forum
NEW ZEALAND

AI FORUM OF NEW ZEALAND DISCUSSION PAPER

The Potential Economic Impacts of AI

LITERATURE REVIEW

CONNECT

PROMOTE

ADVANCE

MEMBER OF THE

TECH ALLIANCE

Published: MAY 2018

aiforum.org.nz

AI Forum of New Zealand Discussion Paper

The Potential Economic Impacts of AI
Literature Review

May 2018

This Artificial Intelligence Forum of New Zealand (AIFNZ) discussion paper provides a detailed literature review of recent global research into the potential economic impacts of artificial intelligence (AI). This literature review was conducted to provide insights for the preparation of the AI Forum's 2018 research report, *Artificial Intelligence – Shaping a Future New Zealand*. A detailed examination of the potential impacts of AI on the New Zealand economy can be found in the report.

EXECUTIVE SUMMARY

The capability of AI to perform increasingly more tasks that were previously considered the sole preserve of humans and expected ongoing continuous improvement has a range of consequences. There is no shortage of reports discussing the potential economic impacts of AI at national and global level or by sector.

This discussion paper explores four key themes arising from a review of the literature on the economic impacts of AI:

- What will be the size of the overall economic impacts?
- Will this really happen any time soon?
- Will we notice?
- What should we do about it?

This literature review also examines the potential economic and labour market effects of the adoption of AI. The main effects are:

- the reduction in costs as computers do more work will generate increased productivity and growth rates.
- the greatest productivity gains are in capital intensive sectors like manufacturing and transport because their operational processes are susceptible to automation.
- although they disagree about the overall scale of the effect, most papers agree that low skilled workers are most at risk from AI because the potential to automate their jobs is higher than for the highly qualified.
- while some jobs may change, many will not be eliminated.
- the technical limitations of automation include perception and manipulation tasks in unstructured environments, creative and social intelligence.
- automation will also produce new jobs for those who create, program, maintain and support AI.

INTRODUCTION

AIFNZ undertook a literature review of AI's potential impacts on the economy, using research tools to isolate significant commentary, important academic and consulting papers. These are in addition to those referenced in the *Artificial Intelligence – Shaping a Future New Zealand* report. The literature is mostly from the last two years, reflecting AI's currency as a topic. Of the 58 papers reviewed, 54 were from 2016 and 2017. Hanson's is the only paper prior to 2013.

This literature review initially explores the potential economic effects of the adoption of AI before delving deeper into the potential labour market effects. This is followed by a brief analysis of potential productivity effects and policy implications.

THE ECONOMIC EFFECTS

The capability of AI to perform increasingly more tasks that were previously considered the sole preserve of humans, and its expected ongoing continuous improvement has a range of consequences.

Hanson (2001) models the wage, population and economic growth consequences of what he calls 'machine intelligence'. The model is best considered as an 'ideal' or 'pure' state, abstracted from reality but helpful to show the main economic effects.

Hanson et al uses a standard economic growth model, modified to introduce computers, which complement human workers and machine intelligence, which eventually replaces them. Machines complement human labour by making them more productive at the jobs they perform, but machines also substitute human labour by taking over human jobs.

The main effects are:

- at first, expensive hardware and software does only the few jobs where computers have the strongest advantage over humans. Eventually, computers do most jobs, due to falling costs and rising capability.
- regarding wages, complementary effects dominate, with human wages rising alongside computer productivity. However, substitution of humans by computers soon dominates and wages fall as fast as computer prices currently do.
- the huge reduction in costs as computers do more work, generates an enormous lift in economic growth with growth rates rising by ten times or more. This occurs because computer technology improves more quickly than technology in general, because machine intelligence becomes a central part of production and because the population of 'machine intelligences' can grow as fast as demand for them.

The overall income effects for humans need not be negative. In Hanson's model, computers are considerably cheaper than human labour, the world can produce the same output at much lower cost. The question becomes how to allocate resources if there is little work to do, or a question of what the displaced humans will do instead?

The model is built on exogenous growth, which means that the rate of technology improvement is determined by assumptions outside the model. Hanson et al shows that an endogenous growth model (where growth builds on growth and rates are determined by equations inside the model) show much bigger impacts. At the extreme they lead to the economic equivalent of the singularity where growth rates skyrocket as intelligent machines learn faster and work more.

