THE BESTSELLING AUTHOR OF *SURVIVING AI*

CALUM CHACE

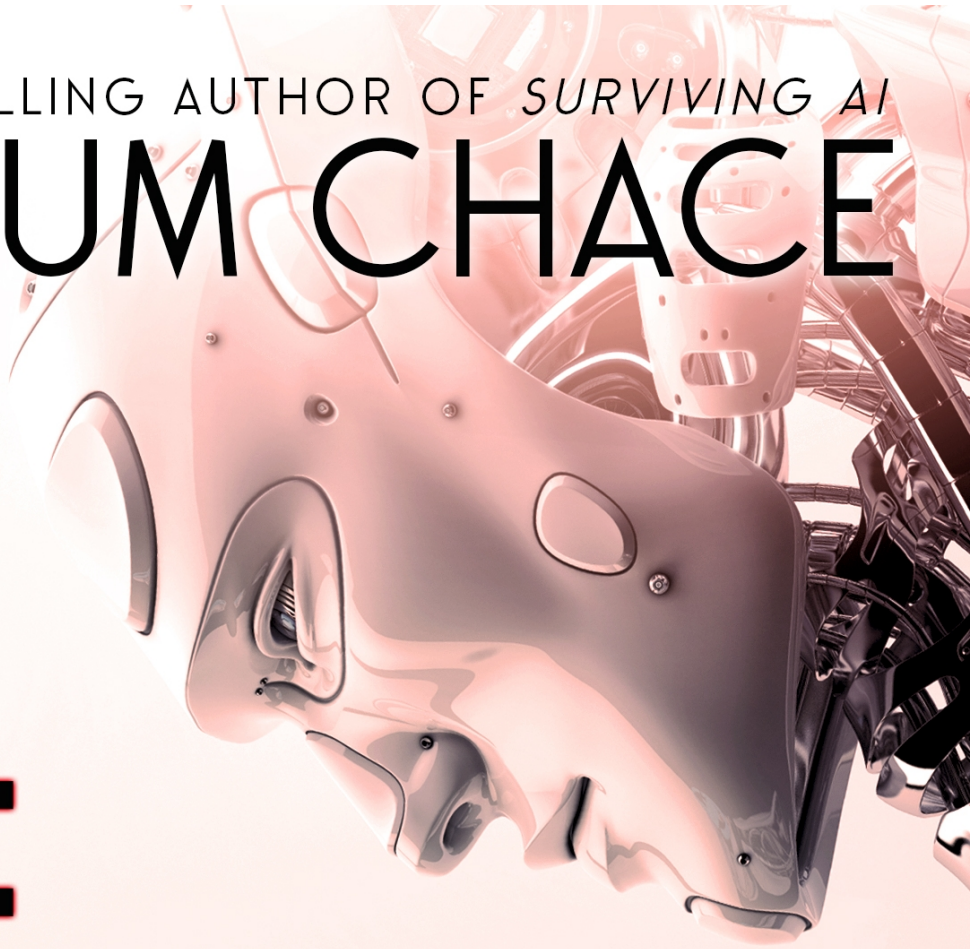

THE ECONOMIC SINGULARITY

Artificial intelligence and the death of capitalism

"Read The Economic Singularity if you want to think intelligently about the future." *Aubrey de Grey*

The Economic Singularity

Artificial intelligence and the death of capitalism

by Calum Chace

Review copy

18 May 2016

Contents

1 Introduction: the Economic Singularity

2 The history of automation
- **2.1** The Industrial Revolution
- **2.2** The Information Revolution
- **2.3** The automation story so far
- **2.4** The Luddite Fallacy

3 Is it different this time?
- **3.1** Prophets of change
- **3.2** Academic and consultancy studies
- **3.3** Crying wolf
- **3.4** AI to date
- **3.5** Exponential future
- **3.6** What people do
- **3.7** Related new technologies
- **3.8** Poster boy: Self-driving vehicles
- **3.9** Who's next?
- **3.10** Jobs or no jobs
- **3.11** What's the problem?
- **3.12** Summary

4 A timeline
- **4.1** Un-forecasts
- **4.2** 2021
- **4.3** 2031
- **4.4** 2041

5 The challenges
- **5.1** Economic contraction
- **5.2** Distribution
- **5.3** Meaning
- **5.4** Allocation
- **5.5** Cohesion

6 Scenarios
- **6.1** No change
- **6.2** Centaurs
- **6.3** Capitalism + UBI
- **6.4** Fracture
- **6.5** Collapse
- **6.6** Protopia

7 Summary and recommendations
- **7.1** The argument
- **7.2** The two singularities
- **7.3** What is to be done?

Comments on The Economic Singularity

A problem that all techno-pioneers face, when “selling” their vision of the future to others (I use quotes because I in no way refer to anything monetary), is to get their audience to focus on the new development in an appropriate context. Above all, this means striking a balance between communicating the significance of the proposed development and setting it within the universe of other developments that are likely to have occurred in the meantime.

The advance of automation, described with great care and accuracy in this book, will almost certainly constitute the substrate within which all other technological developments - be they biomedical, environmental or something else entirely - will occur, and thus within which they should be discussed as regards their value to humanity.

Read "The Economic Singularity" if you want to think intelligently about the future.

Aubrey de Grey – CSO of SENS Research Foundation; former AI researcher

Following his insightful foray into the burgeoning AI revolution and associated existential risks, Calum focuses his attention on a nearer term challenge - the likelihood that intelligent machines will render much of humanity unemployable in the foreseeable future. He explores the arguments for and against this assertion and provides a measured response, acknowledging the risks associated with such a radical shift in our self identity but also outlining the potential significant benefits. Once again he proves a reliable guide through this complex yet fascinating topic.

Ben Medlock, co-founder of Swiftkey, the best-selling app on Android

"It's important that this book and others like it are written. Not because the future will necessarily happen exactly in the way described, but because it's important to be prepared if it does. If automation compels us to shift to a different economic organisation, we better start laying the foundations for the shift right now."

Dr Stuart Armstrong, James Martin Research Fellow at the Future of Humanity Institute, Oxford University

"Chace does a good job of answering the question whether robots will take our jobs. What worries me more though, a bit further down the road, once these robots have become massively intelligent, is whether they may take our lives. Chace covered this issue thoroughly in his previous book, "Surviving AI".

Prof. Dr. Hugo de Garis – author of "The Artilect War", former director of the Artificial Brain Lab, Xiamen University, China

“The jobs of the future don’t exist today and the jobs of today will not exist in the future. Technological Singularity will change everything, but its first manifestation will come in the domain of economics, most likely in the shape of technological unemployment. Calum Chace’s “The Economic Singularity” does a great job of introducing readers of all levels to the future we are about to face. Chace explains what might happen and what we can do to mitigate some of the negative consequences of machine takeover. The book covers unconditional basic income, virtual

environments, and alternative types of economies among other things. Highly recommended."

Dr. Roman V. Yampolskiy, Professor of Computer Engineering and Computer Science, Director of Cybersecurity lab, Author of Artificial Superintelligence: a Futuristic Approach

Unprecedented productivity gains and unlimited leisure—what could possibly go wrong? Everything, says Calum Chace, if we don't evolve a social system suited to the inevitable world of connected intelligent systems.

It's a failure of imagination to debate whether there will be jobs for humans in the automated world, Chace argues - we must look farther and ask how we will organize society when labor is not necessary to provide for the necessities of life. Find an answer, and life improves for all; without one, society collapses. Read this book to understand how social and technological forces will conspire to change the world—and the problems we need to solve to achieve the promise of the Economic Singularity.

Christopher Meyer, author of "Blur", "Future Wealth", and "Standing on the Sun"

It is interesting to listen to our own language. We say things such as "to earn a living", implying that you need to earn the privilege to be alive and to live a moderately enjoyable life. This may be looked upon as strangely in the future as we would now look back and say about slaves that they had to "earn their freedom".

Calum Chace hits the nail on the head in chapter 3 of this extremely timely book. It is probably true that there will be new types of 'jobs' in whatever niches remain best explored by humans in the near future, but we should also consider an entirely different goal for the future.

Who was it in ages past who contributed those things we most remember, over time, as being of great value? It was they who contributed to the arts, the sciences and invention. But who were those people? Throughout the majority of history, these were mainly the people who either did not have to have a 'job' (because they were part of an aristocracy that had a different role to play while being supported by property and subjects), as well as the artists, artisans, philosophers or scientists who were directly supported by those patrons and therefore did not have the need to take a typical 'job'.

It is not the typical jobs that are celebrated as the best of humanity, and therefore it probably should not be our aim to find yet more categories of such jobs. Instead, wouldn't it be much better if a greater proportion of humanity could find the means to engage in preferred and culture-creating activity? With this in mind, it seems to me that it should be our aim to get rid of the need for jobs and employment just for the purpose of survival.

Our strategies for the future should be not about finding new salary jobs, but rather about removing the need for them, and about setting up a better and more advanced social structure. This is where looking at the challenges involved and the path to a successful alternative, as Chace does in chapter 5, is essential.

Where ideas such as a universal basic income (UBI) are concerned, it is useful to keep in mind that the world is not the US. Even if there is some initial antipathy in the US, because of associations between UBI and what might naively be labeled as 'socialist' thinking, the US will not wish to be

left behind if other nations successfully implement the change. The time to dive deeply into the many issues raise in this book, to start a wider conversation about those issues, and to look creatively for the most well-balanced solutions and outcomes, is now.

Randal Koene – founder of carboncopies.org

"The Economic Singularity is fascinating. Calum Chace brilliantly explores the enormous opportunities, and risks, presented to humanity by the rapid advance of technology, and especially artificial intelligence. I couldn't put this book down."

Ben Goldsmith - Menhaden Capital

In his fast-paced new book, Calum Chace explains the challenge facing humanity: to navigate through a dramatic transition which he christens the economic singularity. The culmination of an accelerating wave of automation by robots and AI, this transition threatens to do more than displace employees from the workforce. Unexpectedly, it threatens the end of capitalism itself, and potentially the fracturing of the human species.

Chace compellingly sets out a range of options, before sharing his assessment of the most credible and desirable outcomes, so that we can reach a shared “protopia” rather than a nightmarish “Brave New World” (or worse).

David Wood – chairman, London Futurists

Calum Chace is a best-selling author of fiction and non-fiction books and articles, focusing on the subject of artificial intelligence. His books include “Surviving AI”, a non-fiction book about the promise and the challenges of AI, and “Pandora's Brain”, a techno-thriller about the first superintelligence.

He is a regular speaker on artificial intelligence and related technologies and runs a blog on the subject at www.pandoras-brain.com.

A long time ago, Calum studied philosophy at Oxford University, where he discovered that the science fiction he had been reading since boyhood is actually philosophy in fancy dress.

Also by Calum Chace

Surviving AI

Pandora's Brain

The Internet Startup Bible (co-authored)

The Internet Consumer Bible (co-authored)

For Julia and Alex

THE ECONOMIC SINGULARITY

A Three Cs book

ISBN

First published in 2016 by Three Cs

Printed and bound in Great Britain by

Chapter 1. Introduction: the economic singularity

Accelerating change

In the next few decades, life for most people is going to change in extraordinary ways and at an extraordinary rate. The reason, as usual, is technology.

Human lives and societies can be transformed by religions, cultural memes[i], and the imposition of new economic systems. They can be transformed by the passion and belief of a single great man or woman. But when profound and lasting change takes place it is usually because we found a new way of doing things – a new technology. Thus we name many historical periods after their dominant technology: the iron age, the bronze age, and so on.

When a cluster of related technological innovations come along together they can create sufficient change to merit the title of a revolution. This has happened twice before in human history, with the agricultural and the industrial revolutions, and we are now in the middle of a third, the information revolution.

These revolutions are not overnight affairs: the industrial revolution has been under way for 300 years.[ii] The information revolution is just half a century old,[iii] and in many ways we are nearer to the start than the end. We think the world has changed greatly in the last century, and especially in the last twenty years or so, and indeed it has. But the rate of change is accelerating, and the changes that are coming will dwarf what has happened so far.

Forecasting has always been perilous. Throughout history, most long-term forecasts have been wrong, often blind-sided by the arrival of a new technology like smartphones. But in the coming decades the rate and scale of change will be so great that the future will become mysterious in a new way. So much so that people talk about a coming technological singularity.

The term “singularity” is borrowed from maths and physics, where it means a point at which a variable becomes infinite. The usual example is the centre of a black hole, where matter becomes infinitely dense. When you reach a singularity, the normal rules break down, and the future becomes even harder to predict than usual. In recent years, the term has been applied to the impact of technology on human affairs.[iv]

Superintelligence and the technological singularity

The technological singularity is most commonly defined as what happens when the first artificial general intelligence (AGI) is created – a machine which can perform any intellectual task that an adult human can. It continues to improve its capabilities and becomes a superintelligence, much smarter than any human. It then introduces change to this planet on a scale and at a speed which un-augmented humans cannot comprehend. I wrote about this extensively in my book, “Surviving AI”.

The term “singularity” became associated with a naïve belief that technology, and specifically a superintelligent AI, would magically solve all our problems, and that everyone would live happily ever after. Because of these quasi-religious overtones, the singularity was frequently satirised as “rapture for nerds”, and many people felt awkward about using the term.

The publication in 2014 of Nick Bostrom's seminal book “Superintelligence” was a watershed moment, causing influential people like Stephen Hawking, Elon Musk and Bill Gates to speak out about the enormous impact which AGI will have – for good or for ill. They introduced the idea of the singularity to a much wider audience, and made it harder for people to retain a blinkered optimism about the impact of AGI.

For time-starved journalists, “good news is no news” and “if it bleeds it leads”, so the comments of Hawking and the others were widely mis-represented as pure doomsaying, and almost every article about AI carried a picture of the Terminator. AI researchers and others hastened to warn us (rightly) not to throw the baby of AI out with the bathwater of unfriendly superintelligence, and the debate is now more nuanced.

Technological unemployment and the economic singularity

So for me at least, the term “singularity” no longer seems so awkward. And it seems reasonable to apply it to another event which is likely to take place well before the technological singularity. I call this event “the economic singularity”.

There is a lot of talk in the media at the moment about technological unemployment – the process of people becoming unemployed because machines can do any job that they could do, and do it cheaper, faster and better. There is widespread disagreement about whether this is happening already, whether it will happen in the future, and whether it is a good or a bad thing. This disagreement is natural and inevitable: one of the main features of a singularity is that what lies beyond its event horizon is hard to see[v]. Nevertheless we must try to peer into the hazy future if we are to prepare ourselves for it.

In this book I will argue that technological unemployment is not happening yet (or at least, not much), that it will happen in the next few decades, and that it can be a very good thing indeed - if we prepare for it, and manage the transition successfully.

Naturally, there are challenges. As we will see, a lot of people believe that Universal Basic Income (UBI) is a silver bullet that will solve the problem of technological unemployment. UBI is a guaranteed income paid to all citizens simply because they are citizens. It may take some time for the idea of UBI to be accepted, especially in the USA, where resistance to anything that smacks of socialism is often fierce - almost visceral. Martin Ford's otherwise excellent book “The Rise of the Robots” almost fizzles out at the end because he seems daunted by the scale of the opposition that UBI will face in his home country.

But to my mind, UBI is not the real battle. In Europe we are very comfortable with the idea of a safety net of welfare programmes which prevent the economically unsuccessful from falling into absolute penury. Most American states provide unemployment benefits too, although they usually cease after six months. In fact the US spends more per capita on welfare ($650 in 2011, according to the OECD) than the UK ($610) or Canada ($550).[vi] Unlike some of my American friends, I believe the people of that great country will quickly accept the need for UBI if and when it becomes undeniable that the majority of them are going to be unemployable.

The real problem, it seems to me, is that we will need more than just UBI. We may need an

entirely new form of economy. I see great danger in a world in which most people rub along on handouts while a minority – perhaps a tiny minority – not only own most of the wealth (that is pretty much true already) but are the only ones actively engaged in any kind of economic activity. Given the advances in all kinds of technology that we can expect in the coming decades, this minority would be under immense temptation to separate themselves off from the rest of us – not just economically, but cognitively and physically too. Yuval Harari, author of the brilliant book "Sapiens", says that in the coming century or so, humanity will divide into two classes of people. Rather brutally, he calls them the gods and the useless.[vii]

Capitalism and liberal democracy have served humanity well in the last couple of centuries. I am not convinced they will continue to do so in a post-automation world, but it is no small task to work out what they should be replaced with, and how that can be achieved without turmoil. This sounds like a singularity – an economic singularity.

Chapter 2. The History of Automation

2.1 - The industrial revolution

For a process that began hundreds of years ago, the start date for the industrial revolution is surprisingly controversial. Historians and economists cannot even agree how many industrial revolutions there have been: some say there has been one revolution with several phases, others say there have been two, and others say more.

The essence of the industrial revolution was the shift from manufacturing goods by hand to manufacturing them by machine, and the harnessing of better power sources than animal muscle. So a good date for its beginning is 1712, when Thomas Newcomen created the first practical steam engine for pumping water. For the first time in history, humans could generate more power than muscles could provide - wherever they needed it.

The replacement of human labour by machines in manufacturing dates back considerably earlier, but they were powered by muscles or by wind or water. In the 15th century, Dutch workers attacked textile looms by throwing wooden shoes into them. The shoes were called sabots, and this may be the etymology of the word “saboteur”. A century later, around 1590, Queen Elizabeth (the First) of England refused a patent to William Lee for a mechanical knitting machine because it would deprive her subjects of employment.

In the second half of the 18th century, the Scottish inventor James Watt teamed up with the English entrepreneur Matthew Boulton to improve Newcomen’s steam engine so that it could power factories, and make manufacturing possible on an industrial scale. At the same time, iron production was being transformed by the replacement of charcoal by coal, and “canal mania” took hold, as heavy loads could be transported more cheaply by canal than by road or sea.

Later, in the mid-19th century, steam engines were improved sufficiently to make them mobile, which ushered in the UK's “railway mania” of the 1840s. Projects authorised in the middle years of that decade led to the construction of 6,000 miles of railway – more than half the length of the country's current rail network. Other European countries and the USA emulated the UK's example, usually lagging it by a decade or two.

Toward the end of the 19th century, Sir Henry Bessemer's method for converting iron into steel enabled steel to replace iron in a wide range of applications. Previously, steel had been an expensive commodity, reserved for specialist uses. The availability of affordable steel enabled the creation of heavy industries, building vehicles for road, rail, sea and later the air.

As the 20th century arrived, oil and electricity provided versatile new forms of power and the industrial world we recognise today was born. The changes brought about by these technologies are still in progress.

In summary, we can identify four phases of the industrial revolution:

1712 onwards: the age of primitive steam engines, textile manufacturing machines, and the canals
1830 onwards: the age of mobile steam engines and the railways
1875 onwards: the age of steel and heavy engineering, and the birth of the chemicals industry

1910 onwards: the age of oil, electricity, mass production, cars, planes and mass travel.

From an early 21st-century standpoint, it seems entirely natural that the industrial revolution took off where and when it did. In fact it is something of a mystery. Western Europe was not the richest or most advanced region of the world: there were more powerful empires in China, India and elsewhere. There is still room for debate about whether the technological innovations came about in England at that time because of the cultural environment, the legal framework, or the country's fortuitous natural resources. Fascinating as these questions are, they need not detain us.

2.2 - The information revolution

Even though the industrial revolution is still an on-going process, there is general agreement that we are now in the process of an information revolution. There is less consensus over when it began or how long it is likely to continue.

The distinguishing feature of the information revolution is that information and knowledge became increasingly important factors of production, alongside capital, labour, and raw materials. Information acquired economic value in its own right. Services became the mainstay of the overall economy, pushing manufacturing into second place, and agriculture into third.

One of the first people to think and write about the information revolution and the information society was Fritz Machlup, an Austrian economist. In his 1962 book, *The production and distribution of knowledge in the United States*, he introduced the notion of the knowledge industry, by which he meant education, research and development, mass media, information technologies, and information services. He calculated that in 1959, it accounted for almost a third of US GDP, which he felt qualified the US as an information society.

Alvin Toffler, author of the visionary books *Future Shock* (1970) and *The Third Wave* (1980), argued that the post-industrial society has arrived when the majority of workers are doing brain work rather than personally manipulating physical resources – in other words when they are part of the service sector. Services grew to 50% of US GDP shortly before 1940,[viii] and they first employed the majority of working Americans around 1950.

We have seen that the start and end dates of the economic revolutions (agricultural, industrial and information) are unclear. What's more, they can overlap, and sometimes re-ignite each other.

An example of this overlap is provided by the buccaneers who preyed on Spanish merchant shipping en route to and from Spain's colonies in South America during the 17th century. (Some of these buccaneers were effectively licensed in their activities by the English, French and Dutch crowns, which issued them with “letters of marque”. This ceased when Spain's power declined toward the end of the century, and the buccaneers became more of a nuisance than a blessing to their former sponsors.) When a buccaneer raiding party boarded a Spanish ship the first thing they would look for and demand was the maps. Charts – a form of information which improve navigation – were actually more valuable than silver and gold.[ix]

An example of one revolution re-igniting another is that the industrial revolution enabled the mechanisation of agriculture, causing a second agricultural revolution, making the profession of farming more effective and more efficient. The information revolution does the same, providing farmers with crops that are more resilient in the face of weather, pests and weeds, and allowing them to sow, cultivate and harvest their crops far more accurately with satellite navigation.

Along with the uncertainty about the start date of the information revolution, there is disagreement about how distinct it is from the industrial revolution. The Internet of Things (IoT) is a phenomenon of the information revolution which we will look at in more detail in chapter 3.7. Klaus Schwab, founder and executive chairman of the World Economic Forum which hosts the annual meeting of the global elite in Davos, calls the IoT the fourth industrial revolution.[x] This seems to me to under-state the importance of the IoT, and also to separate it from all the other digital revolutions which comprise the information revolution, including, of course, artificial

intelligence.

2.3 - The Automation story so far

The mechanisation of agriculture

The particular aspect of the industrial and information revolutions which concerns us in this book is automation. Perhaps the clearest example of automation destroying jobs is the mechanisation of agriculture, a sector which accounted for 41% of US employment in 1900, and only 4% by 1970[xi]. (The corresponding figures for the UK are lower in absolute terms, but similar in relative terms: 9% in 1900 falling to 1% in 2000.[xii])

Many of the people who quit farm work moved to towns and cities to take up other jobs because they were easier, safer, or better paid. Many others were forced to find alternative employment because they could not compete with the machines. This process caused much suffering to individuals, but overall, the level of employment did not fall, and society became far richer – both in total and on average. More than one new job was created for every job that was lost.

The reason for this is that as machines replaced muscle power on the farm, humans had other skills and abilities to offer. Factories and warehouses took advantage of our manual dexterity and our ability to carry out a very broad range of activities. Office jobs used our cognitive ability. We turned our hands (often literally) to more value-adding work: you could say that we climbed higher up the value chain.

One-trick ponies

While the mechanisation of agriculture was a good news story for humans, it was less positive for the horse, which had nothing to offer beyond muscle power. 1900 was probably when the US reached "peak horse", with a population of 21m. That number fell to just 3m by 1960[xiii].

Artificial intelligence systems and their peripherals, the robots, are increasingly bringing flexibility, manual dexterity, and cognitive ability to the automation process. One of the big questions addressed in this book is: as computers take over the role of ingesting, processing and transmitting information, will there be anywhere higher up the value chain for humans to retreat to? In other words, can we avoid playing the role of the horse in the next wave of automation? Are we approaching "peak human" in the workplace?

Mechanisation and automation

What went on in farms was mechanisation rather than automation, and the distinction is important. Mechanisation is the replacement of human and animal muscle power by machine power; a human may well continue to control the whole operation. Automation means that machines are controlling and overseeing the process as well: they continuously compare the operation to a pre-set set of parameters, and adjust the process if necessary.

Although the word "automation" was not coined until the 1940s by General Electric,[xiv] this description applies pretty well to the operation of 19th-century steam engines once James Watt had perfected his invention of governors. Automated controllers which were able to modify the operation more flexibly became increasingly common in the early 20th century, but the start-stop decisions were still normally made by humans.

In 1968 the first programmable logic controllers (PLCs) were introduced[xv]. These are rudimentary digital computers which allow far more flexibility in the way an electrochemical process operates, and eventually general-purpose computers were applied to the job.

The advantages of process automation are clear: it can make an operation faster, cheaper, and more consistent, and it can raise quality. The disadvantages are the initial investment, which can be substantial, and the fact that close supervision is often necessary. Paradoxically, the more efficient an automated system becomes, the more crucial the contribution of the human operators. If an automated system falls into error it can waste an enormous amount of resources and perhaps cause significant damage before it is shut down.

Let's take a look at how automation has affected some of the largest sectors of the economy.

Retail and "prosumers"

Retail is a complicated business and there have been attempts to automate many of the processes required to get goods from supplier to customer, and payment from customer to supplier. Demand forecasting, product mix planning, purchasing, storage, goods handling, distribution, shelf stacking, customer service and many other aspects of the business have been automated to varying extents in different places and at different times.

The retail industry has also given us the clearest examples of another, associated phenomenon known as "prosumption". This term was coined in 1980 by the American futurist Alvin Toffler, one of the leading thinkers about the trends we are discussing in this book. At the same time as organisations automate many of their processes, they enlist the help of their customers to streamline their operations. In fact, they get their customers to do some of the work that was previously done for them. The reason why consumers accept this (indeed welcome it) is that the process speeds up, and becomes more flexible – more tailored to their wishes.

Toffler first described this process in *FutureShock* (1970), and in *The Third Wave* (1980) he defined a "prosumer" as a consumer who is also involved in the production process. Where once people were passive recipients of a limited range of goods and services designed or selected by retailers, he foresaw that we would become increasingly involved in their selection and configuration.

Perhaps the simplest example of what he meant is the purchase of gasoline. This dangerous substance was traditionally dispensed by pump attendants, but Richard Corson's invention of the automatic shut-off valve enabled the job to be taken over by customers. Nowadays most consumers in developed countries dispense their own gasoline at self-service pumps. This saves money for the retailer and time for the consumer.[xvi]

Supermarkets have often led the way in automation and prosumption because they are owned by massive organisations with the budgets and the sophistication to invest in the systems needed. Once upon a time, what marketers call fast-moving consumer goods (foods, toiletries, etc.) were requested one at a time by the shopper at a counter and fetched individually by the shopkeeper or his assistant. As these general stores firms grew bigger and more sophisticated they built large stores where shoppers fetched their own items, and presented them for processing at checkouts, like components on a car assembly line. Later on, self-service tills were installed, where shoppers could scan the barcodes of their goods themselves, speeding up the process considerably. Soon, RFID tags[xvii] on goods will enable you to wheel your trolley full of items out of the store and to your car without the fuss of unloading and re-loading them at a checkout.

At each stage of this evolution, the involvement of the consumer in selecting and transporting each item increases, and the requirement for shop staff involvement reduces. This latter effect is disguised because, as society gets richer, people buy many more items, so the store needs more

staff even though their involvement in each individual item is less.

Online shopping is perhaps the ultimate prosumer experience. Consumer reviews replace the retailer's sales force, and its algorithms do the up-selling.

Call centres

Of course, automation and prosumption is not always to the benefit of consumers. In markets where switching costs or partial monopolies dilute the standards-raising effect of competition, companies can save money for themselves in ways which actually make life worse for their customers. We are all familiar with call centres where (for instance) utility companies and banks have automated their customer service operations, obliging frustrated customers to plough through various levels of artificial un-intelligence in order to get their problem resolved. The customer would be much better off if a human picked up the call immediately, but that would cost the companies a lot more money, and they have no incentive to incur that cost.

Things are improving, however, as the AI used in call centres advances. Just as most people choose to withdraw cash from ATMs rather than venture into the bank and wait in line for a human cashier, many call centre operations are now getting good enough at handling or triaging problems that we may soon prefer to deal with the automated system than with a human.

Food service

The automation of service in fast food outlets seems to have been just around the corner for decades. Indeed, elements of it have been a reality for years in Oriental-style outlets like Yo, Sushi!, but it has so far failed to spread to the rest of the sector. There are several reasons for this, including the relatively low labour cost of people working in fast food outlets, and the need for every single purchase to be problem-free, and if not, for there to be a trained human on hand to solve any problem immediately. If a hands-free wash basin fails 5% of the time it is no big deal, but it would be a very big problem if 5% of meals were inedible, or delivered to the wrong customer. Three restaurants in Guangzhou, southern China, which trumpeted their use of robot waiters had to abandon the practice because the machines simply weren't good enough.[xviii]

A combination of factors is poised to overcome this resistance. Increases to the cost of labour caused by rising minimum wage legislation, declining costs of the automated technology, greater cultural acceptance of interacting with machines, and above all, the improved performance of the automated technology. It is increasingly flexible, and it goes wrong less often. McDonald's is one of the many fast food chains that are introducing touch-screen ordering and payment systems in their restaurants, and it is trialling an automated McCafe kiosk in a restaurant in Chicago.[xix] KFC (formerly known as Kentucky Fried Chicken) has a store in Shanghai where customers' orders are taken by a robot equipped with voice recognition software.[xx] Our robot overlords have found a Colonel.

Manufacturing

Car manufacturing has traditionally incurred relatively high labour costs. The work involves a certain amount of physical danger, with heavy components being transported, and metals being cut and welded. It is also a sector where a lot of the operations can be precisely specified and were highly repetitive. These characteristics make it ripe for automation, and the fact that the output (cars) are high-value items means that investment in expensive automation systems can be justified. Around half of all the industrial robots in service today are engaged in car manufacturing.[xxi]

Despite the recession, sales of industrial robots grew at 10% a year from 2008 to 2013, when 178,000 were sold worldwide. Sales in 2014 jumped 29% to 229,000, and the International Federation of Robotics expects the number to jump a further 75% to 400,000 by 2018. China became the biggest market in 2013, installing 37,000 robots compared with 30,000 in the USA[xxii].

Until recently, the industrial robots used in car manufacturing (and elsewhere) were expensive, inflexible, and dangerous to be around. But the industrial robotics industry is changing: as well as growing quickly, its output is getting cheaper, safer and far more versatile.

A landmark was reached in 2012 with the introduction of Baxter, a 3-foot tall robot (6 feet with his pedestal) from Rethink Robotics. The brainchild of Rodney Brooks, an Australian roboticist who used to be the director of the MIT Computer Science and Artificial Intelligence Laboratory, Baxter is much less dangerous to be around. By early 2015, Rethink had received over $100m in funding from venture capitalists, including the investment vehicle of Amazon founder Jeff Bezos. Baxter was intended to disrupt the industrial robots market by being cheaper, safer, and easier to programme. He is certainly cheaper, with a starting price of $22,000. He is safer because his arm and body movements are mediated by springs, and he carries an array of sensors to detect the presence nearby of squishy, fragile things like humans. He is easier to programme because an operator can teach him new movements simply by physically moving his arms in the intended fashion.

Baxter's short life has not been entirely plain sailing. Sales did not pick up as expected, and in December 2013, Rethink laid off around a quarter of its staff. One of the competitors stealing sales from Rethink is Universal Robots of Denmark, a manufacturer of small- and medium-sized robot arms. Universal increased sales to €30m in 2014, and aims to double its revenues every year until 2017.

But Rethink remains well-funded, and in March 2015 it introduced a smaller, faster, more flexible robot arm called Sawyer. It can operate in more environments than Baxter, and can carry out more intricate movements. It is slightly more expensive, at $29,000. Rethink and Universal, along with other companies like the Swiss firm ABB and the German firm KUKI, are making industrial robots more effective, more affordable, and more widespread.

Warehouses

Kiva Systems was established in 2003, and acquired by Amazon in 2012. Kiva produces robots which collect goods on pallets from designated warehouse shelves and deliver them to human packers in the bay area of the warehouse. Amazon paid $775m for the nine-year old company and promptly dispensed with the services of its sales team. Re-named Amazon Robotics in August 2015, it is dedicated to supplying warehouse automation systems to Amazon, which obviously considers them an important competitive advantage.

Secretaries

Most of the examples of automation given above involve manual work. There is one occupation which depends almost entirely on cognitive skills which has been largely automated out of existence: secretaries. In the 1970s, managers had secretaries, and generally did little work on computers themselves. In 1978, secretary was the most common job in 21 US states (41% of them). Today, many managers much of their day staring at computer screens, and secretary is the most common job in only 4 US states (8% of them).[xxiii]

2.4 - The Luddite fallacy

Ned Ludd

A person can have a big impact on society without going to the trouble of actually existing. In 1779, Ned Ludd was a weaver in Leicester who responded to being told off by his father (or perhaps his employer) by smashing a machine. Or maybe he wasn't – the truth is, we don't know. He certainly wasn't the leader of an organised group of political protesters. Nevertheless, in the decades following his alleged outburst, his name was commonly used to take the blame for an accident or an act of vandalism.

As Britain pioneered the industrial revolution in the late eighteenth and early 19th century, many of its people attributed their economic misfortune to the introduction of labour-saving machines. They were no doubt partly correct, although poor harvests, and the Napoleonic Wars against France were also to blame. There was a short-lived phenomenon of organised protest under the banner of Luddism in Nottingham in 1811-13: death threats signed by King Ludd were sent to machine owners.

The government responded harshly, with a show trial of 60 men (many of them entirely innocent) in York in 1813. Machine breaking was made a capital offence. Riots continued sporadically, notably in 1830-31, when the Swing rioters in southern England attacked threshing machines and other property. Around 650 of them were jailed, 500 sent to the penal colony of Australia, and 20 hanged.[xxiv]

The fallacy

The Luddites, and other rioters, were not making a general economic or political observation that the introduction of labour-saving machinery inevitably causes mass unemployment and privation. They were simply protesting against their own dire straits, and demanding urgent help from the people who were obviously benefiting.

It is therefore slightly unfair to them that the term “Luddite fallacy” has become a pejorative term for the mistaken belief that technological development necessarily causes damaging unemployment. (Although, given the hunger they were experiencing, they would probably regard the slur as a very minor irritation.)

The Luddite fallacy pre-dates the industrial revolution, and has taken in quite a few heavyweight thinkers down the years. As long ago as 350BC, the Greek philosopher Aristotle observed that if automata (like the ones said to be made by the god Hephaestus) became so sophisticated that they could do any work that humans do, then workers - including slaves - would become redundant.[xxv]

During the early 19th century, when the industrial revolution was in full swing, most members of the newly-established social science of economics argued that any unemployment caused by the introduction of machinery would be resolved by the growth in overall economic demand. But there were prominent figures who took the more pessimistic view, that innovation could cause long-term unemployment. They included Thomas Malthus, John Stuart Mill, and even the most respected economist of the time, David Ricardo.[xxvi]

The Luddite fallacy and economic theory

The debate can get quite technical, but there are two reasons why it has been correct to reject the Luddite fallacy up until now. The first reason is economic theory: companies introduce machines

because they increase production and cut costs. This increase in supply builds up the wealth in the economy as a whole, and hence the demand for labour.

Say's Law, named after French economist Jean-Baptiste Say, holds that supply creates its own demand, and Say argued that there could not be a "general glut" of any particular goods. Of course we do see gluts in sectors of the economy, but an adherent of Say's Law would argue these are the unintended consequences of interventions in free markets, usually by governments. This law became a major tenet of classical economics, but it was rejected emphatically by British economist John Maynard Keynes, and is not widely accepted today.

But many economists would accept a broader interpretation of the law which states that reducing the cost of a significant product or service will free up money which was previously allocated to it. This money can then be spent to buy more of the item, or other items, thereby raising demand generally, and creating jobs. This assumes, however, that the money freed up is not spent on expensive assets that generate no employment, or invested in companies that employ very few people.

Economists also point out that the Luddite fallacy also depends on a misapprehension about economics called the "Lump of Labour Fallacy", which is the idea that there is a certain, fixed amount of work available, and if machines do some of it then there is inevitably less for humans to do. In fact, economies are more organic and more flexible: they respond to shifts, and innovate to grow. New jobs are created as old ones disappear and the former outnumber the latter.

The Luddite fallacy and economic experience

The second reason to reject the Luddite fallacy hitherto is rather better: history has proved it to be wrong. A great deal of machinery has been deployed since the start of the industrial revolution, and yet there are more people working today than ever before. Put simply, if the Luddite fallacy was correct we would all be unemployed by now.

A study published in August 2015 by the business consultancy Deloitte analysed UK census data since 1871 and concluded that far more jobs have been created than destroyed by technology in that time.[xxvii] Furthermore, the study argued that the quality of the jobs has improved. Where people used to do dangerous and gruelling jobs on the land, and hundreds of thousands used to do the work now done by washing machines, many more Britons are now employed in caring and service jobs. In the last two decades alone there has been a 900% rise in nursing assistants, a 580% increase in teaching assistants, and a 500% increase in bar staff – despite the closure of so many of the country's pubs. (The authors refrained from commenting on the news that the number of accountants has doubled.)

So in the long run, the Luddite fallacy is just that – a fallacy. But in the short run the Luddites had a point. Economists do think that in the first half of the 19th century, wages failed to keep pace with increases in labour productivity. An economist named Arthur Bowley observed in the early 20th century that the share of GDP which goes to labour is generally roughly equal to that which goes to capital,[xxviii] but in the first half of the 19th century, the share of national income taken by profit increased at the expense of both labour and land. The situation changed again in the middle of the century and wages resumed their normal growth in line with productivity. It may be that the slippage in wages was necessary and inevitable to enable enough capital to be accumulated to fuel the investment in technological change.

The period in the early 19th century when wage growth lagged productivity growth is known as the Engels pause, after the German political philosopher Friedrich Engels, who wrote about it in the 1848 "Communist Manifesto", which he co-authored with Karl Marx. The effect ceased at pretty much the same time as he drew attention to it, which may explain why it is not better known.[xxix]

Even in the long run, the picture is not all rosy. A French economist named Gilles Saint-Paul has developed a formula which shows that while demand for unskilled human labour declines, the demand for skilled human capital increases faster. But a side effect can be the increase in income inequality.[xxx]

Is it different this time?

Mechanisation and automation has displaced workers on a huge scale since the beginning of the industrial revolution. It has imposed considerable suffering on individuals, but has led to greater wealth and higher levels of employment overall. The question today is whether that will always be true. As machines graduate from offering just physical labour to offering cognitive skills as well, will they begin to steal jobs that we cannot replace? If the second half of the 19th century saw "peak horse" in the workplace, will the first half of the 20th century see "peak human"? In other words, is it different this time?

3 – Is it different this time?

In this chapter I will argue that the arrival of machine intelligence is also the arrival of a different kind of automation, which spells the end of paid work for many or most people.

We will start in section 3.1 by looking at the most popular books to argue the case for technological unemployment; we will see how they shy away from the logical conclusion of their arguments. We will also hear support for their argument from a couple of unexpected sources. In section 3.2 we will briefly review some of the academic studies, before hearing in section 3.3 from some sceptics: people who think this talk of widespread joblessness is simply the Luddite fallacy at work.

In later sections of this chapter we will explore in more detail the reasons to believe that it really is different this time.

3.1 – Prophets of change

Martin Ford

Martin Ford is the author of perhaps the best book published so far about artificial intelligence causing technological unemployment. "The Lights in the Tunnel" (2009) provoked fierce debate, and his follow-up, "Rise of the Robots" (2015) fleshed out his arguments, and responded to the criticism which the first book attracted. Awarding it the 2015 *Financial Times* and McKinsey Business Book of the Year, Lionel Barber, the *Financial Times*' editor, called it "a tightly-written and deeply-researched addition to the public policy debate … The judges didn't agree with all of the conclusions, but were unanimous on the verdict and the impact of the book."

Ford is well-placed to talk about what technology will do to the world of work. He has a quarter-century of experience in software design, and he lives and works in Silicon Valley, where he runs a software development company. His writing is calm and measured, with an engaging humility.

Exponentials and automation

Ford opens "Rise of the Robots" with a dramatic illustration of the power of exponential increase – the cumulative doubling which is driving digital innovation. He asks us to consider driving a car at five miles an hour, and then doubling our speed 27 times over. The resulting speed would be 671 million miles an hour – fast enough to travel to Mars in five minutes[xxxi]. This, he points out, is the number of doublings that computer power has gone through since the invention of the integrated circuit in 1958. This doubling phenomenon is known as Moore's Law, after one of the founders of the chip manufacturer Intel. We will return to this exponential growth later in this chapter, as understanding it is fundamental to comprehending the scale of the changes that are coming our way.

The book argues that AI systems are on the verge of wholesale automation of white collar jobs – jobs involving cognitive skill such as pattern recognition and the acquisition, processing and transmission of information. In fact it argues that the process is already under way, and that the US is experiencing a jobless recovery from the Great Recession of 2008 thanks to this automation process. Ford claims that middle class jobs in the US are being hollowed out, with average incomes going into decline, and inequality increasing. He acknowledges that it is hard to disentangle the impact of automation from that of globalisation and off-shoring, but he remains convinced that AI-led automation is already harming the prospects of the majority of working Americans.

In fact, since Ford's book was published the US employment figures have improved considerably, and the unemployment rate hovers around 5%, which is considered close to full employment. However, many middle-class Americans do feel squeezed, having been obliged to accept part-time work, or having missed out on wage rises. This suggests that technological unemployment has not yet begun to really bite, but we might be seeing the early warning signs.[xxxii]

Ford pauses to review the prospects for disruption of two sectors of the economy which have so far been relatively unscathed by the digital revolution – education and healthcare. Although there is fierce resistance to the replacement of human activity by AIs in these areas – for instance in essay marking – Ford argues that no industry can ignore for long the benefits of cheaper, faster, more reliable ways of providing their products and services. He goes on to point out that the companies and industries which today are nascent and fast-growing, and tomorrow will be

economic giants, are extremely parsimonious employers of humans. AirBnB, the peer-to-peer rooms rental business, for example, achieved a market cap of $20bn in March 2015 with just 13 employees.

The challenge of UBI

The final chapters of "Rise of the Robots" explore the consequences of the trends which Ford has described. Can an economy thrive and grow if a large minority of people cannot find sufficient work to give themselves and their families a decent life? Would the consequent rise in inequality be economically harmful? More fundamentally, how will these unemployed or under-employed people make ends meet? To Ford, the answer to this last question is clear: governments will need to raise the taxes paid by those who are still working to provide an income for those who are not. But he is acutely aware of the political difficulties that this proposal faces: "American politicians are terrified to even utter the word 'tax' unless it is followed immediately by the word 'cut'."[xxxiii]

In fact, Ford seems daunted by the situation: "The political environment in the United States has become so toxic and divisive that agreement on even the most conventional economic policies seems virtually impossible," he writes. "A guaranteed income is likely to be disparaged as 'socialism'", and "The decades-long struggle to adopt universal health coverage in the United States probably offers a pretty good preview of the staggering challenge we will face in attempting to bring about any whole-scale economic reform."

Ford thinks that most people will probably still be able to find some form of paid employment – just not enough to make a decent living. Unwilling to give up on traditional American ideals like the free market, a capitalist economy and indeed the Protestant work ethic, he advocates a universal basic income of only $10,000 a year - a level low enough to leave the incentive to find work in place. Even so, he is pessimistic about the prospect of persuading his fellow Americans to adopt the idea: "a guaranteed income will probably remain unfeasible for the foreseeable future."

Andrew McAfee and Erik Brynjolfsson

As a pair of MIT professors[xxxiv], McAfee and Brynjolfsson bring academic credibility to their book on AI automation, "The Second Machine Age". They have helped to validate the discussion of the possibility of technological unemployment.

Their book (and their argument) is in three parts. The first part (chapters 1 to 6 inclusive) describes the characteristics of what they call the second machine age. They warn readers that their recitation of recent and forthcoming developments may seem like science fiction, and their prose is sometimes slightly breathless: even tenured professors can get excited about the speed of technological change and the wonders it produces.

The Bounty and the Spread

The second part of the book (chapters 7 to 11) explores the impact of these changes, and in particular two phenomena, which they label "bounty" and "spread". "Bounty" is the "increase in volume, variety and quality, and the decrease in cost of the many offerings brought on by technological progress. It's the best economic news in the world today." This part of the book could have been written by Peter Diamandis, author of "Abundance" and "Bold", and a leading evangelist for the claim that the exponential growth in computer power is leading us towards utopia.

