of exceptions, in every

**Average annual TFP growth in selected OECD countries**

| COUNTRY | 1985-1989 | 1990-1999 | 2000-2007 | 2008-2015 |
|---|---|---|---|---|
| France | 1.9% | 1.0% | 1.0% | 0.0% |
| Germany | 1.4% | 1.4% | 1.0% | 0.4% |
| Italy | 1.5% | 0.6% | -0.1% | -0.4% |
| Japan | 2.9% | 0.8% | 0.8% | 0.6% |
| Korea | 5.2% | 4.1% | 3.1% | 1.7% |
| Spain | 1.5% | 0.3% | -0.1% | -0.1% |
| UK | 1.1% | 1.0% | 1.6% | -0.2% |
| US | 0.7% | 1.0% | 1.3% | 0.5% |
| **20-country average** | **1.6%** | **1.2%** | **1.0%** | **0.1%** |

Source: Economist Intelligence Unit calculations based on OECD data http://stats.oecd.org/Index.aspx?DataSetCode=PDB_GR#

2The Economist Intelligence Unit calculations based on data in The Conference Board "Growth Accounting and Total Factor Productivity, 1995-2015," November 2016, spreadsheet available at https://www.conference-board.org/retrievefile.cfm?filename=TED_2_NOV20161.xlsx&type=subsite


© The Economist Intelligence Unit Limited 2017

country average annual TFP growth in the second half of the 1980s, in the 1990s, and between 2000 and 2007 were all higher than the mean figure for 2008 to 2015. The average figures suggest a slow decline over the first three periods, followed by a shuddering halt to productivity growth since 2008.

Further Conference Board data suggest a similar shift is taking place in much of the developing world. As the chart shows, several key emerging markets saw marked rises in productivity between the late 1990s and the early years of this century, but since 2008 this has tailed off and even gone into reverse. The exception is India, but its relatively low figures for change in productivity indicate it is something of an outlier.

The shift has not gone unnoticed. One influential line of argument in recent years is that the long-term increase in productivity from which developed countries—and particularly the US—have benefited for over a century has either ended or reached a plateau. For example, economist and professor Tyler Cowen argued in his 2011 book *The Great Stagnation*, that the US had taken advantage of all the available low-hanging fruit—both in terms of technology and public policy—that could drive economic growth. As a result, occasional short periods of productivity improvement might still occur, he argues, but these will tend to be focused in specific industries. Slow or non-existent increases in median income and stock market prices instead point to a broader economy that is seeing little productivity improvement. More recently, in 2016, Robert Gordon's *The Rise and Fall of American Growth* asserted that today's innovations are simply not delivering the kind of improved productivity which technology had been doing since the appearance of the electric light bulb.

**TFP growth in selected emerging markets**

| COUNTRY | 1995-2000 | 2001-2007 | 2008-2015 |
|---|---|---|---|
| South Africa | -0.1% | 1.3% | -1.2% |
| China (Official) | 3.0% | 3.6% | 1.4% |
| India | 0.8% | 0.3% | 0.9% |
| Indonesia | -0.4% | 1.7% | 1.2% |
| Argentina | -0.6% | 1.0% | -0.7% |
| Brazil | 0.2% | 0.9% | -0.5% |
| Chile | 0.6% | -0.4% | -1.7% |
| Russian Federation | 1.0% | 5.6% | 0.6% |

Source: Economist Intelligence Unit calculations based on data in The Conference Board "Growth Accounting and Total Factor Productivity, 1995-2015," November 2016, spreadsheet available at https://www.conference-board.org/retrievefile.cfm?filename=TED_2_NOV20161.xlsx&type= subsite

© The Economist Intelligence Unit Limited 2017

The culprit in these arguments is not a lack of technological progress. The phone in your pocket should tell you that's not the case. The issue is whether, like previous waves of new technology, current ones will lift the world's stagnating productivity. Our first scenario projects the likely effect of greater investment in the skills needed to capitalise on the opportunities provided by AI and thereby potentially reverse the slowdown in productivity growth.

## Key assumptions

This scenario differs from the baseline in one key respect: it assumes a higher degree of complementarity between human skills and AI, particularly machine learning, than current trends suggest.

The scenario also assumes that such a shift would result largely from policy changes. Governments would invest more heavily in relevant vocational education for the AI-enabled workplace technologies and tools. The state would at the same time also work to improve general skills across the labour force by putting more money into higher education. Finally, on the assumption that individuals would see the economic value to themselves of being better able to exploit the new technology, the government would expand access to loans and other financing for education.

Another implicit assumption, worth discussing briefly, is that those investing in skills will know broadly which ones will be needed in the emerging technology landscape. The precise mix is less uncertain. Some are obvious. The UK's Royal Academy of Engineering, in a presentation to a House of Commons committee in April 2016, noted that future employees would require a "[challenging] combination of skills... drawing on engineering, computer science, machine learning, mathematics and informatics". Nor can these skills exist in a vacuum, argues Dr Colin Parris, Vice President at GE Software Research: "We need educational policies that allow us to do deep scientific research. If you don't invest to maintain the intellectual horsepower in raw science, it quickly will put you behind the curve, and you never catch up."

More than scientific and technological understanding will be needed, however. The Royal Academy's presentation added that, because it was unlikely any one person would have sufficient understanding of all these areas, "soft skills" like team building and cooperation would also be essential. Ajay Agrawal, Professor of Entrepreneurship at the University of Toronto, goes further. He argues that AI makes prediction based on previous trends much easier, but it does not help with certain fundamental tasks of judgement about the desirability or wisdom of certain ends. Accordingly, he believes that the ability to exercise such judgement will increase in value as prediction can grow more and more automated. As a result, he expects, ironically given current priorities, "a whiplash in the shift in educational emphasis that has been away from the humanities toward engineering back toward


© The Economist Intelligence Unit Limited 2017

the arts and humanities. Topic areas like history, philosophy, and literature will become valuable again, because we'll realize that these subject areas provide the basis for learning what we call judgement."

Meanwhile, for Professor Manuela Veloso, head of the Machine Learning Department at Carnegie Mellon University, the issue will not simply be one of specific skills but of inculcating a whole new mind set. "People are not yet fully aware," she says, "that there are tons of data and that there is no way to benefit from it without AI and machine learning.  Education has to move to the concept of data thinking—it has to be pervasive."

The scenario assumes that public policy will increase overall level of skills in the workforce. The future will show precisely what those are.

In this scenario, the impact of machine learning on productivity is mixed relative to the baseline, which is our official forecast for the ratio of GDP in US dollars to millions of man-hours worked in each of the five economies and Developing Asia as a group.[3] Encouragingly, however, all the countries registered improvements by this measure, although the magnitude of those improvements varies widely.

Australia will reap the greatest benefits in this scenario, with the compound annual growth (CAGR) in productivity increasing to 2.25% against the baseline forecast of just 0.19%. Among its developed world peers, Australia has long lagged in the level of skills in its labour force, although as it continues to make the shift away from reliance on primary industries towards a more services-oriented economy,

## The scenario results
### Changes in productivity in Scenario #1

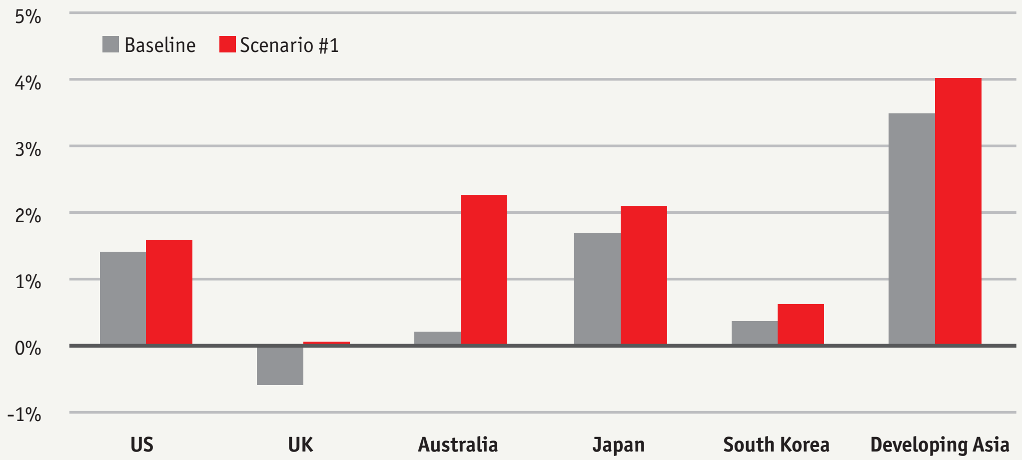


3 For a full explanation of the methodology please see Appendix II

© The Economist Intelligence Unit Limited 2017

some upskilling was already in evidence. Over the forecast period, an acceleration in this trend, coupled with AI, provides a significant boost to the country's productivity.

At the other end of the spectrum are three of Australia's developed world peers: the US, Japan and South Korea. Relative to the baseline, none of the three see major gains from the combination upskilling and AI. One possible explanation is that because these countries are world-leaders in the development of AI, and already have highly-skilled labour forces, at least relative to the rest of the world, productivity improvements are already baked into our baseline forecast in a way that they are not for a country like Australia. Another is that, for all the emphasis on skills, especially in the developed world, upskilling and AI alone are not always and everywhere sufficient to escape secular stagnation.

The UK presents an interesting case in this scenario. For a number of potential reasons, including Brexit and the diminution of the highly productive financial services sector, the baseline forecast is for the country's productivity to be negative between now and 2030. Upskilling and machine learning, while not a means of returning the UK to an era of massive growth in productivity, does stem the decline by maintaining it at current levels. That may not be the most cheering result, but it is better than the alternative.

Developing Asia falls between Australia and the other developed countries covered here. In part, this is due to their strength in our baseline forecast, which is for a CAGR of 3.49%, by far the highest baseline forecast. It's also because taking advantage of upskilling and AI is more difficult for these economies. For one, their labour forces start from a much lower level of overall skills, meaning they have further to go than do the US and Japan, for example. They are also held back by the baseline level of computer capital, which is our proxy for machine learning AI. We'll see how this can be addressed in Scenario 2, the results of which are explained in the next section.

## Productivity in key sectors

The evolving links between productivity, people and machine learning will take a variety of shapes in the key sectors covered by our qualitative scenarios.

### Manufacturing

Manufacturers have for centuries embraced the latest new technology and the current wave will be no exception. One sign among many: the Boston Consulting Group projects that growth in the global installed base of advanced robotics will accelerate from around 2-3% annually in 2015 to roughly 10% per year over the next decade.[4]

Given manufacturing's traditional provision of low-skilled jobs, though, what role will humans play in the new, technology-enabled production?

4Harold Sirkin et al., "The Robotics Revolution: The Next Great Leap in Manufacturing," 23 September 2015, BCG Perspectives, https://www.bcgperspectives.com/ content/articles/lean-manufacturing-innovation-robotics-revolution-next-great-leap-manufacturing/


© The Economist Intelligence Unit Limited 2017

Our qualitative scenario projects that the total number employed in the sector will remain relatively steady, but among the low-skilled there will be job losses. Certain types of employment will disappear completely as AI-enabled machines do the required tasks better; in other cases, because of partial rather than complete automation, some low-skill jobs will remain, but their number will be markedly fewer than in the past. Meanwhile, though, the scenario foresees that for all the employees they cut, companies will take on a similar number of higher skilled individuals able to add additional value by working with the new technology. Although not a universally accepted argument, it is consistent with how past waves of technological change have affected manufacturing. Indeed, the World Economic Forum recently argued that manufacturing jobs "still retain relatively good potential for upskilling, redeployment and productivity enhancement through technology rather than pure substitution."[5]

The potential of the new technology for reshaping manufacturing is profound, with everything from product design, through production, to marketing and shipping already beginning to see substantial signs of incipient change. For instance, Procter & Gamble, a consumer goods company, has cut unplanned down time by 10% to 20% by integrating AI into its operations, leading to better end-to-end supply chain efficiency and enabling humans to do the jobs they are paid to do rather than sitting idle.[6]

Current processes are not only being improved, though; they are being reinvented. To cite just a couple of examples, NotCo, a Chilean company, uses AI software for product development. It breaks down foods into their basic molecules and then uses machine learning algorithms to create vegan alternatives, which in turn are taste-tested by people. It is already marketing a mayonnaise.[7] Meanwhile, AI and robotics can allow increasingly personalised merchandise. Adidas' is building a new Speedfactory which will permit batch sizes of one, as robots manufacture shoes for the specific needs and wants of a given customer. The high level of production automation is not a complete job destroyer. Low paid positions in the mass production of shoes—jobs the company is increasingly finding it hard to fill—may eventually be reduced, but the Speedfactory will also create 160 new jobs.[8] Although the company has not gone into detail about these, they will inevitably include high tech work, such as programming and maintaining the factory's robots. Also, Adidas has noted that humans still do better than machines in certain specialised aspects of shoemaking, such as the final shaping and product design.

5 *The Future of Jobs Employment, Skills and Workforce Strategy for the Fourth Industrial Revolution*, 2016 http://www3.weforum.org/docs/WEF_Future_of_Jobs.pdf.
6 "These factories are crunching production time with artificial intelligence," *Business Insider*, 9 July 2016, http://www.businessinsider.com/sc/artificial-intelligence-change-manufacturing?IR=T.
7 "Chilean start-up that uses AI to reinvent food eyes U.S. deals," *Business Insider*, 5 December 2016, http://www.businessinsider.com/r-chilean-start-up-that-uses-ai-to-reinvent-food-eyes-us-deals-2016-12?IR=T.
8 "Adidas's high-tech factory brings production back to Germany," *Economist* 20 January 2017, http://www.economist.com/news/business/21714394-making-trainers-robots-and-3d-printers-adidass-high-tech-factory-brings-production-back.

© The Economist Intelligence Unit Limited 2017

## BOX OUT: CASE STUDY #1

### GE: Taking the guesswork out of manufacturing

Traditional manufacturing relies heavily on assumptions and approximations based on inevitably imprecise data tempered by experience. When all else fails, sometimes companies have to fall back on educated guesswork. To cite just one example, maintenance schedules for equipment are often based on average wear and tear on a type of machine, not detailed information on how a specific one has been used. The result of such an approach is usually excessive maintenance costs—better safe than sorry—or, in the worst case, unexpected breakdowns.

GE has been looking at ways to reduce the inherent uncertainty in manufacturing. Its solution begins with a straightforward idea which Gartner, an IT research firm, lists as one of its top 10 strategic technology trends for 2017.[9] For any given physical asset, by combining data from automated sensors, detailed operational information, historical information on similar equipment, and the laws of physics, says Dr Colin Parris, vice president of Software Research at GE Software Research, "I can create a digital twin. I don't have to make assumptions. The twin can tell me that here is the exact wear, tear, and damage that will have happened in a certain scenario." Similar to the analytics technology which allows consumer Internet companies to understand the purchasing patterns of each individual customer, the digital twin describes in detail the state of each individual machine, or even part of a machine, in a manufacturing facility.

The first, most obvious use, is more efficient equipment maintenance. Data in the digital twin can make clear the need for early, planned prevention which can take place at times which disturbs productivity least. This is typically much less expensive than the need for emergency repairs which may bring production to a half. Just as important, notes Dr Parris, is reducing unnecessary maintenance. He notes of GE's work servicing transportation equipment—which also now uses digital twins—that avoiding unneeded down time "is a win-win. The customer gets to use the asset more and we do not need to do the servicing."

Maintenance, however, is only the beginning. Dr Parris explains that, just as digital twins of individual parts can be combined into twins of entire machines, these in turn can be combined to create precise digital representations of entire production lines, processes, and even a manufacturing facility

9 Gartner's Top 10 Strategic Technology Trends for 2017," 18 October 2016, http://www.gartner.com/smarterwithgartner/gartners-top-10-technology-trends-2017/


© The Economist Intelligence Unit Limited 2017

as a whole. "Then," he adds, "I can optimise how much more performance I can get from the various parts of the system compared to how much is damaging" to one or more of them. This is not a theoretical optimum, but based on the real-time condition of the equipment in the facility.

With machine learning, for example, those running the factory can see not only if one line may need to reduce production, but also how much load it will be possible to shift to a different line. Such data also affects other parts of the business, Dr Parris adds. "If I know my lines can't deliver a certain output, why should I buy extra components? We use this to go back into the supply chain so we don't tie up free cash flow."

The technology even allows decision-making that optimises costs and benefits across a multi-facility manufacturing operation. "I might see that I need to do something on a different line in a different country," says Dr Parris, but that brings up matters of the cost and transport of the needed materials to that facility, as well as storage issues and even has tax implications. "I can minimise costs and maximise how to benefit" in such a situation by being able to model the process using digital twins and AI.

Even cybersecurity can be improved notes Dr Parris. He explains that GE is creating software called Digital Ghost which uses information from a digital twin to monitor for discrepancies. For example, a given piece of technology will have certain patterns of power usage if it is doing what it is supposed to. If sensors indicated that pattern is not being followed it could be a sign that the controls have been hacked, even if the malware involved has hidden itself from other types of detection.

While still being rolled out, the technology is proving itself. Dr Parris reports that digital twins have at certain facilities "driven improvement of 25% factory throughput and 32% inventory reductions while improving on-time delivery." Informed decisions are more profitable than guesses.

© The Economist Intelligence Unit Limited 2017

## BOX OUT: CASE STUDY #2

### A Coming Revolution: Aravind and the use of AI for retinopathy diagnosis

The Aravind Eye Care System in southern India has grown from an 11-bed clinic set up in the 1970s to a multi-hospital group that, in the year ending March 2016, dealt with 4.7m patients and conducted over 400,000 operations—about 5% of all eye surgeries in the country.

The organisation is famous for its innovation, including a business model that allows cross-subsidisation of surgery costs for the poor so that most patients receive free or subsidised care. Aravind has also been at the forefront of technological innovation in Indian healthcare. Its 57 rural clinics rely on advanced tele-medicine to diagnose patients without the need to travel into the city. Now Aravind is looking at the potential of AI and deep learning to help in its mission.

