

VINTAGE

Contents

Introduction

PART I The Technological Challenge

1 Disillusionment

2 Work

3 Liberty

4 Equality

PART II The Political Challenge

5 Community

6 Civilisation

7 Nationalism

8 Religion

9 Immigration

PART III Despair and Hope

10 Terrorism

11 War

12 Humility

13 God

14 Secularism

PART IV Truth

15 Ignorance

16 Justice

17 Post-Truth

18 Science Fiction

PART V Resilience

19 Education

20 Meaning

21 Meditation

Acknowledgements

Notes

Index

About the Author



Yuval Noah Harari has a PhD in History from the University of Oxford and now lectures at the Hebrew University of Jerusalem, specialising in world history. His books, *Sapiens*, *Homo Deus* and *21 Lessons for the 21st Century* have become an international phenomenon.

ALSO BY YUVAL NOAH HARARI

Sapiens: A Brief History of Humankind
Homo Deus: A Brief History of Tomorrow

Dedicated with love to my husband Itzik, for his unbounded trust and brilliance; to my mother Pnina, for always caring and supporting; and to my grandmother Fanny, for her inexhaustible and selfless joy.

Introduction

In a world deluged by irrelevant information, clarity is power. In theory, anybody can join the debate about the future of humanity, but it is so hard to maintain a clear vision. Frequently, we don't even notice that a debate is going on, or what the key questions are. Billions of us can hardly afford the luxury of investigating, because we have more pressing things to do: we have to go to work, take care of the kids, or look after elderly parents. Unfortunately, history gives no discounts. If the future of humanity is decided in your absence, because you are too busy feeding and clothing your kids – you and they will not be exempt from the consequences. This is very unfair; but who said history was fair?

As a historian, I cannot give people food or clothes – but I can try and offer some clarity, thereby helping to level the global playing field. If this empowers even a handful of additional people to join the debate about the future of our species, I have done my job.

My first book, *Sapiens*, surveyed the human past, examining how an insignificant ape became the ruler of planet Earth.

Homo Deus, my second book, explored the long-term future of life, contemplating how humans might eventually become gods, and what might be the ultimate destiny of intelligence and consciousness.

In this book I want to zoom in on the here and now, but without losing the long-term perspective. How can insights about the distant past and distant future help us make sense of current affairs and of the immediate dilemmas of human societies? What is happening right now? What are today's greatest challenges and choices?

Of course, 7 billion people have 7 billion agendas, and as already noted, thinking about the big picture is a relatively rare luxury. A single mother struggling to raise two children in a Mumbai slum is focused on the next meal; refugees in a boat in the middle of the Mediterranean scan the horizon for any sign of land; and a dying man in an overcrowded London hospital gathers all his remaining strength to take in one more breath. They all have far more urgent problems than global warming or the crisis of liberal democracy. No book can do justice to all of that, and I don't have lessons to teach people in such situations. I can only hope to learn from them.

My agenda here is global. I look at the major forces that shape societies all over the world, and that are likely to influence the future of our planet as a whole. Climate change may be far beyond the concerns of people in the midst of a life-and-death emergency, but it might eventually make the Mumbai slums uninhabitable, send enormous new waves of refugees across the Mediterranean, and lead to a worldwide crisis in healthcare.

Reality is composed of many threads, and this book tries to cover different aspects of our global predicament, without claiming to be exhaustive. Unlike *Sapiens* and *Homo Deus*, this book is not intended as a historical narrative, but rather as a selection of lessons. These lessons do not conclude with simple answers. They aim to stimulate

further thinking, and help readers participate in some of the major conversations of our time.

The book was actually written in conversation with the public. Many of the chapters were composed in response to questions I was asked by readers, journalists and colleagues. Earlier versions of some segments were already published in different forms, which gave me the opportunity to receive feedback and hone my arguments. Some sections focus on technology, some on politics, some on religion, and some on art. Certain chapters celebrate human wisdom, others highlight the crucial role of human stupidity. But the overarching question remains the same: what is happening in the world today, and what is the deep meaning of events?

What does the rise of Donald Trump signify? What can we do about the epidemic of fake news? Why is liberal democracy in crisis? Is God back? Is a new world war coming? Which civilisation dominates the world – the West, China, Islam? Should Europe keep its doors open to immigrants? Can nationalism solve the problems of inequality and climate change? What should we do about terrorism?

Though this book takes a global perspective, I do not neglect the personal level. On the contrary, I want to emphasise the connections between the great revolutions of our era and the internal lives of individuals. For example, terrorism is both a global political problem and an internal psychological mechanism. Terrorism works by pressing the fear button deep in our minds and hijacking the private imagination of millions of individuals. Similarly, the crisis of liberal democracy is played out not just in parliaments and polling stations, but also in neurons and synapses. It is a cliché to note that the personal is the political. But in an era when scientists, corporations and governments are learning to hack the human brain, this truism is more sinister than ever. Accordingly, this book offers observations about the conduct of individuals as well as entire societies.

A global world puts unprecedented pressure on our personal conduct and morality. Each of us is ensnared within numerous all-encompassing spider webs, which on the one hand restrict our movements, but at the same time transmit our tiniest jiggle to faraway destinations. Our daily routines influence the lives of people and animals halfway across the world, and some personal gestures can unexpectedly set the entire world ablaze, as happened with the self-immolation of Mohamed Bouazizi in Tunisia, which ignited the Arab Spring, and with the women who shared their stories of sexual harassment and sparked the #MeToo movement.

This global dimension of our personal lives means that it is more important than ever to uncover our religious and political biases, our racial and gender privileges, and our unwitting complicity in institutional oppression. But is that a realistic enterprise? How can I find a firm ethical ground in a world that extends far beyond my horizons, that spins completely out of human control, and that holds all gods and ideologies suspect?

The book begins by surveying the current political and technological predicament. At the close of the twentieth century it appeared that the great ideological battles between fascism, communism and liberalism resulted in the overwhelming victory of liberalism. Democratic politics, human rights and free-market capitalism seemed destined to conquer the entire world. But as usual, history took an unexpected turn, and after fascism and communism collapsed, now liberalism is in a jam. So where are we heading?

PART I

The Technological Challenge

Humankind is losing faith in the liberal story that dominated global politics in recent decades, exactly when the merger of biotech and infotech confronts us with the biggest challenges humankind has ever encountered.

Disillusionment

The end of history has been postponed

Humans think in stories rather than in facts, numbers or equations, and the simpler the story, the better. Every person, group and nation has its own tales and myths. But during the twentieth century the global elites in New York, London, Berlin and Moscow formulated three grand stories that claimed to explain the whole past and to predict the future of the entire world: the fascist story, the communist story and the liberal story.

The fascist story explained history as a struggle among different nations, and envisioned a world dominated by one human group that violently subdues all others. The communist story explained history as a struggle among different classes, and envisioned a world in which all groups are united by a centralised social system that ensures equality even at the price of freedom. The liberal story explained history as a struggle between liberty and tyranny, and envisioned a world in which all humans cooperate freely and peacefully, with minimum central control even at the price of some inequality.

The conflict among these three stories reached its first critical peak in the Second World War, which knocked out the fascist story. From the late 1940s to the late 1980s the world became a battleground between the two remaining stories: communism and liberalism. Then the communist story collapsed, and the liberal story was left the dominant guide to the human past and the indispensable manual for the future of the world – or so it seemed to the global elite.

The liberal story celebrates the value and power of liberty. It says that for thousands of years humankind lived under oppressive regimes which allowed people few political rights, economic opportunities or personal liberties, and which heavily restricted the movements of individuals, ideas and goods. But people fought for their freedom, and step by step, liberty gained ground. Democratic regimes took the place of brutal dictatorships. Free enterprise overcame economic restrictions. People learned to think for themselves and follow their hearts, instead of blindly obeying bigoted priests and hidebound traditions. Open roads, stout bridges and bustling airports replaced walls, moats and barbed-wire fences.

