

A Report to the Senate Select Committee on the Future of Work and Workers

The Luddite Fallacy Fallacy

About the Authors

James Newton-Thomas - M.App.Sci: Researcher in Artificial Intelligence, for over 30 years. For the last 20 years has been developing automation systems for the mining industry.

Jack Strocchi - B. Comm, Economic consultant, Transport and Construction.

Abstract

This submission addresses the 1st, 3rd, 4th and 5th of the Terms of Reference given although the conclusions drawn are relevant to all. Our goal is to appraise the committee of the very real danger of underestimating the impact that the widespread take up Intelligent Machines (IM) is going to have on the future of work.¹

We approach the problem from first principles, founding it on the theory of computation via Turing, Von-Neumann and Shannon conjoined with the economic theories of Weber and Keynes. From this we derive the practice and application of the automation of capital and the commodification of labour. Our aim is to lay out the worst-case scenario for government as a provider of insurance. Unfortunately, our conclusion is that this worst case scenario is also the most likely given government inaction and the poor advice and projections they have been receiving. In general the economists don't understand the technology and the technologists don't understand the economy. In consequence, the medium term future for many Australians is dire. By this prediction we fall into the class of techno-alarmists, though not Luddites. We state that it is only through recognising and embracing these new technologies and directing them to the common wealth can we provide a prosperous and equitable future for all Australians.

¹ We use the term Intelligent Machine to refer to machines which perform functions that would be deemed to be intelligent if a human were to do them. This includes robots and computers. At some points we use the term Artificial Intelligence (AI) interchangeably.

About the Authors	1
Abstract	1
 TOR#1 - The future earnings, job security, employment status and working patterns of Australians;	 3
Introduction	3
Biology vs Technology	4
The processes of Automation	5
Work	6
The economics of the substitution of labour by intelligent machines	8
The Rise and Fall of the Contingent Workforce	10
Task replacement vs Process Re-engineering	10
Automation Step Change	11
The Tech Giants	11
 TOR#3. the wider effects of that change on inequality, the economy, government and society;	 13
Introduction	13
Polarisation of Income Distribution	13
Fluctuations in the rate of economic growth	13
 TOR#4. the adequacy of Australia's laws, including industrial relations laws and regulations, policies and institutions to prepare Australians for that change;	 15
Introduction	15
1. Throw away the Economists Rear-view mirror	15
2. Reduce Privatisation	16
3. Develop a Transaction Tax on the IoT	16
4. Develop Policies Concerning AI Interaction	17
5. Develop Policies for Data Ownership and Data Rental Rights	17
 Conclusion	 17

TOR#1 - The future earnings, job security, employment status and working patterns of Australians;

Introduction

There is no future of work for humans anywhere. Therefore job security for all workers, from manual to managerial, will be nil and no humans will be competitively employable. Their eventual employment status will be “unemployed”. The pattern of work will be unpaid only. The only economic classification that will survive to provide earnings will be the owners of productive assets: land, raw materials and capital goods.

Capitalism, as the name denotes, is the maximizing the value of the owner's capital stock which is determined via rational cost-benefit calculus. This is facilitated at all levels by computation and which in turn leads to commodification of all marketable goods, including labour and capital. In Australia, in 2015, “labour costs account for more than 60% of corporate expenses”². This provides an obvious target and plenty of scope for economic rationalisation.

That companies will pursue this economic rationale is inevitable. Max Weber³, writing around the turn of the 20th Century, described:

The fate of our times is characterized by rationalization and intellectualization and, above all, by the disenchantment of the world... it means that there are no mysterious incalculable forces that come into play, but rather that one can, in principle, master all things by calculation.

Already now, throughout private enterprise in wholesale manufacture, as well as in all other economic enterprises run on modern lines... rational calculation, is manifest at every stage. By it, the performance of each individual worker is mathematically measured, each man becomes a little cog in the machine,

In the short to medium term, work, as we know it, will be subjected to two major industrial disruptions affecting both the tools we use and the teams we join. These are, firstly, the automation of capital, which will progress from human-operated dumb machines to autonomous intelligent machines. And secondly, the commodification of labour which will degrade from an establishment workforce, contracted under the Standard Employment Relation (SER) and occupying multiskilled jobs, to an increasingly contingent workforce, contracted to third-party agents and performing specialised tasks.

Therefore the future of work largely depends on:

- how far and fast industry can pursue the computational rationalisation of itself through reengineering both machine products and work processes, and

² Tobias Levkovich of Citi, quoted in Sam Ro [DEC 16, 2015], “Labour accounts for 60% of corporate expenses, and it's going up” <https://www.businessinsider.com.au/citi-levkovich-rising-labor-costs-pressure-profit-margins-sp600-2016-2015-12>

³ Weber developed these ideas in the first decade of the 20th century, but the quotes are taken from an anthology. Weber, M. “In his own words” <http://www.faculty.rsu.edu/users/f/felwell/www/Theorists/Weber/words.htm>

- how long and in what fields humans can maintain economic relevance.