“Spread” seems to be a synonym for inequality, although the authors are strangely reluctant to use that word.[xxxv] It is “ever-bigger differences among people in economic success”. This part of the book could have been written by a member of the Occupy movement[xxxvi]. “Spread is a troubling development for many reasons, and one that will accelerate in the second machine age unless we intervene.”

Brynjolfsson and McAfee pose the question whether bounty will overcome the spread. In other words, will we create an economy of radical abundance, in which inequality is relatively unimportant because even though a minority is extraordinarily wealthy, everyone else is comfortably off? Their answer is that current evidence suggests not. Like Martin Ford, they think the American middle class is going backwards financially, and they think this trend will continue unless remedial action is taken.

Hanging onto work

So the third and final part of the book discusses the interventions which could maximise the bounty while minimising the spread. In particular, Brynjolfsson and McAfee want to answer a question they are often asked: “I have children in school. How should I be helping them prepare for the future?”[xxxvii] They are optimistic, believing that for many years to come, humans will be better than machines at generating new ideas, thinking outside the box (which they call “large-frame pattern recognition”) and complex forms of communication. They believe that humans' superior capabilities in these areas will enable most of us to keep earning a living, although they think the education system needs to be re-vamped to emphasise those skills, and downplay what they see as today's over-emphasis on rote learning. They praise the Montessori School approach of “self-directed learning, hands-on engagement with a wide variety of materials … and a largely unstructured school day.” They also have high hopes for digital and distance learning, which use “digitisation and analytics to offer a host of improvements.”[xxxviii]

Brynjolfsson and McAfee offer a series of further recommendations which they say are supported by economists from across the political spectrum: pay teachers more, encourage entrepreneurs, enhance recruitment services, invest in scientific research and infrastructure improvements, encourage immigration by the world's talented migrants, and make the tax system more intelligent.

These seem somewhat unremarkable proposals, and the authors acknowledge that their effectiveness may peter out as the 2020s progress, and machines become even smarter. Looking further ahead, they warn against any temptation to try to arrest the progress in AI, and also against any temptation to move away from the tried and tested economic system of capitalism, which they claim (paraphrasing Churchill's quip about democracy) “is the worst form of [economy] except for all the others that have been tried.”[xxxix]

The authors are very keen on Voltaire's dictum that “work saves a man from three great evils: boredom, vice and need.” They are therefore wary of universal basic income, believing that an absence of work will engender boredom and depression. Instead, they argue for a negative income tax, which incentivises work. With a negative income tax of 50%, if you earn a dollar, the government gives you an additional 50 cents. They cast around for ways to keep us all in work, and rather tentatively suggest a range of exotic schemes, such as a cultural movement to prefer goods made by humans rather than machines.

Other voices

Richard and Daniel Susskind

Father and son team Richard and Daniel Susskind published "The Future of the Professions: How Technology Will Transform the Work of Human Experts" in October 2015. Richard Susskind has impressive credentials, having worked in legal technology since the early 1980s, advised numerous government and industry bodies, and garnered a clutch of honorary fellowships from prestigious universities.[xl] Perhaps even more impressive is that he seems to have retained the respect of his subjects while lambasting them as inefficient, and doomed to extinction.

The Susskinds describe the "grand bargain" whereby members of the professions (lawyers, doctors, architects, etc.) enjoy a lucrative monopoly over the provision of certain kinds of advice in return for policing the standard of that provision. They argue that this bargain has broken down, with many professional services now being available only to the rich and well-connected. They demonstrate how important this is by illustrating the size of the professions. Healthcare in the US alone now costs $3 trillion a year – more than the GDP of the world's fifth-largest country. The combined revenue of the Big Four accounting firms, at $120 billion, is greater than the GDP of the sixtieth-largest country.

Based on 30 years' experience of the legal industry and backed up by extensive research, they paint two scenarios for its mid-term future. The first has professionals working closely with technology, their services enhanced as it improves. The second has most or all of the traditional tasks of professionals carried out by machines. The Susskinds believe that this second outcome is the inevitable one, since what the rest of society cares about is not interaction with humans, but getting our legal, medical and other problems sorted out with the minimum of fuss, risk and expense.

The Susskinds keep their focus on the professions, and refrain from making the obvious read-across to the economy as a whole. As a result, they have little to say about universal basic income, or the possibility of society fracturing. But they do note that once machines have taken on responsibility for most or all the tasks previously carried out by human professionals, big questions will be asked about who should own the machines. They don't provide answers to these questions, although they indicate their preference for some form of common ownership which does not involve the state. In this respect they deserve credit for following the logic of their arguments further than most people writing on the subject.

The book is written with refreshing clarity, precision and felicity of expression - and with such a gloomy message for its audience, that is probably just as well.

Scott Santens

Scott Santens is a writer and a campaigner for Universal Basic Income, based in New Orleans.[xli] He is a moderator of the Reddit Basic Income page, where he maintains a useful FAQ on the subject.[xlii] Self-employed since 1997, towards the end of 2015 he managed to procure a basic income for himself based on pledges from others who support his campaign, via the online giving site Patreon.

Jerry Kaplan

Serial entrepreneur Jerry Kaplan co-founded GO Corporation, which was a precursor to

smartphones and tablets, and was sold to AT&T. He also co-founded OnSale, an internet auction site which pre-dated Ebay, and was sold for $400m. He teaches at his alma mater, Stanford University, and writes books, including one called “Humans Need Not Apply”.

Its message is similar to “The Second Machine Age”: AI has reached a tipping point and is becoming powerfully effective. This will disrupt most walks of life (the computer, he observes, is blind to the colour of your collar), and unless we manage the transition well, the resulting economic instability and growing inequality could be damaging.

Like Ford, Brynjolfsson and McAfee, Kaplan thinks the existing market economy can survive this transition intact.

CGP Grey

Kaplan got the title “Humans Need Not Apply” from a video of the same name[xliii] which appeared on the internet a year before. Posted to YouTube in August 2014 by an Irish-American who goes by the name CGP Grey (his full name is Colin Gregory Palmer Grey[xliv]), the video attracted over 5 million views within a year.

The video is well-produced, engaging and persuasive. It contains plenty of technological eye-candy, and makes its points in punchy sound-bites – ideal for today's short attention spans. Unlike the books described above, it offers no solutions to the problems raised by AI and robotic automation, but - also unlike them – it suggests that capitalism cannot cope with what is coming.

Gary Marcus

A psychology professor at New York University, Gary Marcus has taken an intense interest in artificial intelligence, and where it is leading us. In February 2015 he told a CBS interviewer “Eventually I think most jobs will be replaced, like 75-80% of people are probably not going to work for a living... There are a few people starting to talk about it.”[xlv]

Federico Pistono

Federico Pistono is a young Italian lecturer and social entrepreneur. He attracted considerable attention with his 2012 book “Robots Will Steal Your Job, But That's OK”. A range of eminent people, including Google's Larry Page, were drawn to its optimistic and discursive style. (Google re-named itself Alphabet in October 2015, but most people still call it Google, so in this book I’ll mostly follow that convention.)

After making a forceful case that future automation will render most people unemployed, Pistono argues that there is no need to worry. Much of the book is taken up with musing on the nature of happiness – the word features in the titles of a quarter of its chapters. He is hopeful that we will all discover that the pursuit of happiness through material goods is a fool's errand, and he argues that salvation lies in downsizing. He offers the example of his own family, living in northern Italy. They spend $45,000 a year, but by getting rid of two of their three cars, growing their own food, and generating their own electricity, they can reduce this to $29,000 a year.

He also urges us all to educate ourselves – and encourage everyone else to do likewise – but more for personal fulfilment than in a vain attempt to remain employable.

Two unexpected voices

Andy Haldane

As the chief economist of the Bank of England, Andy Haldane isn't the most obvious person to be found musing about the benefits of universal basic income. But that is exactly what he did in a speech at the Trades Union Congress in November 2015[xlvi]. He wondered whether the displacement effect of automation, whereby jobs are destroyed, might start to outweigh the compensation effect, whereby automation raises productivity sufficiently to generate more demand and thus work.

In his speech, Haldane avoided giving a definitive answer to the question of whether we are nearing "peak human", but he raised many of the concerns explored in this book. He presented an estimate prepared by the Bank of England of the likelihood of automation of the jobs in a range of economic sectors in the UK, adapted from the estimates produced for the US by Frey and Osborne of the Oxford Martin School (of which, more below). The Bank estimated the UK's situation as slightly less alarming than that of the US, but not much. It found that roughly a third of jobs have a low probability of being automated out of existence, another third have a medium probability, and the final third have a high probability. Haldane avoided putting a specific timescale on this, and also avoided saying what would happen after that undisclosed period.

Martin Wolf

As the main financial columnist and associate editor at the *Financial Times*, Martin Wolf is the very epitome of a City establishment figure. He was described by US Treasury Secretary Larry Summers as "probably the most deeply thoughtful and professionally informed economic journalist in the world."[xlvii] Although the credit crunch and subsequent recession have re-kindled his youthful enthusiasm for Keynesian economics, it is still a surprise to read him advocating income redistribution and universal basic income, as he did in this article from February 2014:

"If Mr Frey and Prof Osborne [see below] are right [about automation]… we will need to redistribute income and wealth. Such redistribution could take the form of a basic income for every adult, together with funding of education and training at any stage in a person's life. ... The revenue could come from taxes on bads (pollution, for example) or on rents (including land and, above all, intellectual property). Property rights are a social creation. The idea that a small minority should overwhelming[ly] benefit from new technologies should be reconsidered. It would be possible, for example, for the state to obtain an automatic share in the income from the intellectual property it protects."[xlviii]

3.2 – Academic and consultancy studies

Numerous reports have been written about technological unemployment by academic organisations, consultancies, and think tanks. I have described some of the better-known ones here. Sometimes they reserve judgement or sit on the fence, but as far as possible, I present them in order of increasing scepticism about the proposition of widespread unemployability.

Frey and Osborne

Carl Benedikt Frey and Michael Osborne are the directors of the Oxford Martin Programme on Technology and Employment.[xlix] Their 2013 report "The future of employment: how susceptible are jobs to computerisation?" has been widely quoted. Its approach to analysing US job data has since been used by others to analyse job data from Europe and Japan.

The report analyses 2010 US Department of Labour data for 702 jobs, and in a curious blend of precision and vagueness, concludes that "47% of total US employment is in the high risk category, meaning that associated occupations are potentially automatable over some unspecified number of years, perhaps a decade or two." 19% of the jobs were found to be at medium risk and 33% at low risk. Studies which have extended these findings to other territories have yielded broadly similar results.

The methodology overlays rigour on guesswork. 70 of the jobs were categorised in a brainstorming session, and these categorisations were then extended to the other 632 jobs using calculations which will mystify anyone with only school-level maths, including Gaussian process classifiers - a statistical tool also used in deep learning AI systems. But it would be unfair to criticise the report for lack of rigour. Forecasting is not an exact science; the authors adopted the most scientific approach they could devise, and made no attempt to hide its subjective elements.

As well as sounding the alarm about the possibility of technological unemployment, the report suggests that the "hollowing out" of middle class jobs will stop. A 2003 paper by David Autor (of whom, more below) observed that income has increased for high earners and (albeit less rapidly) for low earners, but stagnated for medium-level earners. Maarten Goos and Alan Manning characterised this hollowing out as the favouring of "Lovely and Lousy jobs". Frey and Osborne argue that in the future, susceptibility to automation will correlate negatively with income and educational attainment, so the Lousy jobs will also disappear. They suggest that people will have to acquire creative and social skills to stay in work, but they don't appear to think that many of us will be able to change the fate that our employment history has assigned to us.

Frey and Osborne followed the 2013 report with another in February 2015, written in collaboration with senior bankers from Citibank. It provides insights into the impact of automation in a number of industry sectors, including stock markets, where the move from trading floors to digital exchanges reduced headcount by 50%. At first glance it is surprising to see bankers suggesting increased taxation to provide income for the unemployed, but it also seems they have little faith in it happening: "Such changes in taxation would seem sensible to us, but they would also be a reversal of the trends of the last few decades." They don't hold out much more hope for their other principal suggested remedy: "education alone is unlikely to solve the problem of surging inequality, [but] it remains the most important factor."

Gartner

Gartner is the world's leading technology market research and advisory consultancy. At its annual conference in October 2014, its research director Peter Sondergaard declared that one in three human jobs would be automated by 2025.[i] "New digital businesses require less labor; machines will make sense of data faster than humans can." He described smart machines as an example of a "super class" of technologies which carry out a wide variety of tasks, both physical and intellectual. He illustrated the case by pointing out that machines have been grading multiple choice examinations for years, but they are now moving on to essays and unstructured text.

The Millennium Project

The Millennium Project was established in 1996 by a coalition of UN organisations and US academic research bodies. Its "2015-16 State of the Future" contained a section on the future of work based on a poll of 300 experts from around the world. Although they mostly thought that technology would impact employment significantly, their collective estimates for long-term unemployment were relatively conservative. They expected global unemployment to reach only 16% in 2030, and just 24% in 2050.

Pew Research Center

The Pew Research Center published a report entitled "AI, robotics, and the future of jobs"[ii] in November 2014. The Center is part of the Pew Charitable Trusts, established in 1948 with over $5bn bequeathed by descendants of the founder of Sun Oil; the Center is the third-largest think tank in the US.

The Center sent a questionnaire to 12,000 selected experts and interested members of the public (mostly but not entirely American), and received 1,900 responses to the question "Will networked, automated, artificial intelligence (AI) applications and robotic devices have displaced more jobs than they have created by 2025?". A slight majority (52%) said no, arguing that technology has always created more jobs than it has destroyed, that it is not advancing fast enough to destroy so many jobs, and that regulatory intervention would stop it if necessary.

The 48% who thought there would be a net loss of jobs believed that the process was already in train, but that it would get much worse, and that inequality would become a severe problem as a result.

Both sides thought that the education system is doing a poor job of preparing young people for the new world of work, and also that the future of employment is not pre-ordained, but is susceptible to good policy.

Fundacion Innovacion BankInter

BankInter, based in Madrid, is one of the largest banks in Spain. In 2003 it established a Foundation to promote the creation of sustainable wealth in Spain through innovation and entrepreneurship. One of the Foundation's main activities is organising the Future Trends Forum, an international think tank which periodically gathers together a group of experts to discuss an important topic, and then produces reports and videos based on the conclusions of those discussions.

In June 2015 I took part in a meeting of the Future Trends Forum entitled "The Machine

Revolution", which addressed "how technological developments (internet, robotics, artificial intelligence, etc.) will boost employment and labour markets in the next decade." Compered by Chris Meyer, author of "Standing on the Sun", the delegates were a mixture of senior government officials from around the world, academic and commercial economists, investors and writers.

When the 34 experts at the meeting were asked whether we thought structural unemployment was a likely result, a slight majority said not. Towards the end of the meeting we each contributed two predictions to a collective timeline, which appears at the end of the report.[lii]

McKinsey

The world's most prestigious management consultancy firm weighed in on the subject of technological unemployment with an article published in its quarterly magazine in November 2015 entitled "Four fundamentals of workplace automation"[liii]. Billed as an interim report of an ongoing research project, its central argument was that instead of asking which jobs can be and will be automated, we should ask which *tasks* will be automated. Few people, it claimed, will find that their entire job disappears, but as much as 45% of the tasks people do at work can be automated with technology that is currently available.

The McKinsey consultants identified 2,000 different "activities" (e.g., greeting customers, demonstrating product features) for a selection of "occupations" (e.g., retail salesperson), and assessed which activities required the 18 "capabilities" which they deem susceptible to automation (e.g., understanding natural language, generating natural language, retrieving information).

They noted that the level of automatability will rise as machines become more capable. For instance, if and when machines equal the median human level of natural language comprehension, then the proportion of tasks which can be automated will rise from 45% to 58%.

At the time of publication, the authors concluded that only 5% of jobs were capable of being fully automated, but 60% of jobs could have 30% of their constituent activities automated. But rather than leading to a 30% headcount reduction, with the other 70% of activities being smeared among the remaining employees, they expect people to become more productive, as machines augment human performance.

Very highly-paid jobs tended to have fewer automatable activities (e.g., 20% for CEOs), but among low- and medium-paid jobs there was a fairly even spread. The consultants found that only 4% of the work activities carried out in the US require creativity at the median human level, and only 29% require a median human level of emotion sensing. Optimistically, they concluded from this that automation will enable humans to do better and more interesting work. Interior designers, for instance "could spend less time taking measurements, developing illustrations and ordering materials, and more time developing innovative design concepts."

Finally, McKinsey suggested that senior managers should increasingly pay close attention to the type, direction and potential of automation within their industry, as it will become a more and more important source of competitive advantage.

A swelling chorus

The reports described above are a selection of the most prominent ones published so far on the

subject of technological unemployment. They are not the only ones, and more are being produced every month – sometimes every week. There is no clear consensus about the likely impact on joblessness of machine intelligence in the coming years and decades. Nevertheless, the theme has an increasingly high profile in the media – it was a focus of the annual gathering of the super-rich and powerful in the ski resort of Davos in January 2016.

In the next section we will see that some people are firmly convinced it is a myth.

3.3 – Crying wolf

In this section we meet a selection of writers who are sceptical about the prospect of technological unemployment. They argue that it is all just a revival of the Luddite Fallacy.

David Autor

David Autor is a professor of economics at MIT. As noted above, he sounded the alarm in a 2003 paper about the "hollowing out" of middle class jobs in the USA – the fact that income has increased for high earners and (albeit less rapidly) for low earners, but stagnated for medium-level earners.

In an interview in October 2015[liv], he gave three reasons why he thinks that some observers have been unduly pessimistic, even hysterical, about the likelihood of job destruction. One is that machines complement and augment humans: they always have, and there is no reason to think that is about to change. The second is that machines increase productivity, which creates wealth, consumption; and demand, which creates more jobs.

The third reason is that humans are creative and ingenious. There are many important businesses and activities now that could not have been imagined 50 years ago. In fact, Autor accuses Martin Ford of arrogance in writing off human ingenuity.

In an article for the Journal of Economic Perspectives (summer 2015) entitled "Why are there still so many jobs?",[lv] Autor forecasts that people will retain a comparative advantage in so-called "human" attributes such as interpersonal interaction, flexibility, adaptivity. He argues that many jobs – like radiologist - combine these with the routine, predictable tasks where computers win. Autor believes it will not be possible to separate these two types of tasks, so humans will continue to carry out the whole bundle.

More generally, Autor is also one of those who believe the rate of change today is over-hyped. He believes that the effect of Moore's Law is substantially muted by regulatory and social frictions which slow down the adoption of new technologies, and he also argues that many technological advances simply don't translate into tangible improvements in the real world. For instance he accepts that his current computer is a thousand times faster than the one he used a few years ago, but he suspects it only makes him 20% more productive. It may be true, he teases, that a new washing machine has more processing power than NASA used to send Neil Armstrong to the moon in 1969, but the washing machine is still not going to the moon.

He derides some of the heralded achievements of AI researchers, arguing for instance that self-driving cars do not emulate human drivers, but instead rely on precise maps of the terrain which have to be prepared before the journey starts. This makes them less flexible than humans, and not fit to be released into the wild without human escorts. We will see how well these arguments stand up later in this chapter.

Although Autor is broadly optimistic about our future, he believes that much depends on the decisions that we take. "If machines were in fact to make human labour superfluous, we would have vast aggregate wealth but a serious challenge in determining who owns it and how to share it." He points out that Norway and Saudi Arabia both enjoy economic abundance (thanks to oil rather than AI), but they use it very differently. Norwegians, he says, work few hours per day and

are generally happy; Saudis import 90% of their labour and nurture terror.

Robin Hanson

Robin Hanson is an associate professor of economics at George Mason University, in Virginia, USA. Like David Autor, Hanson castigates Martin Ford for inappropriate motives, but whereas Autor accuses Ford of arrogance, Hanson alleges dishonesty: "In the end, it seems that Martin Ford's main issue really is that he dislikes the increase in inequality and wants more taxes to fund a basic income guarantee. All that stuff about robots is a distraction."[lvi]

After a few more jibes, Hanson addresses Ford's actual thesis. He starts by admitting that "Ford is correct that, … in the long run, robots will eventually get good enough to take pretty much all jobs. But why should we think something like that is about to happen, big and fast, now?" He attributes four arguments to Ford, and makes short work of the first three. The first is the Frey and Osborne study we reviewed in chapter 3.2, which Hanson dismisses as subjective. The second argument is the decline in labour's share of income since 2000, which Hanson replies could be caused by numerous other factors rather than technological automation. The third argument is the rapid fall in computer prices, which Hanson says has yet to cause any detectable unemployment.

"And then there is Ford's fourth reason: all the impressive computing demos he has seen lately." Hanson is referring, of course, to Google's self-driving cars, real-time machine translation systems, IBM's Watson and so on. Hanson is less impressed by these demonstrations of rapidly improving AI: "We do expect automation to take most jobs eventually, so we should work to better track the situation. But for now, Ford's reading of the omens seems to me little better than fortune telling with entrails or tarot cards."

Having unburdened himself of this cynicism, Hanson proceeds to offer a constructive suggestion. He advocates forecasting by means of prediction markets, where people place bets on particular economic or policy outcomes, like the level of unemployment at some future date. He argues that prediction markets give us a financial stake in being accurate when we make forecasts, rather than just trying to look good to our peers.

Tyler Cowen

A professor at George Mason University and co-author of an extremely popular blog, Tyler Cowen was New Jersey's youngest ever chess champion. He is a man with prodigiously broad knowledge and interests, and although he proposes some key ideas forcefully, there is always some nuance, and he dislikes simplistic and modish solutions. In two recent books, "the Great Stagnation" (2011) and "Average is Over" (2014), he paints a picture of America's future which is slightly depressing, but not apocalyptic. He is alive to the prospect of dramatically improved AI, and the effect it will have on employment. But he does not think widespread permanent unemployment will be one of its results.

For several years, Cowen has championed the claim that the US economy is hollowing out. He expects automation to continue this trend, perhaps to accelerate it. In an article for Politico magazine[lvii], he wrote:

"I imagine a world in which, say, 10 to 15 percent of the citizenry … is extremely wealthy and has fantastically comfortable and stimulating lives, equivalent to those of current-day millionaires,

albeit with better health care. Much of the rest of the country will have stagnant or maybe even falling wages in dollar terms." This grim outlook for the majority is softened because "they will have a lot more opportunities for cheap fun and cheap education [thanks to] all the free or nearly free services that modern technology makes available." But there is a sting in the tail for the real underclass. They, he says, "will fall by the wayside."

Cowen does not expect a universal basic income to be required. Nor does he expect riots. One reason is that the US population will be older: "By 2030, about 19 percent of the US population will be over 65; in other words, we'll be as old as Floridians are today." Floridians are a conservative lot, not given to mayhem. Another is that people will increasingly cluster geographically according to income. Few people in the poorer 85% will live in the hothouse cities of San Francisco and New York, and they will not have the wealth of Manhattan waved in their faces. And perhaps most important, the masses will inure themselves with the opiates of free entertainment and social media.

Geoff Colvin

Geoff Colvin is an editor at Fortune magazine and one of America's most experienced and respected journalists. In August 2015 he published "Humans Are Underrated: What High Achievers Know That Brilliant Machines Never Will." His previous book, "Talent is over-rated" (2006) advanced the proposition that hours of dedicated practice trumps talent in most endeavours, and it was an international best-seller.

His new book accepts that for the first time, technology may be reducing total employment rather than increasing it, but is sceptical for two reasons in particular. First, Colvin argues that because it is so hard to foresee the new types of jobs that are created when economies shift (just as web development and social media marketing were hard to foresee), we under-estimate how many of them there will be.

Colvin believes that skills of deep human interaction – empathy, storytelling, the ability to build relationships - will become far more valuable in the future, and many people will be able to prosper by bringing those skills into the evolving economy. "We're hard-wired by 100,000 years of evolution to value deep interaction with other humans (and not with computers). Those wants won't be changing anytime soon."[lviii]

Crying wolf

"The boy who cried wolf" is one of Aesop's fables, and a commonly-told children's story. The moral usually drawn is that people who earn a reputation for lying are later punished by being disbelieved.[lix] But the story has another lesson for us: a claim that was false in the past may be true in the future, and it can be dangerous to forget that. Automation has been going on for centuries, and past claims that it was causing permanent widespread unemployment have been proven wrong. So far. We should not be complacent when there are good reasons to think that this time, it may be different.

3.4 – AI to date

We now need to spend some time looking at the advances in artificial intelligence which have prompted the discussion we have just witnessed.

If you have read my previous book, "Surviving AI", parts of the next two sections may seem familiar. (And by the way, bless you.) They do contain updates, as the field is moving fast. We get back to new territory in chapter 3.6.

What is AI?

Intelligence is the measure of an agent's ability to achieve goals in a range of different environments.[lx] In both humans and machines, intelligence is not a single, unitary phenomenon. American psychologist Howard Gardner has distinguished nine types of human intelligence: linguistic, logic-mathematical, musical, spatial, bodily, interpersonal, intra-personal, existential and naturalistic. As you read that list, you are probably thinking that you are better at some than others, and that you know other people whose mixture of skills is different. It is the same with artificial intelligence.

Artificial intelligence is simply an intelligence that did not arise naturally by evolution, but was created by humans (or perhaps aliens). Many people think the term is unfortunate and temporary: cars are not called artificial horses, and planes are not called artificial birds. They prefer terms like machine intelligence, or cognitive computing. I sympathise, although for the moment at least, the term artificial intelligence, or AI, is the one understood by the broadest range of people. I will use the terms machine intelligence and artificial intelligence as synonyms.

The value of intelligence

Intelligence is, of course, the distinguishing feature of humans: it is the characteristic which sets us apart from other animals and makes us more powerful than them. And we are much, much more powerful than them. Genetically, we are almost identical to chimpanzees, and our brains are not much heavier than theirs per kilo of body weight. But the difference in structure between our brains means that there are 7 billion of us and only a few hundred thousand of them[lxi]. Their fate depends on our actions, and they are not even aware of that fact.

Our intelligence enables us to communicate, to share information and ideas, and to devise and execute plans of action. It also enables us to develop tools and technology. A single unarmed human would be slaughtered by a mammoth or a lion, but a group of humans working together, or a single human equipped with a rifle, can turn the tables very effectively.

Before we go any further we need to distinguish between intelligence and consciousness, and note that the former does not seem to require the latter. Insects display a level of intelligence – especially collectively – but we have no evidence that they are conscious to any significant extent. Intelligence and consciousness are both more in evidence among mammals, especially primates, but there does not appear to be a straight-line correlation between the two.

Machine intelligence is unlike animal intelligence in that machines can be super-human in very narrow fields – like performing mathematical calculations, or playing chess – but utterly

unintelligent in all other respects, and (so far as we can tell) completely lacking in any degree of consciousness.

We value intelligence highly, since it is the source of our power, but we value consciousness even more. Most humans are happy to kill and eat animals which we deem to have a lower level of consciousness than our own. There is no reason to suppose that humans have attained anywhere near the maximum possible level of intelligence, and it seems highly probable that we will eventually create machines that are more intelligent than us in all respects – assuming we don't blow ourselves up first. We don't yet know whether those machines will be conscious, let alone whether they will be more conscious than us – if that is even a meaningful question.

Artificial General Intelligence (AGI) and Superintelligence

As we noted in chapter 1, the term for a machine which equals or exceeds human intelligence in all respects is artificial general intelligence, or AGI. The day when the first such machine is built will be a momentous one, as the arrival of superintelligence will not be far beyond it. The likelihood of an intelligence explosion is commonly referred to as the technological singularity. This could be an astonishingly positive development for humankind, or a disastrously negative one.

I wrote about this extensively in my previous book, "Surviving AI", and will not cover that ground again here. Suffice to say, we should make strenuous efforts to ensure that if and when we do create the first machines which are destined become superintelligences, we experience a positive outcome rather than a negative one.

Anders Sandberg of Oxford University's Future of Humanity Institute summarised it well by saying that we should aim to become the mitochondria of superintelligence rather than its boot loader. He was referring to Elon Musk's metaphor for how, if we are unwise and / or unfortunate, we could create the thing which destroys us, and saying that we should aim instead for the fate of the prokaryotic cell which was absorbed by another, larger cell and became an essential component of a new, combined, and more complex entity, the first eukaryotic cell.

This book is concerned with the impact of "narrow" AI systems which fall considerably short of AGI.

A quiet revolution

Origins and winters

The science of artificial intelligence got started in 1956 at a conference held at Dartmouth College, in New Hamsphire. Since then it has gone through cycles of optimism and pessimism. Herbert Simon said in 1965 that "machines will be capable, within twenty years, of doing any work a man can do,"[lxii] and two years later Marvin Minksy said that "Within a generation ... the problem of creating 'artificial intelligence' will substantially be solved."[lxiii] These early claims turned out to be ill-founded, and later generations of researchers found their sources of funding dried up in so-called AI winters.

Some leading figures in the field today are worried that a similar fate may befall them because, they say, excessive claims are being made about the capabilities of AI systems today, and what can be achieved in the short term. This seems to me an ungrounded fear. Machine intelligence is the

target of enormous investments – by technology giants like Google and Facebook, by startups, by traditional companies like the automotive manufacturers, and by governments. These investments are being made because machine intelligence delivers results, and they will continue so long as that remains the case. It will only stop if the results stop coming.

Machine learning

In the last few years, the field of AI has undergone a quiet revolution. It goes by the name of machine learning, and a subset called deep learning has proved especially effective at tasks which were previously considered hard problems that were unlikely to be solved for many years to come.

The approach to AI which prevailed in its early days tried to reduce human thought to the manipulation of symbols, such as language and maths, which could be made comprehensible to computers. This became known as symbolic AI, or Good Old-Fashioned AI (GOFAI). Machine learning, by contrast, is the process of creating and refining algorithms which can produce conclusions based on data without being explicitly programmed to do so. The turning point came in 2012 when researchers in Toronto led by Geoff Hinton won an AI image recognition competition called ImageNet.[lxiv] Hinton is a British researcher now at Toronto University and Google, and perhaps the most important figure behind the rise of deep learning as the most powerful of today's AI techniques.

(The word algorithm comes from the name of a 9th-century Persian mathematician, Al-Khwarizmi.[lxv] It means a set of rules or instructions for a person or a computer to follow. It is different from a programme, which gives a computer precise, step-by-step instructions how to handle a very specific situation such as opening a spreadsheet, or calculating the sum of a column of figures. An algorithm can be applied to a wide range of data inputs. A machine learning algorithm uses an initial data set to build an internal model which it uses to make predictions; it tests these predictions against additional data and uses the results to refine the model. The way that some game-playing AIs become superhuman in their field is by playing millions of games against versions of themselves and learning from the outcomes.)

In deep learning, the algorithms operate in several layers, each layer processing data from previous ones and passing the output up to the next layer. The output is not necessarily binary, just on or off: it can be weighted. The number of layers can vary too, with anything above ten layers seen as very deep learning – although in December 2015 a Microsoft team won the ImageNet competition with a system which employed a massive 152 layers.[lxvi]

Deep learning, and especially artificial neural nets (ANNs), are in many ways a return to an older approach to AI which was explored in the 1960s but abandoned because it proved ineffective. While Good Old-Fashioned AI held sway in most labs, a small group of pioneers known as the Toronto mafia kept faith with the neural network approach. They were vindicated when it was discovered that applying them to very large data sets made them surprisingly effective.

How good is today's AI?

Games and quizzes

The game of chess used to be thought of as one of the most challenging intellectual pursuits a person could undertake. (Being rubbish at it, I still do.) It used to be thought that it would take

centuries for machines to become really good at it. That was a long time ago, of course, and we are much wiser now, because as long ago as 1997, IBM's Deep Blue beat Gary Kasparov, the world's best player, in a controversial but conclusive match. Nowadays, humans have no chance against even mid-level chess computers.

With the benefit of hindsight, we know we should have expected this. The rebarbative MIT professor Noam Chomsky observed that a computer winning a chess competition is no more surprising than a forklift truck winning a weightlifting contest. Maybe he was saying this long before Deep Blue beat Kasparov, but if so he was unusual. In truth it has often been hard to forecast what will be hard for machines to do. It turned out to be relatively easy to programme computers to do things that we find very hard, but very hard to teach them how to do things that we find easy, like tying our shoelaces. This is known as Moravec's paradox, after AI pioneer Hans Moravec.[lxvii]

This episode is also a good illustration of another phenomenon, which is that once a computer is able to perform a particular task better than humans, we dismiss it as simple, saying that the next challenge is the really hard one. Until it isn't.

In fact, once a machine is able to perform a particular task, we usually stop calling it artificial intelligence. This is known as Tesler's Theorem, which defines artificial intelligence as that which a machine cannot yet do.

IBM's next bravura AI performance came in 2011, when a system called Watson beat the best human players of the TV quiz game "Jeopardy", in which contestants are given an answer and have to deduce the question. Watson used "more than 100 different techniques … to analyze natural language, identify sources, find and generate hypotheses, find and score evidence, and merge and rank hypotheses." It had access to 200 million pages of information, including the full text of Wikipedia, but it was not online during the contest. The difficulty of the challenge is illustrated by the answer, "A long, tiresome speech delivered by a frothy pie topping" to which the target question (which Watson got right) was "What is a meringue harangue?" After the game, the losing human contestant Ken Jennings famously quipped, "I for one welcome our new robot overlords."[lxviii]

At the beginning of this chapter we noted that intelligence is not a single, unitary skill or process. The fact that Watson is an amalgam – some would say a kludge – of numerous different techniques does not in itself mark it out as different and perpetually inferior to human intelligence. It is nowhere near an artificial general intelligence which is human-level or beyond in all respects. It is not conscious. It does not even know that it won the Jeopardy match. But it may prove to be an early step in the direction of artificial general intelligence.

In January 2016, an AI system called AlphaGo developed by Google's DeepMind beat Fan Hui, the European champion of Go, a board game. This was hailed as a major step forward: the game of chess has more possible moves (35^{80}) than there are atoms in the visible universe, but Go has even more – 250^{150}.[lxix] The system uses a hybrid of AI techniques: it was partly programmed by its creators, but it also taught itself using a machine learning approach called deep reinforcement learning.

(Reinforcement learning is halfway between two other important forms of machine learning: supervised and unsupervised learning. In supervised learning the system is given an example to

follow at each step. In unsupervised learning there are no examples. In reinforcement learning there are rewards for successful steps and penalties for unsuccessful steps. The system has to figure out how to behave according to those signals.[lxx])

A match against the world champion Lee Se-Dol followed in March 2016. Se-Dol was confident, believing it would take a few more years before a computer could beat him. He was genuinely shocked to lose the series four games to one, and observers were impressed by AlphaGo's sometimes unorthodox style of play. AlphaGo's achievement was another landmark in computer science, and perhaps equally a landmark in human understanding that something important is happening, especially in the Far East, where the game of Go is far more popular than it is in the West.

DeepMind did not rest on its laurels. A month after its European Go victory it presented a system able to navigate a maze in a video game without access to any maps, or to the code of the game. Using a technique called asynchronous reinforcement learning, the system looked at the screen and ran scenarios through multiple versions of itself.[lxxi] The ability to navigate by sight, like humans do, will be invaluable for AIs in many real-world applications.

Self-driving vehicles

Another landmark demonstration of the power of AI began inauspiciously in 2004. DARPA offered a prize of $1 million to any group which could build a car capable of driving itself around 150-mile course in the Mojave Desert in California. The best contestant was a converted humvee named Sandstorm which got stuck on a rock after only 7miles.[lxxii] Eight years later, Google's self-driving cars have driven well over a million miles without being responsible for a single serious accident. It is true that they have been rear-ended by human drivers a few times, but this is because they obey traffic regulations, and we humans are not used to drivers doing that. A Google car drove into a bus on Valentine's Day 2016, but the facts of that incident remain somewhat ambiguous.[lxxiii]

The world's automotive manufacturers are now scrambling to master the technologies involved in producing self-driving cars. Toyota is investing billions of dollars in research facilities in Silicon Valley,[lxxiv] and in December 2015, Ford announced a joint venture with Google.[lxxv] Elon Musk, CEO of the startlingly impressive upstart car manufacturer Tesla, said in December 2015 that its first fully autonomous cars would be sold in two years.[lxxvi] Of course it will take longer than that for regulators to catch up, and it will take years for enough of the old, human-piloted cars to be replaced for self-driving cars to deliver the tremendous benefits they are capable of. We will explore this in more detail in chapter 3.8.

Search

We are strangely nostalgic about the future, and we are often disappointed that the present is not more like the future that was foretold when we were younger. 2015 was the 30th anniversary of the 1985 movie "Back to the Future", and it was also the year to which the hero travels at the end of the story. Journalists and commentators complained about the failure of hoverboards and flying cars to arrive, as predicted in the film.

We didn't get hoverboards, but we did get something even more significant. As recently as the late 20th century, knowledge workers could spend hours each day looking for information. Today, less than twenty years after Google was incorporated in 1998, we have something close to

omniscience. At the press of a button or two, you can access pretty much any knowledge that humans have ever recorded. To our great-grandparents, this would surely have been more astonishing than flying cars.

(Some people are so impressed by Google Search that they have established a Church of Google, and offer nine proofs that Google is God, including its omnipresence, near-onmiscience, potential immortality, and responses to prayer.[lxxvii] Admittedly, at the time of writing, there are only 427 registered devotees, or "readers", at their meeting-place, a page on the internet community site Reddit.[lxxviii])

In the early days, Google Search was achieved by indexing large amounts of the web with software agents called crawlers, or spiders. The pages were indexed by an algorithm called PageRank, which scored each web page according to how many other web pages linked to it. This algorithm, while ingenious, was not itself an example of artificial intelligence. Over time, Google Search has become unquestionably AI-powered.

In August 2013, Google executed a major update of its search function by introducing Hummingbird, which enables the service to respond appropriately to questions phrased in natural language, such as, "what's the quickest route to Australia?"[lxxix] It combines AI techniques of natural language processing with colossal information resources (including Google's own Knowledge Graph, and of course Wikipedia) to analyse the context of the search query and make the response more relevant. PageRank wasn't dropped, but instead became just one of the 200 or so techniques that are now deployed to provide answers. Like IBM Watson, this is an example of how AI systems are often agglomerations of numerous approaches.

In October 2015, Google confirmed that it had added a new technique called RankBrain to its search offering. RankBrain is a machine learning technique, and it was already the third-most important component of the overall search service.[lxxx] It is applied to the 15% of searches which comprise words or phrases that have not been encountered before, and converts the language into mathematical entities called vectors, which computers can analyse directly. Microsoft also uses machine learning techniques in its Bing search engine.

In February 2016 Google announced an important change of leadership in its search division: Amit Singhal was replaced by John Giannandrea.[lxxxi] Singhal had overseen the introduction of RankBrain, but was seen as having a bias against applying machine learning techniques to search because it is often impossible to know how the machine has reached its conclusions. Giannandrea has no such reservations: in his previous role he oversaw Google's entire artificial intelligence research activity, including deep learning. This succession is perhaps an allegory of the way that AI is taking over the internet, on the way to taking over everything else.

One of the benefits Google hopes to obtain by increasing its use of AI in search is extra ammunition in its competition with Amazon. Google's competitors in search are not Microsoft's Bing, and certainly not Yahoo. 39% of purchases made online begin at Amazon, compared with 11% at Google.[lxxxii] Improving that ratio is a key aim for the search giant. We have seen before with the relative decline of seemingly invincible goliaths like IBM and Microsoft how fierce and fast-moving the competition is within the technology industry. This is one of the dynamics which is pushing AI forward so fast and so unstoppably.

Image and speech recognition

Deep learning has accelerated progress at tasks like image recognition, facial recognition, natural speech recognition and machine translation faster than anyone expected. In 2012, Google announced that an assembly of 16,000 processors looking at 10 million YouTube videos had identified – without being prompted - a particular class of objects. We call them cats.[lxxxiii] Two years later, Microsoft researchers announced that their system – called Adam – could distinguish between the two breeds of corgi dogs.[lxxxiv] (Queen Elizabeth is famously fond of corgis, so Adam's skill would be invaluable in certain British social circles.)

In February 2015, Microsoft announced that its AI systems could identify an image better than humans according to the tests laid down by ImageNet, the world's top image-recognition competition.[lxxxv] A few days later, Google announced it had done even better.[lxxxvi] Not to be left out, Facebook posted an impressive demonstration video in November 2015.[lxxxvii]

We humans are very good at recognising each other's faces. Throughout history it has been vitally important to distinguish between members of your own group who will help you, and members of rival groups who may try to kill you. A Facebook AI system called DeepFace reached human-level ability to recognise human faces in March 2014, scoring 97% in a test based on a database of celebrity photos called Labeled Faces in the Wild (LFW).[lxxxviii] The following year it announced the ability to recognise faces even when they are not looking towards the camera, with 83% reliability. Google now offers the same functionality to users of Google +.[lxxxix]

These tech giants are sensitive to the privacy concerns that this raises, and are throttling back their offerings so as not to raise alarm. But of course there is no stopping the progress: the genie is well and truly out of the bottle.

In January 2016 Baidu (often described as China's Google) showed off a system called DuLight which uses a camera to capture an image of something in front of you, sends the image to an app on your smartphone, which identifies the object and announces what it is. One application of this is to help blind people know what they are "looking" at.[xc] You can download a similar app called Aipoly for free at iTunes.[xci]

Speech recognition systems that exceed human performance will be available in your smartphone soon.[xcii] Microsoft-owned Skype introduced real-time machine translation in March 2014: it is not yet perfect, but it is improving all the time. Microsoft CEO Satya Nadella revealed an intriguing discovery which he called transfer learning: "If you teach it English, it learns English," he said. "Then you teach it Mandarin: it learns Mandarin, but it also becomes better at English, and quite frankly none of us know exactly why."[xciii]

In December 2015, Baidu announced that its speech recognition system Deep Speech 2 performed better than humans with short phrases out of context.[xciv] It uses deep learning techniques to recognise Mandarin.

Learning and innovating

It can no longer be said that machines do not learn, or that they cannot invent. In December 2013, DeepMind demonstrated an AI system which used a deep learning technique called unsupervised learning to teach itself to play old-style Atari video games like Breakout and Pong.[xcv] These are games which previous AI systems found hard to play because they involve hand-to-eye co-ordination.

The system was not given instructions for how to play the games well, or even told the rules and purpose of the games: it was simply rewarded when it played well and not rewarded when it played less well. As the writer Kevin Kelly noted, "they didn't teach it how to play video games, but how to *learn* to play the games. This is a profound difference."[xcvi]

The system's first attempt at each game was feeble, but by playing continuously for 24 hours or so it worked out - through trial and error - the subtleties in the gameplay and scoring system, and played the games better than the best human player.

The DeepMind system showed true general learning ability. On seeing the demonstration, Google acquired the company for a reported $400m.

Emulating and predicting human cognitive abilities

In June 2015, pictures produced by Google's image recognition neural network fired the public imagination because of their surreal, hallucinogenic properties. In December 2015 a group of AI researchers at MIT published a paper about a system which was able to predict the memorability of images better than humans. The system, called MemNet, reviewed a dataset of 60,000 images, and classified them in 1,000 different ways. It was able to identify why certain images were more memorable than others.

To repeat, these systems are not conscious, and have no imagination. They are neither creating art nor being emotionally affected by the images they process. But that does not matter. They can analyse and process the images in ways that are important to us. And they can do it better than we can.

Is it just Google?

Google, Facebook and the other tech giants pioneered the use of machine learning, and for a while they were pretty much the only organisation with the expertise, the computing resources and the data to implement it. The joke was that machine learning was like teenage sex – everyone talked about it but pretty much no-one did it. That is changing.