The liberal story acknowledges that not all is well in the world, and that there are still many hurdles to overcome. Much of our planet is dominated by tyrants, and even in the most liberal countries many citizens suffer from poverty, violence and oppression. But at least we know what we need to do in order to overcome these problems: give people more liberty. We need to protect human rights, to grant

everybody the vote, to establish free markets, and to let individuals, ideas and goods move throughout the world as easily as possible. According to this liberal panacea – accepted, in slight variations, by George W. Bush and Barack Obama alike – if we just continue to liberalise and globalise our political and economic systems, we will produce peace and prosperity for all.¹

Countries that join this unstoppable march of progress will be rewarded with peace and prosperity sooner. Countries that try to resist the inevitable will suffer the consequences, until they too see the light, open their borders and liberalise their societies, their politics and their markets. It may take time, but eventually even North Korea, Iraq and El Salvador will look like Denmark or Iowa.

In the 1990s and 2000s this story became a global mantra. Many governments from Brazil to India adopted liberal recipes in an attempt to join the inexorable march of history. Those failing to do so seemed like fossils from a bygone era. In 1997, US President Bill Clinton confidently rebuked the Chinese government that its refusal to liberalise Chinese politics puts it ‘on the wrong side of history’.²

However, since the global financial crisis of 2008 people all over the world have become increasingly disillusioned with the liberal story. Walls and firewalls are back in vogue. Resistance to immigration and to trade agreements is mounting. Ostensibly democratic governments undermine the independence of the judicial system, restrict the freedom of the press, and portray any opposition as treason. Strongmen in countries such as Turkey and Russia experiment with new types of illiberal democracies and downright dictatorships. Today, few would confidently declare that the Chinese Communist Party is on the wrong side of history.

The year 2016 – marked by the Brexit vote in Britain and the rise of Donald Trump in the United States – signified the moment when this tidal wave of disillusionment reached the core liberal states of western Europe and North America. Whereas a few years ago Americans and Europeans were still trying to liberalise Iraq and Libya at the point of the gun, many people in Kentucky and Yorkshire have now come to see the liberal vision as either undesirable or unattainable. Some discovered a liking for the old hierarchical world, and they just don’t want to give up their racial, national or gendered privileges. Others have concluded (rightly or wrongly) that liberalisation and globalisation are a huge racket empowering a tiny elite at the expense of the masses.

In 1939 humans were offered three global stories to choose from, in 1969 just two, in 1999 a single story seemed to prevail; in 2019 we are down to zero. No wonder that the liberal elites, who dominated much of the world in recent decades, have entered a state of shock and disorientation. To have one story is the most reassuring situation of all. Everything is perfectly clear. To be suddenly left without any story is terrifying. Nothing makes any sense.

From killing mosquitoes to killing thoughts

The mood of disorientation and impending doom is exacerbated by the accelerating pace of technological disruption. The liberal political system has been shaped during the industrial era to manage a world of steam engines, oil refineries and television sets. It finds it difficult to deal with the ongoing revolutions in information technology and biotechnology.

Both politicians and voters are barely able to comprehend the new technologies, let alone regulate their explosive potential. Since the 1990s the Internet has changed the world probably more than any other factor, yet the Internet revolution was directed by

engineers more than by political parties. Did you ever vote about the Internet? The democratic system is still struggling to understand what hit it, and is hardly equipped to deal with the next shocks, such as the rise of AI and the blockchain revolution.

Already today, computers have made the financial system so complicated that few humans can understand it. As AI improves, we might soon reach a point when no human can make sense of finance any more. What will that do to the political process? Can you imagine a government that waits humbly for an algorithm to approve its budget or its new tax reform? Meanwhile peer-to-peer blockchain networks and cryptocurrencies like bitcoin might completely revamp the monetary system, so that radical tax reforms will be inevitable. For example, it might become impossible or irrelevant to tax dollars, because most transactions will involve only the exchange of information, without a clear-cut transfer of national currency, or any currency at all. Governments might therefore need to invent entirely new taxes – perhaps an information tax, paid on information and *in information* rather than in dollars. Will the political system manage to deal with the crisis before it runs out of money?

Even more importantly, the twin revolutions in infotech and biotech could restructure not just economies and societies but our very bodies and minds. In the past, we humans have learned to control the world outside us, but we had very little control over the world inside us. We knew how to build a dam and stop a river from flowing, but we did not know how to stop the body from ageing. We knew how to design an irrigation system, but we had no idea how to design a brain. If mosquitoes buzzed in our ears and disturbed our sleep, we knew how to kill the mosquitoes; but if a thought buzzed in our mind and kept us awake at night, most of us did not know how to kill the thought.

The revolutions in biotech and infotech will give us control of the world inside us, and will enable us to engineer and manufacture life. We will learn how to design brains, extend lives and kill thoughts at our discretion. Nobody knows what the consequences will be. Humans were always far better at inventing tools than using them wisely. It is easier to manipulate a river by building a dam across it than it is to predict all the complex consequences this will have for the wider ecological system. Similarly, it will be easier to redirect the flow of our minds than to divine what it will do to our personal psychology or to our social systems.

In the past, we have gained the power to manipulate the world around us and to reshape the entire planet, but because we didn't understand the complexity of the global ecology, the changes we made inadvertently disrupted the entire ecological system and now we face an ecological collapse. In the coming century biotech and infotech will give us the power to manipulate the world inside us and reshape ourselves, but because we don't understand the complexity of our own minds, the changes we will make might upset our mental system to such an extent that it too might break down.

The revolutions in biotech and infotech are made by engineers, entrepreneurs and scientists who are hardly aware of the political implications of their decisions, and who certainly don't represent anyone. Can parliaments and parties take matters into their own hands? At present, it does not seem so. Technological disruption is not even a leading item on the political agenda. Thus during the 2016 US presidential race, the main reference to disruptive technology concerned Hillary Clinton's email debacle,³ and despite all the talk about job losses, neither candidate addressed the potential impact of automation. Donald Trump warned voters that the Mexicans and Chinese will take their jobs, and that they should therefore build a wall on the Mexican

border.⁴ He never warned voters that the algorithms will take their jobs, nor did he suggest building a firewall on the border with California.

This might be one of the reasons (though not the only one) why even voters in the heartlands of the liberal West are losing faith in the liberal story and in the democratic process. Ordinary people may not understand artificial intelligence and biotechnology, but they can sense that the future is passing them by. In 1938 the condition of the common person in the USSR, Germany or the USA may have been grim, but he was constantly told that he was the most important thing in the world, and that he was the future (provided, of course, that he was an 'ordinary person' rather than a Jew or an African). He looked at the propaganda posters – which typically depicted coal miners, steelworkers and housewives in heroic poses – and saw himself there: 'I am in that poster! I am the hero of the future!'⁵

In 2018 the common person feels increasingly irrelevant. Lots of mysterious words are bandied around excitedly in TED talks, government think tanks and hi-tech conferences – globalisation, blockchain, genetic engineering, artificial intelligence, machine learning – and common people may well suspect that none of these words are about them. The liberal story was the story of ordinary people. How can it remain relevant to a world of cyborgs and networked algorithms?

In the twentieth century, the masses revolted against exploitation, and sought to translate their vital role in the economy into political power. Now the masses fear irrelevance, and they are frantic to use their remaining political power before it is too late. Brexit and the rise of Trump might thus demonstrate an opposite trajectory to that of traditional socialist revolutions. The Russian, Chinese and Cuban revolutions were made by people who were vital for the economy, but who lacked political power; in 2016, Trump and Brexit were supported by many people who still enjoyed political power, but who feared that they were losing their economic worth. Perhaps in the twenty-first century populist revolts will be staged not against an economic elite that exploits people, but against an economic elite that does not need them any more.⁶ This may well be a losing battle. It is much harder to struggle against irrelevance than against exploitation.