Biology vs Technology

What makes the replacement of mental processes with intelligent machines inevitable is the bounded limits of human biology measured against the boundless potential of machine technology. A human brain can only think as fast as neuro-biology allows, and the human body is constrained by bio-mechanics. An iso-surface representing human capabilities is finite and, whilst we can expand it with better training and tools, there is still a fundamental limit which is not present in networked intelligent machines.

Any attempt to integrate us with machines will eventually run into the problem that we, as defined by our biology, must always remain the weakest link in any such synthesis and thus will be outcompeted by the non-integrated process⁴. In essence, both our internal and external communications are too slow to remain indefinitely competitive. Fibre optic communication is typically 10^9 bit/s compared to the corpus callosum nerve bundle which has a theoretic maximum of 10^{10} bits/s or just on 10 times faster than a single fibre. Nerves however have a much higher signal latency because their signals travel slower. They take a median of 37ms for cross brain communication in an adult human. Fibres are laid in bundles typically more than 10, therefore, we can get information from Canberra to Sydney and back, faster than we can get signals from the left hemisphere to the right hemisphere of our own brain. If the fruits of computation are what we want and will pay for, then there is a lot more scope for it outside of our heads than inside.⁵

It is an “anthropocentric conceit” to imagine there is something special about the human mind, apart from perhaps the hard problem of consciousness. *All* computation is pattern translation and can be reduced to mathematical algorithms and solved, given sufficient computational resources. All mechanical processes can be described mathematically. Therefore *all* industrial work is pattern transformation and is thus computable.⁶

There is no intrinsic reason why intelligent machines cannot replace *all* paid human labour. Sheer scale, speed and connectedness, more than compensate for any current biological efficiency.⁷ All three of these characteristics are subject to exponential improvements and increase the ability to direct computational resources at any specific task.

⁴ Recent press articles quoting Sebastian Thurn, founder of Google X, suggest that AI will enhance us as “super human workers”, accomplishing repetitive tasks whilst leaving humans free to be creative workers. This is not going to happen as humans are the weakest link in the technological chain and if critical will just slow down intelligent machine task execution thus making it less competitive. Steam engines actually did make some workers superhuman 200 years ago.

⁵Data from: Phillips, Kimberley & D. Stimpson, Cheryl & Smaers, Jeroen & Raghanti, Mary Ann & Jacobs, Bob & Popratiloff, Anastas & Hof, Patrick & Sherwood, Chet. (2015). The corpus callosum in primates: Processing speed of axons and the evolution of hemispheric asymmetry. *Proceedings of the Royal Society B: Biological Sciences*. 282. . 10.1098/rspb.2015.1535.

⁶ Alan Turing [1912-1954] showed that essentially, all computation consists of turning patterns into other patterns, and which could be performed by a class of theoretical machine, now called a Universal Turing Machine. Modern computers are homologues of this machine and thus with the the appropriate operational steps and memory, can perform any task we can do.

⁷ Currently, humans still have a general computational edge, as at 2017 our brains are about 100,000 times more energy efficient at computation, pound for pound than computers. Modern supercomputers have about the same computational capability measured in FLOPS as the theoretical maximum of the human brain.

Early AI systems were originally programmed in a God-like, top-down, fashion, to be good at things humans were bad at. This turned out to be surprisingly difficult to do for problems that we found easy and easy to do for problems we found hard⁸. Modern Intelligent Machines now learn in a more human, bottom-up, way. They learn from experience by observing data and through this they are rapidly becoming good at tasks we are good at as a side effect to what we are tasking them for. Eventually they will exceed us in all tasks. The triumph of the latest machine learning paradigms, such as deep learning, represents a technological step change that has dramatically accelerated the process.

As of 2017 intelligent machines are *not* fully substitutable for humans. We know they are similarly computationally capable⁹ however they have not yet learned or been programmed to perform an equivalent array of tasks. Intelligent Machines are definitely substitutable to *some* tasks as is evidenced by their deployment. Given that it takes about a minimum of 10 years to train a human mind to be economically useful and further considering the massive data and social engagement available on the internet, it is likely that a similarly architected neural net can be given equivalent exposure over that time frame and probably much less. Therefore we project 10 years for a fully human substitutable machine mind. This does not mean that all tasks *will* be substituted in this time, rather that they *could* be if there is sufficient economic incentive to do so.

The processes of Automation

There are three exponentially growing processes that have facilitated the phenomenal roll-out of computer technology globally. Whilst described as 'laws' they are actually just observed phenomena with predictive value. The first is the eponymous Moore's Law¹⁰ which is concerned with the computational *speed* of the chips that are the building blocks of modern computers¹¹.