All of Aravind's facilities, including the rural clinics and mobile vans, have the capacity to take images of the back of patients' eyes. These in turn are sent for assessment by the organisation's relatively limited number of expert ophthalmologists who are based in its central hospitals. There they can, among other things, look for signs—in the patients with, or suspected of having, diabetes—of the kind of eye damage which is a common complication of that disease: diabetic retinopathy. If spotted early, this can be treated effectively and inexpensively. If not, blindness can ensue.

In recent trials, in addition to expert human analysis, the photos have been assessed by software which, using deep learning algorithms, had previously trained itself on over 120,000 diabetic eye photographs to spot retinopathy. The AI had already demonstrated good results in controlled tests: these trials were the first in a clinical setting, where automated and manual analyses have been compared side by side.

Dr Kim Ramasamy, chief medical officer at Aravind, explains that the full results of this trial are awaiting publication, so he is not able to describe them in detail. Nevertheless, based on the organisation's experience, it is planning to use the technology on a greater scale.

The most noticeable benefit which the system is likely to bring is speed, says Dr Kim. With human examination of the retinal scans, he explains, the turnaround time could be around two hours, but "once we have automated


© The Economist Intelligence Unit Limited 2017

things, it will be cut to a few seconds. As soon as an image is taken, a report on that person's eyes can be obtained."

Although beneficial, this faster analysis will, Dr Kim believes, revolutionise for the better how diabetic eye care is done. "Currently, the ownership of detecting problems lies with retina specialists who are very few in number. Now that whole ownership can shift to general physicians or diabetologists. Even nurses can do it." This in turn means, he adds, that the capacity to screen for retinopathy will be in the hands of those medical personnel—whom diabetic patients see already for various reasons—rather than ophthalmologists, whom they simply may not take the time to visit. "If you can do this in a physician's office, the whole scenario changes. Now you can screen thousands instead of hundreds in the same amount of time."

Nor will this put humans out of work—quite the opposite. Ophthalmologists are currently too few for the pressing need in India, and with the country often called "the diabetes capital of the world", the amount of retinopathy is not about to decline. Now, though, specialists will be able to spend less time examining retinal images and more time treating the greater number of patients which Dr Kim believes increased screening will identify. Low-skilled jobs will also likely grow in number, he believes: "We need more people to be trained in using the cameras. Anybody can do it."

Currently, however, regulation is a barrier to rapid application of machine learning across eye care in India. "As of today," Dr Kim says, "there are no clear-cut rules for this." Since the current use is for diagnosis, not treatment, and questionable cases are reviewed by a human expert, it is not impeding progress. In the longer term, he believes the problem is soluble: "We are talking with the authorities about how to bring in rules and regulations" for the wider use of AI.

Dr Kim's experience has convinced him that such broader application will occur swiftly. "Once it is deployed for one thing," he says, "more people will find other uses, like detecting glaucoma or age-related macular degeneration early." Indeed, eventually tasks which currently require substantial expertise "will be as simple as checking weight".

© The Economist Intelligence Unit Limited 2017

**Healthcare**

Healthcare has always been a knowledge industry. Expert practitioners draw on an ever-expanding corpus of information to select the relevant pieces of information and apply them in, for example, diagnosing and deciding on treatment for a particular patient with a given ailment.

In doing so, clinicians turn out to be fallible. Across the EU 23% of people surveyed said that they or a family member had experienced a serious medical error. Not surprisingly, such experience correlated with a drop in trust in healthcare.[10] Similarly a 2016 survey of Chinese physicians—a famously overworked group—found 55% admitting to having made a medical mistake in the previous year.[11]

Fortunately, the kind of data analysis and application central to evidence-based medicine is also well suited to AI and machine learning systems. Our qualitative healthcare scenario focuses on the most visible effect which AI is likely to have: a rapid and substantial reduction in medical errors, especially in wealthier countries, and a resultant improvement in overall health and life expectancy.

The impacts on productivity in health and related fields, however, will likely go far deeper. In the pharmaceutical sector, dropping levels of drug discovery have been a concern for some years. AI systems that can comb through medical literature and other health data, however, are beginning to show promise as a source for possible new medications. In 2015, Eve—an AI system at the University of Manchester—used such an approach to find that an existing anti-cancer compound might be useful in combatting malaria, a disease where drug resistance has become an increasing concern.[12]

In healthcare itself, the potential is vast. Productivity here is best conceived of as the improvement, or at least maintenance, of health than as the mere volume of medical interventions. A study by McKinsey, the consultancy, estimated that changes to how care is delivered arising from Big Data discoveries about the optimal prevention strategies, treatments, and providers could reduce total US healthcare system costs by 12-17% without diminishing outcomes.[13]

Healthcare, however, is a field where IT and delivery model innovation is notoriously slow.[14] Nevertheless, already in areas such as diagnosis and supporting treatment decisions, AI systems are coming on stream, even in developing countries. For example, the Aravind Eye Care System—a group of hospitals and clinics in southern India—has been trialing a machine learning system, in cooperation with Google, that spots signs of eye damage related to diabetes. Although the results have not yet been published, they are sufficiently positive that Aravind is rolling out the technology more widely [see case study].

10 Eurobarometer, "Medical Errors," Special report 241, January 2006.
11 Jin Wen et al., "Workload, burnout, and medical mistakes among physicians in China: A cross-sectional study," *BioScience Trends*, 2016, https://www.jstage.jst.go.jp/article/bst/10/1/10_2015.01175/_pdf.
12 Kevin Williams et al., "Cheaper faster drug development validated by the repositioning of drugs against neglected tropical diseases," *Journal of the Royal Society Interface*, 2016.
13 Basel Kayyali et al., "The big-data revolution in US health care," April 2013, http://www.mckinsey.com/industries/healthcare-systems-and-services/our-insights/the-big-data-revolution-in-us-health-care
14 Clayton Christensen et al., *The Innovator's Prescription*, 2009.


© The Economist Intelligence Unit Limited 2017

### Transportation

The driverless vehicle is the machine learning application which has most captured the public imagination. Indeed, under our qualitative scenario for this transportation sector, by 2030 most major modes of public transportation—as well as private automobiles—will be largely or completely controlled by AI-technology. The reasons are many: fewer fatalities arising from human error, more efficient use of existing travel infrastructure, less polluted and less congested cities, among others.

The first steps along this path have already occurred. Numerous companies are working on driverless automobiles. Meanwhile, for freight transport, in April 2016, a fleet of a dozen Daimler and Volvo trucks successfully completed the week-long, nearly autonomous long-distance European Truck Platooning Challenge. Platooning technology allows trucks to travel 1 second apart, eliminating the need for harsh emergency braking and saving on fuel costs.  Meanwhile, Professor Andy Chun of City University of Hong Kong notes that "many cities around the world already have some metro trains and/ or monorails running driverless.  When they first appeared, people were quite concerned about safety. Now, it is just expected, and people don't think twice when boarding one." In fact, they are now unexceptional, with examples as far afield as Kobe—which had the world's first fully automated metro line, Sao Paulo, Vancouver and Paris, to name just a few. Automated public transportation is already moving beyond trains. In 2016, Singapore saw trials of driverless taxis and Ulm, Germany expects to have fully driverless streetcars by 2020.[15]

Although harbingers of substantial future change, too great a focus on such developments misses the larger contribution which AI is making currently and will do increasingly in the near future: helping people operate existing systems more efficiently and safely.  As discussed in a later section, machine learning is letting transport companies use advanced situational and operational intelligence to maintain equipment better.

AI is also helping those driving or piloting vehicles. Autopilots have existed for decades, but are becoming increasingly sophisticated. Researchers at University College, London, for example, are creating one that learns from hundreds of hours of flying simulations by human pilots, so that it is able to react when equipment fails or the environment outside the plane changes. Testing so far has been promising and, to the surprise of the developers, the system was able to learn quickly how to assist with planes other than the kind it was designed for.  Other innovations have already been commercialised: Daimler's freightliner trucks, for example, have predictive cruise control which maintains a safe speed going up and down hills by using GPS, maps and elevation data.[17]

15 World's first driverless taxi trial kicks off in Singapore," *Straits Times*, 26 August 2016, http://www.straitstimes.com/singapore/transport/worlds-first-driverless-taxi-trial-kicks-off-in-singapore; "Green Light for Autonomous Vehicles," *Pictures of the Future*, 8 December 2015, https://www.siemens.com/ innovation/en/home/pictures-of-the-future/research-and-management/innovations-autonomous-vehicles.html
16 "Flight response," *Economist*, 15 September 2016, http://www.economist.com/news /science-and-technology/21707187-artificially-intelligent-autopilot-learns-example-flight-response
17 "Out of road: driverless vehicles and the end of the trucker," 30 March 2017, *Financial Times*, https://www.ft.com/content/2d70469c-140a-11e7-b0c1-37e417ee6c76.

© The Economist Intelligence Unit Limited 2017

# SCENARIO 2: GREATER INVESTMENT IN TECHNOLOGY AND ACCESS TO OPEN DATA

## On the road to automation: technology investment and GDP

Our first scenario dealt with one of the two fundamental contributors to economic activity—labour. This looks at the other—capital, in particular capital investment in information and communication technology (here called "computer capital").

In developed countries covered by this study, the stock of such capital been growing rapidly. We estimate that its value in each—Australia, Japan, South Korea, the UK and the US—rose by between 8% and 12% annually between 2008 and 2016. In some cases, this represents a substantial drop from the figures for the late 1990s and the early 2000s, but it would have been difficult to keep up such levels of growth as the base of computer capital accumulated.

This increase has also been faster than that for other forms of capital. As a result, the proportion of computer capital relative to total capital has steadily risen. In the US, the UK and Australia, for example, our estimate is that between 20% and 30% of the economy's capital is invested in ICT; under our baseline scenario, for the first two countries that figure will be around half by 2030, for Australia over 60%.

The rise of computer capital is much less marked in our two developed Asian countries because of faster increases in investment in other forms of capital. In Japan and South Korea, currently around 10% of economic capital is ICT-related. Developing countries have also seen substantial growth in computer capital investment, although much less steady and from a smaller base. In the countries in our Developing Asia group, for example, by 2016 only 7% of all capital was computer capital.

But is this capital doing any good for the countries in which it is invested? Here the data are difficult to read. On an economy-wide level, it is sometimes hard to see any effect. On the one hand, the output from computer capital has been rising in most economies for which there is data. The rate of that growth, though, has declined markedly. In Australia, for example, in 2008 output attributable to computer capital increased by 13%; by 2016, the rise was just 3.8%. Similar, albeit smaller drops occurred in the US and UK.[18] Part of this results from poor overall GDP growth in the broader economy. More striking, therefore, is that even

18 The Conference Board, "Total Economy Database - Growth Accounting and Total Factor Productivity, 1990-2016," spreadsheet, TCB Original Tab, May 2017, https://www.conference-board.org/retrievefile.cfm?filename=TED_2_MAY20171.xlsx&type=subsite


© The Economist Intelligence Unit Limited 2017

as computer capital makes up a larger proportion of the overall capital pool, its contribution to GDP growth relative to other forms of capital has not been changing markedly in either developed or developing countries.

This is an old problem. In the 1980s, economist Robert Solow famously joked "you can see the computer age everywhere but in the productivity statistics." The debate continues with respected analysts arguing both for and against a marked IT impact on GDP: indeed, a 2015 column in the *Financial Times* called the "impact of technology on long-run growth... one of the great unknowns—perhaps even the greatest—in economics."[19]

Should policy makers, then, be encouraging such investment? The ongoing debate over the economic effect of technology investment has made clear several considerations which indicate that, although hard to discern now, greater investment in AI and machine learning will probably have an important positive impact on GDP in the years to come.

First, traditional measures of output may not pick up a potentially substantial amount of the economic activity which technology permits and thereby possibly underestimating its overall impact.[20] As Professor Agrawal puts it, "we used to recognise billions of dollars of GDP resulting from the manufacturing and sale of film and cameras. Now we take far more photos, but their GDP contribution has gone down to zero, because we don't measure anything that is free or that we share. With AI, labour productivity will go up significantly. The question is whether we will be able to measure it."

Second, existing data gives mixed messages. A 2011 study looking at 62 countries found that ICT capital played a major role in the economic growth of high and upper middle-income states between 2000 and 2006, but had little effect in lower income ones. Similarly, a 159-country study covering of 2000 to 2009 found a correlation between the intensity of ICT use and economic growth, but again this was much higher in wealthier countries.[21] It may be that certain types of economies are readier to gain from the technology.

Finally, a Swedish study suggests that it takes more time for computer capital investment than other types to feed through to GDP figures.[22] The most recent World Economic Forum Global Information Technology Report may provide a partial explanation. It found that those countries most likely to experience a great economic impact from IT also had long traditions of high investment in computer capital. These included Finland, Switzerland, Sweden, Israel, Singapore, the Netherlands and the US. Its scoring of economic impact included items like the extent of business model and organisational change enabled by the technology.[23] Finding ways to use computer capital to reshape companies, and the economy, takes time.

Dr Satyam Priyadarshy, Chief Data Scientist at oil field services firm Haliburton, expects that

19 Gavyn Davies, "The greatest unknown – the impact of technology on the economy," 15 June 2015, https://www.ft.com/content/a89fc212-0ee5-310a-82ec-deb0220a1ec4
20 For an effort to adjust for this problem, see Dave Byrne and Carol Corrado, "ICT Prices and ICT Services: What do they tell us about Productivity and Technology?" Conference Board Economics Program Working Papers Series 16-05, https://www. conference-board.org/pdf_free/workingpapers/EPWP1605.pdf.
21 Ayoub Yousefi, "The impact of information and communication technology on economic growth: evidence from developed and developing countries," *Economics of Innovation and New Technology*, 2011 http://www.tandfonline.com/doi/abs/10.1080/ 10438599.2010.544470; Maryam Farhadi et al, "Information and Communication Technology Use and Economic Growth," PLoS One, 2012, https://www.ncbi.nlm.nih.gov/ pmc/articles/PMC3495961/.
22 Harald Edquista and Magnus Henrekson, "Do R&D and ICT affect total factor productivity growth differently?" *Telecommunications Policy*, 2017, http://www. sciencedirect.com/science/article/pii/S0308596116302531
23 *The Global Information Technology Report*, 2016, http://www3.weforum.org/docs/ GITR2016/ WEF_GITR_Full_Report.pdf.

© The Economist Intelligence Unit Limited 2017

AI's impact will similarly be greatest in the long term. "In the oil and gas industry," he says, "most people believe that with data mining you will have a quick win in terms of return on investment. They are missing that data analytics is about enabling innovation that will give you long term returns." For Dr Parris, this is a familiar pattern. Pointing to experience in healthcare and manufacturing, he notes that "All companies start investing in AI from the point of wanting to reduce inefficiency. In every case, they think there is money on the table. After that comes the transformation stage."

## The potential benefits and risks of open information

Jin Hyung Kim, emeritus professor at the Korea Advanced Institute of Science and Technology, states a basic truth: "The quality of an AI's decision is determined by the quantity and quality of data." Fortunately, on the quantity side at least, as Jerry Kaplan, a futurist and fellow at Stanford University's Centre for Legal Informatics, puts it, "We are swimming in the stuff."

Estimates are inevitably rough, but a frequently cited study by IDC, a market research consultancy, calculated that the world's data is currently doubling every two years.[24] Indeed, Professor Veloso feels that the biggest social challenge around AI is the need for "a change of mentality so people understand that they are producing more data than ever".

Using this data represents a huge opportunity. A 2015 study for the European Commission forecast that opening up just public sector data in EU and EFTA countries would spark direct economic activity of €325bn between 2016 and 2020.[25] This is a fraction of the potential according to McKinsey. In 2013, it estimated that in just seven economic sectors—including several in this study—open public and private data could create value of US$3trn-5trn per year worldwide.[26]

Nor are gains likely to be purely economic. The European Commission study, for example, predicted that open data would reduce traffic fatalities in Europe by 5.5%. Globally, notes Professor Chun, about a million people die from road accidents per year. He expects that by 2030, through the deployment of machine learning, this number might have reached near zero. As discussed in the healthcare qualitative scenario, the possible benefits in terms of lifespan and quality of life are also huge.

Data may even change what it means to be a consumer. Professor Agrawal suggests that people may be able to opt into new kinds of relationships with businesses. "Imagine a company like Amazon that eventually gathers enough data about you that they could really predict what you want with great accuracy. You wouldn't need to go to their website to order things anymore. Instead, they would ship things to you, and you would just return what you don't want." This would not only be convenient for the customers; it would also significantly increase the company's sales by pre-empting customers from purchasing those goods from competing retailers. "In other words," says

24 *The Digital Universe of Opportunities: Rich Data and the Increasing Value of the Internet of Things*, April 2014, https://www.emc.com/leadership/digital-universe/2014iview/ index.htm.
25 *Creating Value through Open Data*, November 2015, https://www. europeandataportal.eu/sites/default/files/edp_creating_value_through_open_data_0.pdf
26 McKinsey Global Institute, *Open data: Unlocking innovation and performance with liquid information*, 2013, http://www.mckinsey.com/business-functions/digital-mckinsey/ our-insights/open-data-unlocking-innovation-and-performance-with-liquid-information


© The Economist Intelligence Unit Limited 2017

Professor Agrawal, "increased prediction accuracy could make the economics of 'predictive shopping' more compelling than traditional shopping, even accounting for the costs associated with more returned items."

At least three bumps exist on the road to this future driven by open data. The first is that not all data is valuable, nor cost-free. Dr Parris makes the point with reference to the GE90 jet engines—each requires 14 sensors. They generate huge amounts of time series data but, he adds, "a lot of it is normal and boring". The truly valuable data comes when discrepancies from the norm occur, but these are already rare. In 2015, for example, the GE90 engine was on around a million flights but there were only 12 significant events. "That's not Big Data," he adds, noting that it needs to be combined with information from other sources. Meanwhile, although declining, the costs of storing and managing vast volumes of data can be huge. Data for data's sake does not necessarily bring value. Indeed, says Dr Parris, "we are still refining out data inputs to cull all the specific data we would like." Quality matters.