Hanson et al does not consider any impacts of AI creating new jobs for humans. Nor does he look closely at the pace at which these effects might take place, or exactly which jobs the computers will do.

PwC (2017) also predicts significant economic impacts. It estimates that AI could add US$15.7 trillion a year to global GDP by 2030, equivalent to a 14 percent boost in today's output.

PWC et al says that $6.6 trillion of this stems from increased productivity due to businesses automating processes and augmenting the capabilities of their existing workforce and $9.1 trillion from an increase in consumer demand, from the availability of higher quality AI-enhanced products and services.

This view was built from a bottom-up analysis of the estimated impacts of the several hundred, "most compelling examples of potential AI applications," supplemented with the results from its General Equilibrium model of the global economy, with parameters adjusted to reflect its view of the impacts of AI on economic interactions.

Gains in business productivity will dominate in the first years of the forecast. As new technologies are adopted, the effects on consumer demand become more important. PwC et al notes that, "AI is still at a very early stage of development overall."

PwC et al recognises that the productivity impacts it expects are very large:

> *"The impact on productivity could be competitively transformative – businesses that fail to adapt and adopt could quickly find themselves undercut on turnaround times as well as costs. They stand to lose a significant amount of their market share as a result."*

The biggest productivity gains are in capital intensive sectors like manufacturing and transport because their operational processes are very susceptible to automation. This includes the productivity impacts of autonomous vehicles.

When it comes to later years and the impacts of new product development, healthcare, the automotive industries and financial services are the sectors with the greatest economic potential.

The impacts are especially large in China (a 26 percent boost to GDP) and North America (15 percent). In China, the importance of manufacturing in the economy is a big influence. In the United States and Canada, advanced technological progress makes implementing AI based improvements to production easier and sophisticated consumers are expected to drive adoption of AI improved products and services more quickly.

PwC et al does not say anything directly relevant to Australia or New Zealand. Oceania is combined with Africa and other Asian markets.

Using a similar method, PwC et al estimates that United Kingdom (UK) GDP could be 10 percent higher in 2030 as a result of AI. It describes AI as "one of the biggest commercial opportunities in today's fast changing economy." Productivity gains in the UK are expected to be quite small (two percent) but the impact of new products and new firms boosting demand is relatively higher (eight percent). This is partly because employment in the UK is less susceptible to automation. PwC et al expects a slightly higher positive impact in the UK, than in Europe because the UK has more technology companies and greater access to the skills required to develop AI technologies.

Also predicting large impacts, Accenture (2016) reports that AI could double the economic growth rate in 12 large developed countries by 2035, based on an economic model that treats AI as a new form of production different from labour and capital.

It recommends policy makers and business leaders treat AI, not as just another enhancer of productivity, but as a tool that can transform how economic growth is created.

The analysis is built on the Hanson et al model and supplemented by research from Frey and Osborne (2013) which identifies what occupations will be affected by AI. It relies on some assumptions about long-run employment, AI adoption and the capacity of countries to absorb AI technologies (which in turn depends on their ICT infrastructure, regulatory framework and investments in the digital economy).

In a further report, Accenture (2017) studied the impact of AI in 16 industries across the same 12 countries. It suggests that AI can boost business profitability by an average of 38 percent by 2035 and lift Gross Value Add (a measure of economic value similar to GDP) by 14 trillion. In addition:

- it expects the largest relative impacts in the ICT sector (from creating new products and services), manufacturing (from more advanced automation) and financial services (from automation of repetitive and mundane tasks).
- the biggest absolute impacts are in manufacturing, wholesale and retail trade, professional services and financial services, reflecting a combination of the impact of AI and the absolute size of those industries.

Chen (2016) takes a different approach and predicts much lower impacts, estimating the cumulative economic impact of AI from 2016 to 2026 as lying between $1.5 and $3 trillion (0.15 to 0.3 percent of global GDP). They generate four ranges for estimated impacts, using the results of previous studies to calculate the economic impacts of AI as if they were the same as computers in general, broadband internet, mobile phones or industrial robotics, and then triangulating to land on a central range. Because of limitations in the previous studies they rely on, some of the estimated economic impacts are for high income economies only.