Remember spam? In the late 2000s there was talk of it crashing the internet but now you rarely see it unless you look at your junk mail box. It was tamed by machine learning. The same is happening with user-generated content (UGC). We like to read the comments on news sites: many of them are dumb, but many are smart and funny. After all, there wouldn't be any point in crowd sourcing if the crowd was all stupid. But some of it needs a grammatical dry-clean to be useful, and the good stuff needs to be surfaced. This is increasingly being done with machine learning, and by companies far down the pyramid from the tech giants. Companies large and small are using machine learning to work out what information to present to their customers and targets at every encounter.[xcvii]

IBM says that its cognitive computing business, which depends heavily on machine learning, now accounts for over a third of its $81 billion annual revenues, and is the main focus for the company's growth. IBM Watson's best-known work today is in the medical sector, but it is also carrying out large-scale projects in food safety with Mars, and in personality profiling for recruitment firms and dating apps.[xcviii]

In December 2015, Elon Musk and Sam Altman, president of the technology incubator Y Combinator announced the formation of a new company called Open AI. They had recruited a clutch of the top machine learning professionals despite the efforts of Google and Facebook to hang onto them with eye-watering financial offers. There is some uncertainty about whether other companies controlled by Musk and Altman (like Tesla and Solar City) will have privileged access to technologies developed at Open AI, but the thrust of the company is to make advanced AI techniques more widely available in the hope that will de-risk them.[xcix]

Because it works, the use of machine learning will continue to grow – fast.

Summary

In the last chapter we saw Robin Hanson's cynicism about "all the impressive computing demos [Martin Ford] has seen lately," and Hanson's conclusion that Ford's "reading of the omens seems … little better than fortune telling with entrails or tarot cards." How impressed you are by demos is of course a personal matter. I side with Ford on this one. But a more important point is that we are still in the early days of AI development. Its rate of improvement is rapid, and what we will see in the next few years and decades will be startling.

3.5 – Exponential future

Big investments, different approaches

The science of artificial intelligence is advancing rapidly, with significant steps announced almost every month. Enormous resources are being devoted to achieving these advances.

Some of the cutting edge work in AI goes on in universities, but much of it happens inside the tech giants on the US West Coast. Four of them - Intel, Microsoft, Google and Amazon - are among the world's top ten R&D (research and development) spenders, with a combined budget in 2015 of $42bn.[c] This equals the entire R&D spend of the UK, both public and private.[ci] IBM, Apple and Facebook are not far behind, and are increasing their R&D spend sharply.

In addition there are around 1,000 startup companies basing their products and services on AI.[cii] But despite all this, it is still early days for the sector. By one count there were over 300 venture capital deals in AI-based companies during 2015, but 80% of them were for less than $5m, and 75% of them were in the US.[ciii]

Openness

In September 2015, Google announced an important change in strategy. Having built a very lucrative online advertising business based on algorithms and hardware which produced better search results than anyone else, it was open sourcing its current best AI software – a deep learning engine called Tensor Flow.[civ] The software was initially licensed for single machines only, so even very well resourced organisations weren't able to replicate the functionality that Google enjoys, but the move was significant. In April 2016 that restriction was lifted.[cv]

In October 2015, Facebook announced that it would follow suit by open sourcing the designs for Big Sur, the server which runs the company's latest AI algorithms.[cvi] Then in May 2016 Google open sourced a natural language processing programme playfully called Parsey McParseFace, and SyntaxNet, an associated software toolkit. Google claims that in the kinds of sentences it can be used with, Parsey's accuracy is 94%, almost as good as the 95% score achieved by human linguists.[cvii]

Open sourcing confers a number of advantages. One is a level of goodwill among the AI community. More importantly, researchers in academia and elsewhere will learn the systems, and be able to work closely with Google and Facebook - and indeed be hired by them. Also, having more smart people working with their systems means there are more smart people making suggestions about improvements and de-bugging.

So far, Apple, the world's largest technology company (and indeed at the time of writing the world's largest company by stock market capitalisation) is an exception to this trend towards open sourcing. There are signs that this is changing, and it may have to: people with academic training are generally more comfortable working in organisations that share their findings, and many of the best people with deep learning experience have significant academic backgrounds.

Extending connectivity

Another example of the enlightened self-interest of the technology giants is Google's and Facebook's initiatives to extend internet access from the current level of 3.2bn people to the remaining 57% of the world's population. Google is experimenting with fleets of helium balloons which are manoeuvred around in the stratosphere (from around 11 miles up, higher than airplanes), providing wi-fi connectivity to mobile devices below. In June 2015, Google signed a deal with the Sri Lankan government to make that country the first in the world to receive blanket wi-fi coverage.[cviii]

The following month, Facebook revealed test flights by scale models of a 42-meter wingspan drone which will beam internet connectivity from the stratosphere to special receivers with lasers and radio. In October 2015 the two companies announced that their respective teams are collaborating.

Exponentials

Doubling up

If you think of artificial intelligence as a car, algorithms are the engine control system, big data is the fuel, and computing power is the engine. (Big data is a term coined in the mid-1990s by John Mashey at Silicon Graphics, a computer firm, to describe very large and growing data sets which could yield surprising insights into a wide range of phenomena.[cix]) The engine is getting more powerful at an exponential rate: its performance is doubling repeatedly. It is impossible to understand the scale of change that we face in the coming years without comprehending the astonishing impact of exponential increase.

Imagine that you stand up and take 30 paces forward. You would travel around 30 yards (or metres, if you are outside Britain and its former colonies). Now imagine that you take 30 exponential paces, doubling the length each time. Your first pace is one metre, your second is two metres, your third is four metres, your fourth pace is eight metres, and so on.

How far do you think you would travel in 30 paces? The answer is, to the moon. In fact, to be precise, the 29th pace would take you to the moon; the 30th pace would bring you all the way back.

That example illustrates not just the power of exponential increase, but also the fact that it is deceptive and back-loaded. Here is another illustration of that. Imagine that you are in a football stadium (either soccer or American football will do) which has been sealed to make it water-proof. The referee places a single drop of water in the middle of the pitch. One minute later she places two drops there. Another minute later, four drops, and so on. How long do you think it would take to fill the stadium with water? The answer is 49 minutes. But what is really surprising – and disturbing – is that after 45 minutes, the stadium is just 7% full. The people in the back seats are looking down and pointing out to each other that something significant is happening. Four minutes later they have drowned.

The fact that exponential growth is back-loaded helps explain another phenomenon, known as Amara's Law, after the scientist Roy Amara. This states that we tend to over-estimate the effect of a technology in the short run and under-estimate the effect in the long run.[cx]

People often talk about the "knee" of an exponential curve, the point at which past progress seems sluggish, and projected future growth looks dramatic. This is a misapprehension. When

you compare exponential curves plotted for ten and 100 periods of the same growth, they look pretty much the same. In other words, wherever you are on the curve, the past always looks horizontal and the future always looks vertiginous.

The author John Lanchester describes how in 1996 the US government started building a new supercomputer to model the behaviour of nuclear explosions. The result was Red, the first machine to process more than a trillion floating point operations per second (a teraflop). It remained the world's fastest supercomputer until 2000, but by 2006 that level of processing was available to schoolchildren in the Sony PS3 gaming computer.[cxi] This is Moore's Law at work.

Moore's Law

In 1965, Gordon Moore was working for Fairchild Semiconductors when he published a paper observing that the number of transistors being placed on a chip was doubling every year. He forecast that this would continue for a decade, which his contemporaries considered extremely adventurous. In 1975 he adjusted the period to two years, and shortly afterwards a Caltech professor named Carver Mead coined the term Moore's Law. In 1968 Moore co-founded Intel, and following an observation by Intel executive David House that the performance of individual transistors was also improving, the Law is generally taken to mean that the processing power of $1,000 of computer doubles every 18 months.

Moore's Law is of course not a law, but an observation which became a self-fulfilling prophecy – a target and a planning guide for the semiconductor industry, and for Intel in particular. Moore's Law celebrated its 50th anniversary in April 2015. Given its importance in human affairs there was remarkably little fanfare.

Exponential curves do not generally last for long: they are just too powerful. In most contexts, fast-growing phenomena start off slowly, pick up speed to an exponential rate, and then after a few periods they tail off to form an S-shaped curve. However exponentials can continue for many steps, and in fact each of us is one of them. You are composed of around 27 trillion cells, which were created by fission, or division – an exponential process. It required 46 steps of fission to create all of your cells. Moore's Law, by comparison, has had 33 steps in the 50 years of its existence.

It would take another two decades for Moore's Law to run through the same number of steps as human cell division. By that time, if the Law holds, the total amount of computing power now available to Google will be available on an ordinary desktop computer. Imagine what we could achieve when every teenager has Google's computing power in her bedroom – and try to imagine how much power Google will have by then!

A number of other technological developments have been observed to progress at an exponential rate, including memory capacity, LEDs (which follow Haitz's Law[cxii]), sensors (where the cost per observation is decreasing exponentially[cxiii]), and the number of pixels in digital cameras.[cxiv]

No more Moore?

In 2015, Intel seemed uncertain about whether its own chip development will keep Moore's Law on track. This is important, as Intel (whose name is a conjunction of **int**egrated **el**ectronics) has

been the world's biggest chip manufacturer since 1991, and is also the world's third-largest R&D spending company (after Volkswagen and Samsung). It has led the miniaturisation of microchips.

In February 2015 Intel updated journalists on their chip programme for the next few years, and the schedule maintained Moore's Law's exponential growth.[cxv] The first chips based on its new 10 nanometre manufacturing process were due to be released in late 2016 / early 2017, after which Intel expected to move away from silicon, probably towards a III-V semiconductor such as indium gallium arsenide.[cxvi] (That 10 nanometres is the distance between the two nearest repeating features on the chip.)

But in July 2015, Intel CEO Brian Krzanich said that it was taking longer for the firm to cut the size of its transistors: "our cadence today is closer to 2.5 years than to 2." At the time of writing, the firm's smallest transistors are the 14 nanometer Skylake model, and the next size down will be the 10 nanometer Cannonlake, due in late 2017, a six-month delay.

Since 2007, Intel had pursued the development of its chips with a "tick-tock" cadence. The tick represented improvements in the manufacturing process, which enabled chip size to be reduced from 45nm to 32nm to 22nm to 14nm. The tock was improvements in the architecture. The new cadence announced by Krzanich was described by one observer as a move from tick-tock to tic-tac-toe, representing process (tic), architecture (tac), and optimisation and efficiency (toe).[cxvii]

However, Mark Bohr, a 37-year Intel veteran and senior fellow in the its processor technology team, argued that taking a longer view, Moore's Law still applied. He is working on the technology to get down to 5 nanometers.[cxviii]

(To provide some context, a human hair is 100,000 nanometers thick, so each hair on your head is 10,000 times thicker than Intel's next release of transistors. Silicon atoms are around 0.2 nanometers across, so a 5nm structure is about 20 atoms wide.)

Of course Intel is not the only game in town. In July 2015, IBM announced that it had a prototype chip of 7 nanometers, using silicon-germanium for some components.[cxix] Manufacturing at scale is very different from prototyping, however, and IBM did not expect to manufacture at scale for two more years. But this would be ahead of Intel's schedule, and the IBM announcement was especially important in heralding a successful move from Deep UltraViolet (DUV) to Extreme UltraViolet (EUV) lithography, which operates at much shorter wavelengths.

3D chips and new architectures

Moore's Law has undergone substantial transitions before. Until 2004, regular increases in the clock speeds of computer chips contributed a large part of their performance improvements. (See here[cxx] for an explanation of clock speeds, if you like that kind of thing.) Over-heating put a stop to this, and instead, chip manufacturers incorporated more processors, or "cores". Modern smartphones may have four cores or even eight, which means the processes they work on have to be broken down into pieces which are operated on in parallel.

However successfully the chip manufacturers prolong it, the existing architecture will reach its end point eventually. Researchers are hard at work on a number of technologies that could replace it. One of these is 3D chips. Placing chips side-by-side delays the signals between them and causes bottlenecks as too many signals try to use the same pathways. These problems can be eased if you

place the chips on top of each other, but this raises new problems. Silicon chips are fabricated at 1,800 degrees Fahrenheit, so if you manufacture one chip on top of another you will fry the one below. If you fabricate them separately and then place one on top of the other, you have to connect them with thousands of tiny wires.

In December 2015, researchers from Stanford announced a new method of stacking chips which they called Nano-Engineered Computing Systems Technology, or N3XT. They claimed this was a thousand times more efficient than conventional chip configurations.[cxxi] They did not give an estimate when mass production of N3XT chips might begin.

Another approach is to combine the memory chips with the traditionally separate processing chips, to reduce the amount of traffic between those two. Another is to design chips specifically to implement neural networks, which is the approach taken by an MIT team that announced the Eyeriss chip in February 2016.[cxxii]

In March 2016, scientists from IBM's TJ Watson Research Centre announced their belief that "resistive processing units" which combine CPU and memory on the same chip could accelerate the processing of machine learning algorithms as much as 30,000 times.[cxxiii] The following month the chip maker Nvidia unveiled its Tesla P100 GPU, optimised for machine learning and boasting huge performance gains.[cxxiv] And if you would like to know what on earth CPUs and GPUs are, go here.[cxxv]

Towards quantum computing

Another much talked-about route to more powerful machines is quantum computing. This is based on the idea that while classical computers use bits (**bi**nary dig**its**) which are either on or off, quantum bits (qubits) can be both on and off at the same time – known as superposition. This enables them to carry out a number of different calculations at once.

Google bought a quantum computer from Canadian company D-Wave in 2013, but was unable to demonstrate to everyone's satisfaction that it actually worked. This changed in December 2015, when Google's engineering director Hartmut Neven announced that its D-Wave computer was 100 million times faster than a traditional desktop computer in a "carefully crafted proof-of-concept problem".

Keeping qubits stable is very hard, but Google thinks it is getting close.[cxxvi] IBM and Microsoft are also bullish about their quantum computing projects.[cxxvii] If they are successful, machine learning will no longer require the massive data sets and extensive training which are necessary today, and computers will edge that bit closer to human-level capabilities.

So what? Where's this all heading?

Moore's Law was an observation which became a target-generator rather than being a description of a fundamental property of the world. Keeping it on track has involved numerous ingenious and unpredictable steps in the past. Commercial imperatives and sheer human inventiveness have managed it so far, and there are plenty of avenues being explored which could maintain that performance.

Using the specific definition arrived at by Gordon Moore and David House in the late 1960s, Moore's Law ceased to hold some years ago. But as Intel's Shekhar Borkar observes, the meaning of Moore's Law as far as users are concerned is broader – namely, that the power of $1,000-worth of computer doubles every couple of years. There are plenty of ways to keep that going, and plenty of incentives too.[cxxviii]

The people who are actually working on the technologies seem determined to maintain the pace of the advance. Moore is more, and more is better. Even the worst case predictions envisage continued rapid improvement in computer processing power, albeit perhaps slower than previously.

In December 2015, Microsoft's chief speech scientist Xuedong Huang noted that speech recognition has improved 20% a year consistently for the last 20 years. He predicted that computers would be as good as humans at understanding human speech within five years.

Geoff Hinton – the man whose team won the landmark 2012 ImageNet competition - went further. In May 2015 he said that he expects machines to demonstrate common sense within a decade.

Common sense can be described as having a mental model of the world which allows you to predict what will happen if certain actions are taken. Professor Murray Shanahan of Imperial College uses the example of throwing a chair from a stage into an audience: humans would understand that members of the audience would throw up their hands to protect themselves, but some damage would probably be caused, and certainly some upset. A machine without common sense would have very little idea of what would happen.

Facebook has declared its ambition to make Hinton's prediction come true. To this end, it established a basic research unit in 2013 called Facebook Artificial Intelligence Research (FAIR) with 50 employees, separate from the 100 people in its Applied Machine Learning team.[cxxix]

So within a decade, machines are likely to be better than humans at recognising faces and other images, better at understanding and responding to human speech, and may even be possessed of common sense. And they will be getting faster and cheaper all the time. It is hard to believe that this will not have a profound impact on the job market.

3.6 – What people do

Jobs and tasks

As we saw in chapter 3.2, consultants from McKinsey pointed out that machines often don't acquire the ability to automate entire jobs in one fell swoop. Instead they become able to automate certain of the tasks which people in those jobs perform. So what are these tasks? What exactly is it that people do for a living?

The economies in the developed world are dominated by services, like finance, health, education, entertainment, retail, transport and so on. In the UK, for instance, service industries account for 78% of GDP, with manufacturing accounting for 15%, construction for 6%, and agriculture less than 1%.[cxxx]

Processing information

In service industries, most tasks involve information: obtaining it, processing it, and passing it on to others. This is also true for many tasks in the manufacturing, construction and agricultural sectors. Obtaining information can involve carrying out research, asking colleagues, looking online or occasionally in books, or coming up with an original idea – which itself usually involves combining two or more ideas from elsewhere.

Processing information can mean checking its accuracy or relevance, determining its importance relative to other pieces of information, making a decision about it or performing some kind of calculation on it. Passing information on is increasingly achieved electronically, for instance by email or online work flow systems.

Obtaining, processing and passing on information can be solitary endeavours, or they can be carried out collaboratively with other people. Almost by definition, the solitary tasks can be carried out by a machine which possesses human-level (or above) ability to understand speech, recognise images, and a modicum of commonsense.

Working with people

Collaboration with other humans is different. Mostly – at least for the time being. It can take many forms: brainstorming with colleagues; preparing for and negotiating a deal which will yield benefit to both sides but maximise your own; pitching an idea to a self-important, unimaginative and prickly boss; coaching a subordinate who has talent, but is also naïve. These appear to be tasks which would be far harder for a machine to emulate.

And indeed they are, but probably not for long. Even now, plenty of interactions with humans can be successfully automated. People seem to prefer withdrawing cash from ATMs than dealing with human cashiers. The centre of gravity of the entire retail industry is shifting online, where consumers generally avoid dealing with humans.

This does not mean that humans are becoming anti-social – far from it. Merely that we like to be able to choose for ourselves when we interact in a leisurely manner with another human, and when we transact some business quickly and efficiently.

Machines are sometimes surprisingly good at tasks which appear at first sight to require a human touch. In chapter 3.10 we will meet Ellie, a machine therapy system developed by DARPA, the research arm of the US military, which has proved surprisingly effective at diagnosing soldiers with post-traumatic stress disorder.

Manual tasks

We noted before that getting machines to do things that we find hard (like playing chess at grandmaster standard) is relatively easy, and getting them to do things that we find easy (like opening a door) is hard. Vivid proof of this was provided by the final round of the DARPA Robotics Challenge, held in June 2015. 25 robots attempted a series of tasks inspired by the rescue missions at the Fukushima nuclear power plant in 2011. None of the robots completed all the tasks, and there was a great deal of hesitation and falling over.

Many jobs involving manual dexterity or the ability to traverse un-mapped territory are currently hard to automate. But as we will see in the next section, that is changing fast.

Tipping points and exponentials

New technologies sometimes lurk for years or even decades before they are widely adopted. 3D printing (also known as additive manufacturing[cxxxi]) has been around since the early 1980s but is only now coming to general attention. Fax machines, surprisingly, were first patented in 1843, some 33 years before the invention of the telephone.[cxxxii]

Sometimes the delay happens because there is at first no obvious application for the inventions or discoveries. Sometimes it is because they are initially too expensive, and engineers have to work on reducing their cost before they can become popular. And sometimes it is because they are simply not good enough when they are first demonstrated by researchers. And sometimes, of course, it is a combination of these factors.

Once it satisfies these conditions, a new technology can take off dramatically, with exciting applications which appear to most people to come from nowhere, when in fact the underlying technology has been known about for a long time.

The applications of deep learning will probably be like that. The technique is a descendant of neural networks, which were first explored in the early days of AI in the mid 20th century. Faster computers, the availability of large data sets, and the persistence of pioneering researchers have finally rendered them effective this decade, leading to “all the impressive computing demos” referred to by Robin Hanson in chapter 3.3, along with some early applications.

But the major applications are still waiting in the wings, poised to take the stage. It won't be long now before machines are decisively better than humans at reading, listening, recognising faces and other images, understanding and processing natural language. And they won't stop at being slightly better than us. They will continue to improve at an exponential rate, or close to it. To say that the impact will be dramatic is an understatement.

Another thing to bear in mind is that to reach the point where technological unemployment forces dramatic change in the way we run our economies does not require everyone to be unemployed

and unemployable. It does not even require a majority to find themselves in that predicament. It just requires a substantial minority to believe that they will be.

Before we proceed to look at some examples of how AI will sweep away many of the jobs we take for granted today, we need to quickly review some of the related technologies which will influence the way that happens.

3.7 – Related technologies

<u>One ring to bind them</u>

Artificial intelligence is increasingly our most powerful technology, and it will increasingly inform and shape everything we do. Its full-blooded arrival coincides with the take-off of a series of other technologies. They are all driven at least in part by AI, and they will all impact the way our societies evolve.

Because they will all unfold in different ways and at different speeds, it is impossible to predict exactly what the impact of these interlacing technologies will be, other than that it will be profound.

<u>The Internet of Things</u>

The Internet of Things (IoT) has been talked about for years – the term was coined by British entrepreneur Kevin Ashby back in 1999.[cxxxiii] Indeed it has been around for long enough to have acquired a selection of synonyms. GE calls it the Industrial Internet, Cisco calls it the Internet of Everything, and IBM calls it Smarter Planet. The German government calls it the Industry 4.0[cxxxiv], the other three being the introduction of steam, electricity, and digital technology. As noted in chapter 2.2, I think this is an unhelpful term, as it shifts the IoT from the information revolution to the industrial one, and it under-states the importance of the information revolution.

My favourite alternative name for the IoT is Ambient Intelligence,[cxxxv] which comes nearest to capturing the essence of the idea, which is that so many sensors, chips and transmitters are embedded in objects around us that our environment becomes intelligent – or at least, intelligible.

When originally conceived, the IoT was based on Radio Frequency Identification tags (RFID), tiny devices about the size of a grain of rice which can be "read" remotely without being visible to the device which "reads" them. The RFID is a passive device, and this concept does not involve any AI.

Later, technologies like Near Field Communication (NFC) were developed, which allow for two-way data exchange. Android phones have been NFC-enabled since 2011, and it powers the Apple Pay system which was launched with the iPhone 6.

The IoT is becoming possible because the component parts (sensors, chips, transmitters, batteries) are becoming cheaper and smaller at – yes – an exponential rate. The technology research company Gartner forecast in December 2013 that 26 billion digitally accessible devices would be installed by 2020, a 30-fold increase within a decade.[cxxxvi] Many of these devices have multiple sensors – smartphones can have as many as 30 each.[cxxxvii]

Looking further ahead, the internet entrepreneur Marc Andreessen predicts that by 2035, every physical item will have a chip implanted in it. "The end state is fairly obvious - every light, every doorknob will be connected to the internet."[cxxxviii]

Making the environment intelligible offers tremendous opportunities. A bridge, building, plane, car or refrigerator with embedded sensors can let you know when a key component is about to fail, enabling it to be replaced safely without the loss of convenience, money, or life which unforeseen failure might have caused. This is known as condition-based maintenance, or

predictive maintenance, and is being pioneered with encouraging results by MTR Corporation which runs Hong Kong's urban transit network.[cxxxix]

The IoT will improve energy efficiency across the economy, as the heating or cooling of buildings and vehicles can be regulated according to their precise temperature, humidity, etc., and the number and needs of the people and equipment using them.

Since its launch in 1990,[cxl] the world-wide web has rendered our lives immeasurably easier, by placing information at our fingertips. The IoT will take that process an important stage further, by dramatically improving the amount and quality of information, and enabling us to control many aspects of our environment. You will be able to find out instantaneously the location and price of any item you want to buy. You will know the whereabouts and welfare of all your friends and family – assuming they don't mind – and the location of all your property: no more lost keys! You will be able to control at a distance the temperature, the volume, the location of things that you own. Your own health indicators can be made available to anyone you choose, which will certainly save many lives.

Like any powerful technology, the IoT will raise concerns, particularly about privacy and security, and we will return to these later. It will also need a set of standards, so that all those semi-intelligent chairs and cars talk the same language. This may come about through government regulation, industry co-operation, or because one player becomes strong enough to impose its standards on everyone else.

Digital assistants

Siri, the digital assistant bundled with more recent releases of the iPhone, is a bit of a joke today, but by 2025 its descendants will be our constant companions, and we will wonder how we ever got along without them. They will be our gateway to the internet, and our invaluable assistants as we navigate our way through the world.

The competition to provide the most useful digital assistant is warm and getting warmer. Siri was the first entrant, but many people think Google's equivalent for android phones is currently better. Microsoft has Cortana, and Amazon has Echo, which operates from an always-on fixed location device rather than a mobile. Facebook is betting on a mixture of AI and human intelligence with its contender, M. There are also numerous smaller players, of which perhaps the most interesting is Viv (from the Latin for "life"), a system developed by the original creators of Siri.[cxli] They span Siri out of a DARPA-funded research project, taking the name from Sigrid, a Scandinavian word meaning both "victory" and "beauty", and sold it to Steve Jobs in 2011.

Artificial intelligences will govern most things in our environment, and something like Siri will be our intermediary, negotiating with and filtering out most of the Internet of Things. Although we may not notice it, this will be a blessed relief. Imagine having to negotiate a world where every AI-enabled device has direct access to you, with every chair and handrail pitching their virtues to you, and every shop screaming at you to buy something. This dystopia was captured in the famous shopping mall scene in the 2002 film "Minority Report", and more laconically in Douglas Adam's peerless "Hitchhiker's Guide to the Galaxy" series, where the Corporation that produces the eponymous guide has installed talking lifts, known as happy vertical people transporters. They are extremely irritating.

Friends?

What generic name will be adopted for these assistants? Most of the essential tools which we use every day have one-syllable names, like phone, car, boat, bike, plane, chair, stove, fridge, bed, gun[cxlii]. Those which have two syllables are often elided or rhyming, like iron and hi-fi. A few, like hoover, are named after the person or company who made the first successful version.

As yet we have no short name for our digital assistants. "Digital personal assistant" and "virtual personal assistant" both capture the meaning but are hopelessly unwieldy. Maybe we'll initialise them, like TVs, and call them DAs, DPAs, or VPAs. Or maybe we'll use the brand name of one of the early leaders, and call them all Siris. Google's chairman Eric Schmidt came up with the interesting idea that we'll find ways to name them after ourselves, and his would be called "not-Eric".[cxliii] Perhaps – and this is my favourite - we'll just call them our "Friends".

Wearables, insideables

At the moment, the vessel which transports the primitive forebears of these essential guides is the smartphone, but that is merely a temporary embodiment. We will surely progress from portables to wearables (Apple Watch, Google Glass, smart contact lenses...) and eventually to "insideables": sophisticated chips that we carry around inside our bodies.

You doubt that Google Glass will make a comeback? The value of a head-up display, where the information you want is displayed in your normal field of vision, is enormous; that's why the US military is happy to pay half a million dollars for each head-up display helmet used by its fighter aircraft pilots.

Apple Watch has been successful because some people will pay good money to simply raise their wrist rather than go to all the bother of pulling their smartphone out of their pocket. How much better to have that hunger for the latest bit of gossip sated, and that essential flow of information about your environment displayed right in front of your eyes with no effort whatsoever?

With regard to "inside-ables", the technology to enable a chip implanted inside you to project imagery into your field of vision is far ahead of where we are now. But it is the next logical step in the process after wearables, and with key aspects of technology advancing at an exponential rate it would be foolish to write it off.

Screens will be everywhere by this time, of course: on tables, walls both interior and exterior, on the backs of lorries so that you can see what is ahead of them.[cxliv] But we will want to carry our own screens around with us, not least because we won't always want other people to see what we are looking at.

We will probably also need to invent a new type of interface to enable us to communicate with our digital assistants. The 2013 movie "Her" is one of Hollywood's most intelligent treatments of advanced artificial intelligence. (I realise that isn't saying much, but Hollywood does frame the way many of us think about future technologies.) The essence of the plot is that the hero falls in love with his digital assistant, with intriguing consequences. Although he uses keyboards occasionally, most of the time they communicate verbally.

There will be times when we want to communicate with our "friends" without making a sound. Portable "qwerty" keyboards will not suffice, and virtual hologram keyboards may take too long to arrive – and they may feel too weird to use even if and when they do arrive. Communication via brain-computer interfaces will take still longer to become feasible, so perhaps we will all have to learn a new interface - maybe a one-handed device looking something like an ocarina[cxlv].

Another way we may communicate with our Friends, and indeed with many of the newly intelligible objects in the Internet of Things is radar. In May 2015 Google posted a video to introduce Soli, a project which embeds a sophisticated radar sensor in a tiny chip. It uses no lenses, and there is nothing to break. It generates a virtual tool in the space above or in front of itself, a way to interpret human intent by tracking the tiniest motion of the human hand and fingers. Soli generates virtual representations of controls we are all familiar with, such as volume knobs, on-off buttons, and sliders.[cxlvi]

Doing business with Friends

Friends will be very big business, and the evolution of their industry will be fascinating. Will it turn out to be a natural monopoly, where the winner takes all? If so, the winner will find itself the subject of intense regulatory scrutiny, and probably of moves to break it up or take it into public ownership. Or will there be a small number of immensely powerful contenders, as in the smartphone platform business, where Apple and Android have the field almost to themselves?

Will we all choose one brand of Friend at an early age, or during adolescence, and stick with it for life, as many people do with smartphones? Doubtless the platform providers will seek to lock us in to that kind of loyal behaviour. Or will we be promiscuous, hopping from one provider to the next as they jostle and elbow each other, taking turns to launch the latest, most sophisticated software?

Robots

The final round of the DARPA Robotics Challenge in June 2015 could have been a triumphal display of engineering prowess and the potency of artificial intelligence. (For what are robots, but the peripherals of AI systems, just as mouses and keyboards are the peripherals of PCs?) Instead, as we noted above, it was a sad affair, with the winning machine taking nearly 45 minutes to complete a series of eight tasks that a toddler could accomplish in 10 minutes. Moreover, the tasks had already been scaled down from the initial targets set in 2012 when it became obvious that none of the teams were going to be able to achieve them.[cxlvii]

But remember the progress in self-driving cars provoked by the DARPA Grand Challenge. In the initial event in 2004 the winning car drove just seven of the 150 miles of the track before crashing. A dozen years later, self-driving cars are demonstrably superior to human drivers in almost all circumstances, and they are closing the remaining gap fast. As far as robotics is concerned, we are at 2004 again. And don't forget the power of exponential improvement.

In chapter 2.3 we met Baxter, a new generation of industrial robot, which is beginning to demonstrate that robots can be flexible, adaptable, and easy to instruct in new tasks. Research teams around the world are teaching robots to do intricate tasks. In October 2015, a consortium of Japanese companies unveiled the Laundroid, a robot capable of folding a shirt in four minutes.[cxlviii] Meanwhile at the University of California, a team developing the Berkeley Robot for

the Elimination of Tedious Tasks (BRETT) spent seven years reducing the time to fold a towel from 20 minutes to 1.5 minutes.[cxlix]

So robots can fold towels - slowly, but it will be a few years before they can carry out efficiently all the tasks that a hotel chambermaid does. What they can already do, however, is automate many of the individual tasks that chambermaids carry out. Half a dozen big-name hotels in California are experimenting with robots that deliver towels and other items to guests' rooms on demand.[cl] Apparently toothpaste is the most-often requested item, presumably to keep all those perfect Hollywood teeth sparkling.

In mid-2015, a team at University of California, Berkeley, announced that by applying deep learning to the problem, they could get robots to screw tops onto bottles and remove nails from wood with a claw hammer, and do so with approximately the same speed and dexterity of a human.[cli]

Researchers are trying out different ways to improve robot performance. Teams at Carnegie Mellon University in Pittsburg and at Google are getting robots to learn about their physical environment by having them simply prod, poke, grasp and push objects around on a table-top, in much the same way that a human child learns about the physical world. Having collected a large data set from this activity, the systems turn out to be better at recognising images from the ImageNet database than systems which have not had the physical training.[clii]

Google's robot army - the dog that didn't bark

In late 2013, Google announced the purchase of no fewer than eight robotics companies. (Since you ask, they are Boston Dynamics - purveyor of the famous Big Dog and Atlas models - Bot and Dolly, Meka, Holomni, SCHAFT, Redwood, Industrial Perception, and Autofuss.) Google also announced that the new division which owned them would be run by Andy Rubin, who created a huge global business with the Android phone platform.

A year later, in October 2014, Andy Rubin left Google to found a technology startup incubator, which prompted observers to remark that Google had been surprisingly quiet about its collection of robot makers. In early 2016, rumours spread that Google was considering selling Boston Dynamics, the creator of Big Dog and Atlas, two of the world's most impressive robots. Google is an experienced acquirer of companies – by the end of 2014 it had acquired 170 of them – and it expects them to make an impact. The hurdle for potential acquisition targets is the "toothbrush test", meaning that their services must be potentially useful to most people once or twice every day. Sooner or later its robotic companies will impress us.

Complicated relationships

It is going to take us humans a while to get used to having robots around. A French company called Aldebaran, which is owned by Japanese firm Softbank, manufactures a robot called Pepper. 120Cm tall and costing around $1,200, they have a limited ability to "read" human emotions and respond appropriately. They have proved extremely popular in Japan, with four batches of 1,000 selling out in less than a minute when they went on sale in September 2015.

The response to Pepper has not been straightforward, however. The manufacturer felt obliged to outlaw any attempt to engage in sex with the robot, and a Japanese man was prosecuted for assaulting one when drunk.[cliii]

A robot called Hitchbot managed to cross Canada from coast to coast in 2015, but was attacked and decapitated in Philadelphia when it tried to repeat the performance in the US.[cliv]

More robots: androids, drones and exoskeletons

It is not clear that robots need to resemble humans closely to perform their tasks, but that doesn't stop researchers from trying to make them. (Robots with human appearance used to be what the word "android" meant before Google appropriated it for phone software.) We are probably quite a few years away from having robots with the verisimilitude of the ones in the film "Ex Machina", or the TV series "Humans", for example. Nadine is a state-of-the-art prototype working as a receptionist at Nanyang Technological University in Singapore. It is humanoid, but doesn't fool anyone who takes a second glance.[clv] Modelled on its inventor, professor Nadine Thalmann, it cannot walk, but it can smile, turn its head, and shake your hand. Its voice is powered by an AI similar to Siri.

Most robots will probably be special-purpose devices, constructed to carry out a very specific task. An example is the Grillbot, a robot the size of a table tennis bat which cleans your barbecue grill, and is otherwise entirely useless.[clvi]

Another form of robot which is taking off fast is drones – flying machines that can be controlled remotely or autonomously. They have a wide range of applications, including taking surreptitious photos of celebrities, taking selfies for life-logging Millennials, and delivering parcels for Amazon. They present a serious challenge for regulators concerned about the impact on more established forms of aircraft. These challenges cannot be dismissed or regulated away: internet-connected drones with powerful sensors and computers on board are quickly becoming essential tools for companies in the utilities and engineering industries, as well as government agencies.[clvii]

Some people argue that exoskeletons are wearable robots. Whether or not that is semantically correct, they will certainly enable one human to do the work of several. At the moment, leading companies in the space like Ekso Bionics[clviii] are focusing on patient rehabilitation systems. But before long similar equipment will be available for people carrying out physically demanding tasks in the military, manufacturing, and distribution.

Virtual Reality

During 2014, many people got their first taste of virtual reality (VR) from Google Cardboard, an ingenious way to let smartphones introduce us to this extraordinary technology. 2016 is widely expected to be the year that VR really takes off, as Facebook's Oculus VR launches Rift, the first VR equipment for consumers that offers high definition visuals and no latency. Latency is a failure of synchronisation between the stimuli from different sources reaching the brain: if your visual experience is out of synch with your other senses, your brain gets confused and unhappy, and can make you feel surprisingly sick.[clix]

When VR is effective it is surprisingly powerful. When the sense data being received by the brain become sufficiently realistic, the brain "flips", and decides that the illusion being presented is the reality.

Google is not giving up on smartphone-based VR. Having sold more than 5m of the cardboard units, it plans to launch a more robust plastic version in 2016, with better sensors and lenses. It will remain considerably cheaper than the Oculus Rift, which will cost hundreds of dollars.[clx]

Augmented reality (AR) is similar to VR except that it is overlaid on your perception of the real world rather than replacing it. It can make elephants swim through the air in front of you, or plant a skyscraper in your back garden. This is handy if you want to remain alert to the threat from dogs and potholes while you are hallucinating swimming elephants. Microsoft's Hololens is the best-known AR brand to date, but great things are expected from a company called Magic Leap, in which Google has a substantial stake.

Insofar as there is any debate about whether VR is going to be an important development, it's between those who think it's going to be huge, and those who think it's going to change everything. Gartner expects two million VR headsets will be sold in 2016, with the volume increasing thereafter – you guessed it - exponentially.[clxi] Digi-Capital, a specialist consultancy, expects VR and AR sales to reach $150bn within five years. It expects AR to account for 80% of that revenue.[clxii] And these are just the projections for the early years – Gartner expects the VR industry will take five years to achieve mainstream adoption.[clxiii]

Digi-Capital's $150bn figure includes software as well as hardware; Goldman Sachs, a bank not given to hyperbole, expects hardware sales alone to reach $80bn in a decade. Interestingly, they expect less than half of that to come from gaming and films, with the rest coming from commercial and professional applications.[clxiv]

Applications of VR and AR

The biggest application in the short term is expected to be video games, which is no small playing field, since gaming has for some time rivalled Hollywood for leadership in global sales of packaged entertainment.[clxv] Judging by the content already being made available for Google Cardboard, people also enjoy ersatz travel, and adventurous experiences. VR versions of Google Street View let you wander around Manhattan until the latency makes you ill, and other developers offer you rollercoaster rides, and adventure sports from skiing to hang gliding.

In the longer term the potential applications are bewildering. Without ever leaving our armchairs we may soon be able to enjoy such realistic simulations of events like sports matches and music concerts that many people will question the value of struggling with transport and crowds to attend the real thing. Of course, the crowd has a lot to do with making the event exciting in the first place, so the organisers of VR events will want to find a way to recreate the effect of being in a crowd. Except that you'll be able to sit next to your friend, who happens to be in a VR rig a couple of continents away at the time.

Education and informal learning is also likely to experience a VR revolution. How much more compelling would it be to learn about Napoleon by experiencing the battle of Waterloo than by reading about it, or listening to a lecturer describe it? How much easier would it be for a teacher to explain the molecular structure of alcohol by escorting her pupils round a VR model of it?

Businesses will find many uses for VR, and because they often have larger budgets than consumers and educational institutions, they may sponsor the creation of the most cutting-edge applications. Computer-aided design environments will become startling places to work, for instance, allowing

architects, designers and clients to explore and discuss buildings in great detail before ground is broken. And who knows what uses the military will find for VR. One frightening thought is that VR could become a powerful and truly terrifying instrument of torture.[clxvi]

Telecommunication will also be taken to a new level. Although audio-only phone calls still predominate, good video-conferencing facilities add enormously to the effectiveness of a long-distance conversation, and the additional step of feeling present in the same space will improve the experience again. Anything which involves your re-location in time or space should be fertile ground for VR.

On the other hand, it is not yet clear whether VR will turn out to be a good medium for movies. In a film, the director wants to direct your attention, and it isn't helpful if half the audience is busy gawping at images or events 180 degrees removed from the focus of the action.

Cynics will point out that new media (TV, video, the internet were all new media in their early days) are established only when users have found ways to apply them to porn, gambling, and then sport. No doubt VR will make its contribution to these areas of human activity, but I'm not going to get sucked into a discussion of what could be achieved with haptic suits – clothing which allows users to experience sensations of heat and touch initiated remotely by someone else.

The death of geography has been declared numerous times, but despite the rise of telephony, digitisation and globalisation, business and leisure travel just keeps on growing. Could the sense of genuine "presence" which good VR confers finally make the old chestnut come true? Will talent continue to be drawn into the world's major cities, or will VR puncture their inflated real estate prices, and smear humanity more evenly across the planet?

Maybe virtual reality can render scarcity less valuable, and less problematic. In the real world, not everyone can live in a beautiful house on a palm-fringed beach, drive an Aston Martin, and be greeted by a Vermeer as they enter their living room. With virtual reality, everyone can – to a fair degree of verisimilitude. As we will see in chapter 5, this might turn out to be extremely important for our overall well-being as a species.

The world-wide web has given us something like omniscience, and virtual reality looks set to give us something like omnipresence. Perhaps all we need now is a technology to give us something like omnipotence.

Related concerns

Powerful new technologies can produce great benefits, but they can often produce great harm. There are four serious concerns about the technologies we have just reviewed: privacy, security, isolation and inequality.

Privacy

AI runs on Big Data, and the Internet of Things will generate it by the bathful. In an intelligent environment, the whereabouts of every citizen is easy to establish, along with who they have met and very possibly what they discussed. Many people are understandably concerned about this information being used and mis-used by all sorts of organisations, including (depending on their political persuasion) governments, corporations, pressure groups and resourceful individuals –

such as jealous spouses. As one group of activists puts it, the net is closing around us and we are increasingly transparent to organisations which are increasingly opaque to us.[clxvii]

Law Enforcement is where we would expect to find the cutting edge uses of this technology. A company called Intrado provides an AI scoring system to the police in Fresno, California. When an emergency call names a suspect, or a house, they police can "score" the danger level of the person or the location and tailor their response accordingly.[clxviii] Optimists would say this is an excellent way to deploy scarce resources. Pessimists would reply that Big Brother has arrived.

Others hope that we can retaliate against this kind of "surveillance on steroids" with "sousveillance". With cameras ubiquitous – including on drones - the actions of those in authority are constrained because they know that their actions are observed and recorded by members of the public. This is already happening with law enforcement, with police officers in the US being prosecuted for harassment in situations where they would previously have been immune from oversight. Some authorities are actively embracing this development, with police officers being required to wear cameras at all times in order to pre-empt false allegations. With cameras on drones, the reach of civilian oversight can be extended so far that some are calling it "Little Brother".[clxix] With the watchers being watched, we may arrive at a balance called "co-veillance".[clxx]

The arms race over data will continue between governments, large organisations, and the rest of us. Hollywood loves the trope of the socially dysfunctional hacker who is smarter, more up-to-date and more motivated than her opposite numbers in the civil service, but perhaps we should not be comforted by that idea. When forced to choose between privacy and the opportunity to share, we generally choose to share. We leave a trail of digital breadcrumbs wherever we go, both in the real world and online, and most of us are careless about it.

In part this is because many of us feel that we have nothing to worry about because we have nothing to hide. But there is a chilling effect on free speech if we start to censor ourselves because we want to stay that way. We might think twice before entering a certain term into a search engine, or hesitate before making friends with someone who is overtly counter-cultural. Recent research shows that we self-censor when we are aware of the possibility that we are being surveilled, even when we know we are saying and writing nothing illegal.[clxxi]

In 2015, the Chinese government provided a chilling demonstration of where this could lead. It is building a "social credit" database of all citizens which ranks them according to their trustworthiness. In a frightening extension of credit scoring systems, the database will incorporate all the financial and behavioural information the government can accumulate, and distil it into a single number, ranging from 350 to 950. A score above 600 qualifies a citizen for an instant loan worth $800. At 650 you can rent a car without leaving a deposit. At 700 you are fast-tracked for a Singapore travel permit. Important jobs will require high scores.

Citizens will earn demerits for reprehensible shopping habits (too many video games? too much wine?) and merits for socially responsible actions, like reporting bad behaviour by others.[clxxii] A particularly scary aspect of the system is that people receive demerits if their friends on social media are marked down.

The system will be compulsory for every Chinese citizen in 2020, and until then, eight pilots are being run by Chinese companies, including Sesame Credit, the financial wing of Alibaba, which is China's version of Amazon.

The US civil liberties pressure group ACLU thinks "China's nightmarish Citizen Scores are a warning for Americans. … There are consistent gravitational pulls toward this kind of behavior on the part of many public and private US bureaucracies, and a very real danger that many of the dynamics we see in the Chinese system will emerge here over time."[clxxiii]

Big Data and AI could enable governments to build apparatus of control which would make Big Brother in George Orwell's "1984" look amateurish.

You're not necessarily safe from this prospect just because you live in a multi-party democracy. In April 2014, Nicole McCullough and Julia Cordray founded Peeple, an app which will enable people to rate each other according to their courtesy and helpfulness. Originally conceived as a way to improve behaviours, it was widely criticised as likely to become a medium for personal attack and bullying. The founders responded by changing the rules so that subjects would have a veto over any comments made about them on the site, although they left open the possibility that users who pay extra could see un-censored inputs.[clxxiv]

Clearly we still have much to learn about how to conduct ourselves individually and collectively in the new world of data tsunamis and massive analytic horsepower.