The second issue is that such quality is not guaranteed. For example, Dr Amy Abernethy, chief medical officer and chief scientific officer at Flatiron Health, a healthcare technology company focusing on accelerating cancer research and improving care, notes that "the biggest barrier [to greater use of machine learning in medicine] is still credible data. Data need to be of sufficient quality that the outcomes of learning are correct." She adds that for many AI applications in medicine, the training data also needs to be labelled accurately. But getting the data required to train AI algorithms is a problem because in medicine much of the information remains in unstructured documents. This is not a permanent problem. She expects that "between now and 2030, data repositories will get better and better, but we are not there yet".

The third issue is that, as with any valuable resource, questions of ownership arise. Unlike other goods, using up information does not necessarily deplete it, but it can have highly negative effects for the owners.

At the individual level, the big question is privacy. The issue goes far beyond the inadvertent release of information which individuals have entrusted to others—a problem on some level even before computers appeared. With AI, the particular worry is what hidden secrets can now be inferred by analysing data which may have been volunteered in other contexts or even passively sensed by Internet of Things devices from the public actions of individuals.

Perhaps the best known exemplar of the issue dates back to a *New York Times* article in 2012. It described a case of the US retailer Target inadvertently revealing the pregnancy of a teenager to her none-too-pleased father by sending maternity-related sales coupons based on the young woman's spending patterns. Target challenged the report's accuracy and others have questioned whether the company's marketing would have been so ham-fisted.[27] Nobody, however,

27 "How Companies Learn Your Secrets," *New York Times*, 16 February 2012, http://www.nytimes.com/2012/02/19/magazine/shopping-habits.html; "Did Target Really Predict a Teen's Pregnancy? The Inside Story," *Kdnuggets*, 7 May 2014, http://www. kdnuggets.com/2014/05/target-predict-teen-pregnancy-inside-story.html.

© The Economist Intelligence Unit Limited 2017

## BOX OUT: CASE STUDY #2

**"Trust Me, I'm a Technology": The need for transparency with AI**

Professor Chun is not alone in believing that "the main barrier to wider adoption of machine learning in general is simply trust. Humans need to trust the technology before they will widely adopt it." Professor Kim adds that currently a level of confidence in AI systems exists, but unfortunately it might not last because it is too often based on a misunderstanding: "Most of the public has been misled by science fiction and believe AI always makes perfect decisions, but this is not true. AI makes mistakes."

The potential sources for these errors are legion. One is flawed or incomplete information to begin with, notes Professor Kim. "A black box making important decisions based on a small and biased data set is very dangerous to our society," he adds. Obviously problematic information is one thing, but a less tractable problem is the ability of AI-enabled machines to absorb widespread biases. A recent article in Science described an AI engaged in learning from the pages across the Internet as part of its training for "word embedding"—a technology that helps with natural language processing. The software came to associate negative words with African Americans and positive ones with those of European descent. This was just one of several problematic links which arose, not from a small sample, but from the technology learning apparently widespread biases.[28] Another fundamental issue is that the algorithms themselves may be flawed, either due to errors in the code or the unrecognised biases of those who created them.[29]

The problem is that, notes Professor Agrawal, "at present, many AI applications are unable to explain why they make the decisions they do". This can lead to frustration even when the technology is apparently getting things right. Deep Patient, software at New York's Mount Sinai hospital, trained itself on 700,000 patient records and is now adept at predicting the onset of schizophrenia—a condition human doctors find very difficult to foresee. Unfortunately, Deep Patient cannot explain what it looks for, so that human clinicians, and their patients, could benefit.[30]

In the context of AI, though, a still bigger need for communication will be between the technology and user. Professor Veloso believes that AI researchers "have focussed too much the on generating solutions to problems rather than how to explain the way solutions are generated in order to enable much better user interaction." Professor Agrawal has similar concerns: "Without the ability to understand why AI is making

28 Aylin Caliskan, "Semantics derived automatically from language corpora contain human-like biases," *Science*, 14 April 2017, http://science.sciencemag.org/content/356/6334/ 183.full
29 Cathy O'Neil, *Weapons of Math Destruction: How Big Data Increases Inequality and Threatens Democracy*, 2016.
30 Will Knight, "The Dark Secret at the Heart of AI," *MIT Technology Review*, 11 April 2017, https://www.technologyreview.com/s/604087/the-dark-secret-at-the-heart-of-ai/


© The Economist Intelligence Unit Limited 2017

decisions, depending on the application it may be difficult to deploy." Indeed, in certain fields, such as medicine, this would run up against millennia of best practice. "As clinicians we are taught to make sure we can understand how a decision was made," says Dr Abernethy, "before we know whether to apply it to the patient in front of us. The first rule of medicine is 'do no harm', but using algorithms that we can't understand can lead to harm."

The signs of growing unease are multiplying, especially as AI algorithms are having a more pronounced impact on decisions of great consequence to humans. The flavour of the debate is captured in the recent headline of an article in the *MIT Technology Review*—no Luddite journal—which reads, "The Dark Secret at the Heart of AI."[31]

Judges and regulators are also taking note. The US Supreme Court is considering whether to hear the appeal of Eric Loomis against a prison sentence for his conviction, following a guilty plea, for eluding police and driving a vehicle without the owner's consent. The court, in assessing the appropriate penalty, relied in part on a AI analysis which projected that Loomis was likely to re-offend. The appeal, however, argues that he was denied due process because he and his lawyers have not been allowed to see the algorithm behind the analysis, which is a private company's intellectual property.[32] Across the Atlantic, meanwhile, the EU's General Data Protection Regulation, which comes into force in 2018, is trumpeted as giving a "right to explanation" of all decisions made by automated or AI systems. A closer analysis, however, suggests that imprecise language within the regulation may limit the ostensible right in practice.[33]

Whatever the precise details, governments are increasingly demanding transparency. This will inevitably be the case in regulated areas like transport and health. Dr Abernethy explains that "if the Food and Drug Administration can't understand how you arrived at a certain conclusion, they can't audit it to be sure it is correct." Without that, drug approval will be impossible. Professor Agrawal predicts "we will soon start classifying applications of AI as ones requiring or not requiring transparency in order to be able to use the decision."

Professor Veloso adds that more than regulatory compliance is at stake. "People will not want to interact with systems that are all black boxes." Indeed, she adds, without transparency about how machines reach their decisions, a widespread and lasting failure of trust in the technology will occur. "This is the biggest challenge" facing AI currently, she concludes.

31 11 April 2017, https://www.technologyreview.com/s/604087/the-dark-secret-at-the-heart-of-ai/

32 "Sent to Prison by a Software Program's Secret Algorithm," *New York Times*, 1 May 2017, https://www.nytimes.com/2017/05/01/us/politics/sent-to-prison-by-a-software-programs-secret-algorithms.html?_r=0

33 Sandra Wachter et al., "Why a right to explanation of automated decision-making does not exist in the General Data Protection Regulation," *International Data Privacy Law*, forthcoming, https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2903469

© The Economist Intelligence Unit Limited 2017

argued that finding out about an unannounced pregnancy was beyond AI. Target had definitely been at least considering a programme to identify pregnant customers for marketing. Now Castlight, a healthcare analytics company, can predict accurately from health insurance data who is likely to be planning a pregnancy, about to become diabetic, or will need back surgery. It will currently reveal such information to the companies it works for only in aggregate for groups—the minimum size is 40 people. It does this to protect individual privacy, but this is a self-imposed ethical restraint, not a technological one.[34]

Pregnancy is but one example. As a 2014 US government report argued, "The challenges to privacy [from AI] arise because technologies collect so much data (e.g., from sensors in everything from phones to parking lots) and analyze them so efficiently... that it is possible to learn far more than most people had anticipated or can anticipate."[35]

Anonymising data is an obvious step, notes Professor Kim, who expects that the technology for this will advance rapidly. Nevertheless, it may not always be possible. As Professor Agrawal points, "some AIs are able to function perfectly well and aggregate data, but in areas like health, that usually require individual-specific recommendations, individual level data is necessary. "How we address privacy will be "a first order issue, "Professor Agrawal concludes.

Much of the response will inevitably be regulatory, issues around which are discussed in the following section. Some of it will also involve personal choices. How we think about privacy also may need to change. Rather than making rigid choices says Dr Parris, "we have to have the sophistication to understand the benefits of privacy and those arising from sharing information, and to trade them off." In practice, people are already doing this. Many are willing to share substantial data about themselves when they perceive a personal, or even societal benefit. Nevertheless, widespread unease also exists about how data on them can be used in aggregate, which in certain circumstances will undermine their willingness to share information[36] Healthcare is an obvious area where consumers might resist greater deployment of AI should there be a scandal surrounding data misuse, but information about individual power usage and travel patterns could also raise important privacy concerns.

An issue analogous to privacy arises for companies if they see information as a source of competitive advantage. Dr Priyadarshi explains of the oil and gas sector, for example, "We've seen some changes but, still today, most companies do not like to share their data, even within their own organisations. It will take a mind-set change in the leadership." The sector is not unusual. In Dr Parris' experience, "initially nobody wants to share their data," because they assume it has value but do not know what it is. He recalls that when he worked as a consultant with financial sector companies, the largest banks would not let his firm access their data. Tier II institutions were in a different situation—"their storage costs were huge" and they did not have the computing capacity for extensive Big Data analysis anyway. To get some value from their data, these

34 "Employers Are Quietly Using Big Data to Track Employee Pregnancies," Forbes, 17 February 2016, http://fortune.com/2016/02/17/castlight-pregnancy-data/.
35 President's Council of Advisors on Science and Technology, *Report to the President: Big Data and Privacy: A Technological Perspective*, 2014, https://bigdatawg.nist. gov/pdf/pcast_big_data_and_privacy_-_may_2014.pdf
36 Lee Rainie and Maeve Duggan, "Privacy and Information Sharing," *Pew Research Center Report,* December 2015, http://www.pewinternet.org/2016/01/14/2016/Privacy-and-Information-Sharing/


© The Economist Intelligence Unit Limited 2017

institutions were willing to allow their information to be aggregated with certain safeguards. As soon as Dr Parris' company revealed that sharing data had economic value, by doing things like spotting which customers were engaged in fraud, the larger banks were willing to do so as well. "That is a maturity curve which industries, including currently manufacturing, have to go through—to understand that this data is costing me money and I can't do anything with it," until they are willing to share it.

If businesses do not become more comfortable sharing data, though, or if the widespread but low-level mistrust of doing so among individuals increases, AI will yield much less of value. As Dr Priyadarshy puts it, "As soon as you ask 'who owns the data?' you can't really do machine learning." In other words, if people are arguing over proprietary control of information, they will not have the necessary mentality to create sufficiently large data pools to truly benefit.

## Key assumptions

Our second scenario looks at the result of positive developments in both the fields of computer capital investment and data openness. It assumes that greater computer capital investment than under the baseline forecast will be the impetus for the broader proliferation of machine learning. It also assumes that human skills growth will keep pace, so that there will be no negative impact on employment.

Driving this accelerated capital investment will be a repeat of the developments accompanied the widespread adoption of ICT across economies: advances in efficiency will push down hardware and software costs, encouraging more investment in machine learning and permitting a wider range of organisations to benefit.

This does not necessarily mean steady, almost linear change: no Moore's Law is likely to describe advances in AI. It will be much more challenging. Mr Kaplan describes the improvements as "continuing but not continuous." Professor Agrawal agrees: "Scientific breakthroughs in this area are not happening in a predictable or linear way. Improvements in performance may be gradual and then see a step function jump." Professor Kim believes we have reached a stage when there will be much more of the latter. "AI technology is developing exponentially," he says. "Each innovation begets another. Deep learning techniques are progressing at the speed of light." The issue for society will be keeping up.

Under this scenario, even as the technology improves, government policies will create an encouraging environment for greater use of AI in two ways. First, new tax credits for machine learning will further drive private capital investment. Governments will also take regulatory steps and spend more to allow greater access to open source data and thereby promote national knowledge communities. They will also take steps to permit sharing of data between countries.

Governments will not, however, be active players in advancing the technology of machine learning: the scenario is not suggesting parallels to Japan's

© The Economist Intelligence Unit Limited 2017

very large state-driven, ultimately unsuccessful Fifth Generation Computer Project of the 1980s. Instead it is about creating an enabling environment where others innovate.

This is consistent with where several interviewees for this study believe the state can contribute most. For example, Mr Kaplan does not "think there is much need for governments to get involved in the technology side." They can, though, facilitate by creating conditions conducive to adoption. For example, he adds that "they can establish rules and regulations that govern how devices, such as self-driving cars, are used. But there is nothing new about that. The same thing happened when the car was introduced and had to compete for resources and space with horses."

It is not the only possible future, however. Professor Kim believes that public R&D, government public health research, and "the innovative application of AI technology for solving social problems" are all very important. As discussed in later sections, China's state-owned enterprises are active in developing smart grids to reshape the power system. Nevertheless, even for those states which focus largely on the regulatory side, much remains to be done, an issue dealt with in greater detail in Scenario 3.

This scenario, though, looks at what would happen in an environment which encourages increased investment in the technology needed for machine learning and where the data needed to take advantage of it is readily available.

This scenario yields the most encouraging results, particularly with regards to economic growth as measured by GDP and against our baseline forecasts for the period covered. For the most part, our current outlook for GDP growth is for it to be subdued in the developed world and to slow from its once great heights in the developing world. Greater investment in computer efficiency (our proxy for AI), along with tax credits and more access to open source data, could provide a substantial boost to economic growth across the five developed countries and Developing Asia as a whole.

Here again Australia is the greatest beneficiary. Our baseline forecast is for Australia's CAGR in GDP to be slightly more than 1% out to 2030. Under this scenario, it rises to 3.74%. Developing Asia is not far behind, however. Its improvement over the baseline is more than 2%, equivalent to a cumulative difference of hundreds of billions of dollars for these countries on an individual basis. As discussed above, this result is the product of the low levels of current and forecast investment in computer capital. In other words, going beyond the baseline is very much in these countries interests.

Elsewhere, the benefits are less marked but nonetheless notable. The UK, South Korea and the US all see their CAGR in GDP rise by more than 1% versus the baseline, putting them all in the range of 2-3% of annual growth that they might not have seen otherwise, the pronouncements of their current leaders notwithstanding. The US and South Korea are at the top of the range, at exactly 3% GDP growth provided they meet this scenario's key assumptions. Japan lags, albeit just, at 2.4%,


© The Economist Intelligence Unit Limited 2017

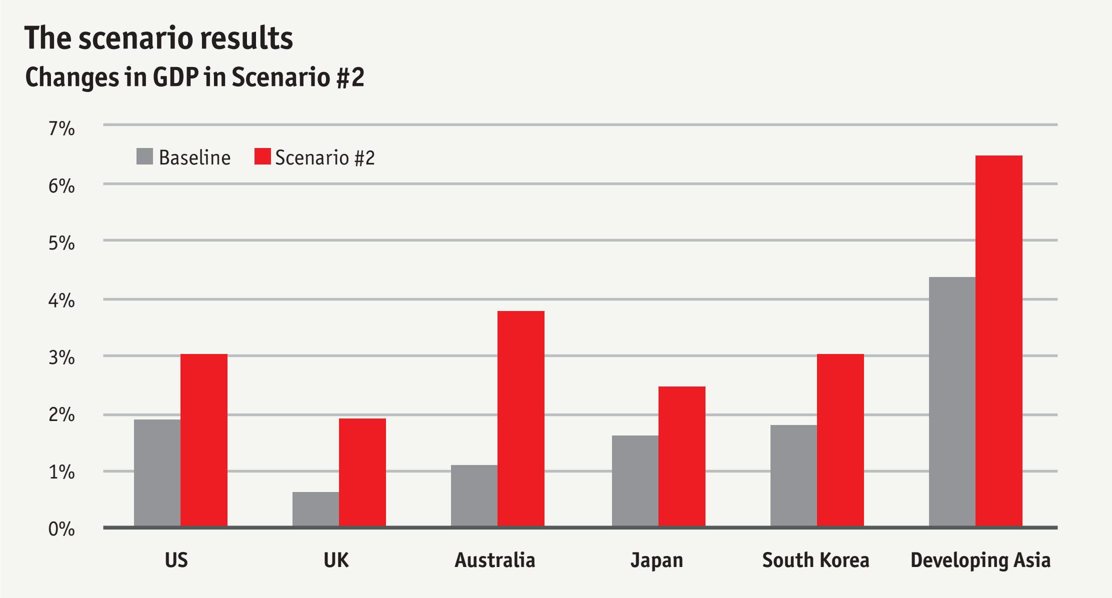


up from the sub-2% growth projected in the baseline. Japan's underperformance relative to South Korea, a country with a similar economic and demographic profile, may be due to its size in terms of population.

These benefits are not just dependent on investment, however. Investment will play a major role, to be sure, but so too will the availability of data, which is a much thornier issue. As our colleagues at The Economist put it in a recent article, the world's most valuable resource is no longer oil, but data.[37] And just as restricting imports and exports of oil affected—and continues to affect—economies throughout the world, restricting data is already having a similar impact and that will only become more pronounced in the years ahead. This scenario assumes, to a certain extent, that the countries covered can get the politics of data right, ensuring an acceptable balance between the privacy concerns of their respective citizens and the needs of the private sector. That outcome appears far from assured at present, however, leaving considerable work to be done.

## Data and capital investment in key sectors

Questions of capital investment and data availability take on particular characteristics in the different sectors which this study covers.

**Manufacturing**

Simple competition will drive many sector firms to invest in AI. Professor Kim explains of Korea, "we are a manufacturing country but traditional industry is not as profitable as before. We need innovation in machine learning in manufacturing to create new value."

37 https://www.economist.com/news/leaders/21721656-data-economy-demands-new-approach-antitrust-rules-worlds-most-valuable-resource

© The Economist Intelligence Unit Limited 2017

The qualitative scenario for this sector draws a distinction between developed and developing world: in the former machine learning turns manufacturing into an employer of highly-skilled, well-paid individuals; in the latter there is little change to the low-skill, low-paid status of factory workers.

The reason is entirely cost-related. The massive growth in manufacturing in the developing world in recent decades was based largely on substantial wage differentials between employees in wealthier and poorer countries. Now, as Professor Agrawal puts it, "When a robot maker is selling an intelligent machine that can potentially replace a human, the first question which the robot company is asked is 'what is the payback period?'"