They also estimate the economic impacts of AI based on the expected economic effects of assumed future industry and private capital investment in AI technologies. This leads to a markedly lower figure for economic impact of only around $500 billion. They present reasons why this method would only capture a portion of the total economic effects of AI. It is common to distinguish between the 'production' effects of technology (the economic impacts from growth in the firms that make the technology) and the 'use' effects (the economic impacts from firms using the technology). The use effects are much larger than the production effects in ICT, i.e. how firms outside the ICT sector use the technology is more significant for economic performance than how quickly firms in the ICT sector itself grow. This is consistent with the previously mentioned analysis. In fact, absolute investments in AI technologies are currently quite small, even if they are growing quickly.

IDC (2017) estimates worldwide revenue for AI systems of $13 billion in 2017, up 60 percent on 2016. Global spending is expected to continue to rise 50 percent a year. Other assessments are consistent with this view, that AI production is small but growing quickly.

THE LABOUR MARKET EFFECTS

There is quite extensive literature on whether, when and how AI will mean the end of human labour.

The Scale of Overall Change

The standard reference is Frey and Osborne et al which estimates that 47 percent of all people employed in the United States of America (USA) in 2010 were working in jobs that could be performed by computers, algorithms and robots in the next 20 years or so. The assessment is built from expert views of the potential future capability of automation that were used to classify 702 occupations in the USA.

The results show a bi-modal distribution, with 47 percent of workers at high risk (more than 70 percent chance of being replaced by a machine) and 33 percent at low risk (under 30 percent chance of replacement), with only 19 percent at medium risk (above 30 percent, but below 70 percent chance of replacement). Frey and Osborne et al view this risk schedule as a timetable with about half of jobs likely to disappear early, followed by a relative plateau while technology continues to improve and then the last third of jobs being replaced some time later.

Several follow up studies applied the risk of automation at the level of occupations to other countries. Compared with the 47 percent figure for the USA, Pajarinen and Rouvinen et al report 35 percent in Finland, Brzeski and Burke et al suggest 59 percent in Germany and Bowles estimates between 45 and more than 60 percent in Europe, with southern European countries having the highest risk.

NZIER (2015) applies Frey and Osborne's estimates of job automation to New Zealand employment data. It finds that 46 percent of the labour force faces a high risk of having their current job automated, similar to Frey and Osborne's estimate for the United States. Part of what generates these large numbers is the methodology used.

Arntz (2016) takes a task based view rather than one based on occupations and comes up with much lower numbers, estimating that, across 21 OECD countries, only nine percent of tasks can be readily automated, with the proportion depending on the structure of the economy from six percent in South Korea to 12 percent in Austria. Arntz et al did not include estimates for New Zealand or Australia.

Oschinski and Wyonch (2016) analyses the occupations at risk in the Canadian labour market. The authors build on Frey and Osborne's method but use updated information on the capabilities of AI and weigh each job that has AI replaceable skills, with the importance of that skill for that job.

They conclude that, although there are many industries with large proportions of employment at high risk of automation, in practice those industries are relatively small employers because they are those that are already highly automated. Only two percent of total Canadian jobs are in industries where more than 75 percent of jobs are at high risk from automation. Nearly 30 percent of jobs are in industries where less than a quarter of employees are at high risk.

The reason for the difference in figures between Arntz, Oschinski and Wyonch, and Frey and Osborne et al is that many jobs include some tasks that computers will struggle to do anytime soon. While these jobs may change, they will not be eliminated by machine intelligence and robotics.

An analysis by Berriman (2017) is an attempt to reconcile Frey and Osborne and Arntz et al. The author essentially concludes that Arntz et al has 'over-egged the pudding'. A task based analysis does reduce the effects markedly from those estimated by Frey and Osborne et al; there are still a great many jobs at risk. They conclude that up to 30 percent of UK jobs could be at high risk of automation by the early 2030s, lower than the USA (38 percent) or Germany (35 percent), but higher than Japan (21 percent).

The Workers Most at Risk from Change

Although they disagree about the overall scale of the effect, most papers agree that low skilled workers are most at risk from AI because the potential to automate their jobs is higher than for the highly qualified. For example:

- Arntz et al say the proportion of jobs at risk for those with only primary education is over 50 percent, compared with less than one percent for those with tertiary qualifications.
- Berriman et al concur that the key difference is education. For those with just GCSE level (equivalent to New Zealand's NCEA Level 1) education or lower, the estimated potential risk of automation is as high as 46 percent in the UK, but this falls to only around 12 percent for those with undergraduate degrees or higher.
- Carey (2016) says the employment effects of automation from AI weigh most heavily on industries that historically were already a focus of concern about job losses due to technological change.
- To those industries of historic focus, we can add some new industries as a result of the new developing capabilities of AI. The Executive Office of the President (2016) estimates that two to three million driving jobs in the USA from a total of 3.7m are threatened by automation. In context the total USA labour force is 160m.