Researchers at Google and Microsoft are experimenting with promising approaches to squaring the circle of protecting privacy while sharing data. Working with Cornell University in New York, Google is trying to enable groups of organisations (e.g., hospitals) to train deep learning algorithms on their own separate data files and then share the outputs of the trained data. They have found that this can work almost as well as combining all the data into one file and using that to train the algorithm.

Microsoft is using a technique called homomorphic encryption to perform analysis on data which is encrypted. It yields encrypted results which can then be decrypted without the sensitive data ever having been available to the analysts in unencrypted form.[clxxv]

Security

In his book "Future Crimes", security expert Marc Goodman sets out in detail how criminals, governments, and organisations use the swelling oceans of data being transmitted about us to steal from us and manipulate us. Cyber crime is probably the fastest-growing type of crime all over the world; much of it goes undetected, and much of what is detected goes unsolved.

Another growing concern about hacking is sabotage. As the internet of things is built out and more and more of our vehicles, buildings and appliances rely on artificial intelligence, the problems that can be caused if their control systems are hacked increase in significance. The possibility of a hacker gaining control of every self-driving car in a city and making them all turn left at the same moment is frightening.

Programmers say that there is no such thing as 100% security: IT systems are designed by humans, and we are fallible. They are also increasingly opaque, and hard to de-bug. An optimist would say that although complex, well-defended systems come under frequent attack, they are rarely successfully hacked. No hacker has yet launched a US nuclear missile, although of course that doesn't mean that it will never happen. Eternal vigilance is the price we must pay to avoid disaster,

and we are not practising it at the moment. Many of us are lax about safeguarding our internet passwords, and many companies' security arrangements also fall far short of best practice.

Policemen say that when they are pursuing a criminal, the criminal needs to be lucky all the time whereas the police only need to get lucky once. But when the criminals are on the offensive, looking for security gaps, the boot is on the other foot.

Inequality

Every time a new technology is launched, people worry that only the rich will have access to it, and there will be a "digital divide" separating the haves and the have-nots. The life experiences and opportunities of the wealthy will diverge unacceptably from those of the rest of us.

So far, while not groundless, this fear has been exaggerated. It is true that in recent years the super-rich have gained more income and wealth than anyone else in most developed economies. (And we're not talking about the 1% here, but the 0.01%.) Meanwhile, there are people who struggle to afford what many would consider the basic necessities of life – although the definition of basic necessity varies greatly between developed countries and elsewhere in the world.

It is also true that the disparity of income between average people in rich countries and average people in the poorest countries is enormous. This disparity, however, is shrinking. And those in America and Europe who protest about the obscene wealth of the 1% in their own countries seem curiously un-troubled by the fact that they themselves are often among the richest 1% of the world's population.

The new technologies which have emerged during the various stages of the industrial revolution have become available to most people in developed economies not long after they were invented. The car, the refrigerator, the washing machine, the TV, the home computer, the smartphone – all have gone through the same cycle. An expensive first version is launched which can be afforded only by the wealthy. It doesn't work very well, and is at least in part a status symbol. Very quickly, the technology improves and the price falls, and pretty soon the great majority of us have one. Next in line for this cycle is virtual reality headgear.

The reason for this is simple economics. Companies make far more money by selling lots of cheap smartphones (for instance) to everyone than by selling a few very expensive ones to the wealthy elite. And in a competitive economy, even if the first company to market is happy to make its money by scalping the rich, other companies will quickly come along to raise the quality and reduce the price. There is no "fridge divide"; why should there be a "digital divide"?

I said this fear was exaggerated "so far". In chapter 5 we will see that there may be more grounds for concern in the not-too-distant future.

Isolation

Parents have long fretted about their teenage children spending long hours in antisocial isolation, hunched over a video game console. Wishing their kids would go outside and kick a ball around instead, they have agonised over a series of scares about the ill effects of video games, which allegedly make kids violent, stop them developing social skills, render them vulnerable to legions

of grooming molesters, and give them impossibly short attention spans. And the blue light of the screen disrupts their sleep.

Meanwhile, the Flynn Effect describes the finding that IQ levels are increasing steadily each generation,[clxxvi] which should not be surprising when you consider the general trends toward less smoking, less drinking, better central heating, better food and better healthcare. And the fact that we are continually learning more about what works in education and what does not.

Humans are intensely social creatures. The need to belong to a tribe – to be accepted by it and perhaps to climb its hierarchy – is programmed deeply into us. Working together in tribes is how we survived in the savannah, surrounded by animals which were stronger and faster, with bigger teeth. Individuals who were cast out of their tribe quickly joined another one or – more likely – got eaten. It would be amazing if in a single bound, one generation of teenagers suddenly freed themselves from this evolutionary programming and isolated themselves in solitary pursuits.

And indeed they haven't. The most popular video games are those which people can play together, and incorporate into their social bonding activities. For teenagers, these activities are just as important as they ever were – and of course it is no less important that their parents be at least slightly appalled by them.

Of course, if and when the day comes when people can plug into utterly compelling virtual reality worlds through a direct neural link, and effectively disappear into the Matrix, things may be different. But unless we have altered our cognitive make-up dramatically by then, my hunch is that we will find a way to make the Matrix social too.

We have explored the state of the art in artificial intelligence, and peered into its likely future, along with related technologies that it drives and will be affected by. We have considered what people do at work. Now it is time to think about the kinds of jobs which wil be automated by these technologies.

We'll start with driving.

3.8 – The poster child for technological unemployment: self-driving vehicles

Why?

The case for introducing self-driving cars is simple and overwhelming: around the world, human drivers kill 1.2 million people a year, and injure a further 20 to 50 million.[clxxvii] Road traffic accidents are the leading cause of death for people aged 15 to 29, and they cost middle-income countries around 2% of their GDP, amounting to $100 billion a year.

90% of these accidents are caused by human error.[clxxviii] Humans become tired, angry, drunk, sick, distracted or just plain inattentive. Machines don't, so they don't cause accidents. To paraphrase Agent Smith in "The Matrix", we are sending humans to do a machine's job.

There is also the wasted time and frustration. We all know that driving can be fun, but not when you're stuck in traffic – perhaps because one of your fellow humans has caused an accident. On average, American commuters spend the equivalent of a full working week stuck in traffic every year – twice that much if they are lucky enough to work in San Francisco or Los Angeles.[clxxix] We drive rather than use public transport because there is no appropriate public transport available, or sometimes because we prefer travelling in our own space. Self-driving cars could give us the best of both worlds, allowing us to read, sleep, watch video or chat as we travel.

Finally, self-driving cars will enable us to use our environments more sensibly, especially our cities. Most cars spend 95% of their time parked.[clxxx] This is a waste of an expensive asset, and a waste of the land they occupy while sitting idle. We will consider later how far self-driving cars could alleviate this problem.

To autonomy and beyond

Self-driving cars, like our artificially intelligent digital assistants, are still waiting to receive their generic name. "Self-driving cars" is the name we are stuck with for the time being, but it is all clunk and no click. At the end of the 19th century it was becoming obvious that horseless carriages were here to stay, and needed a shorter name. *The Times* newspaper adopted "autocar" but the *Electrical Engineer* magazine objected that it muddled Greek (auto) with Latin (car). It argued instead for the etymologically purer "motor-car".[clxxxi] Perhaps we will contract the phrase "autonomous vehicle", and call them "autos".

Some people are going to hate self-driving cars, whatever they are called: petrol-heads like Jeremy Clarkson are unlikely to be enthusiastic about the objects of their devotion being replaced by machines with all the romance of a horizontal elevator. Some people are already describing a person who has been relegated from driver to chaperone as a "meat puppet".[clxxxii]

The US Department of Transport draws a distinction between (partly) autonomous cars and (fully) self-driving cars.[clxxxiii] The former still have steering wheels, and require a human driver to take over when they encounter a tricky situation. Self-driving cars, by contrast, are fully independent, and the steering wheel has been removed to save space. Autonomous cars will probably be merely a staging post en route to the completely self-driving variety.

In fact the US DoT grades cars on a scale from L0, where the driver does everything, to L4, where the car does everything. Google's initial idea was that the first self-driving cars in general use

would be L3, meaning that the human driver should be ready to take over at a moment's notice if anything went wrong, just as airplane pilots are. But the technology proved so reliable that its test drivers became complacent and engaged in "silly behaviour". For instance, one turned round to look for a laptop in the back seat when the car was doing 65 mph. This experience persuaded Google to advocate immediate adoption of L4.[clxxxiv]

<u>The state of the art</u>

Self-driving cars have come a long way since 2004, when the humvee Sandstorm got stuck on a rock seven miles into the first DARPA Grand Challenge, but they are not perfect yet. They struggle with heavy rain or snow, they can get confused by potholes or debris obstructing the road, and they cannot always discern between a pedestrian and a policeman indicating for the vehicle to stop. A self-driving car which travelled 3,400 miles from San Francisco to New York in March 2015 did 99% of the driving itself, but that means it had to hand over to human occupants for 1% of the journey.[clxxxv] With many technology projects, resolving the last few issues is more difficult than the bulk of the project: edge cases are the acid test. Nevertheless, those edge cases are being tackled, and will be resolved.

It is well-known that Google's self-driving cars have travelled well over a million miles in California without causing a significant accident, but what is less well-known is that the cars also drive millions of miles every day in simulators. Chris Urmson, head of the Google project, expects self-driving cars to be in general use by 2020.[clxxxvi]

Sceptics point out that Google's self-driving cars depend on detailed maps. But producing maps for the roads outside California doesn't sound like an insurmountable obstacle, and in any case, systems like SegNet from Cambridge University enable cars to produce maps on the fly.[clxxxvii]

A fully autonomous bus made in France has been serving the centre of the Greek city of Trikala since February 2015. It travels at a top speed of 20 mph along a pre-determined route which is also used by pedestrians, cyclists and cars.[clxxxviii]

In December 2015 Bloomberg reported that Google was preparing to move its self-driving cars unit from its Google X research arm to become a stand-alone business unit within the Alphabet holding company.[clxxxix] At the same time, Elon Musk, CEO of Tesla, remarked that he was revising his estimate of the time when fully automated cars would be available from three years down to two.[cxc] In January 2016 he announced that within about two years, Tesla owners would be able to "summon" their driverless car from New York to pick them up in Los Angeles.[cxci] He claimed that Tesla cars are already better drivers than humans.[cxcii] In April 2016 he went further, claiming that Tesla's autopilot system was already reducing the number of accidents by 50% - where an accident meant an incident where an airbag was deployed.[cxciii]

Ford reported success in January 2016 with tests of its self-driving car in snowy conditions. Unable to determine its location by the obscured road markings, it navigates by using buildings and other above-ground features.[cxciv] In May 2016 an executive in Ford's autonomous vehicle team estimated that the remaining technological hurdles would be overcome within five years, although adoption would of course take longer. He said the amount of computing power each car currently required was "about the equivalent of five decent laptops."[cxcv]

At the time of writing, the only accident which a Google self-driving car might be blamed for happened in February 2016. The car was trying to merge into a line of traffic and expected that a bus which was approaching from behind would give way. It didn't. The car was travelling at 2 mph and no-one was hurt, so no police report was filed to attribute blame officially. The bus driver has declined to comment.[cxcvi]

Of course, just because a product becomes available, that doesn't mean it will be bought, still less that it will comprehensively replace the existing population of products that it is designed to supplant. The rate at which that happens, if it happens at all, depends on a host of factors including regulation, price, design, service support, promotion and PR, and the length of the replacement and upgrade cycle for the product category.

Regulation is an important consideration. Google was disappointed when California's Department of Motor Vehicles (DMV) proposed new rules for self-driving cars in December 2015 which banned vehicles which lacked the capacity for a human to take control. In theory, unco-operative regulators could slow or even stop the arrival of self-driving cars, and there will be powerful lobbies pressing for this. But they can only succeed if all regulators everywhere agree, and work together, and that will not happen – even within the US, never mind globally. In 2015 Google expanded its test driving programme beyond Silicon Valley to Austin, Texas, where the authorities welcome the tech giant's research money and prestige.[cxcvii] In 2016 it added two more cities, Kirkland in Washington and Phoenix in Arizona.[cxcviii]

Several European countries (including the UK) are keen to burnish their credentials as leaders in what will undoubtedly be a massive new industry.

The impact on cities

Enthusiasts for self-driving cars sometimes paint a utopian picture of cities where almost no-one owns a car because communally-owned taxis are patrolling the streets intelligently, anticipating our requirements and responding immediately to our summons. Whereas today, our cars sit idle 95% of the time, squatting like polluting toads on vast acres of city land, in this bright tomorrow they are used efficiently, and the land given over to parking can be returned to pedestrians and useful buildings. Traffic flows smoothly because the cars are in constant communication with each other: they don't bunch into jerky waves and they don't need to stop at intersections.

This is almost certainly an exaggeration. There will still be peak times for journeys, so even if most journeys are undertaken in communal cars, many of them will be parked up during off-peak hours. And traffic will still have to halt at intersections every now and then if pedestrians are ever going to be able to cross the road. Not every pedestrian crossing can have a bridge or an underpass.

Nevertheless, machine-driven cars will be more efficient consumers of road space than human drivers. Traffic conditions are not fixed fates which once imposed can never improve. A congestion charge has significantly reduced traffic flows in London, and the switch to almost-silent hybrid taxis has made walking the streets of Manhattan an even better experience than it used to be.[cxcix] In any case, more efficient road use is not required to justify the introduction of self-driving cars. The horrendous death and injury toll imposed by human drivers is sufficient, together with the liberation from the boredom and the waste of time caused by commuting.

Detroit's response

The car manufacturing industry first experimented with self-driving cars decades ago. From 1987 to 1995, the European Union spent $750m with Daimler Benz and others on the Prometheus project (the PROgramme for a European Traffic of Highest Efficiency and Unprecedented Safety).[cc] There were some impressive technical achievements, but ultimately the project faded. Fortunately, among other things, we have got better at devising acronyms since then.

The automotive industry's response to the implicit challenge from Google and others has been slow and piecemeal. In part this is because the car industry thinks in seven-year product cycles, while the technology industry thinks in one-year cycles at most. Most of the large car companies seem convinced that self-driving technology will be introduced gradually over many years, with adaptive cruise control and assisted parking bedding in during the lifetime of one model, and assisted overtaking being introduced gradually with the next model, and so on. That is far too slow for the tech titans of Silicon Valley. Google, Tesla, Uber and others are racing towards full automation as soon as it can be safely introduced. If Detroit does not join in it may find itself displaced.

In the closing months of 2015, Detroit and its rivals seemed to wake up. Toyota announced a five-year, $1bn investment in Silicon Valley,[cci] Ford announced a JV with Google,[ccii] and BMW's head of R&D declared that in five years, his division had to transition from a department of a mechanical engineering company to a department of a tech company.[cciii] It remains to be seen whether Google will seek to be a supplier of artificial intelligence functionality to a robust and healthy automotive industry, or whether it will follow Tesla's example, and lead the car industry by competing with it. Meanwhile, there are persistent rumours that Apple wants to become a car company too.

<u>Other affected industries</u>

Automotive cover represents 30% of the insurance industry, so a shift to self-driving cars will have a major impact on that industry. The most obvious effect should be a sharp reduction in pay-outs because there will be far fewer accidents. This in turn should mean far lower premiums: bad news for the insurance companies, good news for the rest of us.

Who will take out the insurance policy? When humans drive cars we blame them for any accidents, so they pay for the insurance. When machines drive, does the buck stop with the human owner of the vehicle, the vendor of the self-driving AI system, or the programmer who wrote its code? If the insured parties are Google and a handful of massive competitors, then the negotiating position of the insurance companies will deteriorate sharply from the present situation where they are “negotiating” with you and me.

Warren Buffet ascribes some of his enormous success as the world's best-known investor to his decision to avoid areas he does not understand, including industries based on IT. He has massive holdings in the insurance industry. Unfortunately for him, software is “eating the world”,[cciv] and a large chunk of the insurance industry is about to be engulfed in rapid technological change. Buffet acknowledges that when self-driving cars are established, the insurance industry will look very different, almost certainly with fewer and smaller players.[ccv] It is very hard to say which of today's players will be the winners and losers.

The law of unintended consequences means that we cannot say how the insurance risks will change. Let's hope this never happens, but what if a bug – or a hacker – caused every vehicle in a particular city to turn left suddenly, all at the same time? How does an insurance company

estimate the probability of such an event, and price it? Important issues like this - and the ethical questions we will discuss below – will slow down the introduction of self-driving cars. But they will not stop it. They are capable of resolution, just as we resolved questions about who would build the roads and who would have the right of way in different traffic situations in the decades after the first cars appeared.

People working in insurance companies will certainly not be the only ones affected by the move to self-driving cars. Machines will presumably be programmed not to violate local parking restrictions. That will remove a significant source of income from local authorities: parking charges generate well over $300m a year for the city government of Los Angeles.[ccvi]

Automotive repair shops will still be needed, but their business will shrink as it becomes restricted to maintenance and repairs necessitated by age rather than accidents. Happily, something similar can be said of doctors and nurses.

Programming ethics

Your self-driving car is travelling down the road, minding its own business, when a child, unpredictably, dashes across the street ahead of you. Calculating at super-human speed, it analyses the only three available options: maintain direction, turn right or turn left. Even though it has already applied the brakes far quicker than any human could have, it forecasts (correctly, of course) that these options will result in the death, respectively, of the child, of an innocent adult bystander, or of you, its passenger. Which option should it select? The question will have been answered in advance, even if only by default.

With grim humour, some have suggested that the answer will vary by car. Perhaps a Rolls Royce will always choose to preserve its owner, while a Lada may accord its passenger less respect.

What is happening here is the extension of human control over the world: the arrival of choice. Today, 27% of the victims of accidents are pedestrians and cyclists. What happens to them and the drivers of the cars which hit them is currently decided by the skill of the driver and blind chance. In future we will have the power to affect it, and with increased power comes increased responsibility.

Driving jobs

Clearly, self-driving vehicles will have a huge impact on society, sometimes in surprising ways. What impact will they have on employment? Are they indeed the poster child for technological unemployment?

There are 3.5 million truck drivers in the US alone[ccvii], 650,000 bus drivers[ccviii] and 230,000 taxi drivers[ccix]. How many of these jobs will be lost to machines?

It seems inevitable that machines will drive commercial vehicles. Articulated lorries are driven by professional drivers whose backgrounds are checked and whose working hours and conditions are regulated. They cause fewer accidents per mile driven than cars owned by the likes of you and me. But because they are heavier, when they are involved in accidents they cause much more damage to life and property. It is inconceivable that we will continue to allow humans to do this job for which machines are clearly better suited.

But driving is not the whole of the job. The people who drive trucks, delivery vans, buses and taxis have to deal with the myriad surprises which are thrown at them by life, which is an untidy business at best. If a consignment of barbed wire falls off the back of a truck in front, they will get out and help. They are also often responsible for loading and unloading their vehicles.

Sceptics about technological unemployment could point out that planes have been flying by wire for decades, with human pilots in control for only around three minutes of an average commercial flight. We have yet to dispense with the services of human pilots.

However, a 747 is very different to a truck travelling down Highway 66. A truck is an expensive vehicle, and capable of inflicting severe damage, but commercial planes are on a different scale: they cost many millions of dollars each, and their potential to cause harm was graphically and tragically demonstrated in New York in 2001. Furthermore, those three minutes of human control are in part due to the difficulty of resolving the edge cases we discussed before. In road vehicles if not in planes, we are well on the way towards resolving those.

Consider the process of delivering a consignment from a warehouse to a supermarket or other large retail outlet. Amazon's fleet of Kiva robots show that warehouses are well on the way towards automation. The unloading bays at the retail end are also standardised for efficiency: a system which automates the entire unloading process from a truck into the retailer's receiving area is technically feasible today, and with the exponential improvement in robotics and AI, it won't be long before it is economically feasible as well.

As we have seen, robots are becoming increasingly flexible, nimble and adaptable. They can also increasingly be remotely operated. Most of the situations a driver could deal with on the open road will soon be within the capabilities of a robot which does not need sleep, food or salary. On the rare occasion when human intervention is needed, the gig economy[ccx] can probably furnish one quickly enough.

Once it is economically feasible to replace human drivers with machines, it is a very short step to being economically compelling. Drivers account for 25-35% of the cost of a trucking operation.[ccxi] You can't escape the invisible hand of economics for long. In a free market, once one firm replaces its drivers the rest will have to follow suit, or go out of business. Of course, trade unions and sympathetic governments may try to stop the process in some jurisdictions. They may succeed for a while, but only by rendering their industry uneconomic, and burdening their customers with unnecessary costs which will damage them in turn.

Other governments will take a different approach, and the competitive disadvantage imposed by resistance to change will become apparent. It will not be manifest only in the case of truck drivers, but in all areas of the economy, and regions and countries which do resist will find their living standards declining fast. Over time it will prove unsustainable.

The science fiction writer William Gibson is reported as saying that “The future is already here – it's just not evenly distributed.[ccxii]” In the Yandicoogina and Nammuldi mines in Pilbara, Western Australia, transport operations are now entirely automated, supervised from a centre in Perth, which is 1,200 miles away.[ccxiii] Mining giant Rio Tinto was prompted to take this initiative by economics: the decade-long mining boom caused by China's enormous appetite for raw materials. Drivers earned large salaries in the hazardous and inhospitable environments of these remote

mines, which made the investment case for full automation irresistible.[ccxiv] The economics are going in the same direction everywhere – fast.

The automation of driving will have a major impact on the overall job market. Truck and delivery driver is the most common occupation in 29 US states (57% of them).[ccxv] It will also have the effect of alerting everyone else to the prospect of widespread technological unemployment.

3.9 – Who's next?

Low-income jobs

The Frey and Osborne study we looked at in chapter 3.2 foresaw two waves of automation in the coming decade or two. "In the first wave, we find that most workers in transportation and logistics occupations, together with the bulk of office and administrative support workers, and labour in production occupations, are likely to be substituted by [machines]."[ccxvi] It makes intuitive sense that lower income jobs would be less cognitively demanding, and hence easier to automate.

Food service

Automation is not new to the retail industry. Automated Teller Machines (ATMs) took over the job of dispensing cash in banks years ago, and self-service checkouts are familiar sights in supermarkets. Neither is it new in food service specifically. Few people in major cities now buy their sandwiches bespoke from a human who prepares their food in front of them. Far more buy their lunch in ready-made packages.

Back in 1941, the Automat chain served half a million American customers a day, dispensing macaroni cheese, baked beans and creamed spinach through cubby-holes with glass doors.[ccxvii] The chain declined in the 1970s with the rise of fast-food restaurants serving better-tasting food, such as Burger King and of course McDonalds. But these chains themselves are now discovering the economic appeal of automation.

Chili's Bar and Grill is rolling out a tablet ordering system, and Applebee's began delivering tablets to all its 1,800 restaurants in 2014.[ccxviii] There has been heated political debate about whether these and similar initiatives are prompted by increases in minimum wage levels, but the truth is that AI systems and robots are rapidly becoming more efficient at cost levels which humans can never hope to match.

Customer preference

Cost saving is not the only reason for this kind of automation. In many situations, humans prefer to transact with machines rather than other humans. It can be less time-consuming, and require less effort. It can also make a service available for longer hours, perhaps 24 hours a day. Bank ATMs are the classic example. Another are the automated passport control systems now installed at many airports, which many people opt to use in preference to the manned channels.

A report published in April 2015 by Forrester, a technology and market research company, claimed that 75% of procurement professionals and other people buying on behalf of businesses (i.e., B2B buyers) prefer to use e-commerce and buy online rather than deal with a human sales representative. Once the buyers have decided what they want, the percentage rises to 93%.[ccxix] Forrester pointed out that many vendors were ignoring this fact and obliging customers to speak to a human. This is no doubt partly at least because human sales people are currently much better able to up-sell the buyer, but this is also one of the reasons why buyers prefer e-commerce. Forrester argued that companies which wait too long to offer good e-commerce channels risk losing market share to more digitally-minded competitors.

Call centres

We are still at the very early stages of introducing artificial intelligence to call centres. For many of us, dealing with call centres is one of the least agreeable aspects of modern life. It normally involves a good deal of waiting around, listening to uninspiring hold music, followed by some profoundly unintelligent automated routing, and finally a conversation with a bored person the other side of the world who is reading from a script written by a sadist.

One of the leaders in introducing genuine AI to call centres is Swedbank, one of Sweden's biggest banks, with 9.5m customers and 160,000 employees. It has 700 people working in contact centres, which handle 2m customer calls each year. It has worked with the American software company Nuance to introduce a basic AI called Nina,[ccxx] which learns what customers want and how best to help them by assimilating searches made on the company website and enquiries made at the contact centres.[ccxxi] In December 2015, Nina was handling 30,000 calls a month, and taking care of many of the straightforward transactional calls – like transferring money from one account to another - which were previously clogging up the call centres. The aim is to free up the agents in the contact centres to concentrate on more complicated activities, like taking out a mortgage. But even taking out a mortgage isn't rocket science. Given exponential progress, if Nina can handle transfers today, it will surely be able to handle mortgage applications before long.

Manual work

Occupations requiring physical labour will take longer to automate than clerical and administrative jobs because getting robots to be dextrous and flexible is surprisingly hard. As we saw in chapter 3.7, progress is rapid, but much remains to be done. Manual work in routine, repetitive environments like assembly lines will continue to see machines taking over, but physical labour in unstructured environments like building sites will remain the preserve of human workers for a while longer.

Manufacturing accounts for over a third of China's GDP, and employs more than 100 million of its citizens. Historically, China's competitive strength in manufacturing has been its low wage costs, but this is changing fast: wages have grown at 12% a year on average since 2001, and Chinese manufacturers are embracing automation enthusiastically. As we saw in chapter 2.3, China is now the world's largest market for industrial robots, but it has a long way to go before it catches up with the installed base in more developed countries.

Industrial robots are far from perfect, and manufacturers have under-estimated the progress still required. In 2011 the CEO of Foxconn, a $130bn-turnover Taiwanese manufacturer that is famous for making iPhones, declared a target of installing a million robots by 2014. The robots failed to perform as he hoped, and the actual installation rate has been much slower. But the robots are improving fast.[ccxxii]

The professions

It is certainly not only low-paid, relatively low-prestige service jobs that will be automated. The professions are vulnerable too: lawyers, doctors, architects and journalists. Sometimes accused of being conspiracies against lay people, these are protected occupations, with demanding entry requirements and restrictions on the number of trainees who can join the professions each year. They have commanded prestige and high salaries, but that may be about to change.

Journalists

Nuance, the company behind Swedbank's Nina call centre AI, offers services for journalists, helping them create interviews and articles faster. But Narrative Science, a company established in Chicago in 2010, has an AI system which writes articles without human help. Called Quill, it already produces thousands of articles every day on finance and sports for outlets like Forbes and Associated Press (AP).[ccxxiii] Most readers cannot identify which articles are written by Quill and which by human journalists, and Quill is much faster.

Quill starts with data - graphs, tables and spreadsheets. It analyses these to extract particular facts which could form the basis of a narrative. It then generates a particular plan, or narrative for the article, and finally it crafts sentences using natural language generation software.

A British company called Arria offers the same functionality, but sells mainly to corporations trying to make sense of the tsunami of data which threatens to overwhelm them.[ccxxiv]

In the short term, Quill has not rendered thousands of journalists redundant. Instead it has sharply increased the number of niche articles being written. Newspaper revenues have declined sharply since the turn of the century, as classified ads for jobs, houses and cars migrated online. News services like AP increased the daily quota of articles for each journalist, cut back the number of journalists they employed, and reduced the number of articles they produced on, for instance, the quarterly earnings reports of particular companies. Quill and similar services have enabled them to reverse that decline. AP now produces articles on the quarterly reports of medium-sized companies that it gave up covering in such detail years ago.

Kristian Hammond, founder of Narrative Science, forecast in 2014 that in a decade, 90% of all newspaper articles would be written by AIs. However, he argued that the number of journalists would remain stable, while the volume of articles increased sharply. Eventually, articles could become tailored for particular audiences, and ultimately for each of us individually. For instance, an announcement by a research organisation that inflating your car tyre correctly could reduce your spend on petrol by 7% could be tailored - perhaps with the help of your Digital Personal Assistant - to take into account your particular car, the number of miles you drive each week, and even your style of driving. (Although of course by then you will perhaps not do much of the driving yourself anyway.)

The prediction that the number of human journalists will remain stable sounds reassuring, and indeed you would expect someone marketing an automating technology to say that. Given the exponential improvement of AIs, it is a brave prediction.

TV presenters should also be feeling nervous. In December 2015, Shanghai's Dragon TV featured Xiaoice, (pronounced Shao-ice) an AI weather presenter with a remarkably life-like voice,[ccxxv] based on the Mandarin version of its Cortana digital assistant software. Audience feedback was positive.[ccxxvi]

Other writers

Not everyone who spends their working days crafting crisp sentences is a journalist. They might be PR professionals, or online marketers, for instance. A company called Persado claims that marketing emails drafted by its AI have a 75% better response rate than emails written by human copywriters.[ccxxvii] Citibank and American Express are customers as well as investors.

In January 2016 a researcher at the University of Massachusetts announced an AI which can write convincing political speeches for either of the two main US political parties. The system learned its craft by ingesting and analysing 50,000 sentences from Congressional debates.[ccxxviii]

Two other professions have come under particular scrutiny with respect to their susceptibility to machine automation: the law and healthcare. Let's look at these in turn.

Lawyers

Whatever Hollywood thinks, most lawyers do not spend their days pitting their razor-sharp wits against equally talented adversaries in front of magisterial judges, eliciting gasps of admiration from around the courtroom as they produce the winning argument with a flourish. Most of the time they are reading through piles of very dry material, looking for the thread of evidence which will convict a fraudster, or the poorly drafted phrase which could undermine the purpose of a contract.

Discovery

Many lawyers get a lot of their on-the-job training through the "discovery" process. Known as "disclosure" in the UK, this is a pre-trial process in civil law in which both sides must make available all documents which may affect the outcome of the case. An analogous process takes place in the "due diligence" phase of a corporate merger or acquisition (M&A), in which teams of junior lawyers (and accountants) spend weeks locked away in data rooms, reading through material which can run into millions of documents, looking for something which would clinch the case, or, in the case of M&A work, provide a reason to terminate or renegotiate the deal.

Looking for a needle of fact in a haystack of paper is work more suited to a machine than a man. And although lawyering is a very conservative profession, there are signs that it understands what is coming better than some others. RAVN Systems is the British AI company behind an AI system called Ace, which reads and analyses large sets of unstructured, un-sorted data. It produces summaries of the data, and highlights the documents and passages of most interest according to the pre-set criteria.[ccxxix] When one of the UK's largest law firms started working with Ace, it was regarded as pioneering, experimental, and somewhat risky. Two years later that law firm was promoting its own services to potential new clients on the basis that it knew best how to exploit the advantages of RAVN Ace. Bear in mind that two years is a very short time for anything at all to happen in the legal industry!

Typically, a new client's data will present a new set of challenges. It usually takes a few days to train the system how to read the data, which currently involves human intervention. Once the training is complete, the work proceeds without human involvement, and the system will finish the work much faster than human lawyers could. This means that law firms are having to work out new ways of billing their clients: the old system of hourly rates is under challenge.

Revealing the iceberg

Forward-thinking lawyers are actually excited about the arrival of this sort of automation. Rather than fearing that it will destroy the jobs of junior lawyers, making it impossible for young people to learn the profession, they believe it will increase the amount of cases that can be handled. To illustrate this, imagine a large supermarket chain that wants to know the implication of making a

small change to the employment contracts of all its in-store employees - tens of thousands of people. Previously, its employment law firm would have said that this task could not be undertaken cost-effectively with any degree of rigour. RAVN Ace and systems like it make this kind of work possible, opening up whole new avenues of work for law firms. It is like standing nervously on a body of ice, thinking you are only separated from freezing water by a thin layer, and suddenly discovering that in fact you are standing on an iceberg, with a huge mass of previously unknown solid ice beneath your feet.

As Greg Wildisen, MD of Neota Logic, a firm providing an AI platform for lawyers, puts it, "So many legal questions go 'un-lawyered' today that there is enormous scope to better align legal resources through technology rather than fear losing jobs."[ccxxx]

So in the short and medium term, machine automation of white-collar jobs opens up vast new areas of work that can be undertaken, and doesn't throw the incumbent humans out of work. They are still needed to train the system at the start of a large new assignment, and to process more complicated documents.

But as RAVN Ace and its successors improve – at an exponential rate, of course – they will be able to take on more and more of the sophisticated and demanding aspects of the lawyers' work. No-one can be absolutely sure yet whether this process will hit a wall at some point, leaving plenty of work for humans, or whether it will continue to the point where there are very few jobs left for humans. My own view is that within a few short decades, the machines are coming for most of our jobs.

The short-term explosion of work which happens as the iceberg of latent demand is revealed can give us a false sense of security. The phenomenon of automation leading to job creation is sometimes called the automation paradox.[ccxxxi] But the paradox may turn out to be short-lived.

Forms

Another fairly basic form of legal work is the completion of boilerplate (standard) forms to establish companies, initiate a divorce, register a trademark, request a patent and so on. A company called LegalZoom was established in 2001 to provide these services online, and increasingly, to automate them. LegalZoom now claims to be the best-known law brand in the US,[ccxxxii] and in 2014 the private equity firm Permira paid $200m to become its largest shareholder. Another company, Fair Document, helps clients complete forms for less than $1,000, one-fifth the amount it would have previously cost.[ccxxxiii]

More sophisticated work

At the other end of the spectrum from the "grunt" work of discovery and filling out legal forms, one of the most sophisticated and important jobs that senior and successful lawyers are asked to undertake is to estimate the likelihood of a case winning. The advice is vital as it will determine whether large amounts of money are spent. A team led by Daniel Martin Katz, a law professor at Michigan State University, developed an AI system that analysed 7,700 US Supreme Court cases. It predicted the verdicts correctly 71% of the time.[ccxxxiv]

Another job for experienced lawyers in common law jurisdictions such as the US and the UK is identifying which precedent cases to deploy in support of litigation. A system called Judicata uses

machine learning to find the relevant cases using purely statistical methods, with no human intervention.[ccxxxv]

Will entire sections of the legal industry be automated in the next few years or decades? How about patent lawyers, for instance? Senior patent lawyers are highly skilled and articulate people, but much of the work involved in securing a patent is routine and could perhaps be automated. In November 2015 I took part in a debate at the IMAX cinema in London's Science Museum. The motion was "This House believes that within 25 years, a patent will be applied for and granted without human intervention." Patent lawyers comprised a good part of the audience, and although the motion was vigorously opposed by two senior patent lawyers, the motion was passed. Not exactly turkeys voting for Christmas, but certainly food for thought.

Doctors

Doctors are a scarce resource. Only bright and dedicated people are admitted to the relevant university and post-graduate courses, and these courses demand many years of hard study. Hospitals and local surgeries are organised to maximise the availability of this resource, but some critics argue that they are organised for the benefit of the doctors rather than the patients. In 2015, senior doctor and medical researcher Eric Topol published a book called "The Patient Will See You Now", which he argues should become the mantra for the profession, replacing the current one, which he says is "the doctor will see you now".

Suggesting acerbically that the initials MD stand for Medical Deity, Topol accuses many doctors of being arrogant and paternalistic towards their patients, assuming they are unable to understand the detailed information regarding diagnoses, and withholding information from patients so as not to upset them. He believes that the digital revolution will start to overturn this unsatisfactory state of affairs, as it will place cheap and effective diagnostic tools in the hands of patients.

Better and cheaper diagnostics

In April 2016, researchers at Indiana University announced that a test of open-source machine learning algorithms on 7,000 free-text pathology reports from 30 hospitals yielded equal or better diagnoses than humans had made. The computers were also faster and cheaper.[ccxxxvi]

This sort of technology will become much more widely available. A British startup called Babylon charges customers £5 a month for phone (and videophone) access to a dedicated team of doctors. Before a doctor comes on the line, the patient is triaged by a machine.[ccxxxvii] As AI improves, the role of the human doctor in this process will continuously be reduced.

Smartphones are increasingly able to gather medical data about us, and perform basic analysis. By attaching cheap adapters to their phones, patients can quickly take their blood pressure, sample their blood glucose, and even perform an electrocardiagram. Your breath can be sampled and digitised, and used to detect cancer, or potential heart problems. Your camera's phone can help screen for skin cancers. Its microphone can record your voice, and that data can help gauge your mood, or diagnose Parkinson's disease or schizophrenia.

All this data can be analysed to a certain level within the phone itself, and in many cases that will suffice to provide an effective diagnosis. If symptoms persist, or if the diagnosis is unclear or unconvincing, the data can be uploaded into the cloud, i.e., to server farms run by companies like

Amazon and Google. The heart of diagnosis is pattern recognition. When sophisticated algorithms compare and contrast a set of symptoms with data from millions or even billions of other patients, the quality of diagnosis can surpass what any single human doctor could offer.

Ross Crawford and Jonathan Roberts are professors of orthopaedic research and robotics respectively at Queensland University of Technology. In an article in January 2016,[ccxxxviii] they argued that doctors need to understand that diagnostic services can be made available more cheaply with the assistance of machine intelligence, and reach all the patients who need them, not just those in rich countries who are already manifesting symptoms.

They don't think this will render doctors unemployed. As with the law, there is an iceberg of unmet healthcare needs – needs which automation and machine intelligence can satisfy. Formed in Mumbai in 1996, Thyrocare Technologies is the world's largest thyroid testing laboratory. Its founder, Dr A Velumani, had the insight that 90% of people who could benefit from diagnostic tests were not receiving them because they were too expensive, so the tests were restricted to those already manifesting symptoms of disease. He established Thyrocare to address this latent demand, and now it processes 40,000 samples a day.[ccxxxix]

As this iceberg is revealed, the healthcare industry – like the legal industry - can perform a far better job by reaching many more people at greatly reduced average cost. At first, there will be just as much need for doctors as before: they will continue to carry out the more sophisticated diagnoses, while machines (possibly deployed by less highly-trained people) deliver the routine work. But as we keep observing, the machines are getting smarter at an exponential rate. In time, what is to stop them performing the doctors' other roles as well, and doing them better, faster and cheaper?

Prescribing

If machines can diagnose, can they go on to prescribe, and to fulfil prescriptions? The University of California in San Francisco has installed a robot pharmacist which is reported to have prepared 6,000,000 prescriptions with only one error – a track record which is 60,000 times better than human pharmacists.[ccxl]

Keeping current

As we saw in chapter 3.4, machines are reaching parity with humans in pattern recognition, and will quickly become much better. They are already much better than human professionals at keeping up with new developments in their field.

A human doctor would have to read for 160 hours a week just to keep up with the published medical research. This is clearly impossible for a human, but machines have no such bandwidth restriction. IBM is pushing aggressively into the medical industry with its Watson AI system. According to Samuel Nessbaum of Wellpoint, a private healthcare company, Watson's diagnostic accuracy rate for lung cancer is 90%, which compares favourably with 50% for human physicians.[ccxli]

IBM has come in for criticism for pretending that Watson is a unitary system rather than a kludge of different systems which can be mixed and matched according to need. It is also accused of scaling back its ambition by tackling much smaller projects than the "moonshots" it was originally earmarked for, like curing cancer.[ccxlii] Kris Hammond, the founder of Narrative Science whom we

met when discussing journalists, says that "everybody thought [winning Jeopardy] was ridiculously impossible, [but now] it feels like they're putting a lot of things under the Watson brand name – but it isn't Watson."[ccxliii] In March 2016, DeepMind founder Demis Hassabis went as far as to say that Watson is essentially an expert system as opposed to deep learning one.[ccxliv]

IBM is unfazed by this kind of criticism. It says that Watson is now being used by hundreds of companies to solve particular problems – companies like the Australian energy group Woodside, which used it to review 20,000 documents from 30 years of engineering projects to identify, for instance, the maximum pressure that a certain type of pipeline can withstand. It might be a form of marketing sleight of hand to apply the Watson brand to all these applications, but the company spent a great deal of time and money to create that brand, and it would be unreasonable to expect it not to try and recoup that investment.

That said, IBM is developing a new brand for its commercial AI offering. Celia stands for Cognitive Environments Laboratory Intelligent Assistant, and it seems to be a more user-friendly front end, enabling business analysts, for instance, to interact with it by speech, and by manipulating virtual objects in an augmented reality field.[ccxlv]

And IBM is still pursuing moonshots, in the medical field and elsewhere. As we have noted several times, machine learning is fuelled by data. In October 2015, IBM paid $1bn for Merge Healthcare, a company with 30 billion medical images,[ccxlvi] and $2bn for the digital assets of The Weather Company, to build a weather forecasting service. At the end of the year it unveiled Avicenna, a product of the Watson healthcare business unit designed to help radiologists prioritise which images to review, and help them make diagnoses.[ccxlvii] The interesting question is, how long before at least some of those radiologists turn out to be superfluous to the process.

Operations

You might think that the hands-on physical and frankly messy business of surgical operation will be undertaken by humans rather than machines for the foreseeable future. Probably not. One of the most highly skilled professionals in the emergency suite is the anaesthetist, and Johnson & Johnson has an automated version called Sedasys which, despite fierce opposition from the profession, has FDA approval to provide the anaesthesia in less challenging procedures like colonoscopies. It has carried out thousands of operations in Canada and the USA.[ccxlviii] In March 2016, Johnson & Johnson announced that it was exiting the Sedasys business due to sluggish sales, despite the machine costing $150 per operation whereas a human anaesthetist costs $2,000.[ccxlix] This will certainly not be the last setback in the progress of the machines, but in the long run the economic facts will prevail - although perhaps more slowly in industries where the normal rules of the market do not always apply.

In May 2016 an academic paper announced that a robotic surgeon had out-performed human peers. The Smart Tissue Autonomous Robot (STAR) operated on pig tissue and did the job better, although four times more slowly, than humans operating alone - and also better than humans aided by the semi-robotic Da Vinci system.[ccl]

Education

Teachers are the active ingredient in education – at school level, anyway. Studies have shown repeatedly that the quality of teaching makes an enormous difference to how well a student

performs at school and afterwards. But schools cannot afford enough of them, governments burden them with bureaucracy, and most countries' cultures under-value them.

What happens when the learning of every pupil is monitored minutely by artificial intelligence? When every question she asks and every sentence she writes is tracked and analysed, and appropriate feedback is provided instantly? Teachers will play the role of coach instead of instructor, but as with the other professions, their scope for contribution will shrink.

The beachhead for AI in education is marking, also known as grading. This is the bane of many teachers' lives, and they will welcome an assistant which can relieve them of the duty. A company called Gradescope marks the work of 55,000 students in 100 US universities, marking simple, multiple-choice types of test. It raised $2.6 million in April 2016 to develop its product into complex questions and essays.[ccli] Large corporates which provide education services like Pearson and Elsevier are moving in the same direction.

Towards the end of 2015, 300 students at Georgia Institute of Technology were, unbeknownst to them, guinea pigs in an experiment to see whether they would notice that one of their nine teaching assistants was a robot. Only ever in contact via email, they would ask questions like "Can I revise my submission to the last assignment?" and receive answers back like "Unfortunately there is not a way to edit submitted feedback." None of the students noticed that Jill Watson, named after the IBM Watson system "she" ran on, was in fact an AI.[cclii]

Financial services

The finance sector is an obvious target for machine intelligence, with high-value (and high-priced!) services based on vast amounts of data. Human equity analysts and brokers will increasingly struggle to provide value in the face of competition with machines which can ingest all the relevant data, and never forget any of it. The provision of advice to investors is also migrating to machines, with systems like SigFig incorporating a client's risk appetite and investment style into its algorithms' analysis of low-cost opportunities and recommendations.[ccliii] Similar so-called "robo-adviser" services are available from Betterment, Wealthfront and Vanguard.[ccliv]

These services deploy primitive forms of AI at the moment. According to market research firm Preqin, thousands of hedge funds, managing $200bn of assets, use computer models in most of their trades. But they are using traditional statistical methods rather than AI which learns and evolves. This is changing. Bridgewater, the world's largest hedge fund, hired David Ferrucci away from IBM, where he had project managed the development of the version of Watson which beat Ken Jennings at Jeopardy.[cclv] January 2016 saw the inaugural trades of Aidyia, a hedge fund based in Hong Kong whose chief scientist is Ben Goertzel, one of the leading researchers in artificial general intelligence.[cclvi]

In May 2016, the AHL hedge fund announced that it was stepping up its use of machine learning, having researched it for five years and deployed it experimentally for three. With $19 billion under management, AHL is the largest part of the MAN Group, which in turn is the world's largest publicly-traded hedge fund.[cclvii]

"The human mind has not become any better than it was 100 years ago, and it's very hard for someone using traditional methods to juggle all the information of the global economy in their

head," says David Siegel of Two Sigma, another hedge fund which uses AI. "The time will come that no human investment manager will be able to beat the computer."[cclviii]

"Algo trading" has many critics in financial circles, who point out that they chase spurious correlations (such as the fact that divorce proceedings in Maine have consistently tracked sales of margarine), and that they can move markets in ways that are impossible to follow and are potentially dangerous. But in a financial world become so complex that mere humans can no longer follow it, they may not only be inevitable, but also necessary. David Siegel says "People talk about how robots will destroy the world, but I think robots will save it."