While AI reduces the cost of performing certain functions currently done largely by people, it brings its own expenses. For a manufacturer, the price of a basic, single robot arm is between US$30,000 to US$60,000. That is before any advanced software or any number of essential physical add-ons such as specialised tools, some of which have to be designed and produced for the specific role it will play in a factory. Equipping a large facility can easily grow in cost to millions or tens of millions of dollars.[38] As noted earlier, advanced sensors and even places to store and analyse resultant data bring added expense. In developing countries, even ensuring access to the steady power supply needed for the equipment may bring additional costs. Finally, new employees able to use the technology, while fewer in number, may require skills not available in certain countries. Where such talent is present, it will have to be rewarded with higher wages than those of the workers whose jobs were automated.

In many places where employment costs are high or quality an issue, the balance is already in favour of a shift toward greater use of automation over unskilled labour. Although, as the qualitative scenario indicates, this is more prevalent in wealthy countries, it would be wrong to say it is unique to them. For example, Foxconn, a Taiwanese contract manufacturer of electronics such as the iPhone, reports that in its Chinese factories it plans to automate 30% of jobs by 2020 in response to rising labour costs. In just one factory, it has already cut 60,000 jobs.[39] It is not alone: in mid-2016 the International Federation of Robotics, an industry lobby, predicted that by the end of that year China would become the world's biggest user of industrial robots.[40] Presumably the cost of AI will decline further, making the use of technology more attractive.

On the other hand, with the average minimum wage in Bangladeshi garment factories at 5,300 taka (US$67) per month,[41] it is hard imagine the textile, or similar, industries switching rapidly to AI. Moreover, if automation in certain industries in developing economies leads to job losses, average wages are likely to decline, making the investment needed for machine learning less attractive for other companies.[42]

In short, in manufacturing there will be a marked tension between labour costs and willingness to invest in machine learning. Although more

38 Dmitry Slepov, "The real cost of robotics," *TechCrunch*, 27 March 2016, https://techcrunch.com/2016/03/27/the-real-cost-of-robotics/
39 "Foxconn's Foxbot army close to hitting the Chinese market, on track to meet 30 per cent automation target," *South China Morning Post*, 1 July 2015, http://www.scmp.com/ tech/innovation/article/1829834/foxconns-foxbot-army-close-hitting-chinese-market-track-meet-30-cent.
40 "China's robot revolution," *Financial Times*, 6 June 2016, https://www.ft.com/ content/1dbd8c60-0cc6-11e6-ad80-67655613c2d6
41 "Bangladesh garment factories sack hundreds after pay protests," *Guardian*, 27 December 2016, https://www.theguardian.com/world/2016/dec/27/bangladesh-garment-factories-sack-hundreds-after-pay-protests
42 OECD, *Automation and Independent Work in a Digital Economy*, 2016, https://www.oecd.org/employment/Automation-and-independent-work-in-a-digital-economy-2016.pdf


© The Economist Intelligence Unit Limited 2017

pronounced in developing countries, it is not unique to them. Professor Agrawal notes that, while he sympathises with the motivations of those wishing to raise the minimum wage, "no single policy would be more effective at accelerating the adoption of AI and robots than to increase it".

**Transportation**

In our qualitative scenario for this sector, investment increases but much of this money ultimately comes from governments. The reason is simple: in most of the world transportation relies on public infrastructure, whether roads, railways or ports.

Such investment will almost certainly occur, and not just because AI-enhanced infrastructure is typically greener, safer and more efficient. It can also be less expensive than traditional maintenance, allowing cost conscious governments to reduce their spending by getting the most out of their existing infrastructure. For example, one of the drivers of the UK's shift to so-called Smart Motorways, with sensors providing data to software in order to set dynamic speed limits and permit use of the hard shoulder as an extra lane in certain cases, is that the technology can be retrofitted at about one fifth the price and time needed to plan and create an extra physical lane.[43]

Yet improving individual modes of transport has limited benefits for some cases, particularly in urban areas. For example, in Hong Kong, notes Professor Chun, "the railway lines are already at full capacity during rush hours. The only way AI and machine learning can help shape public rail transportation in the future is to optimize and dynamically balance passenger flow across multiple modes of transportation, such as combining rail with buses or ferries."

Cities are accordingly putting money into technology that can draw on data now available from disparate sources in order to transform transportation networks into a single integrated system for getting people where they want to go as quickly as possible. The western part of the city of Poznan, Poland, to cite one of many examples worldwide, has installed over 200 measurement points that cover current car traffic, parking and public transportation conditions. These are evaluated and the system provides information to travellers on optimal routes to their destination via the Internet and roadside message signs. The data are also used to adjust traffic signals in real time in order to improve flow, especially for public transport. The scheme has reduced traffic volumes and improved its efficiency. Air pollution is also down.[44]

Richer data, meanwhile, is already allowing the first stages of a dramatic improvement in systems safety and reliability through the exploitation of situational and operational intelligence. Situational intelligence is the ability of organisations to monitor where their equipment is, how it is performing over time, and the impact of the interaction of different systems operating within a transport fleet. Advances in data collection and AI now allow early identification of problems and reductions in likely system failures. This can be as simple as knowing exactly where two vehicles are so they do not crash, as with the truck

43 "Hard shoulder scheme cut journey times on motorways," *The Times*, 25 October 2007.
44 United Nations Commission on Science and Technology for Development, "Issues Paper On Smart Cities and Infrastructure," January 2016, http://unctad.org/meetings/en/SessionalDocuments/CSTD_2015_Issuespaper_Theme1_SmartCitiesandInfra_en.pdf

© The Economist Intelligence Unit Limited 2017

## BOX OUT: CASE STUDY #4

**AI proves itself over the long term for Hong Kong's Mass Transit Railway**

Hong Kong's Mass Transit Railway (MTR) runs one of the world's busiest subway systems, with volumes similar to those of New York's and London's. The MTR is also highly efficient, with a 99.9% on-time rate.

A key part of keeping up this level of performance is ongoing maintenance of railway infrastructure and equipment. Scheduling who should be doing what on a such a complex system is no trivial matter. In Hong Kong during a typical week 10,000 engineers are involved and the main window for the majority of activities—the period when no trains are running—lasts only five hours per night.

The scheduling used to be done by human experts. Professor Andy Chun of City University Hong Kong is a consultant with the MTR on the use of AI. He recalls that before its application to this task, "a group of experts would sit down in a long planning session to work out the schedule through discussions, negotiations and mental calculations. As one can imagine, this was a chaotic process and highly error-prone." Mistakes were particularly worrisome, and not just for financial reasons: at an extreme, scheduling errors could endanger engineers' lives.

Now, though, "humans have totally relinquished control to AI, although the final schedule still needs endorsement by humans due to policy and regulations," says Professor Chun. The software sets the schedule, he adds, drawing on what it has learned from "the different fields of railway expertise and existing experts' wealth of knowledge acquired from years of experience". Its model assesses both the urgency of any given maintenance activity as well as the way to maximise use of limited manpower and equipment.

The improvements have been huge. "MTR can carry out more engineering works with fewer resources," reports Professor Chun. Many of MTR's experts now have up to two days a week that were once spent in meetings to devote to other tasks and engineers themselves have an extra half hour—or 10% of their main window of opportunity—to engage in maintenance rather than paperwork. The savings to the company, he estimates, are around US$1m per year.

Other uses of AI and machine learning can boast impressive benefits. What really sets the MTR story apart is how long the


© The Economist Intelligence Unit Limited 2017

organisation has deployed it: the scheduling software entered use in 2004. Since then, the technology has proven itself in two important ways.

The first is reliability. Says Professor Chun, “MTR has been using AI for over a decade now, without a single planning or scheduling error.”

The second is adaptability as the needs of the organisation have changed. In 2007, for example, MTR took over management of Hong Kong’s overland railway system. This brought more complex maintenance challenges than in the past. The MTR uses light rail, for example, but the overland lines use heavy engines. Similarly, engineers working above ground have certain obstacles, such as overland power lines, which are not of concern in MTR tunnels. Moreover, the wider management responsibility brought a dramatic change in scale. Combined with the building of new lines, the size of the system under MTR oversight, as measured by track length, more than doubled between the introduction of AI maintenance scheduling and 2010.

In that year, the company therefore decided to adapt the existing software by expanding its knowledge base and separating it from the AI engine. The latter meanwhile was reconfigured as a scalable AI cloud service.

These technical changes did more than allow the software to oversee Hong Kong’s more complex rail maintenance. MTR’s goal is to become one of the world’s leading rail operators and it is managing an increasing number of rail systems in other cities. It sees the possibility of using the same AI, with a different data set, to schedule maintenance in this growing list of foreign operations.

AI, then, can do more than drive efficiencies in operation intelligence. Over the long term, MTR has shown that it is highly reliable, adaptable, and a potentially important source of competitive advantage.

© The Economist Intelligence Unit Limited 2017

platooning discussed in Scenario 1. Applications can, however, get much more complex: NEC, a Japanese IT company, is working with Singapore's public transport authority to use a range of data, including from bus telematics and human observation, to project whether a bus driver in the city will probably cause an accident in the next three months.[45]

Operational intelligence involves the uses of AI algorithms to improve the process of fleet maintenance, including allowing for preventative maintenance to reduce or eliminate breakdowns

**Energy**

As with urban transportation, the main impact of AI on the power sector is expected to be the transformation of generation, transmission, distribution and use into a single, coherent system.

Since the advent of electricity, centralised grids drawing on large coal, oil, or gas-fired power plants have supplied power. While reliable, these systems are inflexible. Dr Scott Otterson, who works on probabilistic power forecasts at the Frauenhofer Institute for Wind Energy and Energy System Technology, explains that "conventional power generation resources have to be spooled up. Coal and nuclear plants, for example, have a start up and shut down time. You don't want to do that unexpectedly or it gets very expensive." Reserve power plants, brought on line for periods of high demand, are also typically more expensive to operate than those that run constantly—up to ten times costlier per unit output according to a 2004 US government report.[46]

To integrate large amounts of renewable energy, though, says Dr Otterson, power system flexibility is needed. "The system must be ready to rapidly reconfigure itself in response to suddenly missing power plants or to changes in wind or solar production. Smart grid techniques are a way to optimally acquire and control that flexibility," he adds.

These grids will be able to monitor in real time the energy requirements of all customers and draw on the most appropriate range of sources to meet those needs, as well as have dynamic pricing in order to even out demand. For this to function, though, says Dr Otterson, "AI is the alternative. You have millions of solar panels, hundreds of thousands of turbines, and then consumer demand to factor in. AI can pay attention to those huge amounts of data and make models based on them."

A particular attraction of the grids is the promise they hold for integrating large and small sources of renewable energy. A drawback of such types of generation is that their output is weather-dependent and therefore intermittent. Renewable producers also vary in size, from huge offshore wind farms to solar cells on a small house's roof. Adding to the complications, some smaller producers may invest in substantial battery storage which could, in aggregate, also be a source of power for the grid.

Smart grids will be able to price and purchase power from this entire range of sources based on the likely availability given factors such as existing and forecast weather conditions, as well

45 "NEC using artificial intelligence to prevent bus accidents in Singapore," *ZDNet*, 20 March 2017, http://www.zdnet.com/article/nec-using-artificial-intelligence-to-prevent-bus-accidents-in-singapore/

46 United States Government Accountability Office, *Electricity Markets: Consumers Could Benefit from Demand Programs, but Challenges Remain*, August 2004, http://www.gao.gov/new.items/d04844.pdf


© The Economist Intelligence Unit Limited 2017

as current load or any predictable, near-term fluctuations. Where supply—particularly from renewable sources—and demand are not meeting, using the grid companies will be able to do one or a combination of the following: raising prices to dampen demand; using other types of generation, such as gas powered electric plants; or, eventually, even buying energy stored in batteries. Thus, the use of green energy can be maximised and rather than a single large power generator providing a relatively fixed amount to an area, a web of large and small producers and consumers can be kept in balance.

This is the vision, at any rate. The reality is still being worked out. One basic building block, the smart meter, is starting to appear. According to Navigant Research, a clean tech consultancy, China had 348m smart meters installed by late 2016, roughly two-thirds of the world's total.[47] The EU, meanwhile, has since 2009 had a goal of 80% of customers and businesses using smart meters by 2020. Although most countries have lagged in implementation, a few have moved ahead, notably Sweden and Italy, where they are commonplace.[48] Meanwhile, the US states of California and New York are implementing legislation to make energy efficiency a tradable commodity, which opens up potential for fast deployment of such meters. For example, New York's electricity company, ConEdison, plans to install 4.7m of them which would collect an estimated 1.5bn points of near real-time data daily.[49] Overall, estimates of the market for smart grid technology worldwide vary widely, between anywhere from around US$20bn to US$40bn per year.[50]

The next stage will be finding ways to use this data to create smart grids. This is in the nascent stages. A Swiss company Alpiq, for example, is now employing an AI system called GridSense to understand user behaviour for optimal use of energy.[51] On the power production side, Frauenhofer IWES, an applied research organisation in Germany, has been developing open source software to allow consumers to take full advantage of the potential of the smart grid. It has also published research showing that, in theory, it is possible to balance intermittent renewable energy sources.

There will be bumps along the way, however. Early assessments of the roll out of smart meters have shown less benefit to consumers than initially expected.[52] Ultimately, though, Dr Otterson believes that there would be no technological problem in having smart grids by 2030, and the economic issues are not likely to be barriers. The real question is "political", he says. "Smart grids can reduce overall costs and environmental impacts but a reorganisation needs to be done," in order think about producing and using energy in ways that make these societal benefits profitable for smart grid providers.

47 "China Continued to Lead the Global Smart Electric Meter Market through 3Q 2016 with More than 348 Million Smart Meter Installations," press release, 20 December 2016, https://www.navigantresearch.com/newsroom/china-continued-to-lead-the-global-smart-electric-meter-market-through-3q-2016-with-more-than-348-million-smart-meter-installations.
48 European Parliament, "Smart electricity grids and meters in the EU Member States," Briefing, September 2015, http://www.europarl.europa.eu/RegData/etudes/BRIE/2015/ 568318/EPRS_BRI%282015%29568318_EN.pdf; CBI Market Factsheet, "Electronic
Components for Smart Meters in Sweden, Denmark, Italy, UK, France and Spain," https://www.cbi.eu/sites/default/files/market_information/researches/product-factsheet-smart-meters-italy-uk-denmark-sweden-spain-france-electronics-electrical-engineering-2013.pdf.
49 International Energy Agency, *Energy Efficiency Market Report 2016*, 2016, https://www.iea.org/eemr16/files/medium-term-energy-efficiency-2016_WEB.PDF.
50 Markets and Markets, "Smart Grid Market worth 65.42 Billion USD by 2021," http://www.marketsandmarkets.com/PressReleases/global-smart-grid.asp; Zion Research, *Smart Grid... Market: Global Industry Perspective, Comprehensive Analysis, Size, Share, Growth, Segment, Trends and Forecast*, 2014 – 2020, 2015.
51 "The Role Of Artificial Intelligence In Smart Utilities," *energycentral*, 9 December 2016, http://www.energycentral.com/c/iu/role-artificial-intelligence-smart-utilities.
52 European Parliament, "Smart electricity grids and meters in the EU Member States," Briefing, September 2015, http://www.europarl.europa.eu/RegData/etudes/BRIE/2015/ 568318/EPRS_BRI%282015%29568318_EN.pdf.

© The Economist Intelligence Unit Limited 2017

# SCENARIO 3: A POLICY FUMBLE

## Employment: Is it really different this time?

As discussed in an earlier section, labour costs will affect investment in AI. This is just a part of the much wider issue of how AI and machine learning will affect economies and societies.

Currently, only humans can perform certain tasks. AI will enable machines to do some of those. This creates what economists call a "substitute", because anyone wishing to have one of these jobs done can choose between AI and human labour. Such a situation is far from novel. As Professor Chun points out, "Machines and technology have been displacing human workers ever since at least the industrial revolution. The difference is that, instead of blue collar workers, AI will be displacing white collar workers and even experts."

Estimates vary on how many jobs AI will displace, but the numbers are not small. To give just one prominent estimate, in 2016 the OECD projected that 9% of jobs will be completely taken over by technology in the next two decades, with another 25% having around half their tasks automated.[53] The distinction between total and partial displacement may not mean that much. As Stanford's Mr Kaplan puts it: "Whether you make someone more productive or replace them completely, you can do more work with fewer workers. Both put people out of work, but grow the economy."

Professor Kim puts this part of the problem succinctly: "The best benefit of AI is automation, but full automation is not a good story for humans." Fortunately, the ability of the new technology to displace current human labour is only a part of the story. As Professor Kim continues, "Many jobs will be replaced by AI and automation. However, new jobs will be created." Even within existing companies, this is not merely a matter of new IT personnel. Dr Parris of GE, for example, notes that at his firm, in order to make the most effective use of the technology, not only are data experts in higher demand but so are individuals who can explain to IT staff the specific business problems that need attention. There are yet others who can then take the data and fashion it into viable business solutions.

Nor will it be easy to predict which level of seniority and training will see the most job losses and who might gain. In manufacturing, lower skilled jobs look more at risk. In medicine, on the other hand, Dr Abernethy expects that "AI will make it so we can deliver healthcare with a larger cohort of lesser trained individuals," including a greater need for healthcare providers like nurse

53 OECD, "Automation and Independent Work in a Digital Economy," Policy Brief, May 2016, https://www.oecd.org/employment/Automation-and-independent-work-in-a-digital-economy-2016.pdf.


© The Economist Intelligence Unit Limited 2017

practitioners. Similarly, she foresees greater demand for the very experienced and highly trained "individuals who can deal with edge cases and complex situations and individuals who can act as direction setters for further healthcare research including the application machine learning." Those providers, though, who currently deliver relatively straightforward care, such as emergency room physicians, will probably be less in demand, as will certain specialities where AI is already showing promise, such as radiology and pathology.