The liberal phoenix

This is not the first time the liberal story has faced a crisis of confidence. Ever since this story gained global influence, in the second half of the nineteenth century, it has endured periodic crises. The first era of globalisation and liberalisation ended in the bloodbath of the First World War. In the 1930s and early 1940s fascism looked irresistible. From the 1950s to the 1970s it seemed that the future belonged to communism. But every time, liberalism managed to overcome the crisis, learn from its mistakes, and even adopt its rivals' best ideas. In particular, the liberal story learned from communism to expand the circle of empathy and to value equality alongside liberty. Whereas in the late nineteenth century the liberal story cared mainly about the liberties and privileges of middle-class European men, by the late twentieth century it had come to champion the cause of working-class people, women, minorities and non-Westerners.⁷

Yet the current crisis of liberalism is different from the crises of the 1930s and the 1960s. The three major movements that fought one another in the twentieth century all had a coherent vision for the entire human species – be it global domination, revolution or liberation. Today, in contrast, the people who reject the liberal story

economic integration goes hand-in-hand with the freedom of individual consumers – if I can choose from a hundred global brands rather than just three national brands, I enjoy more individual freedom. And so forth. Consequently, if a country wants to enjoy one dish from the liberal set menu – such as economic liberalisation – it has no choice but to take the other dishes too.

What is now common to populist and nationalist movements throughout the world is that even if they describe themselves as ‘anti-liberal’, none of them rejects liberalism wholesale. Rather, they reject the set-menu approach, and want to pick and choose their own dishes from a liberal buffet. Thus Trump is still very much in favour of free markets and privatisation, but he thinks he can have these goods while undermining multilateral cooperation and even free trade. China supports free trade, and its Belt and Road Initiative is one of the most ambitious globalisation projects ever envisioned – but it is far less enthusiastic about free elections. The British Brexiteers uphold democracy and have nothing against individualism, but they don’t like the idea of multilateral cooperation and of giving too much power to international organisations. Viktor Orbán has defined his regime as an ‘illiberal democracy’, arguing that you can have free elections in Hungary without committing yourself to things such as minority rights, diversity and individualism.

The one dish that almost everybody wants, at least in theory, is peaceful international relations. This is the chocolate cake of the liberal buffet. In contrast, the one dish that almost nobody desires – the global celery – is immigration. Even some of the staunchest supporters of democracy, individualism and multilateral cooperation have become decidedly lukewarm about allowing in too many immigrants.

It remains to be seen, however, whether the buffet approach can work. The food analogy might be very misleading. In restaurants, set menus are an arbitrary assemblage of independent dishes. Yet the liberal story always insisted that the liberal system is a living organism made of mutually dependent organs. While you can easily separate the soup from the dessert, you cannot separate the heart from the lungs. Can Trump really promote free markets in the USA while undermining free trade on the global level? Can the Chinese Communist Party continue to enjoy the fruits of economic liberalisation without making any movement towards political liberalisation? Can Hungarians have democracy without personal liberties, or is Orbán’s ‘illiberal democracy’ just a nicer way of saying ‘dictatorship’? Can international peace survive in a world of rising border walls and intensifying trade wars? The buffet approach might well result in the utter breakdown of the liberal system on both the national and international levels.

If this happens, what alternative vision might replace the liberal story? One option might be to completely give up on having a global story of any kind, and instead seek shelter with local nationalist and religious tales. In the twentieth century, nationalist movements were an extremely important political player, but they lacked a coherent vision for the future of the world other than supporting the division of the globe into independent nation states. Thus Indonesian nationalists fought against Dutch domination, and Vietnamese nationalists wanted a free Vietnam, but there was no Indonesian or Vietnamese story for humanity as a whole. When it came time to explain how Indonesia, Vietnam and all the other free nations should relate to one another, and how humans should deal with global problems such as the threat of nuclear war, nationalists invariably turned to either liberal or communist ideas.

But if both liberalism and communism are now discredited, maybe humans should abandon the very idea of a single global story? After all, weren’t all these global stories

– even communism – the product of Western imperialism? Why should Vietnamese villagers put their faith in the brainchild of a German from Trier and a Manchester industrialist? Maybe each country should adopt a different idiosyncratic path, defined by its own ancient traditions? Perhaps even Westerners should take a break from trying to run the world, and focus on their own affairs for a change?

This is arguably already happening all over the globe, as the vacuum left by the breakdown of liberalism is tentatively filled by nostalgic fantasies about some local golden past. Donald Trump coupled his calls for American isolationism with a promise to ‘Make America Great Again’ – as if the USA of the 1980s or 1950s was a perfect society that Americans should somehow recreate in the twenty-first century. The Brexiteers dream of making Britain an independent power, as if they were still living in the days of Queen Victoria and as if ‘splendid isolation’ were a viable policy for the era of the Internet and global warming. Chinese elites have rediscovered their native imperial and Confucian legacies, as a supplement or even substitute for the doubtful Marxist ideology they imported from the West. In Russia, Putin’s official vision is not to build a corrupt oligarchy, but rather to resurrect the old tsarist empire. A century after the Bolshevik Revolution, Putin promises a return to ancient tsarist glories with an autocratic government buoyed by Russian nationalism and Orthodox piety spreading its might from the Baltic to the Caucasus.

Similar nostalgic dreams that mix nationalist attachment with religious traditions underpin regimes in India, Poland, Turkey and numerous other countries. Nowhere are these fantasies more extreme than in the Middle East, where Islamists want to copy the system established by the Prophet Muhammad in the city of Medina 1,400 years ago, while fundamentalist Jews in Israel outdo even the Islamists, and dream of going back 2,500 years to biblical times. Members of Israel’s ruling coalition government talk openly about their hope of expanding modern Israel’s borders to match more closely those of biblical Israel, of reinstating biblical law, and even of rebuilding the ancient Temple of Yahweh in Jerusalem in place of the Al-Aqsa mosque.¹⁰

Liberal elites look in horror at these developments, and hope that humanity will return to the liberal path in time to avert disaster. In his final speech to the United Nations in September 2016, President Obama warned his listeners against retreating ‘into a world sharply divided, and ultimately in conflict, along age-old lines of nation and tribe and race and religion’. Instead, he said, ‘the principles of open markets and accountable governance, of democracy and human rights and international law ... remain the firmest foundation for human progress in this century’.¹¹

Obama has rightly pointed out that despite the numerous shortcomings of the liberal package, it has a much better record than any of its alternatives. Most humans never enjoyed greater peace or prosperity than they did under the aegis of the liberal order of the early twenty-first century. For the first time in history, infectious diseases kill fewer people than old age, famine kills fewer people than obesity, and violence kills fewer people than accidents.

But liberalism has no obvious answers to the biggest problems we face: ecological collapse and technological disruption. Liberalism traditionally relied on economic growth to magically solve difficult social and political conflicts. Liberalism reconciled the proletariat with the bourgeoisie, the faithful with the atheists, the natives with the immigrants, and the Europeans with the Asians by promising everybody a larger slice of the pie. With a constantly growing pie, that was possible. However, economic growth will not save the global ecosystem – just the opposite, it is the cause of the

ecological crisis. And economic growth will not solve technological disruption – it is predicated on the invention of more and more disruptive technologies.