Berriman et al says the risks are highest in transportation and storage (56 percent of jobs at high risk), manufacturing (46 percent), wholesale and retail trade (44 percent), and administration and support services (38 percent). Together, these industries account for over half of the potential job losses. Risks are lower in other large sectors like health and social work (17 percent) or the education sector (nine percent).

In local research, NZIER et al concluded the risk of automation is higher for low skilled jobs than for high skilled jobs. They conclude that 75 percent of labouring jobs are at high risk of automation, but only 12 percent of jobs in professional services. Rural areas have a higher proportion of the industries most affected. Men are also over represented in affected industries.

The Limits to Automation

According to Frey and Osborne et al, the technical limits of computerisation relate to:

- perception and manipulation tasks, especially in unstructured environments.
- tasks that involve creative intelligence.
- work requiring social intelligence, essentially working with and influencing others.

One of the reasons why occupation based assessments of jobs at risk give much higher values than task based assessments is because many jobs involve some amount of working with others. Frey and Osborne et al suggest 98 percent of book-keeping, accounting and auditing clerk jobs are automatable but Arntz et al report that only 24 percent of employees in these roles can do their jobs without interactions with others.

Chui (2016) distinguish three groups of occupational activities by their susceptibility to automation:

- highly susceptible tasks include data collection, data processing, and predictable physical work.
- less susceptible tasks include interactions among stakeholders and unpredictable physical work.
- and least susceptible are managing others and applying expertise.

Oschinski and Wyonch et al includes the following table of skills that challenge AI, therefore operating as barriers to automation.

Table 1– Skills that are hard for AI

Barrier to automation	Definition
Social perception	Being aware of others' reactions and understanding why they react as they do.
Originality	The ability to come up with unusual or clever ideas about a given topic or situation, or to develop creative ways to solve a problem.
Assisting others	Providing personal assistance, medical attention, emotional support or other personal care to others such as co-workers, customers or patients.
Philosophy	Knowledge of different philosophical systems and religions, including their basic principles, values, ethics, ways of thinking, customs, practices and their impact on human culture.
Initiative	Job requires willingness to take on responsibilities and challenges.
Leadership	Job requires willingness to lead, take charge and offer opinions and direction.
Innovation	Job requires creativity and alternative thinking to develop new ideas for and answers to work related problems.
Adaptability and flexibility	Job requires openness to change (positive or negative) and to consider variety in the workplace.
Independence	Job requires developing one's own way of doing things, guiding oneself with little or no supervision and depending on oneself to get things done.

Source: Oschinski, 2017

Other Factors Affecting Uptake

There is an enormous range of economic, legal and regulatory factors that will also affect the extent to which AI replaces human labour. There is uncertainty about the timing of the technical development of AI but also the extent to which other barriers will be overcome. None of the studies reviewed model, for example, the cost of developing and introducing AI into the workplace, the practical issues, or the economic costs and benefits of changes to regulatory approaches that might be required.

Arntz et al conclude that the share of jobs lost will be lower than the share at risk because of the slow speed of technological substitution, because workers will switch to other tasks as some of their work becomes automated and because automation will create new jobs.

McKinsey Global Institute (2017) says nearly a third of tasks in sixty percent of jobs will be automated, but less than five percent of jobs can be completely replaced. They do not expect wide adoption of new working methods until 2030 to 2050, meaning a relatively slow process of change. The limits are related to the economic, organisational and social factors.

> *"Our scenarios suggest that half of today's work activities could be automated by 2055, but this could happen up to 20 years earlier or later depending on the various factors, in addition to other wider economic conditions."*

RSA (2017) cites a range of reasons to think that widespread replacement of low skilled workers in the UK is unlikely including technical limitations to AI and robotics, the fact that most automation is of partial tasks rather than complete jobs (meaning jobs will change, but not be eliminated), the fact that AI and robotics will also create some jobs and complement and empower some existing workers and that rising productivity from AI will lower prices and generate demand for other products.