Such employment churn would also be consistent with past waves of automation. To cite just one example from over the centuries, the same technology which made piecework production of textiles in England unprofitable at the start of the industrial revolution also created factory jobs, in the end employing more people and raising the average standard of living. That did not mean the process was easy. This same change sparked the Luddite protests when skilled textile artisans opposed the changes in working practices and the use of unskilled labour which new machines allowed.

The debate around AI is what, if anything, will be different from the oft-repeated pattern of technology introduction this time around for labour markets.

Techno-pessimists see big trouble ahead. Cowen, in his Great Stagnation, and a follow-on book of 2013, Average Is Over, forecast a future where a small elite able to work with the new technology will benefit greatly but most people will be worse off financially as a result of AI's widespread deployment. Meanwhile, Martin Ford—a futurist and software developer—has two subtitles for different editions of his 2015 book which speak for themselves (*Rise of the Robots: Technology and the Threat of a Jobless Future; Rise of the Robots: Technology and the Threat of Mass Unemployment*).

Others, though, are far more hopeful. They do not see the future as one of pure substitution of technology for human labour, but one where AI can help people do things that were previously very difficult or impossible, making it what economists call a "complementary good". In this case, more technology will mean more jobs for people finding new ways to use it. Authors such as Tom Davenport and Julia Kirby, for example, in their 2016 *Only Humans Need Apply: Winners and Losers in the Age of Smart Machines*, see the potential for AI-enabled technology to augment what humans can currently do, opening the possibility of substantial, widespread economic benefits. In that sense, they see AI as just another in a string of new technologies. Similarly, Professor Kim notes that AI, because it learns from decisions and situations which have already taken place will not be able to devise anything truly new: "AI might look in some cases like it is creating art or music, but it is just imitating. Creativity is the most valuable characteristic of humans that cannot be copied by machine." In the same vein, Professor Veloso notes that "AI will always have limitations." The future will be one of what she calls "symbiotic

© The Economist Intelligence Unit Limited 2017

autonomy" in which the technology and people constantly interact to go beyond what either can do alone.

Even most optimists, however, do not expect the labour market changes to be easy. Professor Veloso argues, "We can't say 'Let's not worry; this is just a transition problem that will be solved in the long run.' The transition will be too fast for that. But it's an inevitable transition and we need to learn how to jointly best handle it so we can maximise the benefits of this fantastic technology."

Appropriate policy—changes in education systems, in particular—will be necessary. Mr Kaplan has argued that the new technology can lead to a new age of affluence and leisure, but reaching this state will depend on addressing labour volatility and income inequality.[54] Ultimately, he adds, "there will be plenty of jobs around but the main effect of technology will be to change the skills people need to accomplish their jobs." For Professor Kim, this is urgent: "To fill these new jobs, you need education. We are in a race between automation and training our people."

Even the less sanguine see the need for policy responses, but usually more far reaching ones. Mr Ford, for example, is sceptical that enough people can be educated to an extent that they could do the work which machines could not, but hopes that policies of wealth redistribution could mean that AI will usher in a bright future over the longer term.

The one common thread is that policy decisions will shape the contours of the inevitably wide ranging effects of the new technology on society.

## Data: Rethinking the rules

Beyond education and skills to enhance employability, government efforts will be key in at least one other area in realising the potential benefits of AI and machine learning.

When asked where governments can make their biggest contribution, Professor Kim's repose was clear: "Data, data, data. Promoting and providing high quality public data is the single most important policy for AI currently and in the near future." Open government data policies are nothing new. In 2013, for example, the G8 countries committed themselves to making as much government data as possible open.[55] Governments worldwide still have far to go, however. The Global Open Data Index, assembled by Open Knowledge International, an NGO, ranks countries on progress in this area on a range of types of data. None are of the kind where secrecy makes much sense, such as things like water quality, government budgets, election results and weather reports. In only the two best performing countries, Australia and Taiwan, were over half of included databases fully open. In all of the G8 countries, that number was between just 20% and 40%.[56]

In addition to releasing their own data, government regulation of how others collect

54 *Humans Need Not Apply: A Guide to Wealth and Work in the Age of Artificial Intelligence*, 2015.
55 *G8 Open Data Charter and Technical Annex*, 18 June 2013, https://www.gov.uk/ government/publications/open-data-charter/g8-open-data-charter-and-technical-annex
56 https://index.okfn.org/


© The Economist Intelligence Unit Limited 2017

and use information will also be important. This will mean finding ways to protect privacy to an appropriate extent in a world where the value of aggregated data is high not only for individuals and companies but for the common good.

Privacy concerns have sparked data regulation of various forms. The EU has played a leading role in this area and in 2016 had, according to Forresters, a technology consultancy, 17 of the 21 most restrictive national data privacy regimes in the world.[57] Also the new *General Data Protection Regulation*, mentioned earlier, will unify the rules across the EU and deepen the requirements which existed under the previous *Data Protection Directive*. In practice, anyone collecting data in Europe needs to be transparent about how they plan to use it and is forbidden from going beyond these plans.

Other countries take a different approach. In the US, for example, data regulation tends to be based on sector. For example, the *Health Insurance Portability* and *Accountability Act* and subsequent regulations have restricted the sharing of private medical information since 1996, but in some other fields no data protection exists at all.

For those wishing to employ AI and machine learning, such regulation is inevitable but not insurmountable. The bigger problem is aggregating information originating in different countries when governments disagree on how to handle privacy. Indeed, the EU has moved from a privacy directive, which mandates member states to adopt legislation of their own that meets certain goals, to a more robust regulation, which imposes common rules across the EU, precisely in order harmonise the legal situation in all its states.

Meanwhile, the ongoing disagreements between the world's two largest markets, the US and the EU, are the clearest example of difficulties arising from regulatory mismatch. The approach to privacy protection clearly differs between jurisdictions and the *Data Protection Directive* forbids transfer of data to countries outside the EU which do not have the same level of protection as in Europe. In 2000, the US and EU agreed to the International Safe Harbour Privacy Principles. American companies which could demonstrate their compliance with these principles could register to be allowed to receive data originating in the EU. Recently, however, in 2015, the European Court of Justice ruled that the Safe Harbour scheme was insufficient. Accordingly, a new arrangement—the EU-US Privacy Shield—is now in place, but has not yet been tested in the courts. Unless some workable arrangement is in place, however, multinationals may not be able to combine their European and American data.

In shaping policy to promote the potential benefits of AI and machine learning, governments will need to work on common approaches to sharing across international borders as much as on enlightened domestic policy. Otherwise, as Professor Agrawal predicts,

57 Chris Sherman et al., *Forrester's 2016 Interactive Data Privacy Heat Map Landscape*: The Data Security And Privacy Playbook, 2016.

© The Economist Intelligence Unit Limited 2017

"variation in privacy regulation will be a basis for competition across countries."

## Key assumptions

Under this scenario, governments fail to act in two key areas, First, policy makers take no action to develop the labour force beyond 2016[58] skill levels. Second, they do not take an active role in the development of national data sharing schemes. As a result, while computer capital increases in line with our baseline forecasts, skills stagnate. Humans' productivity advantage over machines falls, meaning that the latter can perform an ever increasing proportion of tasks within the economy, an ideal recipe for the substitution effect. Hours worked fall due to involuntary unemployment.

In short, this scenario looks at what happens if governments fumble the ball on skills and data.

The results of Scenario 3, which are negative across the board, relative the baseline forecast, stand in sharp contrast to the results of Scenarios 1 and 2, which were all varying degrees of positive. Two of the five economies actually shrink over the forecast period, while another barely grows at all. In the rest, growth is reduced dramatically.

The UK and Australia have the most to lose, according to our model. In the UK, the CAGR of GDP turns from 0.6% to -1.2% and by the end of the forecast period, the UK economy is US$420bn smaller than it was in 2016 and US$670bn less than our baseline, which

**The scenario results**
**Changes in GDP in Scenario #3**

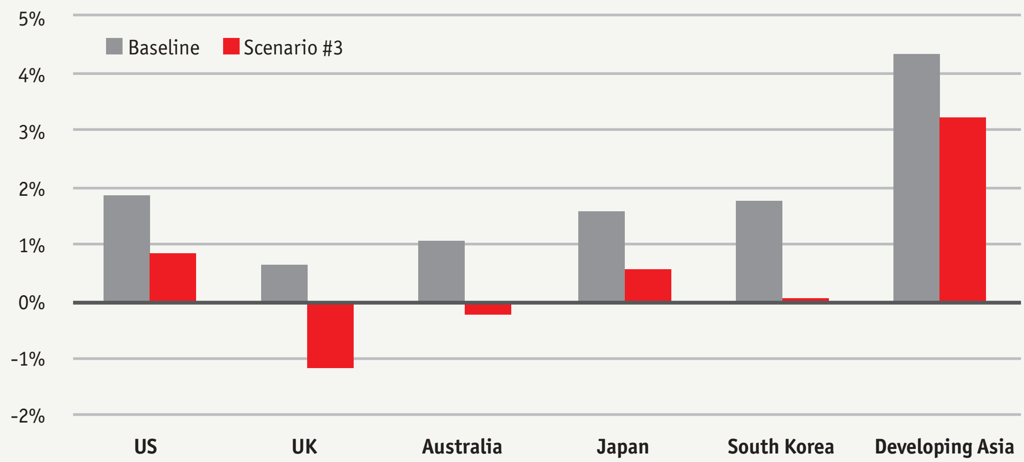


58 2016 being the latest year for which complete data is available.


© The Economist Intelligence Unit Limited 2017

projects modest growth. In Australia GDP will also shrink, albeit at a slower rate of CAGR -0.24%. That still translates to an economy in 2030 that is US$50bn smaller than in 2016 and US$700bn off the baseline forecast.[59] South Korea, which fares only slightly better than the UK and Australia, would see growth fall to 0.02%, with GDP only expanding by US$33m over the next 14 years.

Even in those economies where growth doesn't turn negative or stagnate, the impact is still significant versus the baseline. The US, Japan and Developing Asia all grow around 1% slower in Scenario 3. Given the size of these economies, the impact on overall GDP is dramatic. We currently forecast the US economy to be US$21.9trn in 2030; under Scenario 3, that figure would be US$18.9trn, or US$3trn smaller. In Japan, GDP is US$865bn less than the baseline and in Developing Asia, the lost growth is US$2.7trn.

Together these results should focus minds on just how important it is to get the policies "right" with regards to workforce development and data sharing. While the results of Scenario 1 don't show the benefits of upskilling alone adding a substantial amount of economic growth, Scenario 3 does show that there is great risk in doing nothing. That holds true for data as well. AI runs on and improves because of data. Should governments fail to find to develop and enact policies that assuage privacy and protection concerns while ensuring access to data, the growth potential of AI will be diminished.

## Policy in key sectors

In addition to general policy initiatives, issues of regulation are likely to have specific impacts in each of our study sectors.

### Healthcare

Regulation pervades healthcare, and the advent of AI will do nothing to change that. Indeed, it should not try. As Dr Abernethy says of privacy regulation, "this is fundamental to how we take care of each other."

For the pharmaceutical industry, the most relevant policy consideration for governments—alluded to earlier—will be deciding how to regulate the use and ownership of data. Otherwise, the issues will likely need to be decided by the courts. We have been here before. In 1990, for example, John Moore, a leukaemia patient, sued the University of California for a share of the profits when that institution developed a commercial cell line from cells in Moore's spleen, which had been removed as part of his treatment. In 2006, Washington University had to go to court to prevent a researcher who had relocated to another institution taking a large number of cell samples with him. The researcher had solicited and received written declarations from the patients who had originally donated the cells that this is what they wanted.[60]

In both cases the courts ruled in favour of the respective universities, but the lawsuits highlight that informed consent is not always straightforward, whether for one's DNA or

59 While the model was run using local currency units, the results were converted into US dollars for the sake of comparison. In the case of the UK and Australia, because the pound currently has a higher value against the US dollar than the Australia dollar, the effects appear less dramatic relative to the size of their respective economies.
60 J. Mark Waxman, "Who Owns My Tissue?" *Drug Discovery and Development*, 4 December 2007, https://www.dddmag.com/article/2007/12/who-owns-my-tissue.

© The Economist Intelligence Unit Limited 2017

medical information. Access to data could be undermined by a backlash to what patients consider an invasion of their privacy, or even unduly profiting from something they believed they had freely donated. There also needs to be a clear and reliable legal framework for data ownership, as well as agreements between individual institutions.

In healthcare, meanwhile, as noted earlier, regulators will certainly attach demands for transparency to technology used in treatment, but Dr Abernethy believes that this will be only the first step. Device regulation will need to be transformed in the face of continuous learning. Currently, she notes, regulators approved specific sets of features, but with AI-enabled technology the nature of what those machines will have learned to do is likely to evolve. “It will require regulatory creativity, and willingness on the machine learning side to come up with versions that lock at a particular level for review,” she predicts.

This will be even more important as AI-enabled machines replace clinicians in a variety of ways. In parts of London, for example, in the first half of 2017 Britain’s NHS is trialling Babylon’s chatbot triage service as an alternative to the national medical advice phone line.[61] Even caring roles may be taken on by machines. Advanced research is already taking place in countries such as Japan, which are struggling to cope with ageing populations and shrinking workforces.[62] Such robots, as well as supplying extra labour, could also improve outcomes. Trials with MEDI, a child care robot produced by RxRobots, found that it reduced children’s pain by 50% during medical procedures, and improved vaccination rates by 10%.[63]

These automated or mechanical care providers, however, will need to be regulated just as their human precursors are, in order to insure quality and patient safety. How to do so will be a huge challenge, especially as the line between medical and commercial devices is becoming increasingly blurred through medical applications on smartphones or even virtual reality machines.

Current signs are that governments are sympathetic. Japan’s Robot Revolution Realisation Council (set up to look at various laws which might affect the growth of the country’s robotics industry) has included in its remit smoothing the process of pre-approval for robotic tools under Japan’s medical devices regulations. The European Commission and the US FDA, meanwhile, did relevant research in the early years of this decade, wrestling with questions such as where does wellness or health encouragement tip over into making a phone a medical device. The solution of the FDA—to have a broad definition of what software was a medical device but to announce it would exercise enforcement discretion for the many such apps which it deemed to be low risk—points to the ongoing complexity of that field. As AI increases the power and range of how machines can help with healthcare, these questions are likely to get ever more complicated. For example, when is the Japanese robotic bear bought for a grandmother

61 “The NHS is trialling an AI chatbot to answer your medical questions,” *Wired*, 5 January 2017, http://www.wired.co.uk/article/babylon-nhs-chatbot-app
62 Jon Emont, “Japan Prefers Robot Bears to Foreign Nurses,” *Foreign Policy*, 1 March 2017, http://foreignpolicy.com/2017/03/01/japan-prefers-robot-bears-to-foreign-nurses/
63 “Pediatric Robot Improves Dental Care,” Dentistry Today, 22 Jan 2016, http://www. dentistrytoday.com/news/todays-dental-news/item/686-pediatric-robot-improves-dental-care


© The Economist Intelligence Unit Limited 2017

being a helping hand in moving around the house and when is it being nurse?

**Transportation**

Transportation is seeing some of the most high-profile advances, both current and potential, in the use machine learning. The issues facing policy makers in promoting the safe use of the technology include several questions which are likely to become apparent in other sectors as AI spreads throughout the economy and society.

The most basic challenge is regulation. Progress in coming to grips with autonomous vehicles has been slow. Amendments to the influential Vienna Convention on Road Traffic, which most European states have signed, have allowed certain types of assistance to drivers, such as automated parking and lane changing, since just March of 2016. Driverless cars, or even drivers doing something else while the car drives itself, are still illegal.[64] Meanwhile, a tally kept by Stanford University's Center for Internet and Society reports that, as of May 2017, only five US states have legislation in place allowing the use of autonomous vehicles under certain regulations. It also indicates that proposed legislation has failed to pass in 15 states.[65]

This poor success ratio likely reflects the complexity and novelty of the issues. In order for autonomous vehicles to become the dominant mode of transport, a key regulatory consideration will be, as Professor Agrawal puts it, "who will be liable for what when there are errors." In the event of a fatality, for example, will blame fall on the vehicle owner, its maker, or the manufacturer or designer of relevant software or hardware? "Specifying who in the decision loop is liable for what requires establishing spheres of responsibility," a process which will have to precede mass adoption of the technology, Professor Agrawal continues. He adds that, although potential liability questions will beset any number of uses of AI, its deployment in automobiles, which interact with many people who did not participate in the decision to use the vehicle, makes this issue more salient and also more complex than for, say, machine learning applications in an entirely automated factory or mine.

Liability, however, is only the tip of the regulatory iceberg in transport. Data privacy issues are also relevant for machines that will be know an individual's travel history. Other practical matters will also require revision. For example, how will the requirements for an automobile driver's license, or those for trucks, trains, or aircraft, change as it can increasingly be assumed that some level of automated assistance will be built into the vehicle? Moreover, to what extent will any new user-driver interface need to be standardised?[66] Professor Chun also points out that even a switch from fixed public transport routes to dynamically changing ones based on demand will likely require regulatory approval.

Besides regulation, another key policy issue facing governments will be helping individuals whose jobs are displaced. As noted above, many experts consider this an urgent issue, with

64 "Automated vehicles in the EU," *European Parliament Briefing*, January 2016, http://www.europarl.europa.eu/RegData/etudes/BRIE/2016/573902/EPRS_BRI(2016)573902_EN.pdf
65 "Automated Driving: Legislative and Regulatory Action," *Stanford Law School Centre for Internet and Society Collabowriting Space*, http://cyberlaw.stanford.edu/wiki/ index.php/Automated_Driving:_Legislative_and_Regulatory_Action#State_Bills, accessed 30 May 2017.
66 "Autonomous Car: The regulatory and liability challenges," *Lexology*, 8 April 2016, http://www.lexology.com/library/detail.aspx?g=9f8b0300-d881-4804-8b2d-edc9779c47f7.

© The Economist Intelligence Unit Limited 2017

commercial drivers an oft-cited group likely to need assistance. However, one extended example—US long distance truckers—reveals the potential complexities of setting effective policy in this area.