The liberal story and the logic of free-market capitalism encourage people to have grand expectations. During the latter part of the twentieth century, each generation – whether in Houston, Shanghai, Istanbul or São Paulo – enjoyed better education, superior healthcare and larger incomes than the one that came before it. In coming decades, however, owing to a combination of technological disruption and ecological meltdown, the younger generation might be lucky to just stay in place.

We are consequently left with the task of creating an updated story for the world. Just as the upheavals of the Industrial Revolution gave birth to the novel ideologies of the twentieth century, so the coming revolutions in biotechnology and information technology are likely to require fresh visions. The next decades might therefore be characterised by intense soul-searching and by formulating new social and political models. Could liberalism reinvent itself yet again, just as it did in the wake of the 1930s and 1960s crises, emerging as more attractive than ever before? Could traditional religion and nationalism provide the answers that escape the liberals, and could they use ancient wisdom to fashion an up-to-date world view? Or perhaps the time has come to make a clean break with the past, and craft a completely new story that goes beyond not just the old gods and nations, but even the core modern values of liberty and equality?

At present, humankind is far from reaching any consensus on these questions. We are still in the nihilist moment of disillusionment and anger, after people have lost faith in the old stories but before they have embraced a new one. So what next? The first step is to tone down the prophecies of doom, and switch from panic mode to bewilderment. Panic is a form of hubris. It comes from the smug feeling that I know exactly where the world is heading – down. Bewilderment is more humble, and therefore more clear-sighted. If you feel like running down the street crying ‘The apocalypse is upon us!’, try telling yourself ‘No, it’s not that. Truth is, I just don’t understand what’s going on in the world.’

The following chapters will try to clarify some of the bewildering new possibilities we face, and how we might proceed from here. But before exploring potential solutions to humanity’s predicaments we need a better grasp of the challenge technology poses. The revolutions in information technology and biotechnology are still in their infancy, and it is debatable to what extent they are really responsible for the current crisis of liberalism. Most people in Birmingham, Istanbul, St Petersburg and Mumbai are only dimly aware, if at all, of the rise of artificial intelligence and its potential impact on their lives. It is undoubtable, however, that the technological revolutions will gather momentum in the next few decades, and will confront humankind with the hardest trials we have ever encountered. Any story that seeks to gain humanity’s allegiance will be tested above all in its ability to deal with the twin revolutions in infotech and biotech. If liberalism, nationalism, Islam or some novel creed wishes to shape the world of the year 2050, it will need not only to make sense of artificial intelligence, Big Data algorithms and bioengineering – it will also need to incorporate them into a new meaningful narrative.

To understand the nature of this technological challenge, perhaps it would be best to start with the job market. Since 2015 I have been travelling around the world talking with government officials, business people, social activists and schoolkids about the human predicament. Whenever they become impatient or bored by all the talk of artificial intelligence, Big Data algorithms and bioengineering, I usually need to

mention just one magic word to snap them back to attention: jobs. The technological revolution might soon push billions of humans out of the job market, and create a massive new useless class, leading to social and political upheavals that no existing ideology knows how to handle. All the talk about technology and ideology might sound abstract and remote, but the very real prospect of mass unemployment – or personal unemployment – leaves nobody indifferent.

Work

When you grow up, you might not have a job

We have no idea what the job market will look like in 2050. It is generally agreed that machine learning and robotics will change almost every line of work – from producing yoghurt to teaching yoga. However, there are conflicting views about the nature of the change and its imminence. Some believe that within a mere decade or two, billions of people will become economically redundant. Others maintain that even in the long run automation will keep generating new jobs and greater prosperity for all.

So are we on the verge of a terrifying upheaval, or are such forecasts yet another example of ill-founded Luddite hysteria? It is hard to say. Fears that automation will create massive unemployment go back to the nineteenth century, and so far they have never materialised. Since the beginning of the Industrial Revolution, for every job lost to a machine at least one new job was created, and the average standard of living has increased dramatically.¹ Yet there are good reasons to think that this time it is different, and that machine learning will be a real game changer.

Humans have two types of abilities – physical and cognitive. In the past, machines competed with humans mainly in raw physical abilities, while humans retained an immense edge over machines in cognition. Hence as manual jobs in agriculture and industry were automated, new service jobs emerged that required the kind of cognitive skills only humans possessed: learning, analysing, communicating and above all understanding human emotions. However, AI is now beginning to outperform humans in more and more of these skills, including in the understanding of human emotions.² We don't know of any third field of activity – beyond the physical and the cognitive – where humans will always retain a secure edge.

It is crucial to realise that the AI revolution is not just about computers getting faster and smarter. It is fuelled by breakthroughs in the life sciences and the social sciences as well. The better we understand the biochemical mechanisms that underpin human emotions, desires and choices, the better computers can become in analysing human behaviour, predicting human decisions, and replacing human drivers, bankers and lawyers.

In the last few decades research in areas such as neuroscience and behavioural economics allowed scientists to hack humans, and in particular to gain a much better understanding of how humans make decisions. It turned out that our choices of everything from food to mates result not from some mysterious free will, but rather from billions of neurons calculating probabilities within a split second. Vaunted 'human intuition' is in reality 'pattern recognition'.³ Good drivers, bankers and

replacing all human drivers by computers is expected to reduce deaths and injuries on the road by about 90 per cent.⁸ In other words, switching to autonomous vehicles is likely to save the lives of a million people every year.

Hence it would be madness to block automation in fields such as transport and healthcare just in order to protect human jobs. After all, what we ultimately ought to protect is humans – not jobs. Redundant drivers and doctors will just have to find something else to do.

The Mozart in the machine

At least in the short term, AI and robotics are unlikely to completely eliminate entire industries. Jobs that require specialisation in a narrow range of routinised activities will be automated. But it will be much more difficult to replace humans with machines in less routine jobs that demand the simultaneous use of a wide range of skills, and that involve dealing with unforeseen scenarios. Take healthcare, for example. Many doctors focus almost exclusively on processing information: they absorb medical data, analyse it and produce a diagnosis. Nurses, in contrast, also need good motor and emotional skills in order to give a painful injection, replace a bandage or restrain a violent patient. Hence we will probably have an AI family doctor on our smartphone decades before we have a reliable nurse robot.⁹ The human care industry – which takes care of the sick, the young and the elderly – is likely to remain a human bastion for a long time. Indeed, as people live longer and have fewer children, care of the elderly will probably be one of the fastest-growing sectors in the human labour market.

Alongside care, creativity too poses particularly difficult hurdles for automation. We don't need humans to sell us music any more – we can download it directly from the iTunes Store – but the composers, musicians, singers and DJs are still flesh and blood. We rely on their creativity not just to produce completely new music, but also to choose among a mind-boggling range of available possibilities.

Nevertheless, in the long run no job will remain absolutely safe from automation. Even artists should be put on notice. In the modern world art is usually associated with human emotions. We tend to think that artists are channelling internal psychological forces, and that the whole purpose of art is to connect us with our emotions or to inspire in us some new feeling. Consequently, when we come to evaluate art, we tend to judge it by its emotional impact on the audience. Yet if art is defined by human emotions, what might happen once external algorithms are able to understand and manipulate human emotions better than Shakespeare, Frida Kahlo or Beyoncé?

After all, emotions are not some mystical phenomenon – they are the result of a biochemical process. Hence, in the not too distant future a machine-learning algorithm could analyse the biometric data streaming from sensors on and inside your body, determine your personality type and your changing moods, and calculate the emotional impact that a particular song – even a particular musical key – is likely to have on you.¹⁰

Of all forms of art, music is probably the most susceptible to Big Data analysis, because both inputs and outputs lend themselves to precise mathematical depiction. The inputs are the mathematical patterns of sound waves, and the outputs are the electrochemical patterns of neural storms. Within a few decades, an algorithm that goes over millions of musical experiences might learn to predict how particular inputs result in particular outputs.¹¹

Suppose you just had a nasty fight with your boyfriend. The algorithm in charge of your sound system will immediately discern your inner emotional turmoil, and based on what it knows about you personally and about human psychology in general, it will play songs tailored to resonate with your gloom and echo your distress. These particular songs might not work well with other people, but are just perfect for your personality type. After helping you get in touch with the depths of your sadness, the algorithm would then play the one song in the world that is likely to cheer you up – perhaps because your subconscious connects it with a happy childhood memory that even you are not aware of. No human DJ could ever hope to match the skills of such an AI.