New Jobs and Other Upsides

Some studies look at the benefits of AI for workers whose skills it complements.

"In many cases, machines both substitute for and complement human labour - substituting for workers in routine, codifiable tasks, while amplifying the comparative advantage of workers in problem solving skills, adaptability and creativity. IT, then, should raise earnings in occupations that make intensive use of abstract tasks and among workers who intensively supply them," states Autor (2016).

Automation will also create some jobs for those who create, programme, maintain and support AI and through more general economic effects generate additional wealth and spending that will create more un-automatable jobs.

Agrawal (2016) suggest the core value of 'machine intelligence' is in the dramatic cost reduction of prediction. This will mean that prediction will be used much more, for example, autonomous driving. Human prediction skills will be much less valuable, but that complementary skills (like judgement and choices of action that rest with humans) will become more important.

An Overall View

The aggregate effects of all this labour market change are hard to fathom at this point. It will be a long time before anything becomes evident from backward facing statistics. The World Economic Forum (2016) surveyed 371 leaders in Human Resources (HR) and Strategy at large employers in nine broad industry sectors and asked them what they thought the top technological trend was facing their organisation. Only seven percent of respondents nominated 'Artificial intelligence and machine learning', as compared with 34 percent nominating 'mobile Internet and cloud technology'.

Aris (2017) summarises the overall position as being the difference between optimism and pessimism, with optimists believing that higher labour productivity will increase the demand for work based on the experience of digitisation. Conversely, pessimists believe there will be a surplus supply of labour and falling wages, with economic benefits accruing only to those who control the data, algorithms and intellectual property rights that underlie the rise of the machines.

Ford (2016), for example, is a pessimist believing that humans will largely be replaced by machines. Brynholfsson and McAfee (2014) are more optimistic, focusing on the positive combination of humans and machines and the ability of machines to make human work less arduous, dangerous and demeaning.

Chen et al are reassured by the fact that previous innovations in production have often thought to be 'injurious to the interests of the class of labourers' as economist David Ricardo wrote in the time of the

Industrial Revolution, but the net effect of technological advancement in the last three hundred years has not been a reduction in long term employment but instead a general growth in prosperity.

Similarly, Berriman et al suggests that we are not in an entirely new world with the historical evidence suggesting that AI induced technical change will eventually lead to:

- broadly similar overall rates of employment for human workers, although with different distributions across industry sectors and types of jobs than now.
- higher average real income levels across the country as a whole due to higher overall productivity.
- quite possibly a more skewed income distribution with a greater proportion going to those with the skills to thrive in an ever more digital economy.

The United States National Science and Technology Council et al put the figures for jobs at risk in the USA in the context of the whole economy:

> *"For context, every three months about six percent of jobs in the economy are destroyed by shrinking or closing businesses, while a slightly larger percentage of jobs are added— resulting in rising employment and a roughly constant unemployment rate. The [United States] economy has repeatedly proven itself capable of handling this scale of change, although it would depend on how rapidly the changes happen and how concentrated the losses are in specific occupations that are hard to shift from."*

They also draw a comparison with changes throughout the 20th century in economic activity as machines took over from humans in the agriculture sector:

> *"The anticipated shift in the activities in the labor force is of a similar order of magnitude as the long-term shift away from agriculture and decreases in manufacturing share of employment in the United States, both of which were accompanied by the creation of new types of work not foreseen at the time."*
>
> *"For example, in 1870, almost 50 percent of American employees worked in agriculture, supplying the Nation's food. Today, thanks in large part to technological change, agriculture employs less than two percent of American workers and American food production exceeds domestic demand."*

This does not mean that labour market change would be easy for those affected by them, even if in aggregate the changes are manageable for the economy or for workers overall. They go on to note the effects of job loss:

> *"In recent decades, United States workers who were displaced from their jobs experienced substantial earnings declines. Negative shocks to local economies can have substantial negative and long-lasting effects on unemployment, labor force participation, and wages. Perhaps more significantly, over time, displaced workers' earnings recover only slowly and incompletely. Even ten or more years later, the earnings of these workers remain depressed by 10 percent or more relative to their previous wages. These results suggest that for many displaced workers there appears to be a deterioration in their ability either to match their current skills to, or retrain for, new, in-demand jobs. AI driven automation can act - and in some cases has already acted - as a shock to local labor markets that can initiate long-standing disruptions."*

Significant differences in how change would affect different groups are also noted:

> *"CEA ranked occupations by wages and found that, according to the Frey and Osbourne analysis, 83 percent of jobs making less than $20 per hour would come under pressure from automation, as compared to 31 percent of jobs making between $20 and $40 per hour and four percent of jobs making above $40 per hour."*

Carey et al notes that workers with disabilities and solo parents have particularly high proportions of the low-skilled who will be disproportionately affected.