The second is Metcalfe's Law¹² which is concerned with the *scale* of computational nodes in the network. It predicts that the value of a network rises in proportion to the square of the number of additional users that connect to that network. This is the key behind the growth in the value of companies like Facebook which rely on network externalities.

⁸ Attributed to Hans Moravec in 1988, it came to be known as Moravec's Paradox https://en.wikipedia.org/wiki/Moravec%27s_paradox and played no small part in the two AI winters endured in the latter part of the 20th century.

⁹ The human brain is capable of performing somewhere between 100 Peta - 1 Exo flops. The current supercomputer record holder is the Sunway TaihuLight at the National Supercomputing Center in Wuxi China. It is capable of 125 Peta Flops or slightly over the lower estimate. At current rates of progress a supercomputer will be capable of performing computation in excess of 1 Exaflop (the human brain theoretical maximum) by 2020. Given the brain size of a low-land gorilla is about half of ours to a similar body mass it is likely that human level understanding beyond that of a gorilla, including language and subtracting autonomic bodily functions can be achieved with about half of our mental capacity. The rest runs our bodies and does lowland gorilla stuff. If this estimate is correct, then some supercomputers may already be surpassing the human maximum computational capability. However, having the capability is not the same as deploying it. For that, organisation and/or software is required and which we do not yet have.

¹⁰ Attributed to Gordon Moore in 1965. Moore's Law https://en.wikipedia.org/wiki/Moore%27s_law

¹¹ Moore's Law predicts a doubling in the scale of transistors per chip every two years and hence the computational power of the chip. There is a physical limit to the number of transistors that can be etched into an integrated circuit which has led many to recently posit the end of Moore's Law. These pundits however miss an important reality, Moore's Law has *come to mean* that computers are getting, and will continue get more computationally powerful, doubling in capacity approximately every two years. There are other paradigms and properties apart from transistor density that will allow this trend to continue. For example, within the next decade quantum computers are likely to cause the computational power trend described by Moore's law broken on the upside.

¹² First suggested by Robert Metcalfe c 1980. Metcalfe's Law https://en.wikipedia.org/wiki/Metcalfe%27s_law

The third is Nielsen's Law which is concerned with the *connectivity* of networks. It predicts that bandwidth speed increases by 50% per annum. This is also likely to hold given the planned improvements in bandwidth coming from fibre-optic cabling and low-Earth orbit satellite communications.

These three combined describe and predict the continued exponential rates of growth in the speed, scale and connectivity of computational resources. The network laws are significant with regards labour substitution as they collectively describe how remote computational resources can be focussed on to a local task. ie how work can and will be outsourced to computers. An artificial mind does not have to be in one place, it can be warehoused globally and distributed locally for task performance on demand..

Our ten year prediction to substitutable machine intelligence does not pre-suppose or require any further developments in computing technology. It would come about even if Moore's Law were to stop dead in its tracks tomorrow and chip speed were to stop increasing. Sheer scale and connectivity would carry across the line.

The mechanization of intelligence gives an unbeatable economic advantage to automated industry. Intelligent Machines do not sleep, get bored or distracted by other agendas. Their networks are scalable, their data is replicable and skills near instantaneously upgradeable. They are thus able to replace not just individual workers but entire teams.

Unlike humans however, intelligent machines can pickup or be upgraded from where the last generation left off and do not need to go through the nappies and baby-talk phase of humans. Thus computation engines are becoming to mental work as combustion engines were to manual work.

Work

Work can be decomposed into tasks and it is these tasks (not jobs) that are the targets of automation. Therefore until anthropomorphic robots exist there will be no direct targeting of jobs by intelligent machines¹³, rather it is the tasks that will be targeted.

Task automation is typically a complement to, rather than a substitute for, a worker on the job. As a typical job consists of numerous tasks, in the initial stages of workplace automation a given worker will likely not be made redundant as their job description contains many other tasks which are currently not being automated. In this sense automation is a tool and the workers job description will be modified to integrate with it. Automation creates a ripple effect of job redefinition. As capital will seek to maximise profitability, automation will lead to an increased demand for established multi-skilled (educated) over specialist labour.