You might object that the AI would thereby kill serendipity and lock us inside a narrow musical cocoon, woven by our previous likes and dislikes. What about exploring new musical tastes and styles? No problem. You could easily adjust the algorithm to make 5 per cent of its choices completely at random, unexpectedly throwing at you a recording of an Indonesian gamelan ensemble, a Rossini opera, or the latest K-pop hit. Over time, by monitoring your reactions, the AI could even determine the ideal level of randomness that will optimise exploration while avoiding annoyance, perhaps lowering its serendipity level to 3 per cent or raising it to 8 per cent.

Another possible objection is that it is unclear how the algorithm could establish its emotional goal. If you just fought with your boyfriend, should the algorithm aim to make you sad or joyful? Would it blindly follow a rigid scale of ‘good’ emotions and ‘bad’ emotions? Maybe there are times in life when it is good to feel sad? The same question, of course, could be directed at human musicians and DJs. Yet with an algorithm, there are many interesting solutions to this puzzle.

One option is to just leave it to the customer. You can evaluate your emotions whichever way you like, and the algorithm will follow your dictates. Whether you want to wallow in self-pity or jump for joy, the algorithm will slavishly follow your lead. Indeed, the algorithm may learn to recognise your wishes even without you being explicitly aware of them.

Alternatively, if you don’t trust yourself, you can instruct the algorithm to follow the recommendation of whichever eminent psychologist you do trust. If your boyfriend eventually dumps you, the algorithm may walk you through the official five stages of grief, first helping you deny what happened by playing Bobby McFerrin’s ‘Don’t Worry, Be Happy’, then whipping up your anger with Alanis Morissette’s ‘You Oughta Know’, encouraging you to bargain with Jacques Brel’s ‘Ne me quitte pas’ and Paul Young’s ‘Come Back and Stay’, dropping you into the pit of depression with Adele’s ‘Someone Like You’ and ‘Hello’, and finally aiding you to accept the situation with Gloria Gaynor’s ‘I Will Survive’.

The next step is for the algorithm to start tinkering with the songs and melodies themselves, changing them ever so slightly to fit your quirks. Perhaps you dislike a particular bit in an otherwise excellent song. The algorithm knows it because your heart skips a beat and your oxytocin levels drop slightly whenever you hear that annoying part. The algorithm could rewrite or edit out the offending notes.

In the long run, algorithms may learn how to compose entire tunes, playing on human emotions as if they were a piano keyboard. Using your biometric data the algorithms could even produce personalised melodies, which you alone in the entire universe would appreciate.

It is often said that people connect with art because they find themselves in it. This may lead to surprising and somewhat sinister results if and when, say, Facebook begins creating personalised art based on everything it knows about you. If your boyfriend leaves you, Facebook will treat you to an individualised song about that particular bastard rather than about the unknown person who broke the heart of Adele or Alanis Morissette. The song will even remind you of real incidents from your relationship, which nobody else in the world knows about.

Of course, personalised art might never catch on, because people will continue to prefer common hits that everybody likes. How can you dance or sing together to a tune nobody besides you knows? But algorithms could prove even more adept at producing global hits than personalised rarities. By using massive biometric databases garnered from millions of people, the algorithm could know which biochemical buttons to press in order to produce a global hit which would set everybody swinging like crazy on the dance floors. If art is really about inspiring (or manipulating) human emotions, few if any human musicians will have a chance of competing with such an algorithm, because they cannot match it in understanding the chief instrument they are playing on: the human biochemical system.

Will all this result in great art? That depends on the definition of art. If beauty is indeed in the ears of the listener, and if the customer is always right, then biometric algorithms stand a chance of producing the best art in history. If art is about something deeper than human emotions, and should express a truth beyond our biochemical vibrations, biometric algorithms might not make very good artists. But nor do most humans. In order to enter the art market and displace many human composers and performers, algorithms won't have to begin by straightaway surpassing Tchaikovsky. It will be enough if they outperform Britney Spears.

New jobs?

The loss of many traditional jobs in everything from art to healthcare will partly be offset by the creation of new human jobs. GPs who focus on diagnosing known diseases and administering familiar treatments will probably be replaced by AI doctors. But precisely because of that, there will be much more money to pay human doctors and lab assistants to do groundbreaking research and develop new medicines or surgical procedures.¹²

AI might help create new human jobs in another way. Instead of humans competing with AI, they could focus on servicing and leveraging AI. For example, the replacement of human pilots by drones has eliminated some jobs but created many new opportunities in maintenance, remote control, data analysis and cyber security. The US armed forces need thirty people to operate every unmanned Predator or Reaper drone flying over Syria, while analysing the resulting harvest of information occupies at least eighty people more. In 2015 the US Air Force lacked sufficient trained humans to fill all these positions, and therefore faced an ironic crisis in manning its unmanned aircraft.¹³

If so, the job market of 2050 might well be characterised by human–AI cooperation rather than competition. In fields ranging from policing to banking, teams of humans-plus-AIs could outperform both humans and computers. After IBM's chess program Deep Blue beat Garry Kasparov in 1997, humans did not stop playing chess. Rather, thanks to AI trainers human chess masters became better than ever, and at least for a while human–AI teams known as 'centaurs' outperformed both humans and computers

in chess. AI might similarly help groom the best detectives, bankers and soldiers in history.¹⁴

The problem with all such new jobs, however, is that they will probably demand high levels of expertise, and will therefore not solve the problems of unemployed unskilled labourers. Creating new human jobs might prove easier than retraining humans to actually fill these jobs. During previous waves of automation, people could usually switch from one routine low-skill job to another. In 1920 a farm worker laid off due to the mechanisation of agriculture could find a new job in a factory producing tractors. In 1980 an unemployed factory worker could start working as a cashier in a supermarket. Such occupational changes were feasible, because the move from the farm to the factory and from the factory to the supermarket required only limited retraining.

But in 2050, a cashier or textile worker losing their job to a robot will hardly be able to start working as a cancer researcher, as a drone operator or as part of a human-AI banking team. They will not have the necessary skills. In the First World War it made sense to send millions of raw conscripts to charge machine guns and die in their thousands. Their individual skills mattered little. Today, despite the shortage of drone operators and data analysts, the US Air Force is unwilling to fill the gaps with Walmart dropouts. You wouldn't like an inexperienced recruit to mistake an Afghan wedding party for a high-level Taliban conference.

Consequently, despite the appearance of many new human jobs, we might nevertheless witness the rise of a new 'useless' class. We might actually get the worst of both worlds, suffering simultaneously from high unemployment and a shortage of skilled labour. Many people might share the fate not of nineteenth-century wagon drivers – who switched to driving taxis – but of nineteenth-century horses, who were increasingly pushed out of the job market altogether.¹⁵

In addition, no remaining human job will ever be safe from the threat of future automation, because machine learning and robotics will continue to improve. A forty-year-old unemployed Walmart cashier who by dint of superhuman efforts manages to reinvent herself as a drone pilot might have to reinvent herself again ten years later, because by then the flying of drones may also have been automated. This volatility will also make it more difficult to organise unions or secure labour rights. Already today, many new jobs in advanced economies involve unprotected temporary work, freelancing and one-time gigs.¹⁶ How do you unionise a profession that mushrooms and disappears within a decade?