The RSA et al research takes a different tack, arguing there is no reason at all to bemoan the loss of low skilled, poorly paid and low quality jobs. Writing for the UK, they point out that a third of 'elementary' workers (including waiters and cleaners) have household incomes below the poverty line, so any technology that could improve wages and conditions in those jobs should be welcomed.

THE PRODUCTIVITY EFFECTS

There were two specific studies regarding productivity:

- Accenture et al estimates that AI could boost labour productivity between 11 percent (Spain) and 37 percent (Sweden) by 2035. The difference is driven by the ability of countries to diffuse technological innovation into their wider economic infrastructure.
- McKinsey Global Institute et al suggests automation could raise productivity growth globally by between 0.8 and 1.4 percent a year. This may not sound like much, but New Zealand's average productivity growth in the last forty years is around 1.5 percent a year. These benefits come about from reduced errors and improved quality and speed in production, as well as doing things that humans cannot do easily, and expanding what humans can do at all in some cases.

CEBR (2015) looks at the productivity effects of ICT as a way to get at the potential effects of AI. CEBR et al suggests that ICT has grown from accounting for three percent of labour productivity in office jobs in the 1970s to 12 percent in 2012. At the same time, the sectors they study (professional services, finance and insurance, information and communications, public administration) have grown more quickly from 24 percent of Gross Value Add in 1972 to 43 percent in 2012, and in 2012 accounted for 30 percent of employment in the UK. ICT related labour productivity growth is growing five times faster than overall labour productivity growth.

If computers and the internet have taken so long to have such modest impacts, we would need to isolate something different about AI that means the effects will be larger and will happen more quickly. More generally, predicting big productivity gains from technology is a risky endeavour. Gordon (2016), in a well known book, argues that the advances of the ICT revolution are not as transformative as those of the Industrial Revolution and that they should not be expected to improve productivity in the same way. His view is that the significant jump in productivity during the 1940s was a onetime event.

As Rotman (2016) puts it:

> *"Given impressive advances in artificial intelligence, smart robots, and driverless cars, it's easy to become convinced that we are on the verge of a new technological age. But the troubling reality is that today's advances are having a far from impressive impact on overall economic growth."*

Saniee (2016) suggest we don't yet have enough adoption of the core technologies. They expect there will be a significant productivity jump in the USA around 2030, when the constituent technologies reach more than half of the population. They isolate five core technologies for digital revolution and use proxies for the uptake of each to reflect adoption. The first four of the five technologies are the digital equivalents of the core inventions of the Industrial Revolution.

1. Digital energy networks. Connecting and controlling new power sources, moving them closer to consumption needs and minimizing waste (proxies are: smart electricity meter adoption, renewable energy production).
2. Digital health and sanitation networks. Connecting human health systems and moving diagnosis and treatment to the optimum locations (proxy: household adoption of digital health measurement systems).
3. Digital transportation networks. Connecting and moving people and goods autonomously, improving safety and efficiency (proxy: autonomous vehicle registrations as a proportion of total vehicles).

4. Digital communication networks. Connecting people and systems and moving data and knowledge among them (proxies: smartphone penetration, AI assistant adoption, e-commerce).
5. Digital production networks. Creating local industrial ecosystems to produce and deliver contextual goods and services more efficiently (proxies: 3D printed goods as a proportion of total goods manufactured, establishment of data centres).

Finally, Aghion et al (2017) wonder if in fact the impacts for long run economic activity from AI might end up being modest in any case. They mathematically consider the implications of AI for economic growth in general, treating AI as a new form of automation that offers the possibility of automating non-routine tasks (like driving) or tasks that require high levels of skill (like radiology). They explore the constraint on the economic impacts of AI from Baumol's et al 'cost disease': the limits on overall economic growth that are exposed because sectors with fast growing productivity tend to shrink as a proportion of economy compared with those with slower growing productivity, and conclude that we might end up with a large proportion of tasks automated, but with economic growth being held back by slow growth in areas that cannot be taken over by AI.