What is known as "fintech" is one of the hottest areas for VC investment at the time of writing, and banks are spending considerable amounts of time and energy on working out where the most powerful disruption to their business models will come from, and whether they can do the disrupting themselves rather than be its victims. Banking, especially retail banking, has traditionally been a conservative, slow-moving industry, but the pace is picking up. Goldman Sachs reports that 40% of all US cheques are now processed electronically, despite that service being only four years old.[cclix]

Top analysts are among the highest-paid people in investment banks. They are targeted by a number of fintech companies like Kensho, which sorts through thousands of data sets to produce reports in minutes which would take skilled humans days. For instance, asking the system about Syrian Civil War will generate a report showing its impact on companies, currencies and commodities in as many countries as you like. Kensho's founder, Daniel Nadler, thinks between a third and a half of finance employees will be redundant within a decade.[cclx]

One of the drivers of AI use by financial services firms is the ever-growing and increasingly complex web of compliance requirements imposed by governments and regulators. Systems like IPSoft's Amelia help insurance firms and other financial services companies to navigate this web and make sure the forms and procedures used by staff are up-to-date.[cclxi]

Global banks are regularly fined hundreds of millions of dollars for carrying out illegal or sanctioned trades. Standard Chartered's regulatory costs rose 44% in 2015 to $447m as it was obliged to hire thousands of additional staff to deal with compliance requirements. In March 2016 it announced a major investment in AI systems to oversee its traders' behaviour, and to match their activities against regulatory norms.[cclxii]

It seems that managers in financial services are becoming aware of the potential threat to their livelihood. A survey of 1,700 managers in 17 different industries carried out in the autumn of 2015 by the consulting firm Accenture revealed substantial anxiety about automation. Overall, a third of the managers feared that intelligent machines threatened their jobs, with the level rising to 39% among senior managers. Unsurprisingly, the anxiety was highest (50%) among managers in the technology sector, but it was also high in banking, at 49%.[cclxiii]

3.10 – Jobs or no jobs

The question

In chapter 2 we saw that previous rounds of automation in the industrial revolution did not cause unemployment in the long-term – although the long term is fairly long in this context: the Engels pause that we came across in chapter 2.4 lasted at least a generation (a quarter-century).

In chapters 3.1 to 3.3 we reviewed in brief the arguments of those who believe it is different this time – and those who believe the opposite.

In chapter 3.4 we considered the state of the art of artificial intelligence, and in chapter 3.5 we saw how dramatic the impact of exponential growth can be. (No apologies for repeating that point: it is critical.) Then in chapter 3.7 we reviewed the likely evolution of technologies which are associated with AI, and caused or enabled by it.

Finally, in chapters 3.8 and 3.9 we discussed how various occupations could be automated, starting with the poster child of driving vehicles, and concluding with the privileged elites in the professions.

Now we are ready to address head-on the first great question posed by machine intelligence automation, which is this: Is it different this time? Will the automation of jobs by machine intelligence lead to widespread, lasting unemployment?

For the answer to be negative, we will have to dramatically increase the supply of jobs which are for some reason immune to automation by machine intelligence.

Jobs, not work

It is important at this point to distinguish between jobs and work. Physicists define work as the expenditure of energy to move an object,[cclxiv] but what we mean by it here is the application of energy in pursuit of a project. That energy could be physical, mental, or both. Work could be instigated by an employer, but it also could be purely personal: building or decorating a home, pursuing a hobby, or an unpaid community endeavour.

A job, on the other hand, is always paid labour for the purpose of this discussion. It might be a salaried occupation with a single, stable employer, or it could involve self-employment, or freelance activity. Your job is the way you participate in the economy, and earn the money to buy the goods and services that you need to survive, and enjoy a good standard of living.

If a machine carries out a job, there is no point a human replicating the work it is doing: she will not be paid, so she will have to look for some other way to generate an income.

The gig economy

We saw in chapter 3.2 how consultants at McKinsey noted that jobs can be analysed into tasks, some of which can be automated with current machine intelligence technology, and some of which cannot. This is an important insight and suggests that jobs will be sliced and diced, with some

tasks being automated, and other tasks being retained by the human who previously did the whole job.

Some would argue that this process is already under way. Parts of the economies of developed countries are being fragmented, or Balkanised, with more and more people working freelance, carrying out individual tasks which are allocated to them by platforms and apps like Uber and TaskRabbit.

There are many words for this phenomenon: the gig economy, the networked economy, the sharing economy, the on-demand economy, the peer-to-peer economy, the platform economy, and the bottom-up economy.

Is this a way to escape the automation of jobs by machine intelligence? To break jobs down into as many component tasks as possible, and preserve for humans those tasks which they can do better than machines? Probably not, for at least two reasons. First, it is precarious, and secondly, the machines will eventually come for all the tasks.

Working for yourself can seem an appealing prospect if your current job is a poorly-paid round of repetitive and boring activities. There is freedom in choosing your own hours of work, and fitting them around essential parts of your life like children and hangovers. There is freedom in choosing who you work with, and in not being subject to the arbitrary dictates of a vicious or incompetent boss, or the unfathomable rules and regulations of a Byzantine bureaucracy.

If you are lucky enough to be exceptionally talented, or skilled at a task which is in high demand, then you really can choose how and when you work. But freelancing can have its downsides too. Many freelancers find they have simply traded an unreasonable boss for unreasonable clients, and feel unable to turn down any work for fear that it will be the last commission they ever get. Many freelancers find that in hindsight, the reassurance of a steady income goes a long way to compensate for the 9 to 5 routine of the salaried employee.

Whether or not the new forms of freelancing opened up by Uber, Lyft, TaskRabbit, Handy and so on are precarious is a matter of debate, especially in their birthplace, San Francisco. Are the people hired out by these organisations “micro-entrepreneurs” or “instaserfs” - members of a new “precariat”, forced to compete against each other on price for low-end work with no benefits? Are they operating in a network economy or an exploitation economy? Is the sharing economy actually a selfish economy? Whichever side of this debate you come down on, the gig economy is a significant development: a survey by accounting firm PricewaterhouseCoopers found that as many as 7% of US adults were involved in it.[cclxv]

But our concern here is not whether the gig economy is a fair one. It is whether it can prevent the automation of jobs by machine intelligence leading to widespread unemployment. The answer to that is surely No: as time goes by, however finely we slice and dice jobs into tasks, more and more of those tasks are vulnerable to automation by machine intelligence as it improves its capabilities at an exponential rate.

What tasks, if any, will machines remain unable to automate for the foreseeable future?

Centaurs

A computer first beat the best human at chess back in 1997. Deep Blue was one of the most powerful computers in the world when it beat Gary Kasparov; the match was close and the result was controversial. Today, a programme running on a laptop could beat any human.

But Kasparov claims that a very good human chess player teamed up with a powerful chess computer can beat a second chess computer playing on its own. Humans can undermine the game of a computer by throwing in some surprise moves which don't make much sense in the short term, or by deploying an intuitive strategy. Matches between humans working with computers are called advanced chess, or centaur chess. Kasparov himself initiated the first high-level centaur chess competition in Leon, in Spain, in 1998, and competitions have been held there regularly ever since. Tyler Cowen (one of the sceptics about machine automation that we met in chapter 3.3) explores this form of chess extensively in his book, "Average is Over".

Some people believe this phenomenon of humans teaming up with computers to form centaurs is a metaphor for how we can avoid most jobs being automated by machine intelligence. The computer will take care of those aspects of the job (or task) which are routine, logical and dull, and the human will be freed up to deploy her intuition and creativity. Engineers didn't become redundant just because computers replaced slide rules. Kevin Kelly, founder of Wired magazine, puts it more lyrically: machines are for answers; humans are for questions.[cclxvi]

The trouble is that the intuition and creativity which we humans bring to tasks and jobs is largely a matter of pattern recognition, and machines are getting better at this at an exponential rate. A doctor may be happy to delegate the routine diagnosis of a cold or a flu to a machine which can do it better than she can, if she gets to retain the more interesting and challenging diagnostic work. But what is to stop the machine overtaking the doctor in the more difficult cases as well? The lawyer is in the same boat: the tedious business of sifting through a haystack of documents looking for the needle of evidence is already being outsourced to machines. The more interesting and demanding task of devising a legal strategy is likely to follow suit.

Admittedly, there may need to be some level of human supervision of machine work until the machines acquire a degree of common sense. Before then, the blindly logical thought processes of a machine will not realise when a data glitch or a software bug has generated a bizarre conclusion which is unworkable or dangerous. But as we saw in chapter 3.5, the founding father of deep learning thinks that machines with common sense will appear in a decade or so. This does not mean that they will acquire consciousness, but merely that they will create internal models of the external world which will enable them to appreciate the impacts of glitches and bugs, just as we do.

Machines have already made considerable progress automating routine tasks, and indeed whole occupations where all the tasks are routine. As their performance improves they will increasingly take over tasks and jobs which are non-routine.

In response to a survey published in May 2016, the veteran AI researcher Nil Nilsson suggested laconically that before long, machines would be singing the song Irving Berlin composed for the 1946 Broadway musical "Annie Get Your Gun". The lyric is "Anything you can do, I can do better. I cam do anything better than you."[cclxvii]

The human touch

Some observers think that our salvation from machine intelligence automation lies in our very humanity. Our social skills, and our ability to empathise and to care mean that we carry out tasks in a different way than machines do. Machines are by definition impersonal, the argument goes, and this renders them unsuitable for some types of job.

David Deming, a research fellow at the US National Bureau of Economic Research, believes we are already seeing the implications of this. In a report published in 2015 he claimed that the fastest growth in US employment since as long ago as 1980 has been in jobs requiring good social skills. Jobs requiring strong analytical abilities but no social skills have been in decline – with the implication that they are already being automated.[cclxviii]

Unfortunately, it isn't true that humans want to deal with other humans whenever possible. The first automatic deposit machine, the Bankograph, was installed in a bank in New York in 1960, but it was rejected by its intended customers. Its inventor, Luther Simjian, explained that "The only people using the machines were prostitutes and gamblers who didn't want to deal with tellers face to face," and there were not enough of them to make the machines a worthwhile investment.[cclxix] The first cash dispensing machine, or ATM, was installed in a bank in North London in June 1967. At first, people were again hesitant to use it, but that changed when they realised they no longer had to queue for their cash, and they could access it when the banks were closed (which was most of the time, in those days). Very quickly, people showed a marked preference for the machine over the human bank teller.[cclxx]

Nursing is an occupation long associated with caring people. Images of Florence Nightingale emoting as she nursed the wounded of the Crimean War are deeply ingrained in the profession's self-image. But there is evidence from Japan, Denmark and elsewhere that robots make perfectly acceptable companions for sick people, and are sometimes preferable to their human equivalents. The Paro is a robotic seal developed for use in hospitals. Cute-looking, with big black eyes and covered in soft fur, it contains two 32-bit processors, three microphones, 12 tactile sensors, and it is animated by a system of silent motors. It recharges by sucking on a fake baby pacifier.

The Paro cost $15m to develop; it distinguishes between individual humans, and repeats behaviours which appear to please them.[cclxxi] It has proved especially popular with patients suffering from dementia. As Shannon Vallor, a philosophy professor at Santa Clara University remarked, "People have demonstrated a remarkable ability to transfer their psychological expectations of other people's thoughts, emotions, and feelings to robots."[cclxxii]

So humans are happy to interact with machines more often that we might intuitively expect. Furthermore, machines are much better at understanding humans than we might expect.

Robot therapist

The US Army has a big problem with post-traumatic stress disorder (PTSD) among veterans, not least because soldiers don't like to admit they have it. DARPA funds research at the University of Southern California to develop online therapy services, and the latest result is an online virtual therapist called Ellie.[cclxxiii] She is proving to be better than human therapists at diagnosing PTSD.

There are two reasons for this. First, soldiers feel less embarrassed discussing their feelings with an entity they know will not judge them. In one test, 100 subjects were told that Ellie was

controlled by a human, and another 100 were told that it was a robot. This second group displayed their feelings more openly, both verbally and in their expressions.[cclxxiv]

Secondly, and perhaps more interesting, Ellie gleans most of its information about what is going on inside the soldier's head from his facial expressions rather than from what he says. When talking to a human therapist, the soldier may successfully "sell" the idea that there is nothing wrong, because the human therapist listens closely to what he says, and may miss the subtle facial signals that contradict him. Counter-intuitively, people with depression smile just as frequently as happy people, but their smiles are shorter and more forced. Ellie is superb at catching this.[cclxxv]

Most people would probably agree with David Deming when he says that "Reading the minds of others and reacting is [a skill that] has evolved in humans over thousands of years. Human interaction in the workplace involves team production, with workers playing off of each other's strengths and adapting flexibly to changing circumstances. Such non-routine interaction is at the heart of the human advantage over machines." But we may soon have to re-think that.

It is far from clear that there could ever be enough jobs in the so-called caring professions to employ all the people who would in previous generations have been drivers, doctors, lawyers, management consultants and so on. Especially if machines are muscling into the caring professions too.

Made by hand

Another way that people have suggested the human touch could preserve employment is that we will place a higher value on items manufactured by humans than on items manufactured by machines. It is hard to see much evidence of this in today's world outside some niche areas like hand-made cakes.[cclxxvi] Not many people today buy handmade radios or handmade cars.

There are four reasons why people might prefer products and services made by humans rather than machines: quality, loyalty and variation and status.

If humans produce a better product or provide a better service than machines then other humans will buy from them. But the argument of this chapter is that certainly in many areas, and probably in most, machines will produce goods and services cheaper, better and faster.

Loyalty to our species might be a better defence. "Buy hand-made, save a human!" sounds like a plausible rallying-cry, or at least a marketing slogan. The past is not always a reliable guide to the exponential future, but it is a good place to start, and unfortunately it does not augur well for appeals to loyalty. In the late 1960s, Britain was feeling queasy as the Empire dissolved and Germany's economic power was returning. The "I'm backing Britain" campaign started in December 1967, trying to get British people to buy domestically manufactured products instead of imports. It fizzled out within a few months.[cclxxvii]

Car manufacturing has long been symbolic of a nation's manufacturing virility. In the 1950s, Britain was the world's second-largest car manufacturer after the US, but in the 1960s its designs and build quality fell behind first its European rivals, and then the Japanese. Despite repeated appeals to buy British, sales declined and in 1975 the remaining national manufacturer, British Leyland, was nationalised. It never recovered, and Britain is now home to none of the major global car

brands. (Fortunately it has many innovative and thriving automotive design and component businesses, and it manufactures in record numbers for foreign brands.)

Appealing to people to buy handmade items out of loyalty to one's species may not have a huge economic impact if machine-made items are better quality and much, much cheaper. And in a world of falling employment, most people are going to have to buy as efficiently as they can.

The third reason for buying from humans could be summarised by the phrase "artisanal variation". We like antiques because the patina of age gives them personality: each one is unique. The same goes for the original work of an artist, even if it isn't a Vermeer or a Rubens. But for most people, this is the preserve of luxury items, a few select pieces which we keep on display. Most of our possessions are mass-produced because they are much cheaper, disposable, and we can afford a better lifestyle that way.

We have seen this before, in the second half of the 19th century. With the industrial revolution in full swing, William Morris helped found the Arts and Crafts movement to produce hand-made furnishings and decorations. His concern was to raise quality rather than to reduce unemployment, but in practice he ended up making expensive pieces which only the rich could afford.[cclxxviii]

Some people may choose to buy goods and services from humans rather than machines for reasons of status. But by definition, this could only ever amount to a niche activity, and would not save most of us from unemployability.

New jobs

If machines are going to take a great many, perhaps most, of our existing jobs, can we create a host of new ones – perhaps whole new industries - to replace them? Those who think we can point out that many of the jobs we do today did not exist a hundred years ago. Our grandparents would not have understood what we mean by website builder, social media marketer, user experience designer, chief brand evangelist, and so on. Surely, the argument goes, all these new technologies we have been talking about will throw up many new types of jobs that we cannot imagine today.

As the person probably most responsible for Google's self-driving cars, Sebastian Thrun is a man worth listening to on the subject. He is optimistic: "With the advent of new technologies, we've always created new jobs. I don't know what these jobs will be, but I'm confident we will find them."[cclxxix]

Unfortunately, past experience is (again) not as encouraging as you might think. Gerald Huff is a software engineer working in Silicon Valley, ground zero of the developments we are talking about. Nervous about the prospect of technological unemployment, he carried out a comparative analysis of US occupations in 1914 and 2014. Using data from the US Department of Labour,[cclxxx] he discovered that 80% of the 2014 occupations already existed in 1914. Furthermore, the numbers of people employed in the 20% of new occupations were modest, with only 10% of the working population engaged in them. The US economy is much bigger today than it was in 1914, and employs far more people, but the occupations are not new.

Of course, those of us who argue that it is different this time cannot rely on the historical precedent. It might be different this time in that vast swathes of new jobs *will* be created –

including jobs for averagely-skilled people, not just relatively high-skill jobs like social media marketing. But those who argue that we are falling for the Luddite fallacy cannot argue that history points to everybody getting new types of jobs which are more interesting and safer after a period of adjustment. It doesn't.

If we *were* to create a host of new jobs, what might they be? Maybe some of us will become dream wranglers, guiding each other toward fluency in lucid dreaming. Others may become emotion coaches, helping each other to overcome depression, anxiety, and frustration. Maybe there will be jobs for which we have no words today, because the technology has not yet evolved to allow them to come into being.

It's not hard to imagine that virtual reality will create a lot of new jobs. If it is addictive as enthusiasts think it will be, many people will spend a great deal of their time – perhaps the majority of it – in VR worlds. In that case there will be huge demand for new and better imaginary or simulated worlds to inhabit, and that means jobs.

But does it mean jobs for humans? Although the credit list for the latest superhero blockbuster stretches all the way around the block as it names everyone involved in rotoscoping and compositing the hyper-realistic armies of aliens, the latest CGI technology also makes it possible for two teenagers with a mobile phone to make a film which gains theatrical distribution. Their increasingly powerful software and hardware allows Hollywood directors to conjure visual worlds of such compelling complexity that their predecessors would rub their eyes in disbelief, but it also allows huge quantities of immersive content to be developed by skeleton crews. There will probably always be an elite of directors who are highly paid to push the boundaries of what can be imagined and what can be created, but software will do more and more of the heavy lifting in VR production.

Not for the first time, the games industry shows what is possible. A game called “No Man’s Sky” was announced in 2014 which conjures far more imaginary worlds than you could visit in a lifetime purely by the operation of algorithms and random number generators. You boldly go where no programmer or designer has gone before.[cclxxxi]

Education

It is surprising how many smart people think that education is the answer to automation by machine intelligence. Microsoft CEO Satya Nadella said in January 2016: “I feel the right emphasis is on skills, rather than worrying too much about the jobs [which] will be lost. We will have to spend the money to educate our people, not just children but also people mid-career so they can find new jobs.”[cclxxxii]

Massive Open Online Courses, or MOOCs, are promoted as the way we will all re-train for a new job each time a machine takes our old one. MOOCs are important, and along with flipped lessons, competency-based learning, and the use of Big Data, they will improve the quality of education, and make excellent learning opportunities available to all. (With flipped lessons, students watch a video of a lecture for homework, and then put what they have been told into practice in the classroom. The teacher acts as coach and mentor, a more interactive role than lecturing. Competency-based learning requires students to have mastered a skill or a lesson before they move on the next one; students within a class may progress at different speeds. Big data enables

students and teachers to understand how well the learning process is going, and where extra support is needed.)

Exciting and powerful as these techniques are, they won't protect us from technological unemployment. We have seen that machines are increasingly capable of performing many of the tasks currently carried out by highly educated, highly paid people. The machines aren't just coming for the jobs of bricklayers; they're coming for the jobs of surgeons and lawyers too.

There is a very important postscript to these remarks about education. If we make it through successfully to the new world in which many or most of us are permanently and irrevocably unemployed, then education will be more important, not less. We will need good education to take advantage of our leisured lives even more than we did to survive our working lives.

Entrepreneurs

If machines take over the jobs that are repetitive, humans will look to do things that require creativity, intuition, and pursue counter-intuitive paths. One job title which fits that description is entrepreneur.

In my experience there are two types of entrepreneur. Both are resourceful, determined, and usually of above-average intelligence. The first and most common type is someone who works in an organisation which is doing something poorly. They notice this, and decide to offer a better version. They utilise essential skills and industry know-how acquired while working for the original organisation, and simply improve incrementally on what was being provided there. These people are talented and hard-working, but they also had the good fortune to be in the right place at the right time to spot the opportunity. If they had not been in that position they would have spent their careers working for other people, and because they are hard-working and bright they would probably have make a good fist of that.

The second type is destined to be an entrepreneur whatever circumstances life drops them into. They will never be happy working for someone else. They envision themselves in a future world which looks impossible to anyone else, but they choose to believe it and by dint of sheer force of will they make that future a reality. They will walk through brick walls to make it happen, and will probably be bankrupt more than once. They are charming, astonishingly energetic, and often rather hard to be around. In the words of LinkedIn founder Reid Hoffman, they are people who will happily throw themselves off a cliff and assemble an aeroplane on the way down.[cclxxxiii]

Both types of entrepreneur are rare, and especially the second kind – which may be a good thing for the rest of us. In any case, this is probably not an occupation that is going to save large numbers of people from technological unemployment. The other thing to remember about entrepreneurship as a career is that most startups fail.

Artists

After all these apparently gloomy prognostications, let's close this chapter on a more optimistic note. There is one profession which can probably never be automated until the arrival of an artificial general intelligence which is also fully conscious. That profession is art, and to understand why, it is important to distinguish between art and creativity.

Creativity is the use of imagination to create something original. Imagination is the faculty of having original ideas, and there seems to be no reason why that requires a conscious mind to be at work. Creativity can simply be the act of combining two existing ideas (perhaps from different domains of expertise) in a novel way.

The eminent 19th-century chemist August Kekule solved the riddle of the molecular structure of benzene while day-dreaming, gazing into a fire.[cclxxxiv] True, he had spent a long time before that pondering the problem, but according to his own account, his conscious mind was definitely not at work when the creative spark ignited. You might argue that Kekule's sub-conscious was the originator of the insight, and that a sub-conscious can only exist where there is consciousness, but that seems to me an assertion that needs proving.

Can computers be creative? In mid-2015, Google researchers installed a feedback loop in an image recognition neural network, and the result was a series of fabulously hallucinogenic images.[cclxxxv] To deny that they were creative is to distort the meaning of the word.

Art is something different. Admittedly, this is a personal definition, and perhaps not everyone would agree, but surely art involves the application of creativity to express something of personal importance to the artist. It might be beauty, an emotion, or a profound insight into what it means to be human. (If that disqualifies a good deal of what is currently sold under the banner of art, then so be it - in fact, three cheers.)

To say something about your own experience clearly requires you to have had some experience, and that requires consciousness. Therefore, until a conscious artificial general intelligence (AGI) arrives, AIs can be creative but not artistic. This in turn means that while Donna Tartt and Kazuo Ishiguro are probably OK for a few decades, today's successful genre writers who use stables of assistants to churn out several crime and romance novels each year have chosen the timing of their careers expertly, and their publishers had better find something different to do.

In April 2016, researchers from Microsoft, a Dutch university and two Dutch art galleries created an AI which analysed the way Rembrandt painted. It identified enough of his techniques and mannerisms to enable it to produce paintings in exactly his style – better than any human forger could. They had it design a new picture in Rembrandt's style, of a subject Rembrandt had never worked on, and 3D printed it to capture the old master's technique in three dimensions. Because the machine recognises patterns better than humans can, it may well teach us interesting new things about the way Rembrandt created his masterpieces. But it is not producing art.[cclxxxvi]

3.11 – What's the problem?

The Star Trek economy

It is often said that science fiction tells us more about the present than it does about the future. Most science fiction writers are not actually trying to predict the future, although they may go to considerable lengths to try to make the worlds they create seem plausible. Generally, they are just trying to tell an entertaining story, or maybe use the opportunity that the genre offers to explore something about the fundamental nature of our lives. (At its best, science fiction is philosophy in fancy dress.)

But intentionally or otherwise, science fiction does a very important job for all of us when we think about the future: it provides us with metaphors and scenarios. Many of the most popular science fiction stories present dystopian scenarios: think Terminator, Blade Runner, 1984, Brave New World, and so on. But there are also positive scenarios, and one of the most popular ones is Star Trek.

Set in the 24th century, Star Trek presents a world of immense possibility, of interstellar travel, adventure, and split infinitives. And a world without money or poverty. In the 1996 movie "Star Trek: First Contact", Captain Jean-Luc Picard explains that "Money doesn't exist in the 24th century. The acquisition of wealth is no longer the driving force in our lives. We work to better ourselves and the rest of humanity."

This was not a feature of the original TV series, in which there were quite a few mentions of money, and systems of credit. But before he died, Gene Roddenberry stipulated that there was to be no money in the Federation.[cclxxxvii]

Although there is no money, the people in the later Star Trek stories do compete with each other – for prestige, for approval, for increased responsibility and for career advancement. One of the things that makes James Tiberius Kirk an outstanding starfleet commander is his fiercely competitive nature. He operates in a profoundly meritocratic environment, and will sacrifice a great deal to win.

This is not new. Men and women have always competed for pre-eminence within their tribes and societies, and we are continually applying our ingenuity to work out new ways to do so. Mediaeval knights risked life and limb for honour and glory, and their descendants fought for national self-determination. Today, many people expend considerable sweat and tears – if less blood – to demonstrate their prowess in writing elegant open source software, or edit Wikipedia pages.

Money is not required in Star Trek's United Federation of Planets because energy has become essentially free, and products can be manufactured in so-called Replicators, devices which create useful (including edible) objects out of whatever matter is available.

Another popular science fiction series with a broadly optimistic (if darkly humorous) outlook is the late Iain M. Banks' "Culture" books, set in a distant future when a technologically advanced humanity has colonised swathes of the galaxy, and enjoys mostly peaceful relations with a host of alien civilisations. The humans are kept company and aided by vastly superior and extraordinarily indulgent machine intelligences, and they lead lives of perpetual indulgence. As Banks put it in a 2012 interview, "It is my vision of what you do when you are in a post-scarcity society, you can

completely indulge yourself. The Culture has no unemployment problem, no one has to work, so all work is a form of play."[cclxxxviii]

Abundance

The Star Trek economy is the post-scarcity economy, the economy of radical abundance. In their 2012 book "Abundance: the future is better than you think", Peter Diamandis and Stephen Kotler argue that this world is within reach in the not-too-distant future, thanks largely to the exponential improvement in technology.

Financial Times columnist Martin Wolf urged that we should "enslave the robots and free the poor",[cclxxxix] and who would not welcome such an outcome? Perhaps if we play our cards right, automation by machine intelligence will simply mean that we humans get to spend our long and healthy lives playing, learning, enjoying each other's company, having adventures and fun.

Of course, life is rarely so simple or so easy. In chapter 5 we will explore some of the challenges and hurdles to be overcome. But let's pause for a moment, and beguile ourselves with Richard Brautigan's poetic wish, recorded back in 1967.

"I like to think
(it has to be!)
of a cybernetic ecology
where we are free of our labors
and joined back to nature,
returned to our mammal
brothers and sisters,
and all watched over
by machines of loving grace."[ccxc]

3.12 – Conclusion: yes, it's different this time

It's time to answer the question: is it really different this time? Will machine intelligence automate most human jobs within the next few decades, and leave a large minority of people – perhaps a majority – unable to gain paid employment?

It seems to me that you have to accept that this proposition is at least possible if you admit the following three premises:

1. It is possible to automate the cognitive and manual tasks that we carry out to do our jobs.

2. Machine intelligence is approaching or overtaking our ability to ingest, process and pass on data presented in visual form and in natural language.

3. Machine intelligence is improving at an exponential rate. This rate may or may not slow a little in the coming years, but it will continue to be very fast.

No doubt it is still possible to reject one or more of these premises, but for me, the evidence assembled in this chapter makes that hard.

Counter-arguments

The main argument against the proposition that machines will create widespread structural unemployment is that it hasn't happened in the past. In other words, the proposition is the Luddite fallacy. This is a weak argument at best – akin to saying that we have never sent a person to Mars so we never will. It is also partly false: during the Engels Pause in the first half of the 19th century unemployment rose, and labour's share of national income fell, in the UK at least. But the main point is that the whole question is whether it will be different this time.

People sometimes argue that humans will not become unemployable because there is an inexhaustible well of human wants and needs to fulfil. This doesn't help much if machines can fulfil any need better, faster and cheaper: it will always be economically compelling to build a machine to do the job.

The best argument against the proposition is that humans possess skills which machines will never replicate – or at least, not for many decades. This is inevitably a judgement call, since at the moment, no-one can be certain. My judgement, based on the evidence set out above, is that the skills for which we get paid will be acquired by machines in the next two to four decades.

Raising awareness

2015 was an important year for artificial intelligence. It was the year when our media caught on to the idea that AI presents enormous opportunity and enormous risk. This was thanks in no small part to the publication the previous year of Nick Bostrom's book "Superintelligence". It was also the year when cutting-edge AI systems used deep learning and other techniques to demonstrate human-level capabilities in image recognition, speech recognition and natural language processing. In hindsight, 2015 may well be seen as a tipping point.

Machines don't have to make everybody unemployed to bring about an economic singularity. If a majority of people – or even just a large minority - can never get hired again, we will need a different type of economy.

Furthermore, we don't have to be absolutely certain of this outcome to make it worthwhile to monitor developments and make contingency plans. After all, seeing past the event horizon of a singularity is hard.

As we have seen, and as we will explore in more detail in chapters 5 and 6, the outcome can be fantastic or terrible. To a large extent, it is up to us. If we can't achieve the positive, then perhaps we deserve (as Anders Sandberg put it in chapter 3.4) to be the boot loader for the digital superintelligence rather than its mitochondria.[ccxci]

4. - A timeline

4.1 – Un-forecasts

Three snapshots of a positive scenario

This chapter offers three snapshots of a possible future, one each for 2021, 2031 and 2041. Their purpose is to make the possibility of technological automation seem more real and less academic.

In each one there is a very brief description of the level of automation in a number of industries, and a summary of the impact of that automation on society. Together they depict a positive scenario, with a concluding vision of an economy of radical abundance which has been achieved without massive social dislocation. Why choose a positive scenario? Why not: it's the outcome we should be aiming for.

Before we start, there is an important caveat.

Unpredictable yet inevitable

We know that all forecasts are wrong. The only things we don't know are by how much, and in what direction. The future generally turns out to be not only different to what we expect, but also much stranger. Cast your mind back to 2005. Pretty much everyone thought that cellphones would continue to get smaller, and Facebook was limited to a few thousand universities and schools. Today, just a decade later, larger smartphones are a bit of a thing, and Facebook's valuation has overtaken that of Walmart, the world's largest shopkeeper.[ccxcii] Trying to predict how the world will look in 2031 is like trying to predict the weather on Saturday two months from now. There are just too many variables.

And yet in hindsight what happens appears not only natural, but almost inevitable.

The smartphone is a good example. Pretty much nobody suggested thirty years ago that we would all have telephones in our pockets which would contain powerful artificial intelligences, and which would only occasionally be used for making phone calls. After all, at the time a mobile phone was a fairly hefty device, the size of a small dog. But now that it has happened it seems obvious, logical, and perhaps even inevitable.

Here's why. We humans are highly social animals, and our social habits are facilitated by language. Because we have language we can communicate complicated ideas, suggestions and instructions: we can work together in teams and organise; we can defend ourselves against lions and hostile tribes, we can hunt and kill mammoths, produce economic surpluses, and develop technologies.

It is often said that no species is more savage and more violent than humans. This is no more true than the claim that Americans are more violent than other nationalities because their murder rate is higher than other developed countries. The only reason why humans kill more than other species is that we have more and better weapons.

Humans live cheek by jowl in cities containing millions. This is remarkable: no other carnivorous species can assemble more than a few dozen of its members in a limited space without them

killing each other in rivalry for food, sex, or social dominance. Other species lack our sophisticated ways to communicate and collaborate. Our bigger brains allow us to establish laws and cultural norms which govern the way we interact. We make up stories and agree to believe in them collectively, regardless of whether there is any evidence for them. These stories - about abstract concepts like gods and kingship, nations and ideologies, money and art - give us powerful reasons to cooperate and work together, even to die together.

Non-human primates spend hours every day grooming the other members of their tribe to reassure them that they will not sink their teeth and nails into them. It works, but it is inefficient, and means they cannot readily add new members to their tribe. Humans, by contrast, can walk past complete strangers on a crowded street without a second thought. Our superpowers are communication, and our capacity to sustain mutual belief in things for which we have no evidence. It is thanks to these abilities that we control the fate of this planet and every species on it.

(This means that the old cliché that our dominance is based on our capacity for rational thought is – unlike most clichés – untrue.)

So although it wasn't - and couldn't have been - predicted in advance, in hindsight it is entirely logical that our most powerful technology, artificial intelligence, would first become available to most of us in the form of a communication device.

The way the economic singularity unfolds will probably be like that. Our attempts to forecast the impact of technological unemployment – assuming it arrives – will probably look absurd in hindsight. But when we get there, the outcome will seem not only natural, but perhaps even inevitable.

Not forecasts

I am labouring this point because I want to be clear. The descriptions of a possible future that follow are not predictions. The only thing I am confident of is that the future will not be like this.

Instead, these timelines are intended to serve two functions. First, as I said above, they are a rhetorical device. The arguments in chapter 3.10 that machines will automate our jobs away have been either abstract or fragmentary, and as such, some readers may find them implausible. I'm hoping that the timelines will help make the possible future of an economic singularity seem less academic, less theoretical, and more real.

Secondly, drawing up timelines like these may in some small way help us to construct a valuable body of scenarios. Even when we know the future is unpredictable, it is still worth making plans. There is good sense in the old cliché that failing to plan is planning to fail. If you have a plan, you may not achieve it, but if you have no plan, you most certainly won't.

In a complex environment, scenario development is a valuable part of the planning process. None of the scenarios will come true in their entirety, and many will be completely off the mark. But parts of some of them may approximate some of the outcome. Thinking through how we would respond to a sufficient number of carefully thought-out scenarios could well help us to react more quickly when we see the beginnings of what we believe to be a dangerous trend.

Super un-forecasting

The art of constructing a useful scenario is the same as that of forecasting, which has been extensively studied by Canadian political scientist Philip Tetlock, co-author of the book "Superforecasting: the art and science of prediction." He has found that the best forecasters share a number of traits. First, they treat their views about what will happen as hypotheses, not firm beliefs. If the evidence changes, they change their hypothesis.

Secondly, they look for numerical data. Now we all know that there are lies, damned lies and statistics, and that data is often used in public debate in the same way that a drunk uses a lamp-post: more for support than for illumination. But used carefully and honestly, data is our friend. It is after all the root of the scientific revolution that has lifted most of our species out of poverty and squalor.

Thirdly, they look for context. He cites the example of guests at a wedding, admiring the beauty and grace of the bride and the dashing good looks of the bridegroom, and assuring each other that they will share a long and happy life together. The super-forecaster is a contrarian, noting that around half of all marriages fail, and that the failure rate increases with second and third marriages, especially when one or other partner has a history of infidelity, as with the happy couple today. If she is a tactful super-forecaster, she keeps these thoughts to herself.

Ironically, super-forecasters are often not the people who get listened to in discussions about the future. We tend to pay more attention to those who speak most confidently, and offer clarity and certainty. People who equivocate and offer measured suggestions often don't cut through the noise.

So here goes, with the equivocation minimised.

4.2 – 2021

1. **Transportation**. Numerous cities around the world are experimenting with self-driving cars, but very few are so bold as to omit human drivers altogether. Google's cars can now handle snow and heavy rain, and can distinguish between a pedestrian waving at a friend across the street and a policeman instructing the car to stop. Their sensors are still expensive, but they are getting cheaper quickly. Aspects of self-driving technology are becoming widespread, including intelligent cruise control and assisted parking. Many long-distance commuters let their cars do much of their driving, although watching TV while driving is still illegal.
2. **Manufacturing**. Industrial robots are getting cheaper, and much easier to programme to undertake new tasks. Manufacturers find that the choice between employing another person or buying a new robot is a close thing.
3. **Agriculture**. Farmers are experimenting with robots for both crops and animal husbandry. On a growing number of farms with high-value crops, small wheeled devices patrol rows of vegetables, interrogating plants which don't appear to be healthy specimens, and eliminating weeds with targeted jets of ecologically-sound hot water. Cattle are entirely content to be milked by robots, so the declining population of farm workers no longer has to get up before daybreak every day.
4. **Retail**. The shift towards purchasing goods and services online continues, and there is growing automation within shops. In some supermarkets, shoppers no longer have to unload and re-load their trolleys: the goods are scanned while still inside their baskets. Fewer attendants are required in the checkout area. In fast food outlets, so-called “McJobs” are disappearing as burgers and sandwiches are assembled and presented to customers without being handled by a human.
5. **Construction**. Although some developers are experimenting with pre-fabricated units, most of the cost of a construction project is generated by the variability of conditions on-site, including the foundations. Robots which can handle this unpredictability are still too expensive to replace human construction workers. There are experiments with exoskeletons for construction workers, but these are still expensive.
6. **Technology**. Firms are fighting to recruit and retain machine learning experts; the salaries and bonuses offered were previously unknown outside financial services and professional sports. Sales of wearables are growing, and the successors to Google Glass are out-selling smart watches.
7. **Utilities**. Water companies and power generation and transmission firms are building out fleets of tiny robots and drones which patrol pipes and transmission lines, looking for early warning signs of failure.
8. **Finance**. Retail banking is mostly automated and web-based, and consumer feedback on the quality of service is improving. Wealthy people now get some of their investment advice directly from automated systems, but human investment advisers still serve most of the market. In corporate finance, human advisers show no signs of being replaced, although their back office systems are heavily automated.
9. **Call centres**. Enquiry handling that was offshored to India and then repatriated to home countries is now being offshored again - this time to machines housed in cold climates where the cost of keeping the servers cool are lower.
10. **Media and the arts**. The market for virtual reality apps and shows is booming as Occulus Rift, Meta, and their competitors create demand for latency-free, high resolution content. As usual, porn and sport look like being the killer apps, but there are unexpected hits too, such as “how-to” shows about parenting and relationship enhancement.
11. **Management**. Little change.
12. **Professions**. The tedious jobs which traditionally provided training wheels for accountants and lawyers (“ticking and bashing” for auditors and “discovery” for litigators) are increasingly being handled by machines. Optimists – and sceptics about technological unemployment – point out

that the amount of work for trainee professionals has actually increased, as whole categories of previously uneconomic jobs have become possible, and the machines still need training on each new data set. But more thoughtful practitioners are writing articles in their trade magazines asking how long that will continue, and therefore how tomorrow's partners can learn their trade.

13. **Medical**. In two Scandinavian countries and a handful of US states, AIs are ingesting data sent by patients from their smartphones and carrying out triage. Sometimes they respond with simple diagnoses and treatment recommendations, sometimes they pass the enquiry to a human doctor. Medical professionals and regulators elsewhere are highly critical of these experiments, but the outcome data is impressive. Hospitals in Japan are using robot nurses to great effect, but these are also resisted elsewhere. Pharmaceuticals designed to raise the IQ of adults are in clinical trials.

14. **Education**. Teachers everywhere brandish the empirical evidence that the primary determination of educational outcome is the quality of the teaching. Some teachers are embracing new technologies enthusiastically, especially in competitive environments such as the UK's private school system. Others are resistant.

15. **Government**. There is a worldwide drive to get most government “services” delivered online and cheaper.

There is a great deal of discussion in the media about whether automation will lead to technological unemployment. A growing number of people think that it will, but they are outnumbered by those, including many prominent economists, who continue to deny it.

4.3 – 2031

1. **Transportation**. Most long-distance trucks are now capable of operating autonomously, and many operate routinely without a human on board. In some jurisdictions there are roads which are off-limits to human drivers. Most cars are still privately owned, but many people are experimenting with communally-owned vehicles, which enjoy free parking in most cities. Insurance premiums have plummeted, and fears about self-driving cars being routinely hacked have not been realised. A vocal minority of citizens (which, to general surprise, comprises equal numbers of men and women) are scathing about this arrangement, dubbing the communal cars THEMs, ("tedious horizontal elevator machines").
Many urban deliveries of fast food and small parcels in major cities are now carried out by autonomous drones, operating within their own designated level of airspace. Sometimes the last mile of a delivery is carried out by autonomous wheeled containers. Teenagers delight in "bot-tipping", but with all the cameras and other sensory equipment protecting the bots, it is a risky pastime.
2. **Manufacturing**. Many large factories and warehouses are dark: no light is required because no humans work there. People are becoming a rarity in smaller sites too.
3D printing has advanced less quickly than many expected, as it remained more expensive than mass production. But it is common in niche applications, like urgently required motor parts.
3. **Agriculture**. Farmers are moving heavily into leisure services, as their families and staff are losing their roles to robots.
4. **Retail**. Online shopping reaches 75% of all retail purchases, with a small but growing number of items being 3D printed domestically or in neighbourhood facilities, often with an element of customisation by the consumer. Human shop assistants are starting to be replaced by robots, except in high-margin sectors where they help create an experience rather than simply facilitating straightforward transactions.
5. **Construction**. Human supervision is still the norm for laying foundations, but pre-fabricated (often 3D-printed) walls, roofs and whole building units are becoming common. Robot labour and humans in exoskeletons are increasingly used to assemble them. Drones populate the air above construction sites, tracking progress and enabling real-time adjustments to plans and activities.
6. **Technology**. The first "inside-ables" are appearing, and have been made fashionable by Lord Beckham, the football and fashion magnate. The Internet of Things has materialised, with everyone receiving messages continuously from thousands of sensors and devices implanted in vehicles, roads, trees, buildings, etc. Fortunately, the messages are intermediated by personal digital assistants, which have acquired the generic name of "Friends", but whose owners often endow them with pet names.
New types of relationship and etiquette are evolving to govern how people interact with their own and other peoples' "Friends", and what personalities the Friends should present. Brand loyalty to the companies which provide the best Friends software is fierce.
There is lively debate about the best ways to communicate with their "Friends" and other computers. Most people communicate with them by muttering into implanted microphones, but millions of people are also learning how to use one-handed keyboards which liberates them from traditional keyboards at times when voice is inappropriate. Some believe these new keyboards will be quickly superseded by Brain-Computer Interfaces (BCI), but this has made less progress than its early enthusiasts expected.
A growing amount of entertainment and personal interaction is mediated through virtual reality. It is increasingly rare to see an adolescent in public outside school hours.
Polls suggest that most people now think that artificial general intelligence (AGI – machines which equal or surpass human cognition in all domains) is a serious possibility within a generation or two.

Significant expenditure is flowing into research on how to make sure the outcome is positive, and the moral and religious implications are hotly debated.

7. **Utilities**. In many organisations, most operations are now automated. The main role of humans in these organisations is testing security arrangements. Several hundred people died in two significant hacking incidents – one in the US and one in Europe. This has prompted huge investment in upgraded security arrangements. In another high-profile incident, AI management systems managed the disaster containment and recovery process flawlessly, and much faster than humans could have.