Similarly, human-computer centaur teams are likely to be characterised by a constant tug of war between the humans and the computers, instead of settling down to a lifelong partnership. Teams made exclusively of humans – such as Sherlock Holmes and Dr Watson – usually develop permanent hierarchies and routines that last decades. But a human detective who teams up with IBM's Watson computer system (which became famous after winning the US TV quiz show *Jeopardy!* in 2011) will find that every routine is an invitation for disruption, and every hierarchy an invitation for revolution. Yesterday's sidekick might morph into tomorrow's superintendent, and all protocols and manuals will have to be rewritten every year.¹⁷

A closer look at the world of chess might indicate where things are heading in the long run. It is true that for several years after Deep Blue defeated Kasparov, human-computer cooperation flourished in chess. Yet in recent years computers have become so good at playing chess that their human collaborators lost their value, and might soon become utterly irrelevant.

On 7 December 2017 a critical milestone was reached, not when a computer defeated a human at chess – that’s old news – but when Google’s AlphaZero program defeated the Stockfish 8 program. Stockfish 8 was the world’s computer chess champion for 2016. It had access to centuries of accumulated human experience in chess, as well as to decades of computer experience. It was able to calculate 70 million chess positions per second. In contrast, AlphaZero performed only 80,000 such calculations per second, and its human creators never taught it any chess strategies – not even standard openings. Rather, AlphaZero used the latest machine-learning principles to self-learn chess by playing against itself. Nevertheless, out of a hundred games the novice AlphaZero played against Stockfish, AlphaZero won twenty-eight and tied seventy-two. It didn’t lose even once. Since AlphaZero learned nothing from any human, many of its winning moves and strategies seemed unconventional to human eyes. They may well be considered creative, if not downright genius.

Can you guess how long it took AlphaZero to learn chess from scratch, prepare for the match against Stockfish, and develop its genius instincts? Four hours. That’s not a typo. For centuries, chess was considered one of the crowning glories of human intelligence. AlphaZero went from utter ignorance to creative mastery in four hours, without the help of any human guide.¹⁸

AlphaZero is not the only imaginative software out there. Many programs now routinely outperform human chess players not just in brute calculation, but even in ‘creativity’. In human-only chess tournaments, judges are constantly on the lookout for players who try to cheat by secretly getting help from computers. One of the ways to catch cheats is to monitor the level of originality players display. If they play an exceptionally creative move, the judges will often suspect that this cannot possibly be a human move – it must be a computer move. At least in chess, creativity is already the trademark of computers rather than humans! Hence if chess is our coal-mine canary, we are duly warned that the canary is dying. What is happening today to human–AI chess teams might happen down the road to human–AI teams in policing, medicine and banking too.¹⁹

Consequently, creating new jobs and retraining people to fill them will not be a one-off effort. The AI revolution won’t be a single watershed event after which the job market will just settle into a new equilibrium. Rather, it will be a cascade of ever-bigger disruptions. Already today few employees expect to work in the same job for their entire life.²⁰ By 2050, not just the idea of ‘a job for life’, but even the idea of ‘a profession for life’ might seem antediluvian.

Even if we could constantly invent new jobs and retrain the workforce, we may wonder whether the average human will have the emotional stamina necessary for a life of such endless upheavals. Change is always stressful, and the hectic world of the early twenty-first century has produced a global epidemic of stress.²¹ As the volatility of the job market and of individual careers increases, would people be able to cope? We would probably need far more effective stress-reduction techniques – ranging from drugs through neuro-feedback to meditation – to prevent the Sapiens mind from snapping. By 2050 a ‘useless’ class might emerge not merely because of an absolute lack of jobs or lack of relevant education, but also because of insufficient mental stamina.

Obviously, most of this is just speculation. At the time of writing – early 2018 – automation has disrupted many industries but it has not resulted in massive unemployment. In fact, in many countries, such as the USA, unemployment is at a historical low. Nobody can know for sure what sort of impact machine learning and automation will have on different professions in the future, and it is extremely difficult

the most successful ice-cream vendors in the world are those that the Google algorithm ranks first – not those that produce the tastiest ice cream.

I know this from personal experience. When I publish a book, the publishers ask me to write a short description that they use for publicity online. But they have a special expert, who adapts what I write to the taste of the Google algorithm. The expert goes over my text, and says, ‘Don’t use this word – use that word instead. Then we will get more attention from the Google algorithm.’ We know that if we can just catch the eye of the algorithm, we can take the humans for granted.

So if humans are needed neither as producers nor as consumers, what will safeguard their physical survival and their psychological well-being? We cannot wait for the crisis to erupt in full force before we start looking for answers. By then it will be too late. In order to cope with the unprecedented technological and economic disruptions of the twenty-first century, we need to develop new social and economic models as soon as possible. These models should be guided by the principle of protecting humans rather than jobs. Many jobs are uninspiring drudgery, not worth saving. Nobody’s life-dream is to be a cashier. What we should focus on is providing for people’s basic needs and protecting their social status and self-worth.

One new model, which is gaining increasing attention, is universal basic income. UBI proposes that governments tax the billionaires and corporations controlling the algorithms and robots, and use the money to provide every person with a generous stipend covering his or her basic needs. This will cushion the poor against job loss and economic dislocation, while protecting the rich from populist rage.²³

A related idea proposes to widen the range of human activities that are considered to be ‘jobs’. At present, billions of parents take care of children, neighbours look after one another, and citizens organise communities, without any of these valuable activities being recognised as jobs. Maybe we need to turn a switch in our minds, and realise that taking care of a child is arguably the most important and challenging job in the world. If so, there won’t be a shortage of work even if computers and robots replace all the drivers, bankers and lawyers. The question is, of course, who would evaluate and pay for these newly recognised jobs? Assuming that six-month-old babies will not pay a salary to their mums, the government will probably have to take this upon itself. Assuming, too, that we would like these salaries to cover all of a family’s basic needs, the end result will be something that is not very different from universal basic income.

Alternatively, governments could subsidise universal basic *services* rather than income. Instead of giving money to people, who then shop around for whatever they want, the government might subsidise free education, free healthcare, free transport and so forth. This is in fact the utopian vision of communism. Though the communist plan to start a working-class revolution might well become outdated, maybe we should still aim to realise the communist goal by other means?

It is debatable whether it is better to provide people with universal basic income (the capitalist paradise) or universal basic services (the communist paradise). Both options have advantages and drawbacks. But no matter which paradise you choose, the real problem is in defining what ‘universal’ and ‘basic’ actually mean.

What is universal?

When people speak about universal basic support – whether in the shape of income or services – they usually mean *national* basic support. Hitherto, all UBI initiatives have

been strictly national or municipal. In January 2017, Finland began a two-year experiment, providing 2,000 unemployed Finns with 560 euros a month, irrespective of whether they find work or not. Similar experiments are under way in the Canadian province of Ontario, in the Italian city of Livorno and in several Dutch cities.²⁴ (In 2016 Switzerland held a referendum on instituting a national basic income scheme, but voters rejected the idea.²⁵)

The problem with such national and municipal schemes, however, is that the main victims of automation may not live in Finland, Ontario, Livorno or Amsterdam. Globalisation has made people in one country utterly dependent on markets in other countries, but automation might unravel large parts of this global trade network with disastrous consequences for the weakest links. In the twentieth century, developing countries lacking natural resources made economic progress mainly by selling the cheap labour of their unskilled workers. Today millions of Bangladeshis make a living by producing shirts and selling them to customers in the United States, while people in Bangalore earn their keep in call centres dealing with the complaints of American customers.²⁶

Yet with the rise of AI, robots and 3-D printers, cheap unskilled labour would become far less important. Instead of manufacturing a shirt in Dhaka and shipping it all the way to the US, you could buy the shirt's code online from Amazon, and print it in New York. The Zara and Prada stores on Fifth Avenue could be replaced by 3-D printing centres in Brooklyn, and some people might even have a printer at home. Simultaneously, instead of calling customer services in Bangalore to complain about your printer, you could talk with an AI representative in the Google cloud (whose accent and tone of voice are tailored to your preferences). The newly unemployed workers and call-centre operators in Dhaka and Bangalore don't have the education necessary to switch to designing fashionable shirts or writing computer code – so how will they survive?