A December 2016 White House report projected that, for commercial drivers, the biggest total employment loss would be among heavy truck and tractor-trailer drivers, with between 80% and 100% of the current 1.7m such jobs in the US at risk. Clearly, retraining on a grand scale is in order.[67]

Or is it? The first problem is that the report gave no time scale. This was no oversight: how many jobs will disappear how quickly is difficult to tell. So far, the application of AI to trucks has been to improve drivers' performance or safety, not to replace people. Even platooning, as described in an earlier section, cannot currently be done without a driver at the wheel. One common forecast is that the application of technology to commercial vehicles will bring about evolutionary rather than revolutionary change. In such circumstances, jobs will still be available for humans sitting in the driver's seat for some time to come. Similarly, our qualitative scenario for the transport sector notes that despite dramatic improvements in the automation of aircraft, pilots will still be fixtures in cockpits. At the very least, trucking jobs are likely to disappear more slowly than those of workers in a factory seeing automation, meaning that encouraging younger individuals to avoid a potentially diminishing field of employment may be more appropriate than large-scale retraining of existing drivers.

Another difficulty in devising a response is the current economic role played by employment as a heavy, typically long-distance truck driver in the US. Although clearly some skill is involved—a specialised license is required—becoming a trucker is normally seen as one of the few high-income jobs available to those with little education or technical training.[68] Such individuals, however, have almost by definition already found either uninteresting or too difficult the kind of training in technology subjects mooted by job experts. Moreover, an important reason for the relatively high remuneration from this occupation is not technical expertise but a willingness to work long hours and accept frequent absences which can interfere with a settled home and social life. Even if drivers receive more assistance from technology, as long as people are required to sit in the truck itself, the unsocial conditions of the job—and the premium payment which those willing to put up with them can command—will continue.

In sum, launching any direct retraining programme today for US heavy truck drivers will have to convince a large number of people, many of whom have been unsuccessful at academic subjects, to study in order to acquire the skills needed to take on jobs which, given their other workplace options, will almost inevitably be lower paid than their current employment. All this, while it is uncertain how

67 Executive Office of the President, *Artificial Intelligence, Automation, and the Economy*, December 2016, https://obamawhitehouse.archives.gov/sites/whitehouse.gov/files/documents/Artificial-Intelligence-Automation-Economy.PDF
68 "Out of road: driverless vehicles and the end of the trucker," 30 March 2017, *Financial Times*, https://www.ft.com/content/2d70469c-140a-11e7-b0c1-37e417ee6c76.


© The Economist Intelligence Unit Limited 2017

quickly their existing jobs might disappear, or even if it will before they retire. For governments, knowing to whom, and when, to offer retraining in transport will likely turn out to be as complex as regulation.

### Energy

Professor Kim notes that, in general, governments wishing to create a positive environment for deployment of AI need to focus on cybersecurity. "AI has a weak point in this area." This is especially the case in the energy sector, adds Dr Otterson: "When the entire power system is driven by computers, they really have to work," he says. Data processing and data links are nearly as important as the wires carrying the power, according to Dr Otterson.

The costs of failure in terms of lives and money can be huge. A report by the University of Cambridge's Centre for Risk Studies and the insurance exchange Lloyd's estimated that a virus capable of disabling just 7% of large power generation facilities in the US Northeast would cause US$243bn in damage to that country's economy.[69]

Moreover, Dr Otterson adds "we've had fears of hacking power systems before, but it is smart grids increase the need for data security." Such grids involve far more individuals and companies simultaneously generating and sharing information across the network than traditional generation. While allowing a host of benefits, these connections also invariably greatly increase what a recent MIT report calls the number of "attack surfaces" requiring enhanced security.[70]

Governments have started to work on the issue. The US has issued various requirements for cybersecurity of critical infrastructure and the EU in 2016 created an Energy Expert Cyber Security Platform to provide advice to the European Commission. Nevertheless, much remains to be done: the MIT report argues that issues which still have to be addressed include "cloud security, machine-to-machine information sharing, advanced cybersecurity technologies, outcome-based regulation to avoid prolonged outages and increase system resilience, and international approaches to cybersecurity"—no small agenda.

Another key area for government activity in the energy field is likely to be promoting the collection and sharing of information needed by a range of stakeholders for projects of general public interest. Dr Otterson explains that, in implementing the smart grid, "aside from building a large portfolio of distributed generation, regulatory and data issues are the biggest problems. Deep learning and big data techniques require a huge amount of data which are not available in many cases. For example, you can't get measurements for power flow in some countries, and you can't find regional measurements of power demand." He sees the solution as one of market and regulatory restructuring. "Right now, there is a lack of incentive for that to be shared publicly. Markets are not currently set up

69 University of Cambridge Centre for Risk Studies *The insurance implications of a cyber attack on the US power grid*, 2015.
70 MIT Energy Initiative Utility of the Future, "Cybersecurity White Paper," 15 December 2016," https://energy.mit.edu/wp-content/uploads/2016/12/ CybersecurityWhitePaper_MITUtilityofFuture_-2016-12-05_Draffin.pdf

© The Economist Intelligence Unit Limited 2017

to provide the data needed to do these computations. The requirements for sharing information need to change."

Finally, certain governments may be able to use their existing role in the economy to advance the use of AI in the energy field, with smart grids again furnishing the best example.

The Chinese government has to balance a number of energy-related issues facing the country. On the one hand, demand is increasing, growing by 5% in 2016 which was down from the 8% average of the first three years of this decade. Most of the power comes from fossil fuels, in particular coal, which accounted for just under two-thirds of total generation in 2016.[71] The resultant carbon and other emissions are an issue of increasing salience in a country where urban air quality is notoriously poor. As a result, the government wishes to increase generation of renewable and other low-carbon energy faster than that coming from fossil fuels in order to shift toward a greener overall mix. This includes a large increase in wind and solar power, the generation capacity of which rose 18% and 80% respectively last year.[72] Such a transformation, however, brings the traditional problem of intermittent supply.

The major player in addressing these issues will be the government because the transmission and distribution sectors, although divided into a number of separate companies at the start of the century, largely remain in the hands of state-owned enterprises.

The solution which the government has pursued as part of a long term project—launched in 2009 and integrated into both the last five-year plan and the recently adopted 13th plan—is to build a grid that is both "strong and smart" at the national and local levels. It includes the creation of a huge network of ultra high-voltage (UHV) power lines, allowing transmission from remote sources of solar, wind, hydro and even coal generation to populous areas, as well as rapid exchange of power between regional distributors based on real-time demand. Although the emphasis has been on the transmission element of the transformation, the two large state energy distributors, the dominant State Power Grid Corporation of China (SPGCC), which controls about 80% of the market, and the smaller China South Power Grid Corporation, have been investing in the infrastructure to create smart metering and grids at local levels as well.

The result remains very much a work in progress. Improvements have occurred, but have also fallen short of original hopes. The proportion of non-fossil fuel generated electricity in 2015 was slightly higher than forecast (12% instead of 11.4%), the original projection for 2020 has been revised down from 20% to 15%. Meanwhile, SPGCC had built and put in operation eight UHV lines by the end of last year, covering 11,900 km. On the other hand, it had planned to have 11 such lines complete by that time and looks unlikely to meet its original goal of 37 lines stretching a total of 89,000 km by 2020. Worse still, the first three lines, built between 2009 and 2011, were

71 "Power statistics China 2016: Huge growth of renewables amidst thermal-based generation," *Chinese-European Energy News*, 9 February 2017, http://ceenews.info/en/ power-statistics-china-2016-huge-growth-of-renewables-amidst-thermal-based-generation/; James Lau, "Smart Metering and the Smart Grid in China - Opportunities for EU SMEs," http://www.eusmecentre.org.cn/sites/default/files/files/event/EU SME Centre Webinar - Smart Metering and the Smart Grid in China - Opportunities for EU SMEs.pdf.
72 "Power statistics China 2016: Huge growth of renewables amidst thermal-based generation," *Chinese-European Energy News*, 9 February 2017, http://ceenews.info/en/ power-statistics-china-2016-huge-growth-of-renewables-amidst-thermal-based-generation/


© The Economist Intelligence Unit Limited 2017

carrying only between 21% and 56% of capacity by the end of 2014.[73] Similarly, as noted earlier, China leads the world with 348m smart meters installed by late 2016. On the other hand, most of the installations carried out in 2016 were replacements of existing meters, among which faulty technology and inconsistent use of communication frequencies were common problems.[74]

China, then, has gone some way in showing the world what a nationwide smart (and strong) grid might look like, but still has some way to go before it becomes a reality.

73 "China's under-utilised ultra-high-voltage power lines no silver bullet to rid grid of bottlenecks," *South China Morning Post*, 14 February, 2016, http://www.scmp.com/business/article/1912878/chinas-under-utilised-ultra-high-voltage-power-lines-no-silver-bullet-rid
74 "State Grid to dominate Chinese smart meters deployment," 10 May 2016, *Metering and Smart Energy International,* https://www.metering.com/news/the-state-grid-corporation-of-china-to-dominate-the-deployment-of-smart-meters-ahead-of-other-utilities/

© The Economist Intelligence Unit Limited 2017

# CONCLUSION

The debate over the impact and uses—and misuses—of AI is a necessary one. As we've shown, AI in the form of machine learning offers a great many benefits, although our findings suggest not quite as many and as large as has been claimed, at least not in the near term. Much the same can be said of those warning of AI's severe and deleterious effects on the labour market and society—the spread of AI can and will cause disruption, but that disruption can also be managed and mitigated with the right mix of policies.

One of the key challenges in coming years will be to make sure that discussions around AI and its impact remain grounded in reality so as not to drift into the realm of science fiction. There are a number of approaches to accomplishing this.

**Managing expectations.** Trumpeting AI as a utopian solution to all the world's problems is, in some ways, just as detrimental to the tone of the debate as predicting the eventual enslavement of humanity at the hands of machines. AI can help improve productivity and increase economic output, provided the right conditions. But barring some step-change in the state of the technology (which history teaches us not to rule out), the results of Scenarios 1 and 2 show improvements to be more modest than rhetoric from some quarters seems to project. Left unchecked, outlandish and unfounded claims could simultaneously breed complacency among AI's proponents and policy makers and exacerbate the fears of the doomsdayers.

This applies to the sector and societal-level as well. For all the talk of driverless cars and personalised medicine, AI's impact over the next decade-plus is likely to be in evidence in more quotidian, less media-friendly areas like efficiency gains in manufacturing and energy-usage and rural access to healthcare services. These are important, even if they don't generate headlines.

**Better communication.** There are a number of understanding gaps when it comes to AI, but one of the most important to bridge is between developers, and businesses and government institutions. The former are often only dimly aware of what the latter really need and the latter, in turn, are often only dimly aware of the potential solutions the former could provide. A more robust and frequent exchange of information, capabilities and needs would help remedy this

**Acknowledging the risks.** Although Mr Kaplan's argument about the risks to labour


© The Economist Intelligence Unit Limited 2017

presented by machine learning is not without its critics, it should not be easily dismissed. In fact, being constructive and acknowledging all risks—instead of downplaying or dismissing them altogether—will help to shape the debate about skills and other approaches to employment. The alternative runs the risk of inspiring resignation rather than action.

Data privacy and security in general, and specifically as they relate to AI, are areas that would benefit from a similar approach. The general public only has a dim view on how the data it generates every day is being used and protected. Being forthright on these matters could help assuage their burgeoning concerns; secrecy will only amplify them. Opting for the latter may eventually lead to stricter policies that deprive AI of one of its most important inputs.

**Improving trust and transparency.** As a corollary to the above, efforts to improve public trust are vital and inevitably involve greater transparency and accountability. Cathy O'Neil, the data scientist and author of Weapons of Math Destruction, a book on the uses and misuses of data and algorithms, highlights how many of the decisions that govern society are increasingly being made by algorithms that are not made public. In part this is to protect intellectual property, but it is more complex than that and involves such considerations as security risks, data privacy and how to audit code independent of a specific application.

**Educating the public.** Gaps in knowledge and understanding are more quickly-filled than ever with misinformation and distortion. In the case of AI, there are legitimate concerns about its implications for society, some of which are expressed by scientists and informed business leaders. The public debate, however, has tendency to veer towards more science fiction than science. Given the difficulty of the subject matter, that is somewhat unavoidable. Nevertheless, even though simplification can feel like obfuscation to an expert, efforts still need to be made to explain AI and its uses in a manner that will, at the very least, crowd out the fiction in favour of the science.

Policymakers, for their part, face a number of choices when it comes to addressing AI and its impact on the economy and society.

**Investing in skills and training**. That there is going to be churn in the labour market as a result of AI is widely accepted, as is the attendant need for significant portions of the labour force to undergo education and training throughout their working lives. For some, this will mean a focus on vocational education, a course of studies that is lacking in most countries, including many of those covered in this report. This will need to be addressed.

© The Economist Intelligence Unit Limited 2017

At the same time, it is important for policymakers not to over-adjust towards vocational and STEM education at the expense of liberal arts. As highlighted in the report, observers expect the demand for "soft skills" such as team building, cooperation and critical thinking to increase as AI and machine learning evolve. These skills, as well as the broader concept of good judgement, are not wholly dependent on a liberal arts education, but a basic grounding in history, philosophy and literature will be increasingly important.

Finally, because the combination of skills in demand will be constantly evolving, policymakers, industry executives and educators need to be in frequent communication, lest the infrastructure for skills and training become outdated.

**Dealing with data**. The uses and misuses of data are going to be among the defining issues of the 21st century. For data to be as open and available as it needs to be—both within countries and between them—policymakers need to take steps to assuage the legitimate concerns of their citizens with regards to privacy and security. One key solution is developing rules and regulations that enable and support the use of anonymized data sets.

Another is to make national data privacy schemes inter-operable so that data can flow across borders. To date, this has proven difficult and, as a result, a number of countries have implemented data localisation requirements, among other measures. Should restrictions like these become more prevalent, and cross-border data flows dry up, the impact on technology development and uptake could be considerable.

**Investing in R&D and technology**. Public sector investment in R&D and "deep research" has decreased in many countries. Part of this funding gap has been filled by the private sector, but that could prove unsustainable and if economies are to have the intellectual capacity to capitalise on new technologies, the public sector needs to be more involved.

Related to this, more can be done to incentivise investment in R&D and technology, more broadly. Creating enabling regulations is one approach, as is direct financial support in the form of tax credits.

Advances in technology are accelerating faster than the ability of society and politics to keep up. This is especially true in the field of AI, which has the added complication of being more important and more difficult to comprehend than, say, 5G networks or robotics. Yet, as the age-old saying goes,


© The Economist Intelligence Unit Limited 2017

"With great power comes great responsibility." In the years ahead, developers and users of AI will have great power. The depth and breadth of the benefits will depend on how responsible they are.

© The Economist Intelligence Unit Limited 2017

# APPENDIX 1:

## Australia

### Baseline forecast

The baseline forecast assumes that economic conditions and public policy in Australia correspond to the assumptions underlying our standard model. It sees the Australian economy expanding at a CAGR of 1.03% over the course of the forecast period, bringing the Australian GDP to US$1.4trn in 2030, a nominal increase of US$200bn over the 2016 value. Over the same period, productivity in Australia (measured in billions of dollars per million man-hours worked) is expected to be stagnant, increasing from 0.0598 to 0.062, a CAGR of only 0.19%. Australia is a wealthy country that has long relied on exports from primary industries for incremental growth in the economy—a form of growth that may be increasingly difficult to sustain. This reality is reflected in Australia's modest baseline estimates for GDP and productivity.

### Scenario #1: Greater human productivity through upskilling

Scenario 1 assumes that Australia takes positive policy action to address the threat to human employment posed by machine learning, over and above the public policy actions assumed in the base case. By increasing access to higher education and job training, policy makers in this scenario successfully increase the average skill level of the workforce, thereby ensuring that the complementary effect is stronger than the substitution effect, and that the dominant impact

**Scenario results for the Australia**

| | | 2016 | 2030 (s) | 2030 (b) | Change v. '16 | Change v. baseline | CAGR 2016 - 30 |
|---|---|---|---|---|---|---|---|
| **Scenario #1** | GDP | $1.2trn | $1.9trn | $1.4trn | $700bn | $500bn | 3.11% |
| | Productivity | 0.0598 | 0.0840 | 0.0620 | 0.0242 | 0.0220 | 2.25% |
| **Scenario #2** | GDP | $1.2trn | $2.1trn | $1.4trn | $900bn | $700bn | 3.74% |
| | Productivity | 0.0598 | 0.0920 | 0.0620 | 0.0322 | 0.0300 | 2.88% |
| **Scenario #3** | GDP | $1.2trn | $1.2trn | $1.4trn | -$43bn | -$200bn | -0.24% |
| | Productivity | 0.0598 | 0.0510 | 0.0620 | -0.0088 | -.0110 | -1.07% |

Notes: All GDP in US dollars and rounded from official results
(s) = Scenario result
(b) = The Economist Intelligence Unit baseline forecast
Productivity = output/man hour worked


© The Economist Intelligence Unit Limited 2017

of machine learning technology is to enhance, rather than replace, human labour.

Under Scenario 1, the Australian economy expands at a CAGR of 3.11%, with Australian GDP reaching US$1.9trn by 2030, a figure which represents US$500bn in additional economic activity in 2030 over the baseline forecast. Since Scenario 1 assumes little change in the same number of hours worked in 2030 compared with the base case, this higher level of economic output can be attributed almost entirely to the positive impact on productivity of public policies to upskill workers and increase complementarity between human labour and machine learning technology. In this scenario, productivity increases from 0.0598 to 0.084, a CAGR of 2.25%. Though Australia already has a highly-educated population, it has the most to gain from greater upskilling of its workforce. Under Scenario 1, both GDP and productivity in Australia are predicted to rise by over 2%, compared with the baseline estimate.

**Scenario #2: Greater investment in technology and access to open source data**

Scenario 2 assumes that Australia is successfully able to increase the rate of capital investment in machine learning technology, and hence increase its diffusion and adoption over and above the base case, thanks to the positive impact of public policy initiatives. The scenario assumes efforts will be focused primarily on two main policy interventions: creating new tax incentives to encourage private sector investment in machine learning technology and increasing access to open source data by investing in national knowledge sharing communities; and through the government releasing more of its own data to the public.