Piecemeal automation will cause the aggregate requirement for human labour to decrease with each task automated however initially this will manifest as a lack of new hires and barriers to entry rather than direct retrenchment. The reason for this non-retrenchment is because existing hires have valuable experience (which is another

¹³ A contradiction to this was the job title 'Computer'. These were people that performed mathematical calculations by hand. Digital computers were constructed to supplement and ultimately entirely replaced them... and took their title.

type of multi-skilling). They are already socially integrated, they know the procedures, written and unwritten, and where everything is. New uses will therefore be found for this internally displaced labour, particularly in the improvements in quality of the non-automated tasks. I.e., they will spend more time making pretty reports. The barriers to entry are the beginnings of the true technological unemployment using the term as described by Keynes in 1930¹⁴:

“We are being afflicted with a new disease of which some readers may not yet have heard the name, but of which they will hear a great deal in the years to come—namely, technological unemployment. This means unemployment due to our discovery of means of economising the use of labour outrunning the pace at which we can find new uses for labour.”

This attrition causes a gradual migration of surplus labour to industries where there are more tasks that are currently beyond the reach of intelligent machines. In Australia this has manifested in a reduced share of the labour force employed in the industrial, agricultural and mining sectors, with a corresponding increase in the service industry.

As the computational capabilities of machines increase, more and more tasks become accessible for substitution either by outsourcing or automation. This does not necessarily represent a particular strata or type of task rather, the susceptibility to automation is simply a function of the value of automating it divided by the cost of computational resources required to achieve it. Therefore lawyers and airline pilots who command high wages are more at risk than gardeners or baristas.

This is in stark contrast to the current economic consensus that intelligent machines will target repetitive jobs. Dumb machines, like pumps and spinning jennies, target repetitive jobs, intelligent machines target whatever provides economic reward. There will be a natural perception of this being from low-skilled to high skilled as the low-hanging fruit are plucked but the exponential explosion of computation over the last 50 years means soon that no fruit anywhere are inaccessible.

The actual susceptibility of any task to automation effectively becoming:

$$S_a \propto V_a/R_c$$

Where V_a is the remuneration to be derived from the task after automating and R_c the cost of the computational resources required. The following table has been taken from a recent paper analysing the *recent* impact of automation on occupations ¹⁴. “Impact” describes task substitution rather than job replacement wherein the automation is primarily a complement to labour. However, it is obvious that high value, rather than ‘repetitive’ occupations are being targeted first. This makes sense as high earning individuals are precisely the ones where there is more value in enhancing their performance.

¹⁴ The term technological unemployment was introduced by Keynes in his paper “The Economic Possibilities for our Grandchildren” published in 1930. Keynes at the time saw this as only a temporary malady however it is likely to exist so long as there is exponential technological development.

Most Impacted			Least Impacted	
	Occupations	Scheduled Definition	Occupations	Scheduled Definition
1	Airline Pilots, Copilots, and Flight Engineers	✓	Models	
2	Physicists		Telemarketers	
3	Surgeons	✓	Locker Room, Coatroom, and Dressing Room	
4	Commercial Pilots	✓	Graders and Sorters, Agricultural Products	
5	Air Traffic Controllers		Shampooers	
6	Dentists, General		Maids and Housekeeping Cleaners	
7	Biochemists and Biophysicists		Cleaners of Vehicles and Equipment	
8	Oral and Maxillofacial Surgeons		Slaughterers and Meat Packers	
9	First-Line Supervisors of Fire Fighting and Prevention Workers		Dining Room and Cafeteria Attendants and Bartender Helpers	
10	Microbiologists		Food Servers, Nonrestaurant	

Table 1: List of Occupations most impacted by technology over the last few years¹⁵.

The economics of the substitution of labour by intelligent machines

In any discussion on the likelihood of automation the key economic concept is the elasticity of substitution between labour and capital, commonly referred to by the symbol σ (sigma). Ultimately the value of σ is based on the marginal productivity of a given factor input, with an improvement in productivity increasing that factors substitutability. This has obvious implications for automation as intelligent machines approach human productivity and are therefore eligible to substitute for all human labour.

This σ value is used by economists when writing the standard production function which depicts how changes in the output elasticities cause changes in the labour/capital ratio. The critical value of σ is unity which implies a stable labour/capital ratio as output changes. If the elasticity of substitution is below/above unity (1.0) then factor inputs are gross complements/substitutes. For robot inputs, a value of σ well in excess of unity would portend substantial and irrevocable automation. The implication for economic theory is that the traditional Cobb-Douglas production function, with a σ value of unity and constant returns to scale, is no longer applicable to a robot economy owned by Tech Giants.

¹⁵ Felten, E. W., Raj, M., and Seamans, R. [December 21, 2017] "Linking Advances in Artificial Intelligence to Skills, Occupations, and Industries" <https://www.aeaweb.org/conference/2018/preliminary/paper/EFD8kAG9>

Current research appears to support a value of σ greater than unity for “robotic labour” factor inputs. For the purposes of macroeconomic analysis it is convenient to treat intelligent machines as robotic labour, although they are in legal fact capital goods. De Canio’s¹⁶, using standard economic theory of the production function, came to the grim conclusion that when the value of σ is in the 1.7–2.1 range there will be a depressing effect on wages:

The purpose of this paper is to examine, in a simple model, conditions under which AI will lead to a decrease in aggregate wages, focusing on the role of the elasticity of substitution between human and robotic labor.