8. **Finance**. Retail banking is now fully automated, and investment advice is going the same way. Corporate financiers are in retreat, and their previously stratospheric incomes have fallen sharply.

9. **Call centres**. Almost no humans now work in call centres.

10. **Media and the arts**. All major movies made by Hollywood and Bollywood are now produced in VR, along with all major video games. To general surprise, levels of literacy – and indeed book sales – have not fallen. In a number of genre categories, especially romance and crime, the most popular books are written by AIs.

Major sporting competitions have three strands: robots, augmented humans, and un-augmented humans. Audiences for the latter category are dwindling.

Long-distance communication is massively improved by VR Skype.

Dating sites have become surprisingly effective by requiring their members to provide clothing samples from which they extract data about their smells and their pheromones. The discovery that relationship outcomes correlate closely with these data have slashed divorce rates.

11. **Management**. The ranks of middle management are thinning out. Shareholders are investing heavily in Distributed Autonomous Corporations (DACs), firms consisting of unsupervised AIs which create new business models and strategies and transact with other firms without any humans in the loop.

12. **Professions**. Partners in law firms and accountancy firms are working shorter hours. Human intakes to these firms are dwindling. Most criminal law cases now relate to digital crime: it is more lucrative than physical crime, and easier to avoid surveillance.

13. **Medical**. Opposition to the smartphone medical revolution has collapsed in most countries, and most people obtain diagnoses and routine health check-ups from their “Friends” several times a week. Automated nurses are becoming increasingly popular, especially in elder care.

Several powerful genetic manipulation technologies are now proved beyond reasonable doubt to be effective, but backed by public unease, regulators continue to hold up their deployment.

Cognitive enhancement pharmaceuticals are available in some countries under highly regulated circumstances, but are proving less effective than expected. There are persistent rumours that they are deliberately being engineered that way.

Ageing is coming to be seen as an enemy which can be defeated.

14. **Education**. Data on learning outcomes is steamrollering teachers' resistance to new approaches. Customised learning plans based on continuous data crunching are becoming the norm. Teachers are becoming coaches and mentors rather than instructors. Some schools are experimenting with classroom AIs.

15. **Government**. There is growing pressure to reduce the numbers of politicians and civil servants, as more and more government services are automated. Many jurisdictions are debating the merits of using technology to enable direct democracy, which is being pioneered by Switzerland. Most people are sceptical, fearing the tyranny of the temporary majority. Policemen in most countries record all interactions with members of the public, and public satisfaction levels with them are generally rising.

16. **Charities**. Non-profit organisations are enjoying a surge, thanks to an influx of talent as capable people can't find work elsewhere.

The majority of people in most countries now believe that an economic singularity is coming and that a universal basic income (UBI) will be needed, as before long a very large minority of citizens will be permanently unemployed. Experiments in numerous cities around the world, and a couple of country-scale experiments demonstrate that most people don't succumb to drugs or despair, although a significant minority does, and needs help. There is vigorous debate about how to pay for the UBI.

4.4 – 2041

1. **Transportation**. Humans very rarely drive vehicles on public roads, and few commercial vehicles have human attendants. Young people no longer take driving tests. Motor sports are mostly competitions between self-driving cars. Congestion and parking are no longer a problem. The population of cars has declined dramatically as they are used far more efficiently, and the automotive industry has contracted. Large numbers of dependent businesses (and jobs) are disappearing too, including repair shops and insurance brokers.
2. **Manufacturing**. Almost all factories and warehouses are dark. 3D printing is beginning to look competitive with some forms of mass production.
3. **Agriculture**. Robots do most farm work.
Some countries have large communally-owned agricultural processing concerns which send out meal ingredients on drones in a service described as Netflix for food.
4. **Retail**. Most items are now bought online, and around half of all products sold at retail are 3D printed. Retail outlets on High Streets and city centres are mostly experiential rather than transactional, and mostly staffed by AIs and robots.
5. **Construction**. Robots now carry out most of the work on construction sites.
6. **Technology**. Since AI provides a large proportion of the value in most products and services, there is a major concentration of capital and wealth in the hands of shareholders and key employees in this sector. Its foremost talent is now applied to developing artificial general intelligence and making sure that it is safe for humans. The Internet of Things is all-pervasive, and the environment appears intelligent.
The companies that provide "Friends" have been obliged to make them open-source. Friends are so critical to everyone's lives that being restricted to any one company's walled garden was unacceptable.
7. **Utilities**. Overwhelmingly automated.
8. **Finance**. Overwhelmingly automated.
9. **Call centres**. Unchanged.
10. **Media and the arts**. In sports, robot competitions now generate larger audiences than their human counterparts. The International Olympics Committee de-lists the human versions of around half of all sports.
Haptic body suits combined with VR headsets now provide truly immersive virtual environments. Counselling (by AIs) is required by a section of the population who struggle to maintain the distinction in their minds between reality and VR.
To general surprise, people still read books, but they are very different products now, with holographic illustrations, and often with several alternative story lines developed by their AI authors, which readers can choose between.
Dating sites are now mostly accessed by personal digital assistants. "My Friend likes your Friend" has become a standard opener.
11. **Management**. Many companies now consist of just a few strategists, whose main role is to forecast the optimal business model for the next financial quarter, but they are struggling to keep up with their AI advisers.
12. **Professions**. Accountancy and the law are largely automated.
13. **Medical**. Demand for human doctors is dwindling and professional nursing has been almost entirely automated. Everyone in developed economies has their health monitored continuously by their "Friends". Most people spend a certain amount of time each week visiting family, friends and neighbours who are unwell, just to converse. Sick and disabled people are greatly comforted by their relationships with talking AI companions, some resembling humans, others resembling animals. Significant funds are now allocated to radical age extension research, and there is talk of

"longevity escape velocity" being within reach – the point when each year, science adds a year to your life expectancy. Most forms of disability are now offset by implants and exoskeletons, and cognitive enhancements through pharmaceuticals and brain-computer interface techniques are showing considerable promise.
14. **Education**. The sector has ballooned, with many people now regarding it as recreation rather than work. Most education is provided by AIs.
15. **Government**. Safeguards have now been found to enable direct democracy to be implemented in many areas. Professional politicians are now rare.

Unemployment has passed 50% in most developed countries. Some form of universal basic income or negative taxation is in place everywhere, and most people think the economic singularity has happened.

Radical abundance

Nobody hates their job. People only do work that they enjoy, and most people could not find paid jobs even if they wanted to. Everyone receives a basic income from the state or a non-centralised public organisation, and there is no stigma attached to being unemployed, or partially employed. In most countries the UBI was funded initially by taxes levied on the minority of wealthy people who own most of the productive capital in the economy, and in particular on those who own the AI infrastructure.

In many countries, some of these elites have agreed to transfer the productive assets into communal ownership, either controlled by the state or by decentralised networks operated using blockchain technology (see chapter 6). Those who do this enjoy the sort of popularity previously reserved for film and sports stars.

Some countries mandated these transfers early on by effectively nationalising the assets within their legislative reach, but most retreated from this approach when they realised that their economies were stagnating, as many of their most innovative and energetic people emigrated. Worldwide, the idea is gaining ground that private ownership of key productive assets is distasteful. Most people do not see it as morally wrong, and don't want it to be made illegal, but it is often likened to smoking in the presence of non-smokers. This applies particularly to the ownership of facilities which manufacture basic human needs, like food and clothing, and to the ownership of organisations which develop the most essential technology – the technology which adds most of the value in every industry sector: artificial intelligence.

The gap in income and wealth between rich and poor countries has closed dramatically. This happened in part thanks to a substantial transfer of assets from the West to the rest, but mostly thanks to the adoption of effective economic policies, the eradication of corruption, and the benign impact of technology in the poorer countries.

AIs and robots produce most of the goods and services people need in great abundance and at very low cost. Many products are 3D printed close to where they are consumed, so demand for commercial transportation services has plummeted.

Demand for consumer travel is falling too, as immersive VR provides a close approximation of the experience. This takes the edge off the disappointment many people thought they would feel about not being able to increase their income by working harder in order to buy more luxury goods. Most people accept that they will never own a beach-front property with palm trees

shading the white sand, but they can spend as much time there as they like in convincing VR.

Another concern which has been allayed is that life without work would deprive the majority of people of a sense of meaning in their lives. Just as amateur artists were always happy to paint despite knowing that they could never equal the output of an old master like Vermeer, so people now are happy to play sport, write books, give lectures and design buildings in the knowledge that an AI could do any of those things better than them.

Not everyone is at ease in this brave new world, however. Around 10% of the population in most countries suffers from a profound sense of frustration and loss, and either succumbs to drugs or indulges almost permanently in escapist VR entertainment. A wide range of experiments is under way around the world, finding ways to help these people join their friends and families in less destructive or limiting lifestyles.

Governments and voters in a few countries resisted the economic singularity, seeing it as a de-humanising surrender to machine rule. Although they found economically viable alternatives at first, their citizens' standard of living quickly fell far behind. Several of these governments have now collapsed like the communist regimes of Eastern Europe in the early 1990s, and the others look set to follow – hopefully without violence.

5. - The Challenges

The point of this book so far has been to persuade you that within a few decades, it is likely that many people will be rendered unemployable by machine intelligence. If I have not wholly succeeded in that aim, then I hope you are at least prepared to accept that the possibility is serious enough that we should be thinking about the implications, and what to do about it if it happens.

If I haven't even got you that far, then you're probably about to put this book down. If so, don't throw it away – you might want to come back to it when self-driving vehicles start to make serious impacts on the employment data.

If I have made the case successfully – or if you were persuaded before we started – then welcome to the next stage of the journey. At the end of chapter 4 we saw a rosy scenario in which we are well on the way towards a new type of economy, and the transition has been smooth.

Sadly, life is rarely smooth. There will be challenges. I anticipate five, and I think two of them might cause real harm if we are careless or unlucky. They are: economic contraction, distribution, meaning, allocation, and cohesion. Let's take a look at each of them in turn.

5.1 – Economic contraction

American union boss Walter Reuther recounts a story about a visit he made in the 1950s to a Ford manufacturing plant, where he saw an impressive array of robots assembling cars. The Ford executive who was showing him round asked how Reuther thought he would get the robots to pay union membership fees. Reuther replied that the bigger question was how the robots would buy cars. (The story is usually told with Henry Ford II playing the role of the company executive but it almost certainly wasn't.[ccxciii])

The basic economic problem which this story is supposed to illustrate is that if nobody is earning any money then nobody can buy anything, and even those who do have money and resources can't sell anything. The economy grinds to a halt and everybody starves.

Of course life is never as black-and-white as that. Economies don't go overnight from functioning tolerably well to complete collapse. Even catastrophic decline is less like falling off a cliff and more like tumbling down a slope, with pauses along the way as you hit ledges. But obviously, severe economic contraction is grim, and to be avoided if at all possible.

If, as I have argued, machine intelligence renders more and more people unemployable, then other things being equal, the purchasing power previously exercised by those people will dry up. Their productive output will not be lost - it will just be provided by machines instead of humans. As demand falls but supply remains stable, prices will fall. At first, the falling prices may not be too much of a problem for firms and their owners, as the machines will be more efficient than the humans they replaced, and increasingly so, as they continue to improve at an exponential rate. But as more and more people become unemployed, the consequent fall in demand will overtake the price reductions enabled by greater efficiency. Economic contraction is pretty much inevitable, and it will get so serious that something will have to be done.

But before policy makers are forced to take action to tackle economic contraction, they will be faced by a much more serious problem: what to do about all those people who no longer have a source of income? This is the distribution problem, which is seen by many as the most severe problem raised by the economic singularity. Tackling it successfully will also solve the problem of economic contraction, so we can move right along.

5.2 – Distribution

At the height of the Great Depression in the early 1930s, unemployment reached 25% of the working-age population.[ccxciv] Social security arrangements were primitive then, and developed societies were much poorer than they are today, so that level of joblessness was much harder on people than it is today, when parts of Europe have returned to similar levels overall,[ccxcv] with youth unemployment hitting 50% in some places.[ccxcvi]

The worst levels of unemployment in developed countries today are found in Mediterranean countries like Greece and Spain, where family networks remain strong enough that sons and daughters can be supported for months or even years by fathers and mothers – and vice versa. There are escape valves, too, for the social pressure created by the situation. Economies further north are struggling less, and can absorb the energies and ambitions of many of the unemployed young people from the south.

When self-driving vehicles and other forms of automation render people of all classes unemployed right across the developed world, these safety nets will no longer be available. Articulate, well-connected and forceful middle class professionals will be standing alongside professional drivers and factory workers, demanding that the state do something to protect them and their families.

Universal Basic Income

If and when societies reach the point where we have to admit that a significant proportion of the population will never work again – through no fault of their own - a mechanism will have to be found to keep those people alive. And not just scraping by on the poverty line: they will have to be provided with an income which allows at least the possibility of a decent life by the standards of the societies they live in.

The answer is well-known, and fairly obvious: a universal basic income (UBI), available to all without condition; a living wage which is paid to all citizens simply because they are citizens.

Probably the longest-standing organisation advocating UBI is the Basic Income Earth Network. BIEN was formed as long ago as 1986, and “Earth” replaced “European” in its name in 2004. BIEN defines UBI as “an income unconditionally granted to all on an individual basis, without means test or work requirement." UBI has also been called unconditional basic income, basic income, basic income guarantee (BIG), guaranteed annual income, and citizen's income.

Proponents have argued for various levels of UBI, but in general they choose a level at or around the poverty level in the country of operation. This is partly because they don't think any more would be affordable, or politically acceptable, and partly to ward off criticisms that UBI would make people lazy and unproductive. As mentioned above, this will not be good enough if and when machine intelligence renders most people unable to work for a living. A modern developed society is not sustainable if a majority of its citizens are on the bread line.

UBI is similar to but distinct from the concept of negative income tax (NIT), under which people earning less than a specified amount receive payments. The two systems can be set up to produce the same financial results, but they appeal to different economic and political instincts. UBI involves payments to people who really don't need them, while NIT could stigmatise recipients.

The benefits claimed for UBI address issues which concern both the political left and right. Left-wing proponents see it as a mechanism to eradicate poverty and redress what they view as growing inequality within societies. They sometimes argue that it tackles the alleged gender pay gap, and redistributes income away from capital and towards labour. It has also been held out as a partial solution to the alleged generational theft whereby relatively wealthy pensioners are receiving income generated by taxes on young workers who have no assets, and who may not themselves receive similar benefits in later life because the welfare system looks increasingly unaffordable.[ccxcvii]

Right-wing advocates see UBI as a way to remove swathes of government bureaucracy: abolishing means testing removes the need for the battalions of civil servants who devise and implement it. There would be no incentive for people to game the benefits system, thus reducing government-generated waste and unfairness. They hope it would facilitate a wholesale simplification of tax structures, and perhaps enable a move to a flat tax. And they argue that more lower-income people would go to work because they would no longer be caught in benefit traps which penalise them for raising their income slightly. This would mean fewer children raised in families where nobody works, a particular bugbear of the right.[ccxcviii]

Most current supporters of UBI are on the left, but it has had support from prominent right-wing politicians and economists in the past, notably President Richard Nixon and economists Friedrich Hayek and Milton Friedman.

Experiments

There have been a surprising number of experiments with UBI: the Basic Income page on Reddit lists 25,[ccxcix] and gives potted descriptions of the purpose and outcomes of six of them.[ccc] All the researchers involved reported excellent results, with the subjects experiencing healthier, happier lives, and not collapsing into lazy lifestyles or squandering the money on alcohol or other drugs. Given that, it is curious that none of the experiments have been extended or made permanent.

The declared purpose of many UBI experiments is to investigate the concern that when people receive money for nothing, they stop working. One of the biggest experiments conducted so far, involving all 10,000 people in the small town of Dauphin in Manitoba, Canada, found that the only two social groups which did stop working were teenagers and young mothers, and this was seen as a positive outcome.[ccci]

Of course, people handing in their notice will be of no concern when machines have stolen all our jobs, but a more subtle version of the concern remains: do people in receipt of money for nothing stop doing anything of value? Do they become indolent couch potatoes, watching TV all day long, or collapse into reliance on alcohol and other drugs? Bearing in mind the distinction we made earlier between jobs and work, in a world where intelligent machines have automated most economic activity, the question is not, do people give up jobs, but do they give up work?

Unfortunately, none of the UBI experiments carried out so far constitute a rigorous test. A rigorous test would be universal, randomised, long-term, and basic – in the sense that the income distributed should be enough to live on.[cccii] And so more tests are planned.

In fact, a number of significant UBI experiments are planned or under way at the time of writing. One, in Finland, caused great excitement when it was announced in 2015, but it remains unclear

how broad-based the experiment will be, and what level of income will be paid. The aims are clear, however, and they relate to the right-wing concerns listed above. The Finnish researcher in charge of designing it, Olli Kangas, explains that the UBI experiment is hoping to demonstrate solutions for three problems with the current Finnish benefit system. First, people working part-time (perhaps in the gig economy) receive neither work-based benefits nor unemployment benefits. Second, some people are caught in a benefits trap whereby as their income increases their benefits decrease, which removes their incentive to work more and contribute more to the economy. Third, the existing benefits system is expensive, requiring too many bureaucrats to administer it.

The sample of Finns who are chosen to receive the UBI will be compared with a control sample who are not. Kangas will be exploring their propensity to continue working, their reported happiness and well-being, and any changes in their use of health and social services. He hopes to recruit a substantial sample – perhaps 100,000, which will enable him to detect variations between people of different ages, locations, demographics, and employment histories.[ccciii]

Another interesting UBI experiment is a crowd-funded initiative in Germany, which was launched by Berlin-based entrepreneur Michael Bohnmeyer in 2014.[ccciv] By December 2015, 26 people had been selected by lottery to receive €1,000 (around $1,000) a month, paid for by public donations. Most of the recipients reported that it didn't change their lives enormously, but they felt less stressed, and in many cases were able to embark on creative projects.

There is no shortage of places keen to experiment with UBI. The Dutch cities of Utrecht, Groningen, Wageningen and Tilburg are asking their national government for permission to carry out trials, and a referendum is expected during 2016 in Switzerland. All these initiatives are looking for ways to tackle problems with existing social welfare systems.

We have to go to Silicon Valley to find an experiment specifically designed to explore the impact of UBI in the context of a jobless future when machine intelligence has automated most of what we currently do for a living. Just such an experiment was announced in January 2016 by Sam Altman, president of the seed capital firm Y Combinator, which gave a start in life to Reddit, AirBnB and DropBox. Altman's task is not trivial: he will have to figure out a way to quantify the satisfaction his guinea pigs derive from their UBI, and whether they are doing anything useful with their time.[cccv]

Socialism?

With all these experiments bubbling up, the concept of UBI has become a favourite media topic, but it is controversial. Many opponents – especially in the US - see it as a form of socialism, and the US has traditionally harboured a visceral dislike of socialism. (The strong performance of Senator Bernie Sanders, a self-proclaimed democratic socialist, in the race to become the Democratic Party's candidate in the 2016 Presidential election is a striking departure from this norm of US politics.)

This concern is what seemed to leave Martin Ford somewhat dispirited at the end of his book “The Rise of the Robots”. As we saw in chapter 3, he fears that “guaranteed income is likely to be disparaged as 'socialism'”, and introducing it will be a “staggering challenge”. He is not alone: I have heard similar concerns from a number of thoughtful American friends.

I hope and believe that their fears are over-done. America is, of course, huge – more a continent than a country – so generalisations about it are dangerous. But I do not believe its people are in general un-thinking or malicious. If and when it becomes impossible to deny that a majority (or even a large minority) of its citizens will never do paid work again, and for no fault of their own, I do not believe that the rest will allow them to starve.

The dramatic recent changes in American attitudes towards homosexuality and drugs show how fast opinions there can change, and how far. As recently as 1962, homosexual acts were illegal in every US state, and it was only in 2003 that the federal Supreme Court decision in the *Lawrence v Texas* case invalidated the ban in the last 14 states where it remained unlawful. (Even today, more than a dozen states have yet to repeal or amend their own legislation to reflect this ruling.[cccvi]) And yet in June 2015, the federal Supreme Court ruled that bans on same-sex marriage are unconstitutional, in the case of *Obergefell v Hodges*. According to a Wall Street Journal poll, public support for gay marriage has doubled in the last decade, standing now at 60%.[cccvii]

Attitudes towards the legalisation of cannabis have also undergone a rapid sea change. For years, governments proclaimed a war on drugs, but that policy has clearly failed. Billions of dollars have been spent, and countless lives have been lost, but supply has not been constrained, much less eliminated. Parts of Mexico and other countries where the drugs are grown or routed have become war zones, and hugely powerful criminal organisations have been spawned. Attempts to curb demand have also failed, with tens of thousands of people being criminalised for an activity that harmed no-one. Drugs are dangerous, and their supply should be regulated, but ceding control over that supply to criminal gangs has not proved an enlightened policy. Public opinion in America is swinging rapidly towards that position. In 1969 only 16% of voters polled by Gallup supported legalisation, but now a majority takes that view.[cccviii] Possession of cannabis for personal use is now legal in four states, with the federal government agreeing not to interfere.[cccix]

It is not only America which is experiencing revolutionary changes in social attitudes. Up until 1997, sex before marriage was illegal in China, condemned as “hooliganism”. Nevertheless, a researcher found in 1989 that 15% of citizens had experienced it. The percentage had risen above 70% by 2014. Homosexuality was illegal until 2001 and gay marriage is still not legal. But in 2011, state-owned media began writing positive articles about gay pride marches in Shanghai and elsewhere.[cccx]

These examples show that entrenched societal opinions can and do change, sometimes quickly. If and when machine intelligence renders many of us permanently unemployable, it seems reasonable to expect that opposition to some form of universal basic income will evaporate.

Inflationary?

Opponents of UBI also worry that it will stoke inflation. Other things being equal, a massive injection of money into an economy is liable to raise prices, leading to sudden inflation and perhaps even hyper-inflation. But as campaigner Scott Santens points out, UBI does not necessarily mean an injection of fresh cash into the economy. It would most likely be paid for by increased taxation of the better-off, and by replacing the existing benefits system, together with the bureaucracy which implements it.[cccxi] He also claims that where basic incomes have been introduced, as in Alaska in 1982 and Kuwait in 2011, inflation actually fell.

Unaffordable?

A related objection to UBI which may have more substance is that it is unaffordable. Some argue that it can be funded by raising taxes on the small minority who have become extraordinarily wealthy in recent years. After all, even some of those wealthy people themselves (like Bill Gates and Warren Buffet) have confessed to feeling under-taxed. But experience shows that this can be a losing game. Very wealthy people do sometimes decide to dedicate much of their wealth to charitable causes. Bill Gates (again) and Mark Zuckerberg are obvious examples, and even some of the robber barons of the late 19th century gave fabulous sums to charitable foundations. One of the most successful of those barons was the Scottish-American steel magnate Andrew Carnegie, who endowed some 3,000 municipal libraries, and provided funding for several universities and numerous other organisations before he died. His most famous motto was that "the man who dies rich dies disgraced."[cccxii] But these people generally want to determine for themselves how their wealth is deployed, not least because they believe that they will make better use of it than politicians and bureaucrats.

So even the most generously disposed wealthy people often resist the wholesale appropriation of their assets in the form of taxation. And as demonstrated by the Panama papers scandal that erupted in April 2016, they are well equipped to do so, either by hiring clever lawyers and accountants to find loopholes and dodges, or by shifting themselves and their assets to less demanding jurisdictions.

Furthermore, entrepreneurs and other capable commercial people who are not yet extremely wealthy but aspire to become so may decide to move out of a jurisdiction which raises taxes sharply to pay for UBI. Or if they stay, they may become discouraged and decide against taking the necessary risks and dedicating the necessary time and energy to projects which could achieve their ambition. These people are responsible for much of the dynamism in capitalist countries, and dampening their enthusiasm or incentivising them to move elsewhere can be very damaging to an economy.

This sounds like common sense, but is in fact highly contentious. The political left believes that inequality is a social evil, and argues that taxing the rich does not deter economic activity.[cccxiii] The political right believes that a modicum of inequality is no bad thing, and is anyway inevitable in a thriving economy. It argues that increasing taxes on the rich does deter economic activity, and may actually result in lower government revenues, as the rich look harder for ways to reduce their tax burden.[cccxiv]

The Laffer Curve

Unfortunately, the data is muddy, which enables both sides to marshal apparently convincing arguments. And as is so often the case, the truth lies somewhere between them. We do know that there is a level of taxation beyond which further increases are ineffective, or even self-defeating. The Laffer Curve plots tax rates against the revenue they raise. At 100%, no-one would work, so that is an inefficient rate; 99% would not be much better. Sadly we just don't know for sure what the optimal level is, either in general, or in a specific country at a specific time.[cccxv]

In the UK, the Labour government in 2010 introduced a top rate of tax of 50% for people earning above £150,000. The Conservative government took it down to 45% in 2013, and claimed the result was a sharp revenue increase. The Labour party, of course, claimed the opposite.[cccxvi]

Which side you choose in this debate will be determined largely by your political orientation. Personally, I believe that competing organisations in well-regulated markets are more efficient and effective than monopolistic governments, and I believe that lower tax regimes encourage entrepreneurship. I also think that governments tend to tax their subjects as much as they think they can get away with, which explains why so much of their tax take is achieved through subtle, indirect, and often downright stealthy taxes. Thus a substantial tax increase to fund UBI is likely to be economically damaging.

Fortunately, this debate is rather tangential to my main argument about UBI, which is that we will need to implement it if and when a large minority of people have become unemployable. So if you are on the political left, I don't need to lure you across the parliamentary floor, which is probably a relief to us both.

Before we leave the question of UBI's affordability, we should consider the claim that it can be funded by abolishing existing social benefit arrangements.

Let's kill all the bureaucrats

Channelling Shakespeare,[cccxvii] UBI advocates claim that the massive cost of UBI could be offset by abolishing much or all of the existing benefits systems, along with the legions of bureaucrats who implement them. They offer an enticing vision of a world without means-testing, with no poverty traps, no steely-eyed "advisers" in job centres forcing claimants to apply for unsuitable work, no benefit fraud and no need to game the system.

Unfortunately the world probably won't allow such a nice, tidy outcome. People's needs vary according to their capabilities, their life stage, and their location, among other things. Someone who is disabled might well suffer greatly if their income was equal to that of an able-bodied person in robust health. A single dad with a child may need extra support. People living in London or San Francisco would certainly need more housing benefit than people living in Albuquerque or Auchtermuchty. Having ushered all the bureaucrats out the door thanks to the purifying simplicity of UBI, we would have to apologise and call them right back in again.

The RSA, a British think tank, published a report about UBI in December 2015 which was the result of a year's research and discussions.[cccxviii] It proposed abolishing much of the UK's existing benefits system, and replacing it with a payment of £3,692 for everyone between 25 and 65. This is £307 a month, £71 a week, or £10 a day. The payment amounts to a modest 14% of the average UK wage, which was £26,500 in 2015.[cccxix] People aged between 5 and 25 would receive £2,925, and pensioners would get £7,420. Extra payments would be made for young children.

The RSA estimated the total cost of its proposed system at £280bn, including running costs of £3bn. It claimed that this would be offset by £272bn saved by abolishing most of the existing benefits and pensions infrastructure, including personal income tax allowances and tax relief on pensions payments for higher rate tax payers.

The RSA claimed that families with children and on low wages would be £2,000 to £8,000 better off per year because of the removal of benefit traps. Adjustments required to prevent poorer people being worse off would take the cost to between £10-16bn, around 1% of GDP. This would be funded by taxes on high earners, a group which would also lose income from the changes.

Arguably, the RSA scheme is not a fully-fledged UBI proposal, as payments would taper off for incomes above £75,000, and stop altogether at £100,000. The level of payments are also set at a level which would keep people alive, but would not provide a decent standard of living.

The system is perhaps more an attempt to simplify and streamline the UK's messy and byzantine benefits system. It is also significant that the proposal ignores payments for housing and disability, which are of course substantial, and would require the recall of at least some of those bureaucrats.

Countries are not isolated economic ecosystems. Introducing UBI would significantly affect the competitive position of a country which introduced it, and would have other unintended consequences. A broadly positive article about UBI in the right-of-centre *Daily Telegraph* newspaper speculated that if and when Finland proceeds with its UBI experiment, it will be inundated by economic migrants unless it leaves the EU.[cccxx]

Assets

This review of distribution has focused on income as opposed to wealth. Most people have little wealth, and are therefore dependent on income. A poll published in January 2015 by a US personal finance website[cccxxi] echoed the finding a year earlier by the Federal Reserve[cccxxii] that two-thirds of Americans had savings equal to less than three months income. Half them could not cover an emergency expense of $400 without going into debt. This was aggravated by the recession which began in 2007: the average American family's net worth fell from $136,000 in 2007 to $81,000 in 2013.

Wealth inequality is far more extreme in today's world than income inequality, both globally and within individual nations. It is also less significant. The charity Oxfam created a stir in January 2016 by claiming that the richest 62 people own as much as the poorest 50% of the world.[cccxxiii] The figure may or may not be correct, but it tells us less than it appears to. A young professional in New York living a life of luxury and excess may have no net assets, but it would be perverse to describe her as poor. Furthermore, if the richest billionaires gave their wealth to the poorest half of the world, it would amount to a one-off payment of few hundred dollars each.[cccxxiv]

Nevertheless, if you are one of the lucky minority with substantial net assets, you might be wondering how you will be affected if and when technological unemployment takes hold. Will your house be worth more or less in the new economy? How about your vintage Aston Martin, or your collection of fine wines? Until and unless we move to a completely different kind of economy, it is likely that some of the wealthy people – especially those who control the artificial intelligence which creates most of the added value – will remain wealthy, and perhaps become even more wealthy. Perhaps the prices for Stradivarius violins and prime real estate will continue to rise – for some time at least.

What about the holdings of the much larger number of people in the middle – people who have net assets of a few tens or hundreds of thousands of dollars, perhaps up to as much as a million or two? Unless we switch quickly and smoothly to UBI, it seems likely that the price of assets typically owned by these middle class people, such as suburban houses and mass-produced cars, will slide as their owners try to replace lost income by liquidating their property. This could happen quickly, as people look ahead, see what is coming, and decide to cash in before the slide starts in earnest. Asset prices are notoriously hard to predict because they depend on events which cannot be foreseen, and also upon perceptions about what may happen, and perceptions

about those perceptions. This is another good reason why we should be thinking seriously about these matters sooner rather than later.[cccxxv]

Summary: UBI, but not yet

My conclusion about UBI will probably be unpopular with many, especially those on the political left. It is that the time is not yet ready for a full-fledged UBI, in the sense of a payment made to all citizens with no questions asked, and at a level which affords an acceptable standard of living in the context of the jurisdiction. UBI is a system which requires an economy of abundance, not an economy of scarcity, which is what we still have today. The appropriate system for economies of scarcity is the market (with regulations), because it allocates resources according to what people actually want, not according to what a politician or technocrat thinks they ought to want.

Any government tempted to experiment with a strong version of UBI must carefully consider the possible damage to its country's economic competitiveness, and other international impacts.

FT journalist Tim Harford summed it up well in May 2016, saying that in current circumstances, UBI appeals to three kinds of people: those happy to see the needy receive less income, those happy to see the state balloon (and risk massive capital flight), and those who can't add up.[cccxxvi]

But if and when machines have permanently automated most jobs, we will need to implement some form of UBI. The fears about it being politically unacceptable will probably prove exaggerated, with attitudes changing as the circumstances change.

That being so, if and when it becomes clear that most people are going to be rendered unemployable, it could be helpful to implement a modest form of UBI - something like the RSA scheme mentioned above, perhaps - in order to be able to move quickly to a full-blooded version when automation bites deep.

But UBI will not alone be sufficient to enable us to cope with the end of jobs. The other big problem we will have to tackle is cohesion. We will address that later on in this chapter, but first we should review the alleged problem of how people find meaning in a world without jobs.

5.3 – Meaning

The meaning of life

… is 42, of course.[cccxxvii]

OK, now we've got that out of the way, would you agree with the statement that people's lives need to have meaning in order for them to feel fulfilled, satisfied, and happy? It's certainly true for me, and I'm pretty sure it's true for most of the people I know. It is probably also true of you, or you wouldn't be reading this book.

I have met people who claimed to be pure hedonists – interested only in immediate pleasure. Some of them may even have been telling the truth. But most of us get bored if we feel our lives have no meaning. And not just bored in the sense that you get bored in a queue at a supermarket checkout, but profoundly restless and frustrated. To avoid this feeling we make deep emotional investments in ideas and institutions like family, friendships, work, loyalty to tribes, nations and causes. Deprived of these things, we feel lost and alienated.

Perhaps the most famous quote attributed to the 4th century BC Greek philosopher Socrates is that the unexamined life is not worth living. It is a remarkably strong statement. Why not just say that an unexamined life – a life without philosophy, in other words – is less good than an examined one? Is an unexamined life really worse than death? He made the statement at his trial, when the outcome was a choice between exile and suicide (he chose suicide), so perhaps he was under stress and being hyperbolic. But the claim is usually taken at face value, and perhaps he meant it literally.

It is also an elitist statement. Many people are too preoccupied with making a living, raising a family, escaping drug addiction or whatever immediate challenge they face to indulge in the luxury of philosophical discourse. Are their lives not worth living? You could argue that Socrates and his fellow ancient Athenians had slaves to take care of the menial stuff, but we have labour-saving devices instead, so that's no excuse.

Of course the question of what constitutes a good life, a worthwhile life, a life with meaning is a vexed one, with no simple answers, and probably no single answer. The philosopher John Danaher distinguishes between subjective accounts, which involve *feeling* worthwhile, and objective accounts, which involve helping to *make or do* something worthwhile.[cccxxviii]

Despite not knowing (or at least not agreeing) what a meaningful life is, and despite not spending all that much time in the average day thinking about it, most of us believe we need it. And many of us find it in work. So it's going to be a problem if we stop working.

Or is it?

Meaning and work

Simon Sinek has made a name for himself with books that propound a simple but important truth: if you have a clear purpose which inspires others, you can achieve great things. His best-known saying is "Working hard for something we don't care about is called stress; working hard for something we love is called passion."

You could be forgiven for thinking that a law was passed a few years ago in the US requiring business leaders – and people who want to be business leaders – to talk about their passion for their business. But most people don't feel passionate about their work, even if they pretend they do. In fact, many people are positively alienated by their jobs. They find them meaningless and boring.

Yet even these people usually define themselves by what they do for a living. If you ask someone at a party what they do they are likely to reply that they are an accountant, a taxi driver or an electrician. They are less likely to say that they are the coach of their child's football club, or a cinema-goer, or a reader. No doubt this is partly due to the amount of time that our jobs absorb - but then again we don't define ourselves as sleepers. It also has to do with work being the activity that provides our income, which is why home-based parents often feel sheepish about naming that as their work. (They shouldn't: it is some of the hardest but most rewarding work I've ever done!)

So work helps define us, and it gives many of us purpose. It even gives some of us meaning. So how damaging would it be if we lost it? Unemployed people often struggle with depression, but they are experiencing it in the context of a society where it seems that everyone else has a job. They are also on a lower income than the employed people around them. How bad would it be if everyone else was also unemployed, and receiving a decent income?

Fortunately, there are a couple of places we can look for an answer to that question.

The rich and the old

The agricultural revolution, around 12,000 years ago, created sustainable surpluses of food and other basic resources. This enabled a class of people to stop doing the work that pretty much all humans had done since our arrival on the planet, which was foraging and hunting for food. They became tribal leaders, kings, warriors, priests, traders and so on. Sometimes they spent as much time on these activities as the people who continued to forage and hunt, but sometimes they took time off – deliberately or by happenstance – and engaged in lives of leisure.

In Europe these people became known as aristocrats, from a Greek word meaning the best – originally in a military sense and then a political one. Some aristocrats did jobs: they ran agricultural concerns, they got involved in politics, and in some countries they ran empires. Occasionally they became men (and more rarely, women) of science. Famously, they disdained trade and commerce, regarding those activities as the preserve of the class below them, the middle class.

Many aristocrats did not work – including almost all the female ones. They led lives of leisure. As young men (and in a few cases, young women) they toured classical Greek and Roman sites in the Mediterranean countries. Returning home, they mostly socialised. Their lives revolved around balls, hunts, and visits to their local peers, interspersed with the glamour and tragedy of war, if that was their inclination. This lifestyle was chronicled in the novel, an art form which first acquired its current realistic form in the early 18th century.[cccxxix]

The lives depicted by Jane Austen and her contemporaries may seem tame to modern readers, who have experienced international travel and expect simultaneous global communications. But they were agreeable lives compared to what their poorer contemporaries had to put up with.

Addictions to gambling and drink were a hazard, and of course a minority of this pampered class destroyed themselves and their families with these vices. But this was unusual, and by and large most 18th- and 19th-century European aristocrats seem to have passed their lives without great concern about their lack of meaning. Whether these lives were worthwhile or not, whether or not they had meaning, is probably not for us to judge, but there is no evidence of widespread existential angst among the nobility.

In fact, it is these privileged people who made most of the advances in human thought and art in previous centuries, precisely because they did not need to work for a living, or eke out an existence as subsistence farmers. If they did not produce the memorable work themselves, they often sponsored it by employing talented artisans. So it seems there is much to be said for the ability to be idle.[cccxxx]

The other group we can look to for evidence about the effects of joblessness are retired people. The conventional wisdom used to be that growing old was an almost unmitigated disaster: "Old age ain't no place for sissies", as Bette Davies said,[cccxxxi] although it's obviously better than the only alternative currently available. But starting in the 1990s, researchers began questioning this perception, and found instead that the progress of happiness throughout life is U-shaped. We are at our happiest and most fulfilled when young, we become stressed and discontented in our prime and middle age, and we are happier and more relaxed again when older, despite the onset of physical disabilities and limitations.[cccxxxii] This pattern has been observed across a wide range of societies, and over a substantial period of time.

There are probably numerous causes of this effect, including the relinquishment of responsibility for children, and the acquisition of wisdom to accept what life has thrown at us. But the absence of jobs plays a major role in the lives of the retired. Even if it is not causing the up-tick in happiness, it is at least not preventing it.

Virtually happy

Thus far in human history we have had to find our meaning within the constraints of the three-dimensional world we live in, or in our imaginations. Technology is poised to open up a whole new space for us to explore together – the world of virtual reality. We don't yet know how we will react to this new universe, how we will behave in it, and what it will mean to us. We can be pretty confident that it will have a big impact.

"Diaspora", Greg Egan's novel of the far future, features an environment called the Truth Mines. It is a physical representation of mathematical theorems (albeit in virtual reality) which can seemingly be explored forever without exhausting all the discoveries that can be made. The ability to create virtual worlds that are so convincing to our brains that we almost lose the understanding that they are artificial may well allow us to expand enormously the space within which we find happiness and meaning.

In summary, loss of meaning does not seem likely to be one of the biggest problems that widespread technological unemployment will create.

5.4 – Allocation

The house on the beach

In a world where the majority of people cannot get jobs and are therefore paid a universal basic income, how will we allocate goods and services? At first sight, the answer seems simple and obvious: we will still have money, so we will still have the market. Supply and demand will continue to operate like before.

Let's assume we're at the point where a large majority of people get all their income from UBI. There would have to be adjustments for people looking after children or other dependents, disability and so on, but any extra income allocated for these needs would also be swallowed up by those needs. So in income terms, the society we are discussing would be an extremely egalitarian one.

But this society of egalitarian incomes will have inherited a decidedly un-egalitarian asset base. Houses are the most obvious and the most significant example: they were not all created equal. In a society where everyone's income is more-or-less the same, how will we decide who lives in the nice big house in the posh part of town, and who lives in the small flat with no sound-proofing in the grubby apartment block in the unfashionable suburb?

Will it be like a game of musical chairs? We all work hard to improve our lot, and then when the machines take our jobs, the music stops and we all sit down in the chairs we have arrived at. And just stay there forever. That seems neither fair nor sustainable.

With luck, we will be creating an economy of abundance, in which machines carry out maintenance and improvement works with great efficiency, and hence cheaply. In that case we will set them to revitalising and / or replacing the stock of lower-quality houses. (And cars, and boats, and furniture, and clothes, etc.) But it will take a very long time indeed to build a nice new house for everyone who doesn't start off with one. And even when we have completed that gargantuan task, some houses will still be in much nicer places than others.

This is the allocation problem.

Some scarcity can never be abolished. There is a finite and regrettably small supply of large houses on empty white sand beaches fringed with palm trees leading down to a turquoise sea. Or penthouse apartments on Manhattan's Fifth Avenue. There is a very small supply of Vermeers and Aston Martin DB5s. Do we decide that no-one can own these things? Perhaps we could turn all the nice houses into museums and keep the scarce movable objects on display there, to be visited (and perhaps used) on payment of a fee, or by scheduled appointment. In that case, who will decide what the cut-off point is between a house which people can carry on living in, and one which is too nice to be private property?

VR to the rescue?

At the time of writing, Palmer Luckey and John Carmack are hardly household names, but by the time this book is published they may well be. (In case they're not, they are the key executives of Oculus Rift, which looks set to be the first commercially-available VR equipment to offer a

convincingly immersive user experience.) They talk about a "moral imperative" to make virtual reality available to us all.[cccxxxiii]

Luckey puts it like this: "Everyone wants to have a happy life, but it's going to be impossible to give everyone everything they want. ... Virtual reality can make it so anyone, anywhere can have these experiences." Carmack continues: "you could imagine almost everyone in the world owning [good VR equipment]. … This means that some fraction of the desirable experiences of the wealthy can be synthesized and replicated for a much broader range of people."

Other people have thought about these questions, and not everyone is delighted by the suggestion that VR can assuage the frustration caused by scarcity. Some people think it impossible, and others think it possible but degrading.

The Harvard political philosopher Robert Nozick described a thought experiment back in 1974 featuring an "experience machine" which could recreate any sensation you choose. Your brain is persuaded that the experience is real, which means that you believe it too, but in fact your body is lying in in a flotation tank, deprived of all sensory input while your brain is hooked up to the machine. Philosophers do a lot of their work by investigating their intuitions, and Nozick's intuition was that no-one would use this machine because we value reality too highly. I find it surprising that he came to that conclusion back in 1974, and it would be an even more surprising conclusion to reach today, when so many people spend so much of their lives in simulated realities, albeit only imperfectly simulated. Certainly a great deal of money is being invested by smart people in the belief that we will consume VR avidly. Nozick died in 2002, so he won't have to find out for himself – maybe he would be relieved.

Other critics see the Oculus founders' view of the future as possible but frightening. Ethan Zuckerman is director of the MIT Centre for Civic Media, and thinks that "the idea that we can make gross economic inequalities less relevant by giving [poor people] virtual bread and circuses is diabolical and delusional." Jaron Lanier is a computer scientist and writer who founded VR pioneer VPL Research, and is generally credited with popularising the term virtual reality. He lambasts as "evil" the vision that the rich will become immortal, while "everyone else will get a simulated reality. … I'd prefer to see a world where everyone is a first-class citizen and we don't have people living in the Matrix."

Only time will tell if VR is helpful, or even necessary, in enabling us to live in a world where machines have made humans unemployable. My own guess is that it will play a major role in the lives of most people, and that it will make them more productive, more fun and more fulfilling. As Oculus' John Carmack puts it, "if people are having a virtually happy life, they are having a happy life. Period."

Nevertheless, Zuckerman and Lanier have identified an important problem with the vision. It is not to do with VR so much as the potential separation of our species into two or more divergent camps. We will review this in more detail in the next section.

Algocracy

Decisions about the allocation of resources are being made all the time in societies, on scales both large and small. As argued above, capitalism has proved so effective in raising the living standards of societies which have adopted it (paraphrasing Churchill again[cccxxxiv], it is the worst possible economic system except for all the others) because markets are highly efficient systems for allocating resources in economies characterised by scarcity.

Historically, markets have consisted of people. There may be lots of people on both sides of the transaction (flea markets are one example, eBay is another). Or there may be few buyers and many sellers (farmers selling to supermarket chains) or vice versa (supermarket chains selling to consumers). But typically, both buyers and sellers were humans. That is changing.