Of all the scenarios, including the baseline, Scenario 3 results in the greatest positive impact for the Australian economy. Pursuing policies that encourage greater capital investment in machine learning technology results in significant increases in both GDP and productivity over the base case, due to machine learning both being widespread and having a predominantly complementary impact on the workforce. Australian GDP in this scenario grows at a CAGR of 3.74%, over three and a half times the baseline rate of growth, rising to US$2.1trn in 2030. This represents an increase of US$700trn, fully 50%, over the baseline estimate. Though total hours worked are moderately higher than the baseline in this scenario, the bulk of the increase in GDP comes from productivity, which increases at a CAGR of 2.88%, as opposed to 0.19% in the baseline estimate. As with Scenario 1, Australia is the country which stands to gain the most by far in Scenario 2, with GDP and productivity growth rates that are 2.71% and 2.69% above the baseline estimate, respectively. With a service sector that’s growing in importance and sophistication, Australia is particularly well-placed to benefit from increased capital investment in technology and increased data sharing.

**Scenario #3: Insufficient policy support for structural changes in the economy**

Scenario 3 assumes that policy makers in Australia fail to step up to the economic challenges posed by the development of machine learning technology, perhaps due to continued political instability

© The Economist Intelligence Unit Limited 2017

and short termism in the national government. Without the fillip of new tax incentives for private sector investment assumed in Scenario 2, machine learning technology continues to develop in line with the base case assumptions. The lack of government support for national data sharing, meanwhile, reduces the overall positive economic impact of machine learning, which requires data of both high quality and quantity to be of maximum usefulness. Most crucially, a lack of investment in upskilling the workforce in this scenario results in the substitution effect being stronger than the complementary effect, so that the dominant impact of machine learning technology is to replace workers, rather than to increase their productivity.

The impact of this policy inaction on the Australian economy is profoundly negative, and can be seen in both the GDP and productivity estimates produced by the model. It stands to reason that Australia, as the country which had the most to gain under Scenarios 1 and 2, is also one of the countries to be hit the hardest by the policy paralysis in Scenario 3. GDP is essentially flat in Australia over the forecast period in this scenario, actually contracting slightly at a CAGR of -0.24%, resulting in 200bn fewer dollars of economic activity in 2030 than in the base case. Productivity, meanwhile, heads into negative territory over the period, decreasing at a CAGR of -1.07%. Total hours worked in 2030 increase by 4.2% relative to the baseline estimate, implying that the negative impact on unemployment levels may be less pronounced in Australia than in some of the other developed economies.


© The Economist Intelligence Unit Limited 2017

## Developing Asia

### Baseline forecast

The baseline forecast assumes that economic conditions and public policy in the Developing Asia group of countries correspond to the assumptions underlying our standard model. It sees the economies in Developing Asia expanding at a CAGR of 4.34% over the course of the forecast period, bringing the group's GDP to US$18trn in 2030, a nominal increase of US$8.5trn and almost a doubling of the 2016 value. Over the same period, productivity in Developing Asia (measured in billions of dollars per million man-hours worked) is expected to increase from 0.0027 to 0.0045, a CAGR of 3.49%. The countries in this group have by far the highest expected growth rates of the group, reflecting their status as rapidly developing economies and the fact that they are starting from a much lower level of GDP per capita and productivity than the mature economies.

### Scenario #1: Greater human productivity through upskilling

Scenario 1 assumes that governments in the Developing Asia group of countries take positive policy action to address the threat to human employment posed by machine learning, over and above the public policy actions assumed in the base case. By increasing access to higher education and job training, policy makers in this scenario successfully increase the average skill level of the workforce, thereby ensuring that the complementary effect is stronger than the substitution effect, and that the dominant impact of machine learning technology is to enhance, rather than replace, human labour.

Under Scenario 1, the Developing Asia group of economies expand at a CAGR of 5.04%, with Developing Asia GDP reaching US$20trn by 2030, a figure which represents US$2trn in additional economic activity in 2030 over the baseline forecast. Since this scenario assumes the same

**Scenario results for the Developing Asia**

| | | 2016 | 2030 (s) | 2030 (b) | Change v. '16 | Change v. baseline | CAGR 2016 - 30 |
|---|---|---|---|---|---|---|---|
| **Scenario #1** | GDP | $9.5trn | $20.0trn | $18.0trn | $10.5trn | $2.0trn | 5.04% |
| | Productivity | 0.0027 | 0.0050 | 0.0045 | 0.0023 | 0.0005 | 4.06% |
| **Scenario #2** | GDP | $9.5trn | $24.4trn | $18.0trn | $14.9trn | $6.4trn | 6.47% |
| | Productivity | 0.0027 | 0.0061 | 0.0045 | 0.0034 | 0.0016 | 5.60% |
| **Scenario #3** | GDP | $9.5trn | $15.3trn | $18.0trn | $5.8trn | -$2.7trn | 3.20% |
| | Productivity | 0.0027 | 0.0038 | 0.0045 | 0.0011 | -0.0007 | 2.35% |

Notes: All GDP figures in US dollars and rounded from official results.
(s) = Scenario result
(b) = The Economist Intelligence Unit baseline forecast
Productivity = output/man hour worked

© The Economist Intelligence Unit Limited 2017

number of hours worked in 2030 as in the base case, this higher level of economic output can be attributed entirely to the positive impact on productivity of public policies to upskill workers and increase complementarity between human labour and machine learning technology. In this scenario, productivity increases from 0.0027 to 0.005, a CAGR of 4.06%. Though both GDP and productivity in Scenario 1S show improvement over the base case for Developing Asia, the difference is modest at around 16%. This may be because many of these countries are starting from very low average levels of human capital, in terms of educational attainment and computer literacy, which prevents them from taking full advantage of the potential of machine learning technology.

**Scenario #2: Greater investment in technology and access to open source data**
Scenario 2 assumes that governments in the Developing Asia group of countries are successfully able to increase the rate of capital investment in machine learning technology, and hence increase its diffusion and adoption, over and above the base case, thanks to the positive impact of public policy initiatives. The scenario assumes efforts will be focused primarily on two main policy interventions: creating new tax incentives to encourage private sector investment in machine learning technology, and increasing access to open source data by investing in national knowledge sharing communities; and through the government releasing more of its own data to the public.

Of all the scenarios, including the baseline, Scenario 2 results in the greatest positive economic impact for the economies of Developing Asia. Pursuing policies that encourage greater capital investment in machine learning technology results in significant increases in both GDP and productivity over the base case, due to machine learning both being widespread and having a predominantly complementary impact on the workforce. Developing Asia GDP in this scenario grows at a CAGR of 6.47%, almost 50% higher than the baseline rate of growth, rising to US$24.4trn in 2030. This represents an increase of US$6.4trn, or 36%, over the baseline estimate. Total hours worked are in line with the baseline in this scenario, implying that productivity, which rises to a CAGR of 5.6% over the baseline assumption of a 4.34% CAGR, is to thank for this increase in GDP. This represents a 60% increase over the base case, which, though a solidly positive result, is less of a percentage increase than most of the developed countries experience in the same scenario. This is likely due to the lower cost incentive for employers in Developing Asia to replace workers with automation, as well as the fact that machine learning poses less of a direct threat to workers in manufacturing industries (an important part of many of these economies) than it does to service workers.

**Scenario #3: Insufficient policy support for structural changes in the economy**
Scenario 3 assumes that policy makers in Developing Asia fail to step up to the economic challenges posed by the development of machine learning technology. Without the fillip of new tax incentives for private sector investment assumed in Scenario 22 machine learning technology continues


© The Economist Intelligence Unit Limited 2017

to develop in line with the base case assumptions. The lack of government support for national data sharing, meanwhile, reduces the overall positive economic impact of machine learning, which requires data of both high quality and quantity to be of maximum usefulness. Most crucially, a lack of investment in upskilling the workforce in this scenario results in the substitution effect being stronger than the complementary effect, so that the dominant impact of machine learning technology is to replace workers, rather than to increase their productivity.

This policy inaction has a significant negative impact on the economies of Developing Asia, affecting both the GDP and productivity estimates produced by the model. While Developing Asia is less affected from a relative perspective than the more mature economies, it still loses about a quarter of its GDP growth and about a third of its productivity growth over the forecast period, which fall to CAGRs of 3.2% and 2.35% respectively Scenario 3. This results in US$2.7trn of foregone economic activity in 2030 compared with the base case. Total hours worked in 2030 remain roughly in line with the baseline estimate.

© The Economist Intelligence Unit Limited 2017

## Japan

### Baseline forecast

The baseline forecast assumes that economic conditions and public policy in Japan correspond to the assumptions underlying our standard model. It sees the Japanese economy expanding at a CAGR of 1.57% over the course of the forecast period, bringing the Japanese GDP to US$6.1trn in 2030, a nominal increase of US$1.3bn over the 2016 value. Over the same period, productivity in Japan (measured in billions of dollars per million man-hours worked) is expected to increase from 0.0421 to 0.054, a CAGR of 1.67%. Japan's baseline estimates for growth in GDP and productivity compare favourably with those of other developed economies, especially considering the considerable demographic challenges the country is faced with.

### Scenario #1: Greater human productivity through upskilling

Scenario 1 assumes that Japan takes positive policy action to address the threat to human employment posed by machine learning, over and above the public policy actions assumed in the base case. By increasing access to higher education and job training, policy makers in Scenario 1 successfully increase the average skill level of the workforce, thereby ensuring that the complementary effect is stronger than the substitution effect, and that the dominant impact of machine learning technology is to enhance, rather than replace, human labour.

Under Scenario 1, the Japanese economy expands at a CAGR of 1.96%, with Japanese GDP reaching US$6.4trn by 2030, a figure which represents US$300bn in additional economic activity in 2030 over the baseline forecast. Since this scenario assumes little change in the same number of hours worked in 2030 compared with the base case, this higher level of economic output can be attributed almost entirely to the positive impact on productivity of public policies to upskill workers and increase complementarity between human

**Scenario results for the Japan**

| | | 2016 | 2030 (s) | 2030 (b) | Change v. '16 | Change v. baseline | CAGR 2016 - 30 |
|---|---|---|---|---|---|---|---|
| **Scenario #1** | GDP | $4.8trn | $6.4trn | $6.1trn | $1.6trn | $300bn | 1.96% |
| | Productivity | 0.0421 | 0.0570 | 0.0540 | 0.0160 | 0.0030 | 2.06% |
| **Scenario #2** | GDP | $4.8trn | $6.9trn | $6.1trn | $2.1trn | $800bn | 2.43% |
| | Productivity | 0.0421 | 0.0610 | 0.0540 | 0.0189 | 0.0070 | 2.53% |
| **Scenario #3** | GDP | $4.8trn | $5.2trn | $6.1trn | $400bn | -$900bn | 0.53% |
| | Productivity | 0.0421 | 0.0460 | 0.0540 | 0.0039 | -0.0080 | 0.63% |

Notes: All GDP figures in US dollars and rounded from official results
(s) = Scenario result
(b) = The Economist Intelligence Unit baseline forecast
Productivity = output/man hour worked


© The Economist Intelligence Unit Limited 2017

labour and machine learning technology. In this scenario, productivity increases from 0.042 to 0.057, a CAGR of 2.06%. As with the US and South Korea, the increases in both GDP and productivity over the base case that Japan experiences in Scenario 1 are relatively moderate, most likely due to high starting levels of educational attainment and computer literacy in the Japanese workforce.

**Scenario #2: Greater investment in technology and access to open source data**
Scenario 2 assumes that Japan is successfully able to increase the rate of capital investment in machine learning technology, and hence increase its diffusion and adoption, over and above the base case, thanks to the positive impact of public policy initiatives. The scenario assumes efforts will be focused primarily on two main policy interventions: creating new tax incentives to encourage private sector investment in machine learning technology; and increasing access to open source data by investing in national knowledge sharing communities and releasing more of its own data to the public.

Japan has one of the most highly advanced economies in the world, with very high current levels of investment in computer capital, so it is not surprising that is predicted to make smaller gains in Scenario 2 than the other countries. Nonetheless, the impact of the policy assumptions in Scenario 2 on the Japanese economy is still overwhelmingly positive with both GDP and productivity growth rates more than 50% above the baseline estimates. Pursuing policies that encourage greater capital investment in machine learning technology results in significant increases in both GDP and productivity over the base case, due to machine learning both being widespread and having a predominantly complementary impact on the workforce. The Japanese GDP in this scenario grows at a CAGR of 2.43%—55% higher than the baseline rate of growth—rising to US$6.9trn in 2030. This represents an increase of US$800trn—or 13%—over the baseline estimate. Total hours worked are essentially the same as in the baseline in this scenario, implying that the increase in GDP comes almost exclusively from increases in productivity, which grows at a CAGR of 2.53%, as opposed to 1.67% in the baseline estimate.

**Scenario #3: Insufficient policy support for structural changes in the economy**
Scenario 3 assumes that policy makers in Japan fail to step up to the economic challenges posed by the development of machine learning technology. Without the fillip of new tax incentives for private sector investment assumed in Scenario 2, machine learning technology continues to develop in line with the base case assumptions. The lack of government support for national data sharing, meanwhile, reduces the overall positive economic impact of machine learning, which requires data of both high quality and quantity to be of maximum usefulness. Most crucially, a lack of investment in upskilling the workforce in this scenario results in the substitution effect being stronger than the complementary effect, so that the dominant impact of machine learning technology is to replace workers, rather than to increase their productivity.

The impact of this policy inaction on the Japanese

© The Economist Intelligence Unit Limited 2017

economy is profoundly negative, and can be seen in both the GDP and productivity estimates produced by the model. While GDP still grows in Japan over the forecast period in Scenario 3, it does so at a CAGR of just 0.53%, only about a third of the baseline estimate, resulting in 900bn fewer dollars of economic activity in 2030 than in the base case. Productivity, meanwhile, is sluggish over the period, growing at a CAGR of only 0.63%—a rate which is more than three-fifths lower than the baseline estimate. Total hours worked in 2030 are broadly in line with the baseline estimate. The decline in productivity in Scenario 3 is particularly concerning, given the demographic headwinds facing the Japanese economy. Japan's working age population and overall population are both expected to shrink over the forecast period, leaving fewer workers to support more pensioners. Without high levels of productivity growth, this will lead to an overall fall in Japanese living standards.


© The Economist Intelligence Unit Limited 2017

## South Korea

**Baseline forecast**

The baseline forecast assumes that economic conditions and public policy in South Korea correspond to the assumptions underlying our standard model. It sees the South Korean economy expanding at a CAGR of 1.78% over the course of the forecast period, bringing South Korean GDP to US$1.7trn in 2030, a nominal increase of US$400bn over the 2016 value. Over the same period, productivity in South Korea (measured in billions of dollars per million man-hours worked) is expected to remain relatively flat, increasing from 0.023 to 0.024, a CAGR of 0.35%. South Korea's baseline estimates for growth in GDP and productivity reflect its status as a mature, technologically sophisticated economy.

**Scenario #1: Greater human productivity through upskilling**

Scenario 1 assumes that South Korea takes positive policy action to address the threat to human employment posed by machine learning, over and above the public policy actions assumed in the base case. By increasing access to higher education and job training, policy makers in Scenario 1 one successfully increase the average skill level of the workforce, thereby ensuring that the complementary effect is stronger than the substitution effect, and that the dominant impact of machine learning technology is to enhance, rather than replace, human labour.

Under this scenario, the South Korean economy expands at a CAGR of 2.07%, with its GDP reaching US$1.8trn by 2030, a figure which represents US$100bn in additional economic activity in 2030 over the baseline forecast. Since Scenario 1 assumes only a modest increase in the number of hours worked in 2030 over the base case, this higher level of economic output can be attributed almost entirely to the positive impact on productivity of public policies to upskill workers

**Scenario results for the South Korea**

| | | 2016 | 2030 (s) | 2030 (b) | Change v. '16 | Change v. baseline | CAGR 2016 - 30 |
|---|---|---|---|---|---|---|---|
| **Scenario #1** | GDP | $1.3trn | $1.8trn | $1.7trn | $500bn | $100bn | 1.96% |
| | Productivity | 0.0226 | 0.0250 | 0.0240 | 0.0024 | 0.0010 | 0.63% |
| **Scenario #2** | GDP | $1.3trn | $2.0trn | $1.7trn | $700bn | $300bn | 2.43% |
| | Productivity | 0.0226 | 0.0280 | 0.0240 | 0.0054 | 0.0040 | 1.55% |
| **Scenario #3** | GDP | $1.3trn | $1.3trn | $1.7trn | $139bn | -$400bn | -0.24% |
| | Productivity | 0.0226 | 0.0180 | 0.0240 | -0.0086 | -0.0060 | -1.39% |

Notes: All GDP figures in US dollars and rounded form the official results
(s) = Scenario result
(b) = The Economist Intelligence Unit baseline forecast
Productivity = output/man hour worked

© The Economist Intelligence Unit Limited 2017

and increase complementarity between human labour and machine learning technology. In this scenario, productivity growth, though marginally higher than in the base case, is still disappointing, increasing from 0.023 to 0.025, a CAGR of 0.63%. As with the US and Japan, the increases in both GDP and productivity over the base case that South Korea experiences in Scenario 1 are relatively moderate, most likely due to high starting levels of educational attainment and computer literacy in the South Korean workforce.

**Scenario #2: Greater investment in technology and access to open source data**

Scenario 2 assumes that South Korea is successfully able to increase the rate of capital investment in machine learning technology, and hence increase its diffusion and adoption, over and above the base case, thanks to the positive impact of public policy initiatives. The scenario assumes efforts will be focused primarily on two main policy interventions: creating new tax incentives to encourage private sector investment in machine learning technology, and increasing access to open source data by investing in national knowledge sharing communities; and through the government releasing more of its own data to the public.