THE POLICY IMPLICATIONS

There is relatively limited commentary on formulating policies as this is an emerging technology.

Societal Level Policy

The U.S. National Science and Technology Council et al provides a helpful overview:

> *"Although it is difficult to predict these economic effects precisely with a high degree of confidence, the economic analysis... suggests that policymakers should prepare for five primary economic effects:*
>
> - *positive contributions to aggregate productivity growth.*
> - *changes in the skills demanded by the job market, including greater demand for higher level technical skills.*
> - *uneven distribution of impact, across sectors, wage levels, education levels, job types and locations.*
> - *churning of the job market as some jobs disappear while others are created.*
> - *the loss of jobs for some workers in the short-run, and possibly longer depending on policy responses."*

They are generally optimistic about the effects overall:

> *"Given appropriate attention and the right policy and institutional responses, advanced automation can be compatible with productivity, high levels of employment, and more broadly shared prosperity."*

They also make various policy proposals for investment in and development of AI, to assist affected workers in the transition, to prioritise diversity in STEM and AI in particular, 'to address potential barriers from algorithmic bias.'

Aris et al describes the main societal challenges as making sure that:

- the benefits of digital productivity growth are fairly shared.
- the transition in work is managed as smoothly as possible.
- useful new work is created for those who are displaced.

Accenture et al provides general direction for countries embracing AI including:

- reconsider education to include teaching machines and learning from machines (rather than just learning how to use machines).
- remove laws that might impede the use or adoption of AI enabled societal improvements or putting in place new regulatory frameworks where that is required.
- develop a (or several) code(s) of ethics for AI developers.
- talk up the benefits of AI to avoid too much focus on the economic downsides for humans and work to actively address and pre-empt the downsides of AI for those negatively affected.

RSA et al is also enthusiastic for countries to embrace AI and make active choices for technology that can improve job availability and quality, raise incomes, eliminate prejudice in recruitment, enable gig work (that enables choices, rather than undermining existing labour market protections), and improve living standards in general. Amongst a set of suggested proposals, they encourage venture capitalists and non-profits to invest in benevolent technology that enriches the worker experience, shift the burden of taxation away from labour and towards capital (in the expectation that labour will become less important over time), and draft a blueprint for a UK sovereign wealth fund that would give every citizen a 'technological inheritance' to distribute the economic benefits from moving to embrace AI.

Firm Level Policy

For firms, Accenture et al makes several suggestions, for making the most of AI, including:

- building AI into strategy and leadership.
- integrate oversight of AI's virtual labour into HR and bring AI into workplace training to enable people to use it to learn and share their skills.
- expand the role of the Chief Data Officer beyond security, regulation and governance to include providing data as a crucial part of data supply internally.
- create an open AI culture that is inclusive of new AI employees.
- make automation more intelligent rather than just labour saving.
- measure the impact of AI investments over time.

On the regulatory side, PWC et al expects strong network effects for those firms creating AI technologies. This means that firms that establish an early market advantage in AI enabled services will strengthen their position over time as higher adoption generates more data to use in improving services. This is a virtuous cycle from the firm and country point of view, but one that could lead to monopolisation concerns at an industry level.

Individual Level Policy

Markley-Tower (2017) recommends individuals begin learning more general knowledge and developing their creativity to protect themselves from labour market obsolescence. Similarly, Oschinski et al focuses on teaching critical reasoning and interpersonal skills and an entrepreneurial mindset from an early age so that, by the time students enter the workforce, they will have the necessary skills to succeed, together with increased opportunities for lifelong learning for those already in the workforce.

CONCLUSION

This literature review was prepared by the Artificial Intelligence Forum of New Zealand to provide insights for the preparation of the report, *Artificial Intelligence - Shaping a Future New Zealand (2018).* This paper is a supporting document only. A detailed examination of the potential impacts of AI on the New Zealand economy can be found in the report.

REFERENCES

Accenture (2016), Why Artificial Intelligence is the Future of Growth.

Accenture (2017), How AI Boosts Industry Profits and Innovation.

Agrawal, Aja, Joshua Gans and Avi Goldfarb (2016), The Simple Economics of Machine Intelligence, Harvard Business Review.