The conclusion is that for not-very-high values of the elasticity of substitution between human and robot labor, the proliferation of robots will reduce aggregate wages, ceteris paribus. Starting from the economy at its current position, any σ greater than about 1.9 implies wage decline...any value of σ greater than about 2.8 implies a reduction in the human wage as robotic labor increases.

*The best-fitting distributions..suggest that, if the elasticity of substitution between human and robotic labor is greater than the 1.7–2.1 range, **proliferation of robots will have a depressing effect on human wages.***

What does the future hold for wages and employment as the capabilities of AIs increase? Expansion of AIs’ skill sets (which in the terminology of the paper entails increases in the elasticity of substitution between AIs and humans) is likely to depress wages over time. This will increase measured inequality unless the returns to robotic assets are broadly spread across the population. Yet it is not clear how this spreading of the returns might come about.¹⁷

The economic analysis tends to be targeted to jobs and professions whereas Intelligent Machines are actually being targeted to tasks and it is here that substitution is already occurring, but not at the expense of jobs.

The capitalist drive for efficiency will enact substitution when it is possible. It already does this through migrant labour, offshoring and trade. Given that we have an exponentially increasing unbound trend (Machine Intelligence) approaching a fixed level (us), at some point these lines for any arbitrary task are going to cross. When they do cross, substitution is possible, and given the current dominant economic paradigm, it is also inevitable.

When substitution is possible there is some latency as systems wake up to the fact, after which happens quickly. It becomes a step change even though the apparent changes to that point may have been gradual. The tasks currently done by humans are quite varied

¹⁶ De Canio, S. J. [2016] “Robots and humans – complements or substitutes?” [2016]
https://ac.els-cdn.com/S016407041630043X/1-s2.0-S016407041630043X-main.pdf?_tid=9c6517b2-1544-11e8-b9b9-00000aab0f01&acdnat=1519024686_6ec2c8e0e41d91249eafb4aaaf991a48

¹⁷ De Canio has no doubt that the power of AI systems will increase in the future, thus, further increasing the value of σ .

therefore creating the elasticity which economists speak of and which is why substitution has not all happened at the same time.

We can however expect the rate of substitution to accelerate as the processes that cause it will themselves be automated. ie as computers become more deployed in chip design and manufacture, this in turn leads to faster chips etc. Similarly, improvements in connectivity allow global distribution of tasking during low usage cycles¹⁸. We do not know when full substitution will take place¹⁹, just that it will, and that we have the hardware already in place and it is closer than most government advisors think.

The Rise and Fall of the Contingent Workforce

The drive to efficiency will cause a net attrition of human labour from any existing industrial process. Where there is a degree of specialisation or expertise required in a task that is time finite, there will be a tendency to outsource or employ casual workers, ie, through a contingent workforce. This allows companies to avoid paying for ancillary expenses, such as superannuation, leave, buildings etc. A recent well known example of this is Uber.

In order to outsource tasks effectively they must be specified. This specification in turn is agnostic to who (or what) is actually doing the task so long as the quality of the deliverables are met. This process coupled with massive improvements in communications has allowed offshoring of tasks and in the future, automation. One of the biggest drivers of automation is therefore ISO-9000.

There is an almost sinister element to the use of contingent labour and the task specification it entails. That is, as the specifications are improved and operations of the contingent workforce are monitored, this process itself increasingly provides a defined target and training data for automation²⁰. In this respect the contingent workforce is effectively becoming their own executioners.

Task replacement vs Process Re-engineering

As more and more tasks are automated a new opportunity for efficiency arises. This is process re-engineering where multiple interdependent tasks are re-engineered as a whole to take advantages of technological advances. The process itself is redefined around the capabilities afforded by Intelligent Machines. Process re-engineering allows the removal of infrastructure capital costs required to support human labour and thus the substitution process will accelerate.

Process re-engineering can also allow the ground-up design of new human-less systems. This will cause direct job attrition. The jobs lost to this process are never coming back, furthermore, the re-engineered workplace does not have to be human friendly. A recent significant example of this is the Internet of Things (IoT) which is essentially the currently internet, with machines talking to each other and bypassing human agency.

¹⁸ This is similar in nature to the Buckminster Fuller's suggestion of linking power grids. Diurnal usage patterns means that processes that are slowed because of humans sleeping can have their computational resources diverted to assist operations on the other side of the planet.

¹⁹ Whilst we predict ten years for the capability based on base trends, there is likely to be a few years latency beyond that.