South Korea has high existing levels of computer capital, in line with its status as a world leader in the technology sector, but is also a country in the midst of a transition from a reliance on industrial mass production to higher value-added activities. The policy changes in Scenario 2 enable the country to build on its current strong base, resulting in a significant positive economic impact for the South Korean economy. Pursuing policies that encourage greater capital investment in machine learning technology results in significant increases in both GDP and productivity over the base case, due to machine learning both being widespread and having a predominantly complementary impact on the workforce. South Korean GDP in this scenario grows at a CAGR of 3.0%, which is more than two thirds higher than the baseline rate of growth, rising to US$2trn in 2030. This represents an increase of US$300bn, or almost 18%, over the baseline estimate. Though total hours worked are modestly higher than the baseline in this scenario, majority of the increase in GDP comes from productivity, which increases at a CAGR of 1.55%, over four times the baseline estimate.

**Scenario #3: Insufficient policy support for structural changes in the economy**

Scenario 3 assumes that policy makers in South Korea fail to step up to the economic challenges posed by the development of machine learning technology. Without the fillip of new tax incentives for private sector investment assumed in Scenario 2, machine learning technology continues to develop in line with the base case assumptions. The lack of government support for national data sharing, meanwhile, reduces the overall positive economic impact of machine learning, which requires data of both high quality and quantity to be of maximum usefulness. Most crucially, a lack of investment in upskilling the workforce in this scenario results in the substitution effect being stronger than the complementary effect, so that the dominant


© The Economist Intelligence Unit Limited 2017

impact of machine learning technology is to replace workers, rather than to increase their productivity.

The impact of this policy inaction on the South Korean economy is profoundly negative, and can be seen in both the GDP and productivity estimates produced by the model. GDP growth in South Korea over the forecast period is completely wiped out in Scenario 3, with a CAGR of 0.02%, resulting in 400bn fewer dollars of economic activity in 2030 than in the base case. Productivity, meanwhile, dips deeply into negative territory over the period, with an estimated CAGR of -1.39%. Total hours worked in 2030 increase by almost 2% relative to the baseline estimate, implying at least a partial mitigation of the negative impact on unemployment levels. Along with the US, South Korea is one of the hardest hit of all the countries by the public policy failures in Scenario 3, with falls in GDP and productivity growth rates of 1.76% and 1.74%, respectively.

© The Economist Intelligence Unit Limited 2017

## The UK

### Baseline forecast

The baseline forecast assumes that economic conditions and public policy in the UK correspond to the assumptions underlying our standard model. It assumes that the UK economy will generate lacklustre growth, expanding at a CAGR of just 0.63% over the course of the forecast period, raising the UK's GDP to US$2.8trn in 2030, a nominal increase of US$300bn over the 2016 value. Productivity in the UK (measured in billions of dollars per million man-hours worked) is expected to decline from 0.045 to 0.041 over the same period, a CAGR of -0.62%. The falling productivity assumed in the base case is in line with recent performance in the UK, where the country's struggle with poor productivity has long been one of economic policy makers' top concerns. It is also one of the main factors behind the disappointing GDP growth predicted for the UK in the baseline estimate. Increased economic uncertainty and negative repercussions on trade and investment related to the UK's impending exit from the EU are another factor contributing to low growth in the UK's baseline estimate. In the base case, the UK is expected to underperform every other country in the analysis, both in terms of GDP and productivity growth.

### Scenario #1: Greater human productivity through upskilling

Scenario 1 assumes that the UK takes positive policy action to address the threat to human employment posed by machine learning, over and above the public policy actions assumed in the base case. By increasing access to higher education and job training, policy makers in Scenario 1 successfully increase the average skill level of the workforce, thereby ensuring that the complementary effect is stronger than the substitution effect, and that the dominant impact of machine learning technology is to enhance, rather than replace, human labour.

### Scenario results for the UK

| | | 2016 | 2030 (s) | 2030 (b) | Change v. '16 | Change v. baseline | CAGR 2016 - 30 |
|---|---|---|---|---|---|---|---|
| **Scenario #1** | GDP | $2.5trn | $3.1trn | $2.8trn | $600bn | $300bn | 1.29% |
| | Productivity | 0.0450 | 0.0453 | 0.0410 | 0.003 | 0.002 | 0.04% |
| **Scenario #2** | GDP | $2.5trn | $3.4trn | $2.8trn | $900bn | $600bn | 1.94% |
| | Productivity | 0.0450 | 0.0498 | 0.0410 | 0.0048 | 0.088 | 0.68% |
| **Scenario #3** | GDP | $2.5trn | $2.1trn | $2.8trn | -$300bn | -$700bn | -1.20% |
| | Productivity | 0.0450 | 0.0310 | 0.0410 | - 0.004 | - 0.0110 | -2.41% |

Notes: All GDP figures in US dollars and rounded from the official results
(s) = Scenario result
(b) = The Economist Intelligence Unit baseline forecast
Productivity = output/man hour worked


© The Economist Intelligence Unit Limited 2017

Under Scenario 1, the UK economy expands at a CAGR of 1.29%, with the UK's GDP reaching US$3.1trn by 2030, a figure which represents US$300bn in additional economic activity in 2030 over the baseline forecast. Since Scenario 1 assumes little change in the same number of hours worked in 2030 compared with the base case, this higher level of economic output can be attributed almost entirely to the positive impact on productivity of public policies to upskill workers and increase complementarity between human labour and machine learning technology. In this scenario, productivity is essentially flat, increasing over the period at a CAGR of 0.04%. While still disappointing, this represents a significant improvement over the base case, which assumes that UK productivity will decline.

**Scenario #2: Greater investment in technology and access to open source data**

Scenario 2 assumes that the UK is successfully able to increase the rate of capital investment in machine learning technology, and hence increase its diffusion and adoption, over and above the base case, thanks to the positive impact of public policy initiatives. The scenario assumes efforts will be focused primarily on two main policy interventions: creating new tax incentives to encourage private sector investment in machine learning technology, and increasing access to open source data by investing in national knowledge sharing communities; and through the government releasing more of its own data to the public.

Of all the scenarios, including the baseline, Scenario 2 results in the greatest positive impact for the UK economy. Pursuing policies that encourage greater capital investment in machine learning technology results in significant increases in both GDP and productivity over the base case, due to machine learning both being widespread and having a predominantly complementary impact on the workforce. The UK's GDP in this scenario grows at a CAGR of 1.94%, three times higher than the baseline rate of growth, rising to US$3.4trn in 2030. This represents an increase of US$600bn, or 21.4%, over the baseline estimate. Total hours worked are roughly the same as in the baseline in this scenario, implying that the increase in GDP is almost entirely a function of productivity, which increases at a CAGR of 0.68%, as opposed to the baseline estimate, in which it falls by 0.62%.

**Scenario #3: Insufficient policy support for structural changes in the economy**

Scenario 3 assumes that policy makers in the UK fail to step up to the economic challenges posed by the development of machine learning technology. Without the fillip of new tax incentives for private sector investment assumed in Scenario 2, machine learning technology continues to develop in line with the base case assumptions. The lack of government support for national data sharing, meanwhile, reduces the overall positive economic impact of machine learning, which requires data of both high quality and quantity to be of maximum usefulness. Most crucially, a lack of investment in upskilling the workforce in this scenario results in the substitution effect being stronger than the complementary effect, so that the dominant

© The Economist Intelligence Unit Limited 2017

impact of machine learning technology is to replace workers, rather than to increase their productivity.

The impact of this policy inaction on the UK economy is profoundly negative, and can be seen in both the GDP and productivity estimates produced by the model. GDP actually decreases by a CAGR of -1.20% in the UK over the forecast period in Scenario 3, resulting in a level of economic activity in 2030 that is US$400bn below 2016 levels, and US$700bn below the baseline estimate. Productivity, meanwhile, already expected to be negative in the base case, drops to a staggering -2.41% CAGR over the period. Total hours worked in 2030 remain higher than in 2016, but fall relative to the baseline estimate. Together these grim numbers imply that the UK will experience significant periods of economic recession and an overall drop in living standards over the course of the forecast period. Of all the countries in the analysis, the UK has the most to lose from poor public policy, with falls in GDP and productivity growth rates of 1.83% and 1.79% respectively.


© The Economist Intelligence Unit Limited 2017

## The US

### Baseline forecast

The baseline forecast assumes that economic conditions and public policy in the US correspond to the assumptions underlying our standard model. It sees the US economy expanding at a CAGR of 1.84% over the course of the forecast period, bringing US GDP to US$21.9trn in 2030, a nominal increase of US$5.3trn over the 2016 value. Over the same period, productivity in the US (measured in billions of dollars per million man-hours worked) is expected to increase from 0.059 to 0.073, a CAGR of 1.40%. The US' baseline estimates for growth in GDP and productivity reflect its status as a highly technologically mature economy with an ageing native population and moderate levels of immigration. Though not high by historical standards, they compare positively with other developed economies such as the UK, Australia and Japan.

### Scenario 1: Greater human productivity through upskilling

Scenario 1 assumes that the US takes positive policy action to address the posed to the labour market by machine learning, over and above the public policy actions assumed in the base case. By increasing access to higher education and job training, policy makers in Scenario 1 successfully increase the average skill level of the workforce, thereby ensuring that the complementary effect is stronger than the substitution effect, and that the dominant impact of machine learning technology is to enhance, rather than replace, human labour.

Under Scenario 1, the US economy expands at a CAGR of just over 2%, with US GDP reaching US$22.5trn by 2030, a figure which represents US$641bn in additional economic activity in 2030 over the baseline forecast. Since this scenario assumes the same number of hours worked in 2030 as in the base case, this higher level of economic

**Scenario results for the US**

| | | 2016 | 2030 (s) | 2030 (b) | Change v. '16 | Change v. baseline | CAGR 2016 - 30 |
|---|---|---|---|---|---|---|---|
| **Scenario #1** | GDP | $16.6trn | $22.5trn | $21.9trn | $5.8trn | $641.0bn | 2.04% |
| | Productivity | 0.059 | 0.075 | 0.073 | 0.016 | 0.002 | 1.59% |
| **Scenario #2** | GDP | $16.6trn | $25.9trn | $21.9trn | $9.3trn | $4.0trn | 3.00% |
| | Productivity | 0.059 | 0.086 | 0.073 | 0.027 | 0.014 | 2.56% |
| **Scenario #3** | GDP | $16.6trn | $18.8trn | $21.9trn | $2.6trn | $3.1trn | 0.84% |
| | Productivity | 0.059 | 0.063 | 0.073 | 0.04 | -0.010 | 0.40% |

Notes: All GDP figures in US dollars and rounded from official results
(s) = Scenario result
(b) = The Economist Intelligence Unit baseline forecast
Productivity = output/man hour worked

© The Economist Intelligence Unit Limited 2017

output can be attributed entirely to the positive impact on productivity of public policies to upskill workers and increase complementarity between human labour and machine learning technology. In this scenario, productivity increases from 0.059 to 0.075, a CAGR of 1.59%. At just under 3%, the increases in both GDP and productivity over the base case that the US experiences in Scenario 1 are relatively modest, and are the lowest of the six countries in this analysis. This may be due to relatively high starting levels of educational attainment and computer literacy in the US workforce.

**Scenario #2: Greater investment in technology and access to open source data**
Scenario 2 assumes that the US is successfully able to increase the rate of capital investment in machine learning technology, and hence increase its diffusion and adoption over and above the base case, thanks to the positive impact of public policy initiatives. The scenario assumes efforts will be focused primarily on two main policy interventions: creating new tax incentives to encourage private sector investment in machine learning technology,  increasing access to open source data by investing in national knowledge sharing communities; and through the government releasing more of its own data to the public.

Of all the scenarios, including the baseline, Scenario 2 results in the greatest positive impact for the US economy. Pursuing policies that encourage greater capital investment in machine learning technology results in significant increases in both GDP and productivity over the base case, due to machine learning both being widespread and having a predominantly complementary impact on the workforce. The US GDP in this scenario grows at a CAGR of 3%, more than 60% higher than the baseline rate of growth, rising to US$25.9trn in 2030. This represents an increase of US$9.3trn, or 18%, over the baseline estimate. Though total hours worked are modestly higher than the baseline in this scenario, the lion’s share of the increase in GDP comes from productivity, which increases at a CAGR of 2.56%, as opposed to 1.4% in the baseline estimate.

**Scenario #3: Insufficient policy support for structural changes in the economy**
Scenario 3 assumes that policy makers in the US fail to step up to the economic challenges posed by the development of machine learning technology, perhaps due to the political gridlock that plagues the US federal government. Without the fillip of new tax incentives for private sector investment assumed in Scenario 2, machine learning technology continues to develop in line with the base case assumptions. The lack of government support for national data sharing, meanwhile, reduces the overall positive economic impact of machine learning, which requires data of both high quality and quantity to be of maximum usefulness. Most crucially, a lack of investment in upskilling the workforce in this scenario results in the substitution effect being stronger than the complementary effect, so that the dominant impact of machine learning technology is to replace workers, rather than to increase their productivity.

The impact of this policy inaction on the US


© The Economist Intelligence Unit Limited 2017

economy is profoundly negative, and can be seen in both the GDP and productivity estimates produced by the model. While GDP still grows in the US over the forecast period in Scenario 3, it does so at a CAGR of 0.84%, less than half of the baseline estimate, resulting in 3.1trn fewer dollars of economic activity in 2030 than in the base case. Productivity, meanwhile, largely stagnates over the period, growing at a CAGR of 0.4%. Total hours worked in 2030 remain higher than in 2016, but fall relative to the baseline estimate.

© The Economist Intelligence Unit Limited 2017

# APPENDIX 2: METHODOLOGY

**Introduction**

Machine learning is a new technological promise. Policy makers' decisions in the next decade will determine the impact of machine learning on society. On one hand, proactive efforts to harness the power of machine learning are capable of driving shared growth and greater productivity. On the other, policy actions or inactions that fail to consider the macroeconomic relationship between humans and machines may amplify the negative effects of machine learning.

We identify three main levers available to policy makers: **upskilling, investment in computer capital** and **access to open data**. These policies have varying impacts on the substitutability and complementarity of machines and humans, which has a consequence on the impact on output and productivity. We developed a baseline and three scenarios to quantify the economic impact of machine learning on five countries (US, UK, Australia, Japan and South Korea) and developing Asia (as a grouping).

**The model**

Our estimates are based on a model developed by Hanson (2001). Hanson developed an exogenous growth model where machines can complement human labour when they are more productive at tasks and jobs they perform (complementing human labour), but they can also take over human jobs (substituting human labour). There are various factors such as computer prices, human labour and skills that can lead to complementarity or substitutability effects. An overview of the model is as follows:

$$\ln Y = \alpha + ß \ln K + \rho \ln (a^*H + M)$$

Where:
Y is output (GDP)
Y is output (GDP)
a is the human productivity advantage
β and ρ is are the output elasticities of capital and labour, respectively
K is other forms of capital such as plants and human capital
H is human labour
M is computer capital

**Scenario 1: Greater human productivity advantage through upskilling (optimistic scenario)**

In the first scenario, we modify our baseline scenario to accommodate a higher degree of complementarity between labour and machine learning than is presently anticipated. Greater upskilling through more tertiary education drives an increase in the human productivity advantage (a) over machine learning as compared to baseline levels. Changes in the regulatory environment result in more vocational education, where the curriculum has been realigned to complement machine learning. Access to finance for tertiary education also expands, enabling more people to reap the



© The Economist Intelligence Unit Limited 2017

returns to more schooling (the OECD estimates each additional year of education has a ROI of 3.7%). Under this scenario, the investment in computer capital continues to grow at the same rate as in the baseline scenario. Overall, both GDP and productivity were found to improve for all countries in the forecast period (2017-2030) relative to our baseline forecast.

**Scenario 2: Greater investment in technology and access to open source data (optimistic scenario)**

In this scenario, greater computer capital investment is the impetus for the broader proliferation of machine learning relative to our baseline forecast, assuming machine learning's complementary effect dominates its substitution effect. Changes in the regulatory environment enable investment in greater access to open source data to facilitate national knowledge communities. New tax credits for investment in computer capital drive greater private sector adoption of machine learning. Continued advances in computing efficiency drive down hardware costs, inviting further investment in machine learning software.  Computer capital grows at a country-specific incremental growth rate on top of our baseline forecast. As in the first scenario, both GDP and productivity improve for all countries in the forecast period (2017-2030) relative to our baseline forecast.

**Scenario 3: Disproportionately increasing economic investment in technology (negative scenario: substitution effect)**

Scenario 3 assumes machine learning's substitution effect dominates the complementary effect. In the forecast period, policy makers do not take action to develop the labour force beyond 2016 skill levels, nor do they take an active role in the development of national data sharing schemes. Computer capital progresses in line with our baseline forecasts, but labour skill stagnates. The human productivity advantage over machines falls, while the share of tasks machines can perform rises compared to baseline levels of both variables. Such an apathetic regulatory environment provides the ideal conditions for a dominant substitution effect. Machine learning substitutes labour as hours worked fall due to involuntary unemployment. Overall, under this scenario, GDP and productivity are significantly lower for all countries in the forecast period (2017-2030) relative to the baseline scenario.

© The Economist Intelligence Unit Limited 2017

© The Economist Intelligence Unit Limited 2017

© The Economist Intelligence Unit Limited 2017

**LONDON**
20 Cabot Square
London
E14 4QW
United Kingdom
Tel: (44.20) 7576 8000
Fax: (44.20) 7576 8500
E-mail: london@eiu.com

**NEW YORK**
750 Third Avenue
5th Floor
New York, NY 10017, US
Tel: (1.212) 554 0600
Fax: (1.212) 586 0248
E-mail: newyork@eiu.com

**HONG KONG**
1301 Cityplaza Four
12 Taikoo Wan Rd
Taikoo Shing
Hong Kong
Tel: (852) 2585 3888
Fax: (852) 2802 7638
E-mail: hongkong@eiu.com

**SINGAPORE**
8 Cross Street
#23-01 PWC Building
Singapore 048424
Tel: (65) 6534 5177
Fax: (65) 6428 2630
E-mail: singapore@eiu.com

**GENEVA**
Rue de l'Athénée 32
1206 Geneva Switzerland
Tel: (41) 22 566 2470
Fax: (41) 22 346 9347
E-mail: geneva@eiu.com



Whilst every effort has been taken to verify the accuracy of this information, neither The Economist Intelligence Unit Ltd. nor the sponsor of this report can accept any responsibility or liability for reliance by any person on this report or any of the information, opinions or conclusions set out herein.