Aris (2017), How Disruptive will Automation Be?

Artnz, Melanie, Terry Gregory and Ulrich Zierahn (2016), The Risk of Automation for Jobs in OECD Countries: A Comparative Analysis, OECD Social, Employment and Migration Working Papers, No 189.

Autor, David (2016), The Shifts - Great and Small - in Workplace Automation Frontiers, MIT Sloan Management Review.

Berriman, Richard (2017), Will Robots Steal our Jobs? The Potential Impact of Automation on the UK and Other Major Economies.

Bowles (2014), The Computerisation of European Jobs, Bruegel, July 2014.

Brzeski and Burke (2015), cited in Arntz et al.

Brynholfsson and McAfee (2014), The Second Machine Age.

Carey, David (2016), Adapting to the Changing Labour Market in New Zealand, OECD Economics Department Working Papers, No 1420.

CEBR (2015), O2 Individual Productivity Index: How ICT Technology Drives the UK Office Economy.

Chen, Nicholas, Lau Christensen, Kevin Gallagher, Rosamond Mate, Greg Rafert (2016), Global Economic Impacts Associated with Artificial Intelligence.

Chui, Michael, James Manyika and Mehdi Miremadi (2016), Where Machines Could Replace Humans - and Where They Can't (yet), McKinsey Quarterly.

Executive Office of the President (2016), Artificial Intelligence, Automation and the Economy, National Science and Technology Council, United States. December 2016.

Ford, Martin (2016), Rise of the Robots: Technology and the Threat of a Jobless Future

Frey, Carl Benedikt and Michael A Osborne (2013), The Future of Employment: How Susceptible are Jobs to Computerisation.

Gordon (2016, The Rise and Fall of American Growth: The U.S Standard of Living since the Civil War.

Hanson (2001), Economic Growth Given Machine Intelligence.

IDC (2017), Worldwide Spending on Cognitive and Artificial Intelligence Systems Forecast to Reach $12.5 Billion This Year, New IDC Spending Guide.

McKinsey Global Institute (2017), A Future That Works: Automation, Employment and Productivity, January 2017.

McKinsey Global Institute (2017), Artificial Intelligence, the Next Digital Frontier? June 2017.

NZIER (2015), Robot Nation? The Impact of Disruptive Technologies on Kiwis, NZIER Insight, 55-2015.

Pajarinen, Mika and Petri Rouvinen (2014), Computerization Threatens one Third of Finnish Employment.

PWC (2017), Sizing the Prize, What's the Real Value of AI for your Business and how can you Capitalise?

PWC (2017), The Economic Impact of Artificial Intelligence on the UK Economy.

Rotman, David (2016), Dear Silicon Valley: Forget Flying Cars, Give us Economic Growth.

RSA (2017), The Age of Automation: Artificial Intelligence, Robotics and the Future of Low-Skilled Work, September 2017.

Saniee, Iraj, Sanjay Kamat, Subra Prakash and Marcus Weldon (2017), Will Productivity Growth Return in the new Digital Era? An Analysis of the Potential Impact on Productivity of the Fourth Industrial Revolution.

World Economic Forum (2016), The Future of Jobs: Employment, Skills and Workforce Strategy for the Fourth Industrial Revolution, January 2016.

This literature review was built on research conducted by Aaron Schiff (independent economist) and Hayden Glass of Sapere (an independent economic consultancy), on behalf of AIFNZ.

Learn more about the AI Forum of New Zealand at www.aiforum.org.nz

The Artificial Intelligence Forum of New Zealand (AIFNZ) sets out to raise the level of awareness and capabilities of Artificial Intelligence (AI) in New Zealand. The rapid development of AI technologies presents major opportunities and challenges for our country: from creating world leading AI businesses, nurturing a pool of talented AI engineers, applying AI technologies to our agriculture, Government, manufacturing and service industries to holding a meaningful national debate on the broader implications for society, New Zealand needs to actively engage with AI now in order to secure our future prosperity. The Forum brings together citizens, business, academia and the Government, connecting, promoting and advancing the AI ecosystem to help ensure a prosperous New Zealand.

AI Forum of New Zealand
L1 Building C, 14-22 Triton Drive, Auckland 0632, New Zealand
Ph +64 9 394 7693
www.aiforum.org.nz