²⁰ An example of this use is the crowd sourcing system Mechanical Turk where an increasing amount of work is now explicitly nominated as AI training. Classification of images, writing etc.

Automation Step Change

The magnitude of effect on the labour force of a single task substitution depends on how ubiquitous that task is. Some tasks substitutions will have such a significant effect that they will rapidly escalate the eventual performance of process re-engineering. We propose that a major substitution event will occur when computers pass the Turing Test, that is, they will become conversationally indistinguishable to humans²¹. This test is not some abstract milestone as many claim; rather, it represents the point at which AI can be applied to any task describable.

Once computers can understand speech the revolution will then begin to eat its children. STEM sciences will suffer, particularly IT professionals who will be amongst the first to go as the skill of talking to computers is essentially what they now do. This inflection point with its consequential assault on employment in the STEM fields will likely be reached well within the next generation.

The Tech Giants

There is race underway between the richest and smartest companies on earth to position themselves as global cloud-based AI service providers. The most visible competitors in this race are the North American and East Asian tech giants including Google, Apple, Amazon, Alibaba, Baidu, Tencent, Facebook, Samsung and Microsoft.

These companies, which have dominated IT, are now changing their focus from digital 'bits' to analog 'its'. Their proximate aim is to position themselves between customers and suppliers so as to capture and mine proprietary data for marketing purposes. Ultimately they aim to become the economy-wide IT broker and out-sourceable AI contractor through the privatisation of the IoT, which will become the medium for the robot economy.

It no coincidence that the Tech Giants have launched virtual assistants (such as Siri, Google, Bixby, Cortana and Alexa) as the heralds of their flagship services. These can be best thought of as interns, paid for by us, tasked with learning how to speak whilst gathering as much information as they can.

This strategy gives the tech giants two unassailable advantages:

1. They have direct contact with consumers
2. They get massive free distributed data to train and improve their AIs particularly in Natural Language Processing (NLP)

The Tech Giants, having more data and computational resources than anyone else are rationalisation agents *par excellence*, and use both institutional and instrumental systems to subject work to computation internally and sell computation for arbitrary tasks externally. It is their computational systems that will condition the extent and pace of automation across the entire world. As they substitute intelligent machines for human

²¹ This does not mean the computer needs to lie, ie if asked "Are you a computer", it would presumably answer truthfully. Rather, it implies the conversation is seamless and useful. It is important to note that in the same vein that we build aircraft that outperform birds yet do not flap their wings, an AI may be constructed differently to the human brain but still outperform it.

labour via the contingent workforce, these companies will increasingly *become* the economy.

There is nothing stopping them from continuing their monopolisation of the tech sector given the inherent economies of scale.

On the production side they benefit from their huge server farms and massive data sources. Likewise the Chinese Tech Giants (Baidu, Alibaba and Tencent) appear entrenched given they have the guiding support of the Chinese Government and a huge domestic market. Moreover Tech Giants have further strengthened their position by developing the strategy of buying up promising tech start-ups which prevents future competition.

On the consumption side they benefit from network externalities that exponentially increase the value of each incremental member of the network. Consumers tend to stay with a network which has familiarity. Furthermore, as the AIs powering the smart assistants improve they can further this process by tailoring themselves to become increasingly complementary and useful to the individual.

TOR#3. the wider effects of that change on inequality, the economy, government and society;

Introduction

The automation of capital and the commodification of labour will inflict tremendous industrial disruption in itself, given that people spend half their waking lives, earn most of their income, and build a large part of their social lives, at work. But this is only an initial effect, profound and widespread inequality and instability will follow. Specifically, there will be increased polarisation of income distribution and dramatic fluctuations in the rate of economic growth.

These effects are macroeconomic owing to the relationship between the production of output on the one hand and both the distribution and expenditure of income on the other. The macroeconomic effect is reinforced as automation and commodification becomes widespread.

Polarisation of Income Distribution

Automation on wide-spread scale will result in “*distributional carnage*”.²² Looking first at the macroeconomics of income distribution we can see two sources of increased class division arising directly out of the automation-induced change in the production function, sometimes called capital-biased technological change.

Firstly, there will be a winnowing of numbers *within* the capitalist class as economies of scale generate a winner-take-all payoff to the owners of the Tech Giants. Big Fish will eat the Little Fish.

Secondly, there will be an increased polarisation *between* the the capitalist class and everyone else who will be forced into a gradually deteriorating labour market. The labour market itself will stratify into establishment workers who enjoy the Standard Employment Relation and contingent workers who will effectively work at piece rates. As automation withers the establishment workforce away, the fraction of labour employed in contingent tasks will increase until it too succumbs.

A further adverse effect on labour income will be the glut of labour as capitalists and workers made redundant flood the market seeking employment in the contingent workforce. Wages will be low and exploitation high.