If AI and 3-D printers indeed take over from the Bangladeshis and Bangalorians, the revenues that previously flowed to South Asia will now fill the coffers of a few tech-giants in California. Instead of economic growth improving conditions all over the world, we might see immense new wealth created in hi-tech hubs such as Silicon Valley, while many developing countries collapse.

Of course, some emerging economies – including India and Bangladesh – might advance fast enough to join the winning team. Given enough time, the children or grandchildren of textile workers and call-centre operators might well become the engineers and entrepreneurs who build and own the computers and 3-D printers. But the time to make such a transition is running out. In the past, cheap unskilled labour has served as a secure bridge across the global economic divide, and even if a country advanced slowly, it could expect to reach safety eventually. Taking the right steps was more important than making speedy progress. Yet now the bridge is shaking, and soon it might collapse. Those who have already crossed it – graduating from cheap labour to high-skill industries – will probably be OK. But those lagging behind might find themselves stuck on the wrong side of the chasm, without any means of crossing over. What do you do when nobody needs your cheap unskilled labourers, and you don't have the resources to build a good education system and teach them new skills?²⁷

What then will be the fate of the stragglers? American voters might conceivably agree that taxes paid by Amazon and Google for their US business could be used to give stipends or free services to unemployed miners in Pennsylvania and jobless taxi-drivers in New York. However, would American voters also agree that these taxes

should be sent to support unemployed people in places defined by President Trump as ‘shithole countries’?²⁸ If you believe that, you might just as well believe that Santa Claus and the Easter Bunny will solve the problem.

What is basic?

Universal basic support is meant to take care of basic human needs, but there is no accepted definition for that. From a purely biological perspective, a Sapiens needs just 1,500–2,500 calories per day in order to survive. Anything more is a luxury. Yet over and above this biological poverty line, every culture in history defined additional needs as ‘basic’. In medieval Europe, access to church services was seen as even more important than food, because it took care of your eternal soul rather than of your ephemeral body. In today’s Europe, decent education and healthcare services are considered basic human needs, and some argue that even access to the Internet is now essential for every man, woman and child. If in 2050 the United World Government agrees to tax Google, Amazon, Baidu and Tencent in order to provide basic support for every human being on earth – in Dhaka as well as in Detroit – how will it define ‘basic’?

For example, what does basic education include: just reading and writing, or also composing computer code and playing the violin? Just six years of primary school, or everything up to a PhD? And what about healthcare? If by 2050 medical advances make it possible to slow down ageing processes and significantly extend human lifespans, will the new treatments be available to all 10 billion humans on the planet, or just to a few billionaires? If biotechnology enables parents to upgrade their children, would this be considered a basic human need, or would we see humankind splitting into different biological castes, with rich superhumans enjoying abilities that far surpass those of poor *Homo sapiens*?

Whichever way you choose to define ‘basic human needs’, once you provide them to everyone free of charge, they will be taken for granted, and then fierce social competitions and political struggles will focus on non-basic luxuries – be they fancy self-driving cars, access to virtual-reality parks, or enhanced bioengineered bodies. Yet if the unemployed masses command no economic assets, it is hard to see how they could ever hope to obtain such luxuries. Consequently the gap between the rich (Tencent managers and Google shareholders) and the poor (those dependent on universal basic income) might become not merely bigger, but actually unbridgeable.

Hence even if some universal support scheme provides poor people in 2050 with much better healthcare and education than today, they might still be extremely angry about global inequality and the lack of social mobility. People will feel that the system is rigged against them, that the government serves only the super-rich, and that the future will be even worse for them and their children.²⁹

Homo sapiens is just not built for satisfaction. Human happiness depends less on objective conditions and more on our own expectations. Expectations, however, tend to adapt to conditions, including to the condition of *other people*. When things improve, expectations balloon, and consequently even dramatic improvements in conditions might leave us as dissatisfied as before. If universal basic support is aimed at improving the objective conditions of the average person in 2050, it has a fair chance of succeeding. But if it is aimed at making people subjectively more satisfied with their lot and preventing social discontent, it is likely to fail.

To really achieve its goals, universal basic support will have to be supplemented by some meaningful pursuits, ranging from sports to religion. Perhaps the most

successful experiment so far in how to live a contented life in a post-work world has been conducted in Israel. There, about 50 per cent of ultra-Orthodox Jewish men never work. They dedicate their lives to studying holy scriptures and performing religious rituals. They and their families don't starve partly because the wives often work, and partly because the government provides them with generous subsidies and free services, making sure that they don't lack the basic necessities of life. That's universal basic support *avant la lettre*.³⁰

Although they are poor and unemployed, in survey after survey these ultra-Orthodox Jewish men report higher levels of life satisfaction than any other section of Israeli society. This is due to the strength of their community bonds, as well as to the deep meaning they find in studying scriptures and performing rituals. A small room full of Jewish men discussing the Talmud might well generate more joy, engagement and insight than a huge textile sweatshop full of hard-working factory hands. In global surveys of life satisfaction, Israel is usually somewhere near the top, thanks in part to the contribution of these jobless poor people.³¹

Secular Israelis often complain bitterly that the ultra-Orthodox don't contribute enough to society, and live off other people's hard work. Secular Israelis also tend to argue that the ultra-Orthodox way of life is unsustainable, especially as ultra-Orthodox families have seven children on average.³² Sooner or later, the state will not be able to support so many unemployed people, and the ultra-Orthodox will have to go to work. Yet it might be just the reverse. As robots and AI push humans out of the job market, the ultra-Orthodox Jews may come to be seen as the model of the future rather than as a fossil from the past. Not that everyone will become Orthodox Jews and go to the yeshivas to study the Talmud. But in the lives of all people, the quest for meaning and for community might eclipse the quest for a job.

If we manage to combine a universal economic safety net with strong communities and meaningful pursuits, losing our jobs to the algorithms might actually turn out to be a blessing. Losing control over our lives, however, is a much scarier scenario. Notwithstanding the danger of mass unemployment, what we should worry about even more is the shift in authority from humans to algorithms, which might destroy any remaining faith in the liberal story and open the way to the rise of digital dictatorships.

Liberty

Big Data is watching you

The liberal story cherishes human liberty as its number-one value. It argues that all authority ultimately stems from the free will of individual humans, as it is expressed in their feelings, desires and choices. In politics, liberalism believes that the voter knows best. It therefore upholds democratic elections. In economics, liberalism maintains that the customer is always right. It therefore hails free-market principles. In personal matters, liberalism encourages people to listen to themselves, be true to themselves and follow their hearts – as long as they do not infringe on the liberties of others. This personal freedom is enshrined in human rights.

In Western political discourse the term ‘liberal’ is sometimes used today in a much narrower partisan sense, to denote those who support specific causes like gay marriage, gun control and abortion. Yet most so-called conservatives also embrace the broad liberal world view. Especially in the United States, both Republicans and Democrats should occasionally take a break from their heated quarrels to remind themselves that they all agree on fundamentals such as free elections, an independent judiciary and human rights.