Algorithms, as we saw in chapter 3, are sets of rules or instructions for a computer to follow. A model is built and refined by testing with data. (Lots of data!) In situations with well-defined processes and target outcomes, algorithms can be extremely efficient decision-makers. We see this in our everyday lives. ATMs are operated by algorithms – no human responds to the instruction you type into the screen. They are efficient and effective, and most of us prefer to draw our cash from ATMs rather than queueing at a counter to have a human provide the service.

Algorithms now take many decisions which were formerly the responsibility of humans. They initiate and execute many of the trades on stock and commodity exchanges. They manage resources within organisations providing utilities like electricity, gas and water. They govern important parts of the supply chains which put food on supermarket shelves.

With the rapidly increasing volume of data flowing from tiny sensors embedded in buildings, machinery, vehicles and all kinds of products, and with the ever-improving performance of machine learning systems thanks to Moore's Law, algorithms are getting better and better at their jobs. They are getting better, cheaper and faster than humans. As well as taking our jobs they are taking our decisions.

In his 2006 book "Virtual Migration", Indian-American academic A. Aneesh coined the name "algocracy" for this phenomenon.[cccxxxv] The difficulty with it has been explored in detail by the philosopher John Danaher, who sets the problem up as follows. Legitimate governance requires transparent decision-making processes which allow for involvement by the people affected. Algorithms are often not transparent and their decision-making processes do not admit human participation. Therefore algorithmic decision-making should be resisted.[cccxxxvi]

Danaher thinks that algocracy poses a threat to democratic legitimacy, but does not think that it can be, or should be, resisted. He thinks there will be important costs to embracing algocracy and we need to decide whether we are comfortable with those costs.

Of course many of the decisions being delegated to algorithms are ones we would not want returned to human hands – partly because the machines make the decisions so much better, and partly because the intellectual activity involved is deathly boring. It is not particularly ennobling to be responsible for the decision whether to switch a city's street lights on at 6.20 or 6.30, but the decision could have a significant impact. The additional energy cost may or may not be offset by the improvement in road safety, and determining that equation could involve collating and analysing millions of data points. Much better work for a machine than a man, surely.

There are many decisions which machines could also make better than humans, but we might feel less comfortable having them do so. The allocation of new housing stock, the best date for an important election, the cost ceiling for a powerful new drug, for instance. Arguments are probably going to become increasingly commonly and increasingly vehement over which decisions should be made by machines, and which by humans.

5.5 – Cohesion

The scenario of "the Gods and the Useless"

As mentioned in chapter 1, at the end of his July 2015 TED talk,[cccxxxvii] the author of "Sapiens", Yuval Harari, makes a seemingly throw-away comment about humanity devolving into two classes: the gods and the useless. The audience laughs at this brutal assessment, but I suspect Harari is deadly serious.

Imagine a society where the great majority of people lead lives of leisure, their income provided by a beneficent state, or perhaps a gigantic charitable organisation. They are not rich, they don't travel first class or frequent expensive restaurants, and they don't own multiple houses. But they have no pressing needs and in fact they want for little: they enjoy socialising, learning, sports, exploration, and much of this is carried out in virtual worlds which are almost indistinguishable from reality.

A small minority of people in this society do have jobs. Their work is pleasurable and intellectually simulating, and not stressful. It involves monitoring and occasionally guiding or re-setting the performance of the machines which run their society – machines which they own.

Let's say that this elite minority is generous towards the majority which lives outside their gated communities, and which does not visit the luxurious resorts they migrate between, and does not travel with them on their private heli-jets. They are effectively benign rulers, although both camps refrain from putting it like that.

Now, in this future world, all members of the species of homo sapiens are changing. They are using new technologies to enhance themselves both cognitively and physically. They use smart drugs, exoskeletons and genetic technologies, among others. Maybe they have engineered themselves to need less sleep.[cccxxxviii]

Everyone has access to these technologies, but the elite has privileged access. They get them sooner, and this could be vitally important. I argued in chapter 3.7 that concerns about a "digital divide" were exaggerated. Companies make much more money by selling lots of relatively cheap cars and smartphones to almost everyone than they ever could by selling just a handful of diamond-encrusted versions to the super-rich.

But we mustn't forget that technology is advancing at an accelerating rate. In the future society we are envisioning, important breakthroughs in physical and cognitive enhancement are announced every year, then every month, then every week. As artificial intelligence gets better and better it fuels this improvement – even though it (the AI) is still narrow AI, and far from becoming human-level, artificial general intelligence, or AGI.[cccxxxix]

It may become hard or even impossible to disseminate these cognitive and physical improvements quickly enough to avoid a profound separation between those with privileged access to them and the rest of us.

So the elite will change faster than the rest. As the two groups lead largely separate lives, the widening gap may not be apparent to the majority, but the elite surely will know about it. They will decide that they must draw attention away from the fact, and they will take precautions to prevent attack, in case the majority should become aware of and resentful about what is happening. They will surround their gated communities with discreet machines which possess astonishingly powerful defensive – and offensive – capabilities. They will keep themselves more and more to themselves, meeting members of the majority less and less often. When they do

meet, it will almost always be in virtual reality, where their avatars (their representations in the VR environment) do not betray the widening gulf between the two types of humans.

After a while, humans will evolve into two different species: the gods and the useless.

Brave New World

This is not your average science fiction dystopia. The scenario most commonly offered up by Hollywood is that technology has pulled back the curtain which concealed the truly evil nature of the capitalism system, and mankind has fallen into a form of high-tech slavery, where the rich descendants of company CEOs and scheming tycoons and financiers brutalise an impoverished and oppressed majority. The movie "Elysium" is just one of many dreary examples of this tired old cliché. In it, as in so many others, society has actually regressed from capitalism into a sort of techno-feudalism. Any viewer who is half-awake is wondering, since machines can do all the work, what is the point of enslaving humans?

(It is curious that the left-leaning culture prevalent in Hollywood impels it to issue these tirades against capitalism, when Hollywood studios are themselves formidable exponents of the capitalist art. And it is curious that Hollywood stars who demand millions of dollars to make a single film complain that corporations are fuelled by nothing but greed.)

A far more interesting scenario is Aldous Huxley's "Brave New World", which is a very subtle piece of world-building. Almost everyone is content in the society which has been developed after an appalling military conflict, but it is clear to the reader that humanity has lost something important, and most of it has regressed to an almost infantile condition. Yet when a talented outsider arrives, he is unable to devise a way to improve the system, or to accommodate himself to it. He instinctively feels – and the reader is encouraged to agree – that his own life has shown him there is a better way to live, but he is unable to articulate or maintain it.

In Huxley's story, humanity has achieved a stable equilibrium. Only a tiny minority of humans are aware of what has been sacrificed in order to achieve a society of such docility and acquiescence. Regular sex and a powerful drug called soma are the opiates of the masses. (Huxley wrote the book in 1931, before the advent of rock and roll, so he couldn't anticipate Ian Dury's hedonistic mantra that sex and drugs and rock and roll were all his brain and body needed.[cccxl])

Brave New World is certainly not intended as a blueprint. Even so, it is calm, and the spectre of economic and social collapse seems to have been abolished - at least for the time being. It would be foolish of us to take even that much for granted. A society comprising gods and the useless might turn out to be inherently unstable. In a full-on conflict between them it seems likely that the gods would have the means to protect themselves, but at what cost?

Will capitalism remain fit for purpose?

Private property is an essential feature of capitalism, and in particular, the private ownership of the means of production, exchange and distribution. In market economies, most people earn their living by selling their labour – their time and their physical and intellectual skills.

People called entrepreneurs hire workers and combine their labour with the other major element of the capitalist economy – capital, which consists of money, machinery, land, buildings and intellectual property. The entrepreneurs use the labour and the capital to develop and make products and services which they hope will be bought by enough people to turn a profit. The capital is put at risk during this process, and the profit is a reward and an incentive for the entrepreneurs and the owners of the capital (known as capitalists) who took that risk.

The picture is complicated, because we are all capitalists today. Much of the capital deployed by entrepreneurs is owned by financial institutions like insurers and pension companies, and ownership of these institutions is very widely distributed through pension plans. Many people are also shareholders in the companies they work for, thanks to employee share ownership schemes.

In a capitalist economy where most people work, individuals who start off without any capital can acquire it by saving some of what they are paid for their labour, and by starting companies themselves. Countries where this is easy and celebrated tend to be better off than countries where entrepreneurship is discouraged, or hampered by corruption, over-regulation, or lack of infrastructure.

Similarly, people who start off with capital can lose it, by bad luck or poor judgement. It is not uncommon for families to go from rags to riches and back to rags within three generations.

There are countries today where this kind of social mobility (both upward and downward) is very limited, and the elites are entrenched. Arguably, nowhere in the world has sufficient economic social mobility, but in many countries it does exist to some degree.

As we have seen, after the economic singularity there will be two major differences. First, if the prediction that most people will be unemployable comes true, then there will be pretty much no traffic – no economic and social mobility - between the elite and the rest. This will be the case all over the world. If you can't do paid work, it is very hard to accumulate capital. Second, if the rate of technological progress continues to accelerate, the elite may avail themselves of the means of cognitive and physical enhancement to diverge from the majority, both physically and cognitively.

The obvious but difficult remedy for this is to end the institution of private property. The means of production, exchange and distribution would be placed into some kind of collective ownership to prevent the possibility of social and species fracture.

As we saw in chapter 3.1, this conclusion is rejected by the two most popular books published so far about technological unemployment. I share their inclination, and it makes me extremely uncomfortable. I was in business for 30 years before becoming a full-time writer and speaker, and I remain convinced of the largely positive effects of a regulated market economy with a welfare safety net.[cccxli] It seems to me that capitalism is the best economic system we know of for a society where humans do the work.

But I fear that capitalism may not be so well-suited for an economy of abundance, where machines do the work, where most people are unemployed, and where technology is changing the species quickly. I am not pushing this argument hard. Not only does it make me uncomfortable, but one of the fundamental characteristics of a singularity is that it is even harder than usual to predict the future when there is an event horizon in the way.

It is by no means certain that abandoning capitalism and private property is the only way to avoid fracture and collapse. It may be possible for all kinds of people to live in harmony in a society where a minority gets paid to work and owns most of the economy's assets, while everyone else lives happily on a universal basic income. It is not impossible to imagine a benign technocracy in which the majority of people really don't care who owns what, because they are wholly satisfied with their abundant supply of material and digital goods and services.

Unfortunately, every time I try to envisage this world, the picture degrades into a variation on the theme of "Brave New World" – or worse. Perhaps this is simply a failure of my imagination. I hope so.

If it is true that we need to move away from capitalism, we have two major jobs on our hands. First, we need to determine what that new economy should look like. Second, we need to work out how to transition from the economy we have to the economy we need. This will not be easy. Humans dislike change, and as always, there will be winners and losers. The losers may not take their losses calmly.

The scenario of the gods and the useless is not the only possible outcome of technological unemployment. The next chapter explores half a dozen of the most plausible scenarios. By assessing the likelihood and utility of each scenario, and understanding how to achieve or avoid them, we may be able to achieve the most positive outcome of the economic singularity.

6. - Scenarios

6.1 – No Change

In a July 2015 interview with Edge, an online magazine, Pulitzer Prize-winning veteran New York Times journalist John Markoff lamented the deceleration of technological progress - in fact he claimed that it has come to a halt.[cccxlii] He reported that Moore's Law stopped reducing the price of computer components in 2013, and pointed to the disappointing performance of the robots entered into the DARPA Robotics Challenge in June 2015 (which we reviewed in chapter 3.7).

He claimed that there has been no profound technological innovation since the invention of the smartphone in 2007, and complained that basic science research has essentially died, with no modern equivalent of Xerox's Palo Alto Research Centre (PARC), which was responsible for many of the fundamental features of computers which we take for granted today, like graphical user interfaces (GUIs) and indeed the PC.

Markoff grew up in Silicon Valley and began writing about the internet in the 1970s. He fears that the spirit of innovation and enterprise has gone out of the place, and bemoans the absence of technologists or entrepreneurs today with the stature of past greats like Doug Engelbart (inventor of the computer mouse and much more), Bill Gates and Steve Jobs. He argues that today's entrepreneurs are mere copycats, trying to peddle the next "Uber for X".

He admits that the pace of technological development might pick up again, perhaps thanks to research into meta-materials, whose structure absorbs, bends or enhances electromagnetic waves in exotic ways. He is dismissive of artificial intelligence because it has not yet produced a conscious mind, but he thinks that augmented reality might turn out to be a new platform for innovation, just as the smartphone did a decade ago. But in conclusion he believes that "2045 ... is going to look more like it looks today than you think."

It is tempting to think that Markoff was to some extent playing to the gallery, wallowing self-indulgently in sexagenarian nostalgia about the passing of old glories. His critique blithely ignores the arrival of social media and much else, and dismisses the basic research that goes on at Google X, DeepMind, the Human Brain Project and elsewhere.

Nevertheless, Markoff does articulate a fairly widespread point of view. Many people believe that the industrial revolution had a far greater impact on everyday life than anything produced by the information revolution. Before the arrival of railroads and then cars, most people never travelled outside their town or village, much less to a foreign country. Before the arrival of electricity and central heating, human activity was governed by the sun: even if you were privileged enough to be able to read, it was expensive and tedious to do so by candlelight, and everything slowed down during the cold of the winter months.

But it is facile to ignore the revolutions brought about by the information age. Television and the internet have shown us how people live all around the world, and thanks to Google and Wikipedia, etc., we now have something close to omniscience. We have machines which rival us in their ability to read, recognise images, and process natural language. And the thing to remember is that

the information revolution is very young. What is coming will make the industrial revolution, profound as it was, seem pale by comparison.

As we stressed at the start of chapter 4, the future is unknown, and all predictions are perilous. But the idea that the world will be largely unchanged three decades hence seems the least plausible of the scenarios set out in this chapter.

6.2 – Racing with the machines

Centaurs

We came across the idea of centaurs in chapter 3.10. The notion comes from Gary Kasparov, the chess grandmaster who was beaten by IBM's Deep Blue computer in 1997. He noticed that (for the time being at least) a human teamed up with a machine can beat the best chess-playing computer, and he called the combination a centaur. It has since become a metaphor for the hope that humans will not be rendered economically redundant by intelligent machines, but instead will work ever more closely with them, always bringing some special human magic to the combination.

Icebergs

As we saw in chapter 3.9, this hope that we can be centaurs gains some credence from the early experience with AI systems in what lawyers call discovery, or disclosure. After a few hours or days of training, deep learning systems are dramatically quicker and better than humans at ploughing through huge piles of documents, looking for particular pieces of information. It was feared that these systems would remove the need for junior lawyers, but instead it turned out that the machines brought a huge new class of work into the realm of possibility. Projects that would previously have been uneconomic were now feasible, and the junior lawyers are still needed to carry out the initial training. This has been called the iceberg phenomenon: it was thought the junior lawyers were standing on thin ice, but it turned out instead that they were standing on top of a massive bulk of newly available work. Their position began to look secure again.

A similar phenomenon is likely to be observed with medical diagnosis. Again in chapter 3.9, we saw that cheap attachments to smartphones will soon enable us all to be tested far more often and far more cheaply than at the moment. The AI resident on your smartphone – or servers it accesses in the cloud – will assess your blood pressure, blood glucose, your breath, your voice and many more indicators, and deliver instant verdicts in most cases. Initially at least, the upshot will be that we will have much better information about our medical condition, and we will still need doctors to carry out the more significant tests.

Creativity and caring

Sceptics about technological unemployment also argue that human creativity will remain in demand, as will human empathy, which they see as a pre-requisite in caring professions like nursing. Senior people in the technology companies which are developing AI often make this argument. Microsoft's principal researcher Jonathan Grudin, for instance, says that "technology will continue to disrupt jobs, but more jobs seem likely to be created. There is no shortage of things that need to be done and that will not change."[cccxliii] It is not surprising to hear these arguments from executives in businesses which are transforming themselves into AI companies: they would presumably feel very uncomfortable if they thought their work was hastening an economic crisis.

But while technology company executives sound breezy about the prospects for continued unemployment as machine intelligences get smarter, some of the academic authors who broadly agree with them sound more tentative. In chapter 3.1 we saw that in their book "The Second Machine Age", MIT professors Erik Brynjolfsson and Andrew McAfee believe that for many years to come, humans will be better than machines at generating new ideas, and complex forms of

communication. They think that capitalism should be defended and retained, but they sound less confident about what will happen in the medium term. They argue for an overhaul of the US education system, but they don't sound convinced that will be enough, and they speculate that a negative income tax may eventually become necessary.

Tyler Cowen, whom we encountered in chapter 3.3 as the author of "Average is Over", is certainly not breezy in his assessment of the outlook, nor is he tentative. He is confident that UBI will not be needed, and he does not expect riots. But his prognosis is lugubrious. He foresees 10-15% of the population being extremely wealthy, and the rest getting by on incomes which are stagnant at best, but putting up with it because many of them are too old to riot, and they are pacified by the excellent cheap entertainment that technology provides.

We addressed these arguments in chapter 3.10, and we need not go over them in detail again now. In summary, the danger is that the icebergs will stop growing after a while, and the machines will need less and less training to tackle each one. The 5m Americans who currently earn a living by driving are not all going to become computer programmers or custodians of the AIs, and it is hard to believe that we can all become professional artists, nurses and therapists.

6.3 – Capitalism + UBI

Martin Ford, author of not one but two books on technological unemployment, also believes that capitalism can survive the "rise of the robots", as he calls it in the title of his second book. He thinks it will be a struggle for many people, though, so he urges policy-makers to start thinking about introducing a Universal Basic Income (UBI) of around $10,000 a head when the time is right. He believes it will be extremely challenging to introduce this policy in the US, given the ingrained hostility to socialism there.

But as we saw in the last chapter, there are reasons to suppose that this challenge may evaporate if technological unemployment bites severely. And as we also saw in the last chapter, it may be replaced by challenges which are much harder to overcome.

6.4 – Fracture

This is the “gods and the useless” scenario which we discussed in chapter 5.5 about the challenge of cohesion.

6.5 – Collapse

Civilisation is fragile. Any schoolchild can name some great empires which collapsed: the Romans, the Maya, the Persians. The ancient Egyptians managed to rise and fall several times during their extraordinary 3,000-year history.

We also know how fragile civilisation is from two famous episodes in experimental psychology. In 1961, Yale psychologist Stanley Milgram recruited students from that elite university, and told them to administer mild electric shocks to incentivise a stranger who was supposed to be learning pairs of words. The shocks were fake, but the students did not know this, and an extraordinary two-thirds of the students were prepared, when urged on by the experimenter, to deliver what appeared to be very painful and damaging doses of electricity.[cccxliv] The experiment has been replicated numerous times around the world, with similar results.

Ten years later, Stanford psychology professor Philip Zimbardo, a school friend of Milgram's, ran a different experiment in which students were recruited and arbitrarily assigned the roles of prisoners and guards in a make-believe prison. He was shocked to see how enthusiastically sadistic the students who were chosen to be guards became, and he was obliged to terminate the exercise early.[cccxlv] This experiment has also been replicated numerous times.

Our 21st century global civilisation seems pretty robust. We have just gone through what is frequently described as the worst recession since the Great Depression of the 1930s, and for the great majority of people, the experience was nothing like as awful as those terrible years, which did so much to set up the disastrous carnage of World War Two.

But history and experimental psychology demonstrate that we cannot afford to be complacent. If the argument of this book is correct, we are about to embark on a journey towards a new type of economy which we have not yet designed. Unless we are careful, there will be plenty of opportunities for mis-steps, misunderstandings, and downright mischief by populists and demagogues.

If technological unemployment arrives in a rush, and we are not prepared, a lot of people will lose their incomes quickly, and governments may not move fast enough to avert drastic collapses in asset prices as people sell their belongings to make ends meet. If the introduction of UBI is slow or botched in some countries, the resulting economic crises could lead to their governments being overthrown by irresponsible or foolish leaders. We must hope this does not happen in any countries with significant stocks of nuclear weapons.

6.6 – Protopia

Utopia, Dystopia, Protopia

Kevin Kelly is a writer, and the founding editor of *Wired* magazine. He has been called the most interesting man in the world.[cccxlvi] I'm not sure whether he enjoys the burden of that appellation, but he does produce a lot of interesting ideas. One of his good ones is Protopia.

Too much of today's thinking about the future is dystopian, and that is partly because too many people fail to realise just how much progress homo sapiens has made in the last few centuries and decades. It is natural and indeed helpful for our species to be discontented: if we weren't discontented, we probably wouldn't struggle to make the world a better place. But it can lead to misunderstandings.

Many people think that all politicians are corrupt, and that all corporations are run by Bond villains who are greedy and bent on world domination. Most of us could think of some group, clique or tribe that we are suspicious, fearful or disdainful of. But the truth is that in most of the world, today is the best time there has ever been to be alive. Most people in developed countries today live better than kings and queens did a couple of centuries ago. We live longer, eat better, have better healthcare, and inconceivably better access to information and entertainment than previous generations. Of course everything could go horribly wrong tomorrow. There might even be an iron law of nature that when civilisations reach a certain stage they either blow themselves up, or create machines to do it for them. But from where we stand today, there is no reason to believe that. It seems more likely that the future is open, and potentially very good indeed.

Utopian visions of the future are less common, but also problematic. A future in which life has to all intents and purposes become perfect sounds sterile and boring. It is also highly improbable: the more we learn about the universe, the more we discover that we don't know, so it seems unlikely the universe will one day stop presenting us with puzzles and challenges. Perhaps this is why the two best-known literary descriptions of utopias, Thomas More's "Utopia" and Voltaire's "Candide", are essentially critiques of the societies they lived in rather than recipes for an ideal future one.

So it is refreshing to read Kelly saying this: "I am a protopian, not a utopian. I believe in progress in an incremental way where every year it's better than the year before but not by very much—just a micro amount. I don't believe in utopia where there's any kind of a world without problems brought on by technology. Every new technology creates almost as many problems that it solves." But crucially, it gives us "a choice that we did not have before, and that tips it very, very slightly in the category of the sum of good."[cccxlvii]

Collective ownership

What might a society look like which has passed successfully through the event horizon of the economic singularity? Machines are producing virtually all the goods and services that we would otherwise pay people to produce, and the society has avoided the twin traps of fracture and collapse. Everyone shares in the bounty that the machines are producing, and almost everyone has found meaning and joy in their lives without jobs. There is no enforcement of a rigid equality of personal outcomes across the lives of everyone in this society, but there is also no increasingly entrenched divergence between those with access to all the latest technologies and those without.

In chapter 5.5, we confronted the possibility that this society has felt obliged to abandon our powerful attachment to the concept of private property, and has moved to some form of collective ownership of the means of production, exchange and distribution. In other words, socialism.

Those of us who are convinced that the free market economy, suitably regulated, is an ingenious system that has demonstrably created the best standards of living that humans have ever enjoyed will find this hard to swallow. Certainly, the idea takes some getting used to.

Advocates of capitalism who have thought this far ahead have suggested that their favourite economic model is the best system for an economy of scarcity, but that perhaps it won't be appropriate for an economy of abundance. The problem highlighted by the scenario of "the gods and the useless" reveals a different problem. Scarcity hasn't disappeared: it has changed, and become more dangerous. The new scarcity is the privileged access to an accelerating flow of powerful new enhancement technologies. The danger is that the elite which enjoys this privileged access will rapidly become a separate species - a dominant species. Both history and common sense suggest that is unlikely to work out well for everyone else. The new scarce resource - the privileged access to the cascade of new technologies – is even more valuable and powerful than any scarce resource that we value today.

When normal people read about the lives of billionaires and movie stars, we often think they live in a different world. But the distance which separates them from us is tiny compared to the gulf which could open up between the AI-owning elite and the unemployed majority in a world which passed through the economic singularity while retaining private property.

In chapter 4, I described a timeline for a successful navigation of the economic singularity. It ended up with the AI-owning elites transferring their assets into collective ownership, and being hailed as heroes and heroines for doing so. How might this work in practice?

Decentralisation

Part of the genius of the market economy – the reason why it is so effective - is that decisions are taken by the people best qualified to take them. The market enables (indeed obliges) each of us to provide truthful signals about what we do and don't want, what we do and don't value. We buy this car and not that car because we prefer it (given our budgetary constraints) and there is no doubt that we are providing a correct signal because we are spending our own money. The decisions taken by the market overall about how many of each car to make are the aggregate of all these signals.

By contrast, when decisions are made in a centralised, planned economy, somebody is guessing about what is wanted and needed at every level below them. However good their data collection system, and however well-intentioned they are, they will always be out of date and they will often be just plain wrong. There is also a very good chance that corruption will set in, because it is so easy for that to happen. With apologies to Lord Acton,[cccxlviii] power corrupts, absolute power corrupts absolutely, and corruption is central to centralised planning.

Common ownership can work well in small communities, such as families, tribes, and small villages. But as soon as a society attains any level of size and sophistication, the bonds of kinship weaken and individuals start to claim ownership over land and property. The society becomes

regulated by power structures which begin as means of self-defence and evolve into expressions of ambition.

If (and it is a big "if") surviving the economic singularity and avoiding fracture means ending the system of private ownership, how can this be done without falling into the unwelcome embrace of an over-mighty state and centralised planning?

The answer just might be the blockchain.

Blockchain

People have gone mad trying to understand how the blockchain works, never mind trying to explain it. Its most famous application is Bitcoin, the world's first completely decentralized digital currency.[cccxlix] In just a few years, the Bitcoin "economy" has grown larger than the economies of some countries. The value of a Bitcoin has fluctuated wildly, hitting a peak of $1,216 in November 2013.

The insights which made Bitcoin possible were published in 2008 under the pseudonym Satoshi Nakamoto, and the blockchain is at the heart of it. The blockchain is a public ledger which records transactions. The clever bit is that the ledger is completely trustworthy despite having no central authority, like a bank, to validate it. It is trustworthy in that you can have full confidence that if someone gives you a Bitcoin, then you do own that Bitcoin: the person who gave it to you will not be nipping off to spend the same piece of currency elsewhere, even though it is entirely digital.

This confidence arises because transactions are recorded in blocks which are added to the chain by people (or rather computer algorithms) called miners. These miners are working continuously on mathematical problems whose solutions are hard to find but easy to verify. A problem is solved ("mined") roughly every ten minutes, and each solution creates a block. The new block is added to the chain, and incorporates the transactions made since the last block was added to the chain. Your transaction is published on the blockchain's network as soon as it is agreed, but it is only confirmed, and hence reliable, when a miner has incorporated it into a block.

Satoshi Nakamoto's innovation solved a previously intractable challenge in computer science known as the Byzantine General's Problem. Imagine a mediaeval city surrounded by a dozen armies, each led by a powerful general. If the armies mount a co-ordinated attack, their victory is assured, but they can only communicate by messengers on horseback who visit the generals one by one, and some of the generals are untrustworthy. The blockchain provides a way for each general to know that a message calling for an attack at a particular time is genuine, and has not been fabricated by a dishonest general before it reached him.[cccl]

Digital currency is only one of the possible applications of blockchain technology. It can register and validate all sorts of transactions and relationships. For instance, it could be used to manage the sale, lease or hire of a car. When you take possession of a car, it could be tagged with a cryptographic signature, which would mean that you are the only person who could open and start the car.[cccli]

The revolutionary benefit of the blockchain is that all kinds of agreements can be validated without setting up a centralised institution to do so. By removing the need for a central intermediary, the

blockchain can reduce transaction costs, and it can enhance privacy: no government agents need have access to your data without your permission.

Most importantly, for our present purposes, the blockchain may make possible the decentralised ownership and management of collective assets.

Collective ownership

Imagine a future in which it is apparent to many people that we are heading towards the scenario of the gods and the useless. The elite few who own the machines are as uncomfortable as the rest of us about this – or at least a sizeable number of them are. They do not want to hand their assets over to a government organisation, as they believe this would simply swap one potentially dangerous elite for another one.

But they realise that if the scenarios of the gods and the useless becomes reality, they will end up as pariahs, feared and perhaps hated by the rest of their species. This outcome might well be grim for the useless, but it would be unpleasant for the gods as well.

I don't believe the common meme that rich people are all bad, greedy and selfish. I have known quite a few, and worked for some of them. They seem to me to be the same mix of good and bad, greedy and generous as the rest of us. They tend to be smart and hard-working, but otherwise they are pretty normal, which is to say, that curious human blend of similar and different, happy and sad, predictable and unpredictable.

When machines do all the paid jobs they will be enormously efficient. Many goods and services will be effectively free. Others will be cheap and getting cheaper, but new products and services will be being invented all the time, and those which require considerable resources will be more expensive. Everyone will have a certain amount of UBI, and the way they spend it will give the machines a set of pricing signals to enable them to allocate resources sensibly. Life should be good and getting better for most people. Protopia could be within reach.

It strikes me as entirely plausible that in this scenario, the smallish minority which owns most of the assets when the game of musical chairs stopped – including notably the AI – would prefer to throw in their lot with the rest of us rather than hide behind heavily fortified gates, outcasts from the rest of their species.

It would be a non-trivial project to work out in detail how the assets could be transferred into universal common ownership, validated by the blockchain, and managed in a decentralised fashion. And it is certainly not a forgone conclusion that the rich minority would endorse it. But I suspect it will turn out to be our best way forward.

Chapter 7. Summary and recommendations

7.1 - The argument

Automation and unemployment

This book has argued that improvements in machine intelligence over the next few decades are going to make it impossible for most humans to earn a living. It concluded that we would be wise to devote some resource to working out how to deal with this development – indeed that we would be foolish not to.

In chapter 2 we looked back at the industrial revolution, and saw how concerns that automation would lead to permanent mass unemployment turned out to be unfounded. (The Engels pause was lengthy, but not permanent.) Instead, automation raised productivity and output across the economy. The unfounded concerns became known as the Luddite fallacy.

In chapter 3.1 and 3.2 we reviewed claims that this time is different. In the information revolution, mankind's third great wave of transformation, machines are increasingly able to out-perform humans in cognitive tasks. This might put humans in the predicament that horses were placed in by the industrial revolution.

In 1900, 40% of American workers were employed in agriculture, and that has now fallen below 3%. The farm workers found better jobs elsewhere in the economy, sometimes in occupations which their parents could not have imagined. But horses didn't. 1900 was "peak horse" in America, with about 25m of them working on farms; now there are fewer than 3m. The difference between the horses and the humans is that when machines took over the muscle jobs, humans had something else to offer: our cognitive, emotional and social abilities. Horses had nothing else to offer, and their population collapsed.

In chapter 3.3 we heard the response that this is merely a revival of the Luddite fallacy, and the rest of chapter 3 explored this debate in detail.

We considered whether technological automation is biting yet, and saw that the US economy has added a significant number of jobs since the banking crisis, although incomes have stagnated-. It is hard to disentangle the effects of technology on the job market from the effects of globalisation and other effects, but it does not seem likely that machine intelligence is creating widespread unemployment yet.

In chapter 3.4 we reviewed the state of the art in artificial intelligence, and in chapter 3.5 the exponential rate at which it is improving. We saw that machines are still a long way from becoming artificial general intelligences (AGIs) – computers with all the cognitive abilities of an adult human, including our flexibility. Even today's smartest machines are still artificial narrow intelligences: they may be superhuman at playing video games, or Go, but they cannot tie a pair of shoelaces or sell you a house, and they are not even aware that they are playing a game.

The fact is that machines don't need to become AGIs to displace most of us from our jobs. They simply have to become better than us at what we do for a living. Because they are overtaking us at many forms of pattern recognition, including image recognition, speech recognition, and natural

language processing, they are in the process of overtaking us. And of course, once a machine can do your job, it will quickly be able to do it faster, better and cheaper than you can. Machines don't eat, sleep, get drunk, tired or cranky. And unlike human brains, their abilities continue to improve at an exponential rate.

In chapter 3.6 we asked what capabilities people bring to the workplace, and in chapter 3.7 we reviewed the related technologies which are being introduced alongside AI, including personal digital assistants, robotics, virtual reality, augmented reality, and the internet of things. We discussed the concerns which these technologies raise about privacy, security, isolation and inequality.

In chapters 3.8 and 3.9 we discussed how widespread unemployment might arise in practice, first among drivers, as self-driving vehicles take over their function, and then in a range of other industries, including occupations like retail and sales, and the professions.

In chapter 3.10 we assessed the claim that fears about unemployment are overdone because we will invent new kinds of jobs as machines take over the old ones. We reviewed the idea that we can race with the machines instead of against them, becoming "centaurs", and occupying ourselves with the "icebergs" of new work which machines have made possible.

But at the end of it all, we concluded that these are only likely to be temporary respites. Whatever jobs we invent, the machines will take over most of them as well. In the medium term, a large minority of people - perhaps the majority – will not be able to earn a living through work.

The upside

In chapter 3.11 we saw how this does not have to be bad news. In fact in can be extremely good news. Some people are lucky enough to love their jobs, and find fulfilment in them. For many more people, work is simply a way to generate an income. It may provide a purpose, but it does not provide meaning. A world in which machines do all the boring work could be wonderful. They could be so efficient that goods and services could be plentiful, and in many cases free. Humans could get on with the important business of playing, relaxing, socialising, learning and exploring. Surely this is what we should be aiming for.

Chapter 4 provided a timeline – emphatically not a forecast – of a scenario in which humanity makes a successful transition towards this enormously positive outcome.

Challenges and scenarios

Of course there are challenges, and we turned to these in chapter 5, looking in turn at concerns about economic contraction, distribution, meaning, allocation, and cohesion.

In chapter 5.2 we explored the idea of universal basic income (UBI), a payment made to all citizens which allows them to live fulfilling lives when they are no longer able to find paid jobs. We heard that thoughtful people in America are concerned that the traditional antipathy to socialism will prevent the introduction of UBI, but we concluded that once it becomes obvious that plenty of sensible, diligent and hard-working people can no longer afford to keep themselves and their families, any such opposition will quickly fade. Nevertheless, the fact that people are concerned does underscore the need for a public debate about what is coming towards us.

In chapter 5.3 we discussed the fear that unemployed humans will find their new lives hollow, lacking in meaning, and perhaps even boring. We concluded that this too is unfounded. For centuries, aristocrats in most countries didn't work for a living, and in many societies they viewed work as a demeaning activity, to be avoided by "people of quality". Some of them got into trouble with drink, drugs and gambling, but only a small minority. Most of them seem to have led contented lives, however questionable we might find the economic systems they operated in.

Likewise, retirement is rarely considered a disaster in developed countries. Even though most of us only get to enjoy it when we are past our prime, most retirees find enough projects and pastimes to keep themselves busy and at peace. Numerous surveys have found that happiness is U-shaped: we are at our most content during childhood and retirement, and it is probably no coincidence that these are the periods in our lives when we don't work for a living. If we retired when still in our primes, we would be even better equipped to enjoy our lives of leisure.

In chapter 5.4 we considered whether virtual reality might help resolve the problem of how to allocate rare goods and services in a world where incomes are hard to vary.

In chapter 5.5 we tackled what may turn out to be the biggest challenge raised by the economic singularity: cohesion. We asked whether capitalism, and in particular the institution of private property, will be as suitable for the post-work world as it has been during the industrial revolution. This an uncomfortable discussion for people like me who believe that a sensibly-regulated market economy has been enormously beneficial for humanity. Along with the Enlightenment and the consequent scientific revolution, capitalism has made our time the best era to be born human, bar none.

But a world where almost no-one works, and an elite owns the intelligent machines, is going to be a world of fantastic and entrenched inequality. Inequality is often over-estimated as a contemporary social evil, but this post-economic singularity world will also be one where advancing technology makes available radical enhancements to our physical and cognitive performance. These enhancements will come along faster and faster, and groups with privileged access to them may start to diverge from everyone else, and become a separate species. The author Yuval Harari has referred to this scenario in the chilling phrase, "the gods and the useless". The "Brave New World" depicted by Aldous Huxley in 1931 might be one of the least bad outcomes of this scenario.

If the post-economic singularity world needs a different type of economy, then we need to start thinking now about what that might be - and also how to get there. The damage that could be caused by an uneven or violent transition to the new world could be immense.

In chapter 6 we pulled these threads together with a review of half a dozen potential scenarios. Chapter 6.1 presented and dismissed the idea that technological progress has slowed almost to a halt, and there is nothing to worry about. Chapter 6.2 rehearsed the hope that we can race with the machines by becoming centaurs and enjoying the icebergs of new work. Chapter 6.3 offered the idea that unemployment will grow, but can be accommodated by UBI.

Chapter 6.4 reprised the scenario of the gods and the useless, and chapter 6.5 reminded us that civilisation is fragile, and that a poorly-planned transition towards a new economy could be hazardous. Chapter 6.6 adopted Kevin Kelly's term Protopia for a successful transition, and

suggested that the blockchain might turn out to be the mechanism to administer society's collectively owned assets, notably its artificial intelligence.

7.2 - The two singularities

In my previous book, “Surviving AI”, I wrote at length about the challenge and the opportunity presented by the technological singularity, the moment when (and if) we create an artificial general intelligence which continues to improve its cognitive performance and becomes a superintelligence. Ensuring that we survive that event is, I believe, the single most important task facing the next generation or two of humans – along with making sure we don’t blow ourselves up with nuclear weapons, or unleash a pathogen which kills everyone.

If we secure the good outcome to the technological singularity, the future of humanity is glorious almost beyond imagination. As DeepMind co-founder Demis Hassabis likes to say, humanity’s plan for the future should consist of two steps: first, solve artificial general intelligence, and second, use that to solve everything else. “Everything else” includes poverty, illness, war and even death itself.

The stakes in the economic singularity are not so high (which is why I tackled it second.) If we find ourselves in the “gods and the useless” scenario, or if our societies collapse as we fail to transition from modern capitalism to something more suitable for the new world, it is unlikely that every human will die. (Not impossible, though, as someone might initiate a catastrophic nuclear war.) Civilisation would presumably regress, perhaps drastically, but our species would survive to try again. Trying again is something we are good at.

On the other hand, if it is coming at all, the economic singularity is coming sooner than the technological singularity. No-one knows how long it will take to build an artificial general intelligence, but it looks tremendously hard. It is probably only a matter of time, but that time may well be quite a few decades. The economic singularity is likely to be with us in two or three decades – perhaps not in the sense that a majority of people will be unemployable by then, but in the sense that it will be obvious and undeniable that it is going to happen. Asset prices may collapse at that point.

Relatively speaking, then, the technological singularity is more important but less urgent, while the economic singularity is less important but more urgent.

7.3 - What is to be done?

Relinquishment won't work

Impressed by the dangers attending the two singularities, you might think it would be a good idea to call a halt to further development of artificial intelligence research, either permanently, or simply for long enough to allow us to work out how to ensure that both events are beneficial. Unfortunately this is impossible.

First, we would not know what research to pause. Improvements in the performance of AI comes from many directions: chip design and manufacture, algorithm development, the accumulation and statistical analysis of data, to name but three. Unless we could arrest pretty much all scientific and technological research, we could not be sure that someone, somewhere, was not working on something which will advance AI.

Second, the incentive to develop and deploy a better AI than the competition is literally irresistible. For companies like Google and Facebook, who are leading the way in AI research, it is a matter of critical commercial performance in the short term, and of economic survival in the medium and long term. For military commanders it is quite literally a matter of life and death. Even if by some miracle all the world's leading politicians could be gathered together and persuaded to sign a joint declaration that all AI research will stop, they would not abide by it. We can all agree that North Korea would cheat, but who would be so naïve as to think that their own government would not do the same?

If it is possible to create an artificial general intelligence, it will be created – and it will be created as soon as it becomes possible. The same applies to the technologies required to render most humans unemployable.

Monitoring

There are at least four permanently-established organisations studying the risk to humanity posed by the potential arrival of superintelligence.[ccclii] There is only one that I know of which is studying the future of automation and technological unemployment.[ccliii] There should be more.

In chapter 4 we explored how hard it is to make accurate forecasts, but failing to keep a lookout for approaching dangers (and opportunities) is foolish. The view that most people will be rendered unemployable by machine intelligence within the next few decades is probably a minority opinion at the moment, but a great many people are uncertain about what will happen, and the case argued in this book and elsewhere is surely at least plausible enough to be worth watching out for. We should be employing economists and others to monitor the available data for signs of technological unemployment, and devising new ways to detect it.

The economist Robin Hanson thinks that machines will eventually render most humans unemployed, but that it will not happen for many decades, probably centuries. Despite this scepticism, he proposes an interesting way to watch out for the eventuality: prediction markets. People make their best estimates when they have some skin in the forecasting game. Offering people the opportunity to bet real money on when they see their own jobs or other peoples' jobs being automated may be an effective way to improve our forecasting.[ccliv]

Planning

We don't have sufficient information to draw up detailed plans for the way we would like our economies and societies to evolve. But we can, and probably should, be doing detailed scenario planning.

Scenario planning has been practised by military leaders since time immemorial. It was given the name by Herman Kahn, who wrote narratives about possible futures for the US military while working for the RAND Corporation in the 1950s. (His suggestion that a nuclear war might be both winnable and survivable made him one of the inspirations for Dr Strangelove in the classic 1964 movie.[cccv]) Scenario planning was adopted by Shell Oil after it (along with the rest of the industry) was disastrously wrong-footed by the rise of the oil cartel OPEC in the 1970s.[cccvi]

Scenario planning is more art than science, but it can be a valuable discipline. When we commit our thoughts about a possible future to paper we are forced to consider them rigorously. Institutes consisting of smart people doing this work could make a valuable contribution.

An informal version of this is the daily business of futurists and futurologists, people who are often viewed with scepticism by the wider public. Perhaps that will change – in fact, perhaps futurology will come to be seen as a mission-critical profession. Science fiction writers also have an important role, in providing vivid metaphors and warnings.

The role of the tech giants

Google, Facebook, Amazon, Microsoft, IBM and Apple are shaping the new world we are moving into, along with their Chinese counterparts Baidu, Alibaba and Tencent. Their motivation is partly commercial: they understood sooner than anyone else that artificial intelligence and related technologies will increasingly provide most of the world's economic value. They are moving aggressively to dominate the AI space, and competing fiercely with each other for talent and market positions.

Although I have no privileged access, it seems to me that many of the leading figures in this industry are also motivated by something else: a belief that the future will be better than today, and an impatience to make it arrive faster.

It is ironic, then, that these companies are often reluctant to talk about their vision. In particular, they shy away from discussing AI. It is understandable: every time they talk to a journalist about AI, the resulting article is accompanied by a picture of the Terminator. It must be enormously frustrating to be working hard to conjure a happier, richer, safer world, when all you get in response is talk about existential risks.

It may be understandable, but it is also dangerous. The idea that artificial intelligence is improving quickly is now firmly in the public mind. When self-driving cars become common, smartphones are capable of sensible conversations, and domestic robots can carry out many of our domestic chores, people will increasingly ask where it is all heading. In the absence of optimistic answers, they will gravitate towards the bad ones, and Hollywood has given us plenty of those.

We need potent new memes, illustrating the current benefits and the future promise of AI. The tech giants are creating this new world; even if only for their own self-preservation, it would be a good idea for them to explain how it is capable of being a glorious new world.

What should I study?

The question that young people (and their parents) naturally ask about the economic singularity is, how can I best prepare for the economy that we are moving towards? It's an important question for me, too: at the time of writing my son is 15.

The obvious answer is to study computing. Computers are at the heart of the changes sweeping the world in the information revolution, so it has to be valuable to understand how they work and what they can and cannot do. If possible, study machine learning, and in particular, deep learning. It seems a safe bet that these powerful techniques will remain important for years to come. Carrying out maintenance, supervision or development of hardware of software may keep you in a job for longer than most.

In the long run, however, if the argument of this book is correct, we are probably all unemployed. It may well be an advantage for a while to be rich, but if we manage the transition successfully that may become less important and less worthwhile. And if we don't … well, let's just say we have to.

Beyond the economic singularity you're going to want to have as rich an interior life as possible, so give yourself as broad an education as you can. Studying your own and other people's languages will give you insights into how our minds work. Studying sciences will give you insights into how the world works. And studying the humanities will give you insights into how societies work. All of these should make what could be a very long life an interesting one.