At the limit, there will just be owners of capital and an unemployed workforce.

Fluctuations in the rate of economic growth

Automation on a widespread scale would create conditions for unpredictable high amplitude oscillations in the rate of economic growth. Traditionally, macroeconomics

²²Berg, A., Buffie, E. and Zanna, F. [December 2017] “Should We Fear the Robot Revolution? (The Correct Answer is Yes)” <https://pdfs.semanticscholar.org/96a2/1932a7427de71fd29da6bd95176841d17a96.pdf>

devolves into two areas, studying changes in both the (supply-side) structural and (demand-side) cyclical rate of economic growth. Both of these will be profoundly affected by automation, although making predictions about growth rates during periods of industrial disruption is a mugs game without a crystal ball.

On the structural supply side, it is clear that the industrial disruption caused by automation has the potential to massively increase the long term structural rate of economic growth, given that a robotic IoT could massively extend the production possibility frontier. The traditional neoclassical model of economic growth, employing the Cobb-Douglas production function²³, no longer applies to an economy dominated by Intelligent Machines. It is equally clear that there is great potential for technological unemployment in the short term as displaced workers face fewer employment opportunities and little scope to upgrade their human capital to compete with intelligent machines. The situation is chaotic and therefore unclear as to which effect will dominate over the medium term. Government decisions will be a critical determiner.

On the cyclical demand side, the forced redundancy of more than even 10% of the workforce²⁴ would have a massively depressing effect on labours expenditure and therefore aggregate demand. The functional income-expenditure multiplier²⁵, which implies that one person's income is another person's expenditure, would go into reverse. An income-expenditure multiplier in excess of one implies that this depressing effect on expenditure will be larger than the initial income loss. Estimates of the value of the income-expenditure multiplier in Australia vary but 1.5 was used by the IMF in estimating the Obama administration's fiscal stimulus.²⁶ Workers have a higher marginal propensity to consume²⁷ than owners, so the multiplier effect will be larger among workers just made redundant.

At the very least a demand-side shock of this scale raises the prospect of prolonged recession - we may in fact be seeing early signs of this if the notion of "secular stagnation" has any validity. This could easily turn into a financial depression as a functional liquidity crisis can rapidly morph into a financial solvency crisis when toxic mortgages are tossed onto the asset fire sale.

²³ See the Solow-Swan model https://en.wikipedia.org/wiki/Solow%E2%80%93Swan_model

²⁴ This figure is based on the automation of tasks, which mostly leaves establishment jobs in place. Arntz, M., T. Gregory and U. Zierahn [2016], "The Risk of Automation for Jobs in OECD Countries: A Comparative Analysis", <http://www.oecd-ilibrary.org/docserver/download/5jlz9h56dvq7-en.pdf?expires=1519141690&id=id&accname=guest&checksum=8A3F4764CCE17DCC26884725715A66EB>

²⁵ The theory was introduced by Keynes in 1936 but formalised by Samuelson and Hansen in 1939 into the Multiplier-accelerator model https://en.wikipedia.org/wiki/Multiplier-accelerator_model

²⁶ Quiggin, J. [26 August 2013] <http://johnquiggin.com/2013/08/26/fiscal-multipliers-and-employment-wonkish/>

²⁷ The multiplier is calculated using the formula $1 / (1 - MPC)$ where MPC is the Marginal Propensity to Consume an additional unit of income.

TOR#4. the adequacy of Australia's laws, including industrial relations laws and regulations, policies and institutions to prepare Australians for that change;

Introduction

No country is prepared for wide-spread technological unemployment since the end of labour as an industrial input is an unprecedented change in social organization. The chief problem for government is recognising that there is an immediate problem. Unfortunately the economics profession, which is tasked with warning and advising government officials of looming economic problems, is largely persuaded that automation is not a problem or that it is one that can mostly be left to the free play of market forces. This is wrong.

There are numerous things an informed Government can do to lessen the coming shocks. Herein is a brief list.

1. Throw away the Economists Rear-view mirror

The consensus view of the economics profession is that automation of capital has not greatly increased productivity, and the commodification of labour has not disrupted employment patterns. Bob Solow, back in 1987, dubbed the absence of marked automation-induced improvements in growth as the "productivity paradox"²⁸ More recently Australian economists²⁹ have made much the same point about, what we might call, the "employment paradox", namely the stable ratio of full-time employment and the puzzling absence of overall growth in the contingent workforce. This, despite the fact that, to paraphrase Solow, "you can see gig workers everywhere but the statistics".

The problems with the economists essentially neoclassical interpretation is that it is parochial, backward looking and superficial. It is parochial because evidence on the growth of the contingent labour force in the US is clear.³⁰ It is backward looking because it relies on historical evidence which is myopic to the imminent step-change in the workplace and workforce, driven by both exponential improvements in digital technology and the increasing monopolisation of industry by the Tech Giants..