Test yourself. Do you think people should choose their government rather than blindly obey a king? Do you think people should choose their profession rather than be born into a caste? Do you think people should choose their spouse rather than marrying whomever their parents select? If you answered ‘Yes’ to all three questions, congratulations, you are a liberal.

In particular, it is vital to remember that even conservative heroes such as Ronald Reagan and Margaret Thatcher were great champions not only of economic freedoms but also of individual liberties. In a famous interview in 1987, Thatcher said that ‘There is no such thing as society. There is [a] living tapestry of men and women ... and the quality of our lives will depend upon how much each of us is prepared to take responsibility for ourselves.’¹

Thatcher’s heirs in the Conservative Party fully agree with the Labour Party that political authority comes from the feelings, choices and free will of individual voters. Thus when Britain needed to decide whether it should leave the EU, Prime Minister David Cameron didn’t ask Queen Elizabeth II, the Archbishop of Canterbury or the Oxford and Cambridge dons to resolve the issue. He didn’t even ask the Members of Parliament. Rather, he held a referendum in which each and every Briton was asked: ‘What do you *feel* about it?’

$b \times c \times d = \text{ahh!}$

Biological knowledge multiplied by Computing power multiplied by Data equals Ability to Hack Humans.

We can already witness how this formula works in the field of medicine. The most important medical decisions in our life rely not on our feelings of illness or wellness, or even on the informed predictions of our doctor – but on the calculations of computers which understand our bodies much better than we do. Within a few decades, Big Data algorithms informed by a constant stream of biometric data could monitor our health 24/7. They could detect the very beginning of influenza, cancer or Alzheimer's disease, long before we feel anything is wrong with us. They could then recommend appropriate treatments, diets and daily regimens, custom-built for our unique physique, DNA and personality.

People will enjoy the best healthcare in history, but for precisely this reason they will probably be sick all the time. There is always something wrong somewhere in the body. There is always something that can be improved. In the past, you felt perfectly healthy as long as you didn't sense pain or you didn't suffer from an apparent disability such as limping. But by 2050, thanks to biometric sensors and Big Data algorithms, diseases may be diagnosed and treated long before they lead to pain or disability. As a result, you will always find yourself suffering from some 'medical condition' and following this or that algorithmic recommendation. If you refuse, perhaps your medical insurance would become invalid, or your boss would fire you – why should they pay the price of your obstinacy?

It is one thing to continue smoking despite general statistics that connect smoking with lung cancer. It is a very different thing to continue smoking despite a concrete warning from a biometric sensor that has just detected seventeen cancerous cells in your upper left lung. And if you are willing to defy the sensor, what will you do when the sensor forwards the warning to your insurance agency, your manager and your mother?

Who will have the time and energy to deal with all these illnesses? In all likelihood, we could just instruct our health algorithm to deal with most of these problems as it sees fit. At most, it will send periodic updates to our smartphones, telling us that 'seventeen cancerous cells were detected and destroyed'. Hypochondriacs might dutifully read these updates, but most of us will ignore them just as we ignore those annoying anti-virus notices on our computers.

The drama of decision-making

What is already beginning to happen in medicine is likely to occur in more and more fields. The key invention is the biometric sensor, which people can wear on or inside their bodies, and which converts biological processes into electronic information that computers can store and analyse. Given enough biometric data and enough computing power, external data-processing systems can hack all your desires, decisions and opinions. They can know exactly who you are.

Most people don't know themselves very well. When I was twenty-one, I finally realised that I was gay, after several years of living in denial. That's hardly exceptional. Many gay men spend their entire teenage years unsure about their sexuality. Now imagine the situation in 2050, when an algorithm can tell any teenager exactly where he is on the gay/straight spectrum (and even how malleable that position is). Perhaps

the algorithm shows you pictures or videos of attractive men and women, tracks your eye movements, blood pressure and brain activity, and within five minutes ejects a number on the Kinsey scale.⁶ It could have saved me years of frustration. Perhaps you personally wouldn't want to take such a test, but then maybe you find yourself with a group of friends at Michelle's boring birthday party, and somebody suggests you all take turns checking yourself on this cool new algorithm (with everybody standing around to watch the results – and comment on them). Would you just walk away?

Even if you do, and even if you keep hiding from yourself and your classmates, you won't be able to hide from Amazon, Alibaba or the secret police. As you surf the Web, watch YouTube or read your social media feed, the algorithms will discreetly monitor you, analyse you, and tell Coca-Cola that if it wants to sell you some fizzy drink, it had better use the advertisement with the shirtless guy rather than the shirtless girl. You won't even know. But they will know, and such information will be worth billions.

Then again, maybe it will all be out in the open, and people will gladly share their information in order to get better recommendations – and eventually in order to get the algorithm to make decisions for them. It starts with simple things, like deciding which movie to watch. As you sit down with a group of friends to spend a cozy evening in front of the TV, you first have to choose what to see. Fifty years ago you had no choice, but today – with the rise of view-on-demand services – there are thousands of titles available. Reaching an agreement can be quite difficult, because while you personally like science-fiction thrillers, Jack prefers romantic comedies, and Jill votes for artsy French films. You may well end up compromising on some mediocre B-movie that disappoints all of you.

An algorithm might help. You can tell it which previous movies each of you really liked, and based on its massive statistical database, the algorithm can then find the perfect match for the group. Unfortunately, such a crude algorithm is easily misled, particularly because self-reporting is a notoriously unreliable gauge for people's true preferences. It often happens that we hear lots of people praise some movie as a masterpiece, feel compelled to watch it, and even though we fall asleep midway through, we don't want to look like philistines, so we tell everyone it was an amazing experience.⁷

Such problems, however, can be solved if we just allow the algorithm to collect real-time data on us as we actually watch movies, instead of relying on our own dubious self-reports. For starters, the algorithm can monitor which movies we completed, and which we stopped watching halfway through. Even if we tell the whole world that *Gone with the Wind* is the best movie ever made, the algorithm will know we never made it past the first half-hour, and we never really saw Atlanta burning.

Yet the algorithm can go much deeper than that. Engineers are currently developing software that can detect human emotions based on the movements of our eyes and facial muscles.⁸ Add a good camera to the television, and such software will know which scenes made us laugh, which scenes made us sad and which scenes bored us. Next, connect the algorithm to biometric sensors, and the algorithm will know how each frame has influenced our heart rate, our blood pressure and our brain activity. As we watch, say, Tarantino's *Pulp Fiction*, the algorithm may note that the rape scene caused us an almost imperceptible tinge of sexual arousal, that when Vincent accidentally shot Marvin in the face it made us laugh guiltily, and that we didn't get the joke about the Big Kahuna Burger – but we laughed anyway, so as not to look stupid. When you force yourself to laugh, you use different brain circuits and muscles

This ebook is copyright material and must not be copied, reproduced, transferred, distributed, leased, licensed or publicly performed or used in any way except as specifically permitted in writing by the publishers, as allowed under the terms and conditions under which it was purchased or as strictly permitted by applicable copyright law. Any unauthorized distribution or use of this text may be a direct infringement of the author's and publisher's rights and those responsible may be liable in law accordingly.

VINTAGE
20 Vauxhall Bridge Road,
London SW1V 2SA

Vintage is part of the Penguin Random House group of companies whose addresses can be found at global.penguinrandomhouse.com.



Copyright © Yuval Noah Harari 2018

Cover painting from the series *We Share Our Chemistry with the Stars*, Marc Quinn, oil on canvas © & courtesy Marc Quinn studio; Design Suzanne Dean

Yuval Noah Harari has asserted his right to be identified as the author of this Work in accordance with the Copyright, Designs and Patents Act 1988

First published in Vintage in 2019
First published in hardback by Jonathan Cape in 2018

penguin.co.uk/vintage

A CIP catalogue record for this book is available from the British Library

ISBN 9781473554719