The most important generations

Every generation thinks the challenges it faces are more important than what has gone before. They can't all be correct. American journalist Tom Brokaw bestowed the name "the greatest generation" on the people who grew up in the Great Depression and went on to fight in the Second World War. As a "baby boomer" myself, I certainly take my hat off to that generation.

Speaking at the United Nations in 1963, John F Kennedy said something which would not sound out of place today: "Never before has man had such capacity to control his own environment, to end thirst and hunger, to conquer poverty and disease, to banish illiteracy and massive human misery. We have the power to make this the best generation of mankind in the history of the world - or make it the last."[cccLvii]

Today's rising generation is the Millennials, born between the early 1980s and the early 2000s. They are also known as Generation Y, and the one after them, born from the early 2000s to the early 2020s, is provisionally called Generation Z. Let's hope that is not prophetic.

The Millennials and Generation Z have been born at the best time ever to be a human, in terms of life expectancy, health, wealth, access to education information and entertainment. They have also been born at the most interesting time, and the most important. Whether they like it or not, they have the task of navigating us through the economic singularity of mass unemployment, and then the technological singularity of super-intelligence. If they succeed, humanity's future is

almost incredibly good. If not, it could be bleak. It will fall largely to them to plot the course, adjust it where necessary, avoid the rocks and the cries of the Sirens, and bring the ship safely home.

Acknowledgements

I am enormously grateful to the following people, who have given their time and energy to support this book, and in many cases to provide constructive criticism on its earlier drafts. I have learned a lot from their insights, and the book is much better for them. All errors, omissions and solecisms are of course my fault, not theirs.

Adam Jolly, Adam Singer, Aubrey de Grey, Ben Medlock, Ben Goldsmith, Chris Meyer, Clive Pinder, David Wood, Gerald Huff, Hugo de Garis, Jeff Pinsker, Jim Muttram, Justin Stewart, Kenneth Cukier, Malcolm Myers, Peter Fenton O'Creevy, Peter Monk, Randal Koene, Roman Yampolskiy, Stuart Armstrong, Will Gilpin, William Charlwood, William Graham.

As always, my profoundest thanks go to my partner Julia, who is my adviser, my cheerleader, and my kindest but most penetrating critic. My hugely talented designer Rachel Lawston has again produced a cover which (I think) looks great and answers the brief perfectly.

i Memes are ideas or beliefs which spread from person to person to become pervasive within a culture.

ii Many textbooks place the start of the industrial revolution in the second half of the 18th century, but I like the argument that Thomas Newcome's creation of the first practical steam engine in 1712 provides the best origin story.

iii There is no general agreement about when the information revolution started. In his 1962 book "The Production and Distribution of Knowledge in the United States", the Austrian economist Fritz Machler suggested that with 29% of GDP accounted for by the knowledge industry, it had begun.

iv The term was first applied to human affairs back in the 1950s by John von Neuman, a key figure in the development of the computer. The physicist and science fiction author Vernor Vinge argued in 1993 that artificial intelligence and other technologies would cause a singularity in human affairs within 30 years. This idea was picked up and popularised by the inventor and futurist Ray Kurzweil, who believes that computers will overtake humans in general intelligence in 1929, and a singularity will arrive in 2045. https://en.wikipedia.org/wiki/Technological_singularity

v The event horizon of a black hole is the point beyond which events cannot affect an outside observer, or in other words, the point of no return. The gravitational pull has become so great as to make escape impossible, even for light.

vi http://fivethirtyeight.com/features/universal-basic-income/?utm_content=buffer71a7e&utm_medium=social&utm_source=plus.google.com&utm_campaign=buffer

vii At the end of this video: http://bit.ly/1MtEqNb

viii https://www.minnpost.com/macro-micro-minnesota/2012/02/history-lessons-understanding-decline-manufacturing

ix http://blogs.rmg.co.uk/longitude/2014/07/30/guest-post-pirate-map/

x https://www.weforum.org/pages/the-fourth-industrial-revolution-by-klaus-schwab

xi http://www.ers.usda.gov/media/259572/eib3_1_.pdf. Employment in agriculture declined in absolute terms as well, from 11.7m in 1900 to 6.0m in 1960. http://www.nber.org/chapters/c1567.pdf

xii_www.ons.gov.cuk/ons/rel/census/2011-census-analysis/170-years-of-industry/170-years-of-industrial-changeponent.html

xiii http://www.americanequestrian.com/pdf/us-equine-demographics.pdf

xiv https://en.wikipedia.org/wiki/Automation#cite_note-7

xv M. A. Laughton, D. J. Warne (ed), *Electrical Engineer's Reference book*

xvi http://www.oleantimesherald.com/news/did-you-know-gas-pump-shut-off-valve-was-invented/article_c7a00da2-b3eb-54e1-9c8d-ee36483a7e33.html

xvii Radio frequency Identification tags. They can take various forms – for instance, some have inbuilt power sources, while others are powered by interacting with nearby magnetic fields, or the radio waves which interrogate them.

xviii http://www.businessinsider.com/three-chinese-restaurants-fired-their-robot-workers-2016-4

xix https://www.illinoispolicy.org/mcdonalds-counters-fight-for-15-with-automation/

xx http://www.eater.com/2016/5/5/11597270/kfc-robots-china-shanghai

xxi http://www.ehow.com/about_4678910_robots-car-manufacturing.html

xxii http://www.ifr.org/industrial-robots/statistics/

xxiii http://www.npr.org/sections/money/2015/02/05/382664837/map-the-most-common-job-in-every-state

xxiv http://www.nationalarchives.gov.uk/education/politics/g5/

xxv http://jetpress.org/v24/campa2.htm

xxvi Ricardo originally thought that innovation benefited everyone, but he was persuaded by Malthus that it could suppress wages and cause long-term unemployment. He added a chapter called "On Machinery" to the final edition of his book "On the Principles of Political Economy and Taxation".

xxvii http://www.theguardian.com/business/2015/aug/17/technology-created-more-jobs-than-destroyed-140-years-data-census

xxviii https://en.wikipedia.org/wiki/Bowley%27s_law

xxix http://www.economics.ox.ac.uk/Department-of-Economics-Discussion-Paper-Series/engel-s-pause-a-pessimist-s-guide-to-the-british-industrial-revolution

xxx http://press.princeton.edu/titles/8659.html

xxxi This depends on the two planets being pretty much as close as they ever get.

xxxii http://fortune.com/2015/11/10/us-unemployment-rate-economy/

xxxiii This and the other quotes in this paragraph and the next one are from Chapter 10: Toward a New Economic Paradigm.

xxxiv Brynjolfsson is the director of the MIT Center for Digital Business and McAfee is a principal research scientist there.

xxxvThe word "inequality" crops up 42 times in the book, including in the titles of sources, but the authors never explicitly connect it with "spread".

xxxvi The loosely-organised protest organisation that sprang up after the 2008 credit crunch to campaign against inequality.

xxxviiChapter 12: Learning to Race with the Machines: Recommendations for Individuals.

xxxviii
Chapter 13: Policy Recommendations.

xxxixChapter 14: Long-Term Recommendations.

xl http://www.susskind.com/

xli http://www.scottsantens.com/

xlii
https://www.reddit.com/r/BasicIncome/ and https://www.reddit.com/r/basicincome/wiki/index

xliiihttps://www.youtube.com/watch?v=7Pq-S557XQU.

xliv https://www.youtube.com/watch?v=C5MVXdg6nho.

xlv http://www.cbsnews.com/videos/how-technology-may-change-our-labor-and-leisure/

xlvi http://www.bankofengland.co.uk/publications/Pages/speeches/2015/864.aspx

xlvii https://newrepublic.com/article/69326/call-the-wolf

xlviii http://www.ft.com/cms/s/0/dfe218d6-9038-11e3-a776-00144feab7de.html#axzz3stkJb1V2

xlix The Programme was established in January 2015 with funding from Citibank, one of the largest financial institutions in the world. The Oxford Martin school was set up as part of Oxford University in 2005, as an institution dedicated to understanding the threats and opportunities facing humanity in the 21st century. It is named after James Martin, a writer, consultant and entrepreneur, who founded the school with the largest donation ever made to the university – which was no mean feat given that Oxford was founded 1,000 years ago, and is the oldest university in the world (after Bologna in Italy).

l http://www.computerworld.com/article/2691607/one-in-three-jobs-will-be-taken-by-software-or-robots-by-2025.html

li http://www.pewinternet.org/2014/08/06/about-this-report-and-survey-2/?beta=true&utm_expid=53098246-2.Lly4CFSVQG2lphsg-KopIg.1&utm_referrer=https%3A%2F%2Fwww.google.co.uk%2F

lii https://www.fundacionbankinter.org/web/fundacion-bankinter/ficha-documento?param_id=173404#_48_INSTANCE_av33_%3Dhttps%253A%252F%252Fwww.fundacionbankinter.org%252Fweb%252Fglobal-site%252F-%252Fthe-machine-revolution%253F

liii http://www.mckinsey.com/insights/business_technology/four_fundamentals_of_workplace_automation

liv http://www.socialeurope.eu/2015/10/the-limits-of-the-digital-revolution-why-our-washing-machines-wont-go-to-the-moon/

lv https://www.aeaweb.org/articles.php?doi=10.1257/jep.29.3.3

lvi https://reason.com/archives/2015/03/03/how-to-survive-a-robot-uprisin

lvii http://www.politico.com/magazine/story/2013/11/the-robots-are-here-098995

lviii http://www.forbes.com/sites/danschawbel/2015/08/04/geoff-colvin-why-humans-will-triumph-over-machines/2/

lix http://www.eastoftheweb.com/short-stories/UBooks/BoyCri.shtml

lx German academic Marcus Hutter, and Shane Legg, co-founder of DeepMind

lxi http://www.savethechimps.org/about-us/chimp-facts/

lxii The Shape of Automation for Men and Management by Herbert Simon, 1965

lxiii

lxiv http://www.wired.com/2016/01/microsoft-neural-net-shows-deep-learning-can-get-way-deeper/

lxv http://www.etymonline.com/index.php?term=algorithm

lxvi http://www.wired.com/2016/01/microsoft-neural-net-shows-deep-learning-can-get-way-deeper/

lxvii Moravec wrote about this phenomenon in his 1988 book "Mind Children". A possible explanation is that the sensory motor skills and spatial awareness that we develop as children are the product of millions of years of evolution. Rational thought is something we have only been doing for a few thousand years. Perhaps it really isn't hard, but just seems hard because we are not yet optimised for it.

lxviii https://www.youtube.com/watch?v=Skfw282fJak

lxix http://futureoflife.org/2016/01/27/are-humans-dethroned-in-go-ai-experts-weigh-in/

lxx http://www.nervanasys.com/demystifying-deep-reinforcement-learning/

lxxi https://www.newscientist.com/article/2076552-google-deepmind-ai-navigates-a-doom-like-3d-maze-just-by-looking/

lxxii http://www.popsci.com/scitech/article/2004-06/darpa-grand-challenge-2004darpas-debacle-desert

lxxiii https://www.theguardian.com/technology/2016/mar/09/google-self-driving-car-crash-video-accident-bus

lxxiv http://www.wsj.com/articles/toyota-to-invest-1-billion-in-artificial-intelligence-firm-1446790646

lxxv
http://www.forbes.com/sites/chunkamui/2015/12/23/5-reasons-why-automakers-should-fear-googles-partnership-with-ford/

lxxvi
http://electrek.co/2015/12/21/tesla-ceo-elon-musk-drops-prediction-full-autonomous-driving-from-3-years-to-2/

lxxvii http://www.thechurchofgoogle.org/Scripture/Proof_Google_Is_God.html

lxxviii
https://www.reddit.com/r/churchofgoogle/

lxxix The answer, if you're searching from England, is to fly over Asia.

lxxx http://searchengineland.com/faq-all-about-the-new-google-rankbrain-algorithm-234440?utm_campaign=socialflow&utm_source=facebook&utm_medium=social

lxxxi http://www.wired.com/2016/02/ai-is-changing-the-technology-behind-google-searches/

lxxxii http://www.thedrum.com/opinion/2016/02/08/why-artificial-intelligence-key-google-s-battle-amazon

lxxxiii http://www.wired.com/2012/06/google-x-neural-network/

lxxxiv They are the Pembroke and the Cardigan Corgi. http://research.microsoft.com/en-us/news/features/dnnvision-071414.aspx

lxxxvhttp://image-net.org/challenges/LSVRC/2015/index#news

lxxxvi
http://www.eetimes.com/document.asp?doc_id=1325712

lxxxvii
https://youtu.be/U_Wgc1JOsBk?t=33

lxxxviiihttp://news.sciencemag.org/social-sciences/2015/02/facebook-will-soon-be-able-id-you-any-photo

lxxxix
http://www.computerworld.com/article/2941415/data-privacy/is-facial-recognition-a-threat-on-facebook-and-google.html

xc http://www.wired.com/2016/01/2015-was-the-year-ai-finally-entered-the-everyday-world/

xci At the time of writing, April 2016, Aipoly is impressive, but far from perfect.

xcii http://www.bloomberg.com/news/2014-12-23/speech-recognition-better-than-a-human-s-exists-you-just-can-t-use-it-yet.html

xciii
http://www.forbes.com/sites/parmyolson/2014/05/28/microsoft-unveils-near-real-time-language-translation-for-skype/

xcivhttp://www.technologyreview.com/news/544651/baidus-deep-learning-system-rivals-people-at-speech-recognition/#comments

xcv https://youtu.be/V1eYniJ0Rnk?t=1

xcvihttp://edge.org/response-detail/26780

xcvii http://techcrunch.com/2016/03/19/how-real-businesses-are-using-machine-learning/

xcviii http://www.latimes.com/business/technology/la-fi-cutting-edge-ibm-20160422-story.html

xcix http://www.wired.com/2016/04/openai-elon-musk-sam-altman-plan-to-set-artificial-intelligence-free/

c http://www.strategyand.pwc.com/global/home/what-we-think/innovation1000/top-innovators-spenders#/tab-2015

ci
2013 data: http://www.ons.gov.uk/ons/rel/rdit1/gross-domestic-expenditure-on-research-and-development/2013/stb-gerd-2013.html

cii http://insights.venturescanner.com/category/artificial-intelligence-2/

ciii
http://techcrunch.com/2015/12/25/investing-in-artificial-intelligence/

civ http://www.wired.com/2015/11/google-open-sources-its-artificial-intelligence-engine/

cv https://www.theguardian.com/technology/2016/apr/13/google-updates-tensorflow-open-source-artificial-intelligence

cvi http://www.wired.com/2015/12/facebook-open-source-ai-big-sur/

cvii The name Parsey McParseFace is a play on a jokey name for a research ship which received a lot of votes in a poll run by the British government in April 2016. http://www.wsj.com/articles/googles-open-source-parsey-mcparseface-helps-machines-understand-english-1463088180

cviii Assuming you don't count the Vatican as a proper country. http://www.ibtimes.co.uk/google-project-loon-provide-free-wifi-across-sri-lanka-1513136

cix https://setandbma.wordpress.com/2013/02/04/who-coined-the-term-big-data/

cx http://www.pcmag.com/encyclopedia/term/37701/amara-s-law

cxi http://www.lrb.co.uk/v37/n05/john-lanchester/the-robots-are-coming

cxii Haitz's Law states that the cost per unit of useful light emitted decreases exponentially

cxiii
http://computationalimagination.com/article_cpo_decreasing.php

cxiv
http://www.nytimes.com/2006/06/07/technology/circuits/07essay.html

cxv . http://arstechnica.com/gadgets/2015/02/intel-forges-ahead-to-10nm-will-move-away-from-silicon-at-7nm/

cxvi
. The “III-V” refers to the periodic table group the material belongs to. Transistors made from these semiconductors should consume far less power, and also switch much faster.

cxvii http://www.extremetech.com/extreme/225353-intel-formally-kills-its-tick-tock-approach-to-processor-development

cxviii http://www.nextplatform.com/2015/11/26/intel-supercomputer-powers-moores-law-life-support/

cxix http://www.theguardian.com/technology/2015/jul/09/moores-law-new-chips-ibm-7nm

cxx Clock speed, also known as clock rate or processor speed, is the number of cycles a chip (central processing unit, or CPU) performs each second. Inside each chip is a small quartz crystal which vibrates, or oscillates, at a particular frequency. It takes a fixed number of oscillations, or cycles, to perform the instructions that a chip is given. One cycles is one Hertz, and today’s chips operate in gigaHertz (GHz), billions of cycles per second. As other aspects of chip designs diverge, clock speed is no longer a reliable measure of a chip’s effective performance.

cxxi http://www.popularmechanics.com/technology/a18493/stanford-3d-computer-chip-improves-performance/

cxxiihttp://gadgets.ndtv.com/science/news/mit-builds-low-power-artificial-intelligence-chip-for-smartphones-799803

cxxiiihttp://www.engadget.com/2016/03/28/ibm-resistive-processing-deep-learning/

cxxivhttp://arstechnica.com/gadgets/2016/04/nvidia-tesla-p100-pascal-details/

cxxvCPU stands for Central Processing Unit. They are general purpose processors which can carry out many kinds of computation, but are not necessarily optimised for any of them. GPU stands for Graphics Processing Unit, and as the name suggests, they were originally designed for displaying graphics in video games. They are very good at taking huge quantities of data and carrying out the same operation over and over again. It turns out that machine learning benefits from their particular capabilities. CPUs and GPUs are often deployed in tandem.

cxxvihttp://www.technologyreview.com/news/544421/googles-quantum-dream-machine/

cxxvii
http://www.technologyreview.com/news/537041/ibm-shows-off-a-quantum-computing-chip/

cxxviiihttp://www.nature.com/news/the-chips-are-down-for-moore-s-law-1.19338

cxxixhttp://fortune.com/facebook-machine-learning/

cxxx 2013 data: http://www.ons.gov.uk/ons/dcp171778_315661.pdf

cxxxi This has given rise to the term “subtractive manufacturing” for the traditional forms of manufacturing. This method of naming is rather splendidly called a retronym.

cxxxii https://en.wikipedia.org/wiki/Fax

cxxxiiihttp://www.rfidjournal.com/articles/view?4986

cxxxiv
http://www.vdi-nachrichten.com/Technik-Gesellschaft/Industrie-40-Mit-Internet-Dinge-Weg-4-industriellen-Revolution

cxxxv
Coined by another British entrepreneur, Simon Birrell:
https://www.linkedin.com/in/simonbirrell

cxxxvihttp://www.gartner.com/newsroom/id/2636073

cxxxvii http://singularityhub.com/2016/02/09/when-the-world-is-wired-the-magic-of-the-internet-of-everything/

cxxxviii
http://www.telegraph.co.uk/technology/internet/12050185/Marc-Andreessen-In-20-years-every-physical-item-will-have-a-chip-implanted-in-it.html

cxxxix http://www.information-age.com/it-management/strategy-and-innovation/123460379/trains-brains-how-artificial-intelligence-transforming-railway-industry

cxl http://home.cern/topics/birth-web

cxli http://www.theguardian.com/technology/2016/jan/31/viv-artificial-intelligence-wants-to-run-your-life-siri-personal-assistants

cxliiNot an everyday object outside the USA, of course

cxliiihttp://www.bloomberg.com/news/articles/2016-01-11/google-chairman-thinks-ai-can-help-solve-world-s-hard-problems-

cxlivThis is actually a great idea, which is being trialled in Argentina at the time of writing: http://www.telegraph.co.uk/motoring/motoringvideo/11680348/Transparent-trucks-with-rear-mounted-Samsung-safety-screens-set-to-save-overtaking-drivers.html Of course it may be less valuable when cars drive themselves and their human occupants don't look at the road.

cxlv An ocarina is a wind instrument about the size of a fist. First introduced to Europeans by the Aztecs, it looks like a toy submarine.

cxlvi https://www.youtube.com/watch?v=0QNiZfSsPc0

cxlviihttp://www.popsci.com/darpa-robotics-challenge-was-bust-why-darpa-needs-try-again

cxlviiihttp://uk.businessinsider.com/laundroid-japanese-robot-folds-laundry-2015-10

cxlix
http://www.npr.org/sections/money/2015/05/19/407736307/robots-are-really-bad-at-folding-towels

cl http://www.techinsider.io/savioke-robot-butler-in-united-states-hotels-2016-2

cli http://www.kurzweilai.net/the-top-ai-breakthroughs-of-2015

clii http://www.nextgov.com/emerging-tech/2016/05/robots-are-starting-learn-touch/128065/

cliii http://www.theguardian.com/world/2015/sep/28/no-sex-with-robots-says-japanese-android-firm-softbank

cliv https://www.theguardian.com/technology/2015/aug/03/hitchbot-hitchhiking-robot-destroyed-philadelphia

clv http://www.telegraph.co.uk/news/science/science-news/12073587/Meet-Nadine-the-worlds-most-human-like-robot.html

clvi http://techcrunch.com/2016/01/07/the-grillbot-is-a-robot-that-cleans-your-grill/#.w9z87m:Hd0d

clvii http://singularityhub.com/2016/02/29/drones-have-reached-a-tipping-point-heres-what-happens-next/

clviiihttp://intl.eksobionics.com/

clix Your brain is wired to make you see things before you hear them as it knows that light travels faster than sound. Thus the brain can tolerate audio lagging video, but is much less tolerant of video lagging audio. This is known as Multi-Modal Perception.

clx http://www.ft.com/cms/s/0/b33d75fe-cc5a-11e5-be0b-b7ece4e953a0.html#axzz3znOxP8QH

clxi http://www.kitguru.net/peripherals/anton-shilov/gartner-two-million-vr-headsets-to-be-sold-in-2016/

clxii
http://www.digi-capital.com/news/2015/04/augmentedvirtual-reality-to-hit-150-billion-disrupting-mobile-by-2020/#.VoV65vmLRD8

clxiii
http://uk.businessinsider.com/virtual-reality-on-gartner-hype-cycle-2015-8

clxivhttp://techcrunch.com/2016/01/30/how-the-growth-of-mixed-reality-will-change-communication-collaboration-

and-the-future-of-the-workplace/

clxvThe games industry is much bigger than Hollywood if you stop measuring move income at the box office. If you add in DVD and other "windows", plus merchandising, it is hard to say. https://www.quora.com/Who-makes-more-money-Hollywood-or-the-video-game-industry

clxvi https://versions.killscreen.com/we-should-be-talking-about-torture-in-vr/

clxviihttp://www.tomdispatch.com/post/175822/tomgram%3A_crump_and_harwood%2C_the_net_closes_around_us/

clxviiihttps://www.washingtonpost.com/local/public-safety/the-new-way-police-are-surveilling-you-calculating-your-threat-score/2016/01/10/e42bccac-8e15-11e5-baf4-bdf37355da0c_story.html

clxixhttp://www.newyorker.com/tech/elements/little-brother-is-watching-you

clxxhttp://www.wired.com/2014/03/going-tracked-heres-way-embrace-surveillance/

clxxi https://www.washingtonpost.com/news/the-switch/wp/2016/03/28/mass-surveillance-silences-minority-opinions-according-to-study/

clxxiihttp://www.bbc.co.uk/news/world-asia-china-34592186

clxxiii http://www.computerworld.com/article/2990203/security/aclu-orwellian-citizen-score-chinas-credit-score-system-is-a-warning-for-americans.html

clxxivhttp://www.theguardian.com/technology/2015/oct/06/peeple-ratings-app-removes-contentious-features-boring

clxxv https://www.technologyreview.com/s/601294/microsoft-and-google-want-to-let-artificial-intelligence-loose-on-our-most-private-data/?utm_source=Twitter&utm_medium=tweet&utm_campaign=@KyleSGibson

clxxviThe Flynn Effect: http://www.bbc.co.uk/news/magazine-31556802

clxxviiWHO "Global Status Report on Road Safety 2013: supporting a decade of action

clxxviiihttp://www.japantimes.co.jp/news/2015/11/15/business/tech/human-drivers-biggest-threat-developing-self-driving-cars/#.Vo7D5fmLRD8

clxxixhttp://www.theatlantic.com/business/archive/2013/02/the-american-commuter-spends-38-hours-a-year-stuck-in-traffic/272905/

clxxxhttp://www.reinventingparking.org/2013/02/cars-are-parked-95-of-time-lets-check.html

clxxxihttp://www.etymonline.com/index.php?term=autocar

clxxxiihttp://www.digitaltrends.com/cars/audi-autonomous-car-prototype-starts-550-mile-trip-to-ces/

clxxxiiihttp://www.nhtsa.gov/About+NHTSA/Press+Releases/U.S.+Department+of+Transportation+Releases+Policy+on+Automated+Vehicle+Development

clxxxiv http://www.reuters.com/investigates/special-report/autos-driverless/

clxxxvhttp://www.wired.com/2015/04/delphi-autonomous-car-cross-country/

clxxxvihttp://recode.net/2015/03/17/google-self-driving-car-chief-wants-tech-on-the-market-within-five-years/

clxxxviihttp://techcrunch.com/2015/12/22/a-new-system-lets-self-driving-cars-learn-streets-on-the-fly/

clxxxviiihttp://cleantechnica.com/2015/10/12/autonomous-buses-being-tested-in-greek-city-of-trikala/

clxxxixhttp://www.bloomberg.com/news/articles/2015-12-16/google-said-to-make-driverless-cars-an-alphabet-company-in-2016

cxc http://electrek.co/2015/12/21/tesla-ceo-elon-musk-drops-prediction-full-autonomous-driving-from-3-years-to-2/

cxci http://venturebeat.com/2016/01/10/elon-musk-youll-be-able-to-summon-your-tesla-from-anywhere-in-2018/

cxcii https://www.washingtonpost.com/news/the-switch/wp/2016/01/11/elon-musk-says-teslas-autopilot-is-already-probably-better-than-human-drivers/

cxciii http://electrek.co/2016/04/24/tesla-autopilot-probability-accident/

cxcivhttp://www.bbc.co.uk/news/technology-35280632

cxcv http://www.zdnet.com/article/ford-self-driving-cars-are-five-years-away-from-changing-the-world/

cxcvi http://www.reuters.com/investigates/special-report/autos-driverless/

cxcviihttp://www.wired.com/2015/12/californias-new-self-driving-car-rules-are-great-for-texas/

cxcviii http://www.reuters.com/investigates/special-report/autos-driverless/

cxcix It has been suggested that electric cars should make noises so that people don't step off the pavement in front of them. A friend tells me he would like his to make a noise like two coconuts being banged together, in homage to the scene in Monty Python and the Holy Grail where King Arthur, unable to afford a horse, has a camp follower fake the noise of one with coconuts.

cc http://www.pcmag.com/article2/0,2817,2370598,00.asp

cci http://www.nytimes.com/2015/11/06/technology/toyota-silicon-valley-artificial-intelligence-research-center.html?_r=0

ccii https://www.yahoo.com/autos/google-pairs-with-ford-to-1326344237400118.html

cciii http://uk.businessinsider.com/bmw-says-cars-with-artificial-intelligence-are-already-here-2016-1?r=US&IR=T

ccivhttp://www.wsj.com/articles/SB10001424053111903480904576512250915629460

ccv http://fortune.com/2014/05/04/6-things-i-learned-at-buffetts-annual-meeting/

ccvihttp://www.thenewspaper.com/news/43/4341.asp

ccvii. http://www.alltrucking.com/faq/truck-drivers-in-the-usa/

ccviii . http://www.bls.gov/ooh/transportation-and-material-moving/bus-drivers.htm

ccix
. http://www.bls.gov/ooh/transportation-and-material-moving/taxi-drivers-and-chauffeurs.htm

ccx See chapter 3.10.

ccxihttp://www.joc.com/trucking-logistics/truckload-freight/driver-wage-hikes-could-raise-truckload-pricing-12-18-percent_20150325.html

ccxii. *The Economist*, December 4, 2003

ccxiii
http://www.abc.net.au/news/2015-10-18/rio-tinto-opens-worlds-first-automated-mine/6863814

ccxiv
http://www.mining.com/why-western-australia-became-the-center-of-mine-automation/

ccxv http://www.npr.org/sections/money/2015/02/05/382664837/map-the-most-common-job-in-every-state

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

ccxviihttps://en.wikipedia.org/wiki/Horn_%26_Hardart#Automated_food

ccxviii
http://www.computerworld.com/article/2837810/automation-arrives-at-restaurants-but-dont-blame-rising-minimum-wages.html

ccxixhttp://blogs.forrester.com/andy_hoar/15-04-14-death_of_a_b2b_salesman

ccxxhttp://www.nuance.com/for-business/customer-service-solutions/nina/index.htm

ccxxi
http://www.zdnet.com/article/swedbank-humanises-customer-service-with-artificial-intelligence-platform/

ccxxii https://www.technologyreview.com/s/601215/china-is-building-a-robot-army-of-model-workers/

ccxxiiihttp://www.theguardian.com/technology/2014/sep/12/artificial-intelligence-data-journalism-media

ccxxivhttp://www.arria.com/

ccxxvhttps://www.youtube.com/watch?v=HXKDnqM9Ulw

ccxxvi
http://www.chinadaily.com.cn/china/2015-12/24/content_22794242.htm

ccxxviihttp://persado.com/

ccxxviiihttp://www.techtimes.com/articles/127526/20160126/ai-politics-how-an-artificial-intelligence-algorithm-can-write-political-speeches.htm

ccxxixhttp://www.ravn.co.uk/

ccxxxhttp://www.legalweek.com/legal-week/sponsored/2434504/is-artificial-intelligence-the-key-to-unlocking-innovation-in-your-law-firm

ccxxxihttp://linkis.com/www.theatlantic.com/SoE5e

ccxxxiihttp://www.legalfutures.co.uk/latest-news/come-americans-legalzoom-gains-abs-licence

ccxxxiii
https://www.fairdocument.com/

ccxxxivhttp://msutoday.msu.edu/news/2014/using-data-to-predict-supreme-courts-decisions/

ccxxxvhttp://uk.businessinsider.com/robots-may-make-legal-workers-obsolete-2015-8

ccxxxvi http://www.kurzweilai.net/machine-learning-rivals-human-skills-in-cancer-detection

ccxxxviihttp://uk.businessinsider.com/deepmind-cofounders-invest-in-babylon-health-2016-1

ccxxxviiihttp://singularityhub.com/2016/01/18/digital-diagnosis-intelligent-machines-do-a-better-job-than-humans/?utm_content=bufferb9e5d&utm_medium=social&utm_source=twitter.com&utm_campaign=buffer

ccxxxixhttp://forbesindia.com/article/hidden-gems/thyrocare-technologies-testing-new-waters-in-medical-diagnostics/41051/1

ccxlhttp://www.ucsf.edu/news/2011/03/9510/new-ucsf-robotic-pharmacy-aims-improve-patient-safety

ccxlihttp://www.qmed.com/news/ibms-watson-could-diagnose-cancer-better-doctors

ccxliihttp://www.ft.com/cms/s/2/dced8150-b300-11e5-8358-9a82b43f6b2f.html#axzz3xL3RoRdy

ccxliii
http://www.ft.com/cms/s/2/dced8150-b300-11e5-8358-9a82b43f6b2f.html#axzz3xL3RoRdy

ccxliv http://www.theverge.com/2016/3/10/11192774/demis-hassabis-interview-alphago-google-deepmind-ai

ccxlvhttp://qz.com/567658/searching-for-eureka-ibms-path-back-to-greatness-and-how-it-could-change-the-world/

ccxlvihttp://www.forbes.com/sites/peterhigh/2016/01/18/ibm-watson-head-mike-rhodin-on-the-future-of-artificial-intelligence/#24204aab3e2922228b9c30cc

ccxlvii
http://www.dotmed.com/news/story/29020

ccxlviiihttp://www.wsj.com/articles/SB10001424052702303983904579093252573814132

ccxlix http://www.outpatientsurgery.net/outpatient-surgery-news-and-trends/general-surgical-news-and-reports/ethicon-pulling-sedasys-anesthesia-system--03-10-16

ccl http://www.wired.co.uk/news/archive/2016-05/05/autonomous-robot-surgeon

ccli https://www.edsurge.com/news/2016-04-18-gradescope-raises-2-6m-to-apply-artificial-intelligence-to-grading-exams

ccliihttp://www.wsj.com/articles/if-your-teacher-sounds-like-a-robot-you-might-be-on-to-something-1462546621

ccliiihttps://www.sigfig.com/site/#/home

ccliv http://www.nytimes.com/2016/01/23/your-money/robo-advisers-for-investors-are-not-one-size-fits-all.html?_r=0

cclv http://www.bloomberg.com/news/articles/2015-02-27/bridgewater-is-said-to-start-artificial-intelligence-team

cclvihttp://www.wired.com/2016/01/the-rise-of-the-artificially-intelligent-hedge-fund/

cclviihttps://next.ft.com/content/c31f8f44-033b-11e6-af1d-c47326021344 (Paywall)

cclviiihttp://www.ft.com/cms/s/0/5eb91614-bee5-11e5-846f-79b0e3d20eaf.html#axzz3zEmSvuZs

cclixhttps://itunes.apple.com/gb/podcast/exchanges-at-goldman-sachs/id948913991?mt=2&i=361020299

cclx http://uk.businessinsider.com/high-salary-jobs-will-be-automated-2016-3

cclxihttp://www.fiercefinanceit.com/story/will-regulatory-compliance-drive-artificial-intelligence-adoption/2016-01-05

cclxii http://www.liverpoolecho.co.uk/news/business/liverpool-fc-sponsor-standard-chartered-11104215

cclxiiihttp://www.cnbc.com/2015/12/30/artificial-intelligence-making-some-bosses-nervous-study.html

cclxivAssuming the work is happening on Earth. Wikipedia offers a more general but less euphonious definition: "Work is the product of the force applied and the displacement of the point where the force is applied in the direction of the force."

cclxvhttp://www.wsj.com/articles/can-the-sharing-economy-provide-good-jobs-1431288393

cclxvi https://www.edge.org/conversation/kevin_kelly-the-technium

cclxvii https://www.singularityweblog.com/techemergence-surveys-experts-on-ai-risks/

cclxviiihttp://uk.businessinsider.com/social-skills-becoming-more-important-as-robots-enter-workforce-2015-12

cclxixhttp://www.history.com/topics/inventions/automated-teller-machines

cclxx http://www.theatlantic.com/technology/archive/2015/03/a-brief-history-of-the-atm/388547/

cclxxihttp://www.wsj.com/articles/SB10001424052748704463504575301051844937276

cclxxii http://kalw.org/post/robotic-seals-comfort-dementia-patients-raise-ethical-concerns#stream/0

cclxxiiihttp://viterbi.usc.edu/news/news/2013/a-virtual-therapist.htm

cclxxivhttp://observer.com/2014/08/study-people-are-more-likely-to-open-up-to-a-talking-computer-than-a-human-therapist/

cclxxvhttp://mindthehorizon.com/2015/09/21/avatar-virtual-reality-mental-health-tech/

cclxxvihttp://www.handmadecake.co.uk/

cclxxviihttp://www.bbc.co.uk/news/magazine-15551818

cclxxviiihttp://www.oxforddnb.com/view/article/19322

cclxxixhttp://www.ft.com/cms/s/2/c5cf07c4-bf8e-11e5-846f-79b0e3d20eaf.html#axzz3yLGlrr1J

cclxxxhttp://www.bls.gov/cps/cpsaat11.htm

cclxxxi https://en.wikipedia.org/wiki/No_Man%27s_Sky

cclxxxiihttp://www.ft.com/cms/s/2/c5cf07c4-bf8e-11e5-846f-79b0e3d20eaf.html#axzz3yLGlrr1J

cclxxxiii http://www.inc.com/john-brandon/22-inspiring-quotes-from-famous-entrepreneurs.html

cclxxxivhttp://www.uh.edu/engines/epi265.htm

cclxxxvhttp://googleresearch.blogspot.co.uk/2015/06/inceptionism-going-deeper-into-neural.html

cclxxxvi http://www.bbc.co.uk/news/technology-35977315

cclxxxviihttp://fee.org/freeman/the-economic-fantasy-of-star-trek/

cclxxxviiihttps://www.wired.co.uk/news/archive/2012-11/16/iain-m-banks-the-hydrogen-sonata-review

cclxxxixhttp://www.ft.com/cms/s/0/dfe218d6-9038-11e3-a776-00144feab7de.html#axzz3yUOe9Hkp

ccxchttp://www.brautigan.net/machines.html

ccxci As noted in chapter 3.4, Anders Sandberg is James Martin Fellow at the Future of Humanity Institute at Oxford University. He was referring to Elon Musk's warning that we might be the boot loader for a digital superintelligence, meaning that we create it and then disappear. Anders suggests that a better fate would be what happened to a prokaryotic cell which was absorbed by another, larger cell and became an essential component of a combined, more complex entity, the first eukaryotic cell.

ccxciihttp://money.cnn.com/2015/06/23/investing/facebook-walmart-market-value/

ccxciiihttp://quoteinvestigator.com/2011/11/16/robots-buy-cars/

ccxcivhttp://thegreatdepressioncauses.com/unemployment/

ccxcv
http://www.statista.com/statistics/268830/unemployment-rate-in-eu-countries/

ccxcvi
http://www.statista.com/statistics/266228/youth-unemployment-rate-in-eu-countries/

ccxcviihttp://www.scottsantens.com/

ccxcviiihttp://www.economonitor.com/dolanecon/2014/01/27/a-universal-basic-income-conservative-progressive-and-libertarian-perspectives-part-3-of-a-series/

ccxcixhttps://www.reddit.com/r/BasicIncome/wiki/index#wiki_that.27s_all_very_well.2C_but_where.27s_the_evidence.3F

ccc https://www.reddit.com/r/BasicIncome/wiki/studies

ccci http://basicincome.org.uk/2013/08/health-forget-mincome-poverty/

cccii http://fivethirtyeight.com/features/universal-basic-income/?utm_content=buffer71a7e&utm_medium=social&utm_source=plus.google.com&utm_campaign=buffer

ccciii http://www.fastcoexist.com/3052595/how-finlands-exciting-basic-income-experiment-will-work-and-what-we-can-learn-from-it

ccciv http://www.latimes.com/world/europe/la-fg-germany-basic-income-20151227-story.html

cccv http://www.vox.com/2016/1/28/10860830/y-combinator-basic-income

cccvi https://en.wikipedia.org/wiki/Sodomy_laws_in_the_United_States#References

cccvii http://blogs.wsj.com/washwire/2015/03/09/support-for-gay-marriage-hits-all-time-high-wsjnbc-news-poll/

cccviii http://www.huffingtonpost.com/2009/05/06/majority-of-americans-wan_n_198196.html

cccix http://blogs.seattletimes.com/today/2013/08/washingtons-pot-law-wont-get-federal-challenge/

cccx http://www.bbc.co.uk/news/magazine-35525566

cccxi https://medium.com/basic-income/wouldnt-unconditional-basic-income-just-cause-massive-inflation-fe71d69f15e7#.3yezsngej

cccxii http://streamhistory.com/die-rich-die-disgraced-andrew-carnegies-philosophy-of-wealth/

cccxiii http://www.forbes.com/sites/greatspeculations/2012/12/05/how-i-know-higher-taxes-would-be-good-for-the-economy/#5b0c080b3ec1

cccxiv http://taxfoundation.org/article/what-evidence-taxes-and-growth

cccxv https://en.wikipedia.org/wiki/Laffer_curve

cccxvi http://www.bbc.co.uk/news/uk-politics-26875420

cccxvii A minor character in Shakespeare's Henry VI called Dick the Butcher has the memorable line, "First thing we do, let's kill all the lawyers." It seems Shakespeare was not fond of lawyers: http://www.spectacle.org/797/finkel.html

cccxviii https://www.thersa.org/action-and-research/rsa-projects/public-services-and-communities-folder/basic-income/

cccxix http://www.icalculator.info/news/UK_average_earnings_2014.html

cccxx http://www.telegraph.co.uk/finance/economics/12037623/Paying-all-UK-citizens-155-a-week-may-be-an-idea-whose-time-has-come.html

cccxxi http://www.marketwatch.com/story/most-americans-are-one-paycheck-away-from-the-street-2015-01-07

cccxxii http://www.federalreserve.gov/econresdata/2014-economic-well-being-of-us-households-in-2013-executive-summary.htm

cccxxiii http://www.theguardian.com/business/2016/jan/18/richest-62-billionaires-wealthy-half-world-population-combined

cccxxiv http://www.bbc.co.uk/news/magazine-26613682

cccxxv I'm indebted to Dr Justin Stewart, an investor, for prodding me to address the issue of assets more closely.

cccxxvi http://timharford.com/2016/05/could-an-income-for-all-provide-the-ultimate-safety-net/

cccxxviiIn case you only recently arrived on this planet, that was a reference to the sainted Douglas Adam's "Hitchhiker's Guide to the Galaxy" series. If you haven't read it, I recommend that you put this book down and read that one instead. I won't be offended. But please come back here afterwards.

cccxxviiihttp://philpapers.org/archive/DANHAT.pdf

cccxxixThe novel is sometimes said to have originated in the early 18th century, but in fact it is a much older art form. What happened then was that writers began publishing books which described life as they actually saw it. https://en.wikipedia.org/wiki/Novel#18th_century_novel

cccxxx I am indebted to AGI researcher Randal Koene for this observation.

cccxxxihttps://en.wikiquote.org/wiki/Bette_Davis

cccxxxiihttp://www.economist.com/node/17722567

cccxxxiiihttp://www.wired.com/2016/02/vr-moral-imperative-or-opiate-of-masses/

cccxxxiv See chapter 3.1

cccxxxvhttp://heather.cs.ucdavis.edu/JIntMigr.pdf

cccxxxvihttp://philosophicaldisquisitions.blogspot.co.uk/2014/01/rule-by-algorithm-big-data-and-threat.html

cccxxxviihttps://www.ted.com/talks/yuval_noah_harari_what_explains_the_rise_of_humans/transcript?language=en

cccxxxviiihttp://motherboard.vice.com/read/sleep-tech-will-widen-the-gap-between-the-rich-and-the-poor

cccxxxix Covered in detail in my previous book, "Surviving AI".

cccxlhttps://en.wikipedia.org/wiki/Sex_and_drugs_and_rock_and_roll

cccxliI am that terrible old cliché: a socialist student whose left-wing views did not long survive contact with the world of work. As a trainee BBC journalist writing about Central and Eastern Europe long before the Berlin Wall fell, I soon realised how fortunate I was to have grown up in the capitalist West. I didn't expect to be heading back in the other direction in later life.

cccxliihttps://edge.org/conversation/john_markoff-the-next-wave

cccxliiihttp://uk.pcmag.com/robotics-automation-products/34778/news/will-a-robot-revolution-lead-to-mass-unemployment

cccxlivhttps://en.wikipedia.org/wiki/Milgram_experiment

cccxlvhttp://www.prisonexp.org/

cccxlvihttp://fourhourworkweek.com/2014/08/29/kevin-kelly/

cccxlviihttps://www.edge.org/conversation/kevin_kelly-the-technium

cccxlviii http://history.hanover.edu/courses/excerpts/165acton.html

cccxlixhttp://mercatus.org/sites/default/files/Brito_BitcoinPrimer.pdf

ccclhttp://www.dugcampbell.com/byzantine-generals-problem/

ccclihttp://www.economistinsights.com/technology-innovation/analysis/money-no-middleman/tab/1

ccclii: The Machine Intelligence Research Institute (MIRI) in Northern California, The Future of Humanity Institute (FHI) and the Centre for the Study of Existential Risk (CSER) in England's Oxford and Cambridge respectively, and the Future of Life Institute (FLI) in Massachussetts.

ccliii The Oxford Martin Programme on Technology and Employment, part of Oxford University and part-funded by Citibank. There have been numerous less permanent initiatives, such as the 2015-16 research programme of Fundacion Bankinter, one of Spain's leading banks. (Disclosure: I am involved in that one.)

ccclivhttp://archive.fortune.com/magazines/fortune/fortune_archive/2003/09/15/349149/index.htm

ccclvPaul Boyer, 'Dr. Strangelove', a chapter in "Past Imperfect: History According to the Movies" edited by Mark C. Carnes

ccclvihttp://s05.static-shell.com/content/dam/shell/static/public/downloads/brochures/corporate-pkg/scenarios/explorers-guide.pdf

ccclviihttp://www.jfklibrary.org/Research/Research-Aids/Ready-Reference/JFK-Quotations.aspx