It is superficial because it does not see that changes in both production processes and consumption products are masked by quantitative statistics. In the production process, task substitution is already underway and is entirely consistent with stable full-time employment patterns. Thus skilled workers are complementing their labour with automated systems. Whilst unskilled workers are supplementing their income with moonlight jobs in the contingent workforce. For consumption products, it is clear that much of the increase in productivity is reflected in qualitative improvements in

²⁸ "You can see the computer age everywhere but in the productivity statistics."Robert Solow, "We'd better watch out", New York Times Book Review, July 12, 1987, page 36

²⁹ Borland, J. and Coelli, M., Are robots taking our jobs? [November 2017]
<http://onlinelibrary.wiley.com/doi/10.1111/1467-8462.12245/full>

³⁰ Katz, L. F. and Alan B. Krueger, A. B., [2016] "The Rise and Nature of Alternative Work Arrangements in the United States, 1995-2015 [2016] https://scholar.harvard.edu/files/lkatz/files/katz_krueger_cws_resubmit_clean.pdf

performance rather than quantitative increases in the value of output. In fact, many digital products are subject to price deflation which depresses productivity.

2. Reduce Privatisation

As Intelligent machines increasingly move into management any private sector inventiveness or inefficiency will disappear. Public sector firms are also less likely to tolerate exploitation and other forms of unfairness.

Therefore there is increasingly less advantage to the country to encourage privatisation.

3. Develop a Transaction Tax on the IoT

The IoT has the potential to grow exponentially larger than the existing economy. It will however be very difficult to tax. One possible solution would be to develop a block-chain based transactional unit for the IoT which would allow the government to derive revenue from a kind of transaction tax. This would be accomplished by issuing units to itself based solely on the economic activity which is directly reflected in the blockchain. This specie would be similar to a cryptocurrency in that it would provide a degree of transactional anonymity but as it is government backed would not suffer the same speculative price fluctuation. Nor would it be mineable. The tax would be infinitesimally small per transaction but as the IoT grows could allow a non-evadeable source of revenue.

4. Develop Policies Concerning AI Interaction

AI systems as deployed by the large tech companies will have a significant amount of information both pertaining to any given individual and any association of individuals.

They will therefore be able to engage in social engineering to an unprecedented degree and scale. As AI is automated and learns itself there will be a legislative deniability between those owning or commissioning the AI and the actions of that AI as it seeks to manipulate people for the financial gain of its commissioners. The likely primary use of this will be in sales and marketing where AI sales bots could cheaply and massively target people or interest groups. Similar to what spam does now but far more engagingly. Elderly people would be particularly vulnerable. Legislation should therefore be preemptively developed to reduce the negative impacts of this asymmetry.

Some overarching principles of such legislation could include:

1. An AI must identify itself as an AI
2. An AI must always be truthful
3. An AI must always identify its purpose when conversing with a human.
4. An AI must always disclose the source of any information given by it.
5. An AI must retain a transcript of all conversations between itself and humans and which is freely and only available to the human participant in the conversation and law enforcement.
6. A human must not pretend to be an AI

5. Develop Policies for Data Ownership and Data Rental Rights

Data is valuable and people are effectively giving it away. The government should develop policies of education and ensure privacy and other rights are not being sacrificed as well as exploring the right to impose some kind of a data acquisition tax or minimum rights for the benefit of all Australians who participate in the generation of data.

Conclusion

We have here discussed the accelerating and unstoppable introduction of Artificial Intelligence (AI) which is just the final phase of the industrial revolution which began in 18th century; the process in which labour is systematically substituted by more productive capital.

We give no suggestion that the future with ascendant AI will necessarily be a bad place. The replacement of humans as a source of labour holds forth the tantalising prospect of a World of Plenty. However, getting to this future is fraught with the catastrophic risks of industrial disruption, financial depression and social division.

Economic disruption is painful. It was traumatic for the displaced peasants leaving the lands they had worked for generations. Similarly it will be traumatic for those leaving their factories and offices, still lumbered with their mortgages. The difference is that the former, although they did not know it at the time, had greener pastures to plow, their children, new industries to found. We, the next displaced group, do not. It is ironic that at the start of the industrial revolution people thought they had no-where to run but actually they did, and now, as the industrial revolution nears its end they think they do, but really don't.

We have become so used to change that we do not fear it as our forefathers did.

The disruption poses a formidable challenge to public officials in their capacity as both representatives of their constituents interests and as ministers of public monies. Unlike global warming, there is no possibility of deferring this problem.

Essentially, we believe the Luddites were correct in recognition of the threat, wrong in the response, but 200 years ahead of their time. We must do better.

Thank you