Contents

Preface to the Paperback Edition

This book is about one of the greatest economic challenges of our time: the threat of a world where there is not enough well-paid work for everyone to do, because of the remarkable technological changes on the horizon. It was written with a sense of urgency since, in my view, we are not yet taking this threat seriously enough. But nobody could have predicted how, only a few months after it was first published, a global pandemic would bring economic life as we knew it to an end and make the ideas and concerns in here more urgent than ever.

At the time of this writing, the COVID-19 pandemic has been with us for about six months. When the pandemic began, the hope was that it would be a short-lived crisis. Economies would need to be temporarily placed in a sort of suspended animation, but once the virus had passed – in a matter of weeks, it was originally thought – we would swiftly return to economic life as usual. As we now know, this initial hope turned out to be completely misplaced. The virus is here to stay for some time. The frantic policy firefighting of the first few months of the crisis has given way to far more long-lasting interventions. And the economic consequences of the pandemic have been more destructive than most of us had first imagined. From April to June 2020, for instance, the US suffered the sharpest collapse in output since the Second World War. The UK lost almost 18 years of growth in a matter of months.[1]

At the core of this economic meltdown has been the labour market. Work was already precarious in many parts of the world before the pandemic began, marked by stagnating wages, rising insecurity, pockets of unemployment, and declining participation. COVID-19 pushed it off a cliff: in many countries hit hard by the virus, like the

US and the UK, joblessness has surged to extraordinary levels. As the pandemic took hold, in other words, we found ourselves unexpectedly thrust into a world with much less work – not because that work has been automated, but because the measures that we were forced to adopt in response to the virus (lockdowns, social distancing, self-isolation, and so on) completely decimated the demand that so many jobs rely upon.

As a result, we have had to confront even sooner the challenges that concern me in this book. Andrew Yang, the 2020 US presidential candidate who focused on the threat of job displacement, made this point well: 'Apparently I should have been talking about a pandemic instead of automation,' he wrote on Twitter.[2] To be clear, the threat of technological unemployment in the future has not diminished: on the contrary, there are reasons to think that the threat is now greater than before. But the pandemic has also given us a frightening preview of what this future might look like, and an insight into the immensity of the challenges that we will have to face when it arrives.

A GLIMPSE OF THE FUTURE

As we shall see in this book, the fundamental difficulty that lies ahead is a distributional one. Technological progress may make us collectively more prosperous than ever before, but how are we to share out that prosperity when our traditional way of doing so – paying a wage for the work that people do – is less effective than in the past? And this, of course, is precisely the economic problem that has dominated in 2020. Overnight, vast numbers of workers around the world woke up to suddenly find themselves without a job and an income.

What should be done? I argue that in a moment like this the state must take on a far larger role in sharing out prosperity in society, through what I call the 'Big State'. The pandemic has now proven that there is no credible alternative. Different countries have adopted slightly different schemes, but all of them involve a vastly bigger state providing an income to those without work. Indeed, ideas which only a few months ago were viewed by some as outlandish – a basic income, for instance – have swiftly become commonplace in almost all corners

of political conversation. To provide support to the unemployed, and to prop up the economy more generally, the US has already borrowed more than five times what it did at the height of the financial crisis of 2007–8; the UK is on track in 2020 to set a peacetime borrowing record.[3]

Aside from sharing out prosperity, there are two other big challenges we can expect to face in a world with less work, both of which have little to do with economics. One of these is the growing power of a small handful of large technology companies, or Big Tech. Here, too, the pandemic offers a glimpse of the future: it is a conspicuous feature of the COVID-19 economic landscape that such companies have done particularly well. At one point during this crisis, just five of them accounted for more than 20 per cent of the worth of the entire S&P 500 index, comprised of 500 large companies listed on US stock exchanges.[4] Apple alone was worth more than all the companies in the London Stock Exchange's FTSE 100 index combined.[5]

My concern in this book, though, is far less with the tech companies' economic power – great and growing though it may be – than with their *political* power, and the impact they may have on issues of liberty, democracy and social justice in the future. So it is important to note, for instance, how debates about data privacy and security have quietly disappeared from public discussion since the pandemic began. A 'do whatever it takes' mentality took hold at the start of this crisis, with many countries permitting CCTV surveillance footage, smartphone location data, and credit-card purchasing history, among much else, to be collected, sifted, sorted and studied on a huge scale in an effort to control the virus. The threat may have required this. But in time we must ensure that the new political power we have granted to Big Tech, and the heightened ability to shape how we all live together in society that comes with it, is properly scrutinized and reined in if need be.

The final challenge that we will face in a world with less work, I argue, is finding meaning in life. It is often said that work is not simply a source of income, but a source of purpose as well. And so, if employment dries up, then where will that sense of direction come from? My own view is that the relationship between work and meaning is actually far murkier than is commonly supposed: many people

do not get a strong sense of purpose from their jobs today, and our relationship to work has looked wildly different at other moments in history. The pandemic has strengthened that belief. Yes, there have been awful stories of those who lost work and felt a sense of devastation that could not be explained by the loss of an income alone; but there are also many accounts of those who felt quite the opposite, a sense of relief in being freed from jobs that were simply not worth the wages they provided.

But what will people actually *do*, if they do not have to work for a living? I fear we do not yet have good answers to this question. In a world like ours, where work sits at the centre of our lives, it is very difficult to imagine how we might spend our time any differently. Our struggles in this pandemic-induced world have shown that. We can point to some telling changes in spending patterns over the last few months: the UK, for instance, has suffered significant shortages of flour, timber and bedding plants as people took up baking, DIY and gardening to fill their spare time; the US has suffered from similarly disruptive spikes in demand. But there have also been unfamiliar public conversations about bigger issues: work–life balance, the value of family and community, the merits of city life, how best to spend our time in idleness, how to maintain our mental health in tough times. (Depression among UK adults almost doubled at the start of the pandemic; text messages to a US government mental health hotline rose almost 1,000 per cent.[6]) That these conversations feel so novel, and that the conclusions can at times seem so provisional and unsatisfying, strengthens my sense that the all-consuming nature of our traditional working lives has distracted us from these big questions until now.

THE INCREASING THREAT OF AUTOMATION

While the pandemic offers a preview of the problems that a world with greater automation will have to grapple with – issues regarding the distribution of prosperity, the power of Big Tech, and the search for meaning – it is also probably hastening the arrival of that world.

One reason for this is that many countries around the globe are

now in severe recession, and evidence from the past suggests that when economies slow down automation can pick up. Around the beginning of the twenty-first century, for example, jobs (as a proportion of total employment) for secretaries, clerks, salespeople and the like shrivelled up as new technologies began to take on these workers' tasks and displace them from their roles. In the book, I explore exactly why such 'middle-skilled' jobs were lost, while both high-paid and low-paid workers increased their share of employment. But what matters for thinking about our current situation is that, at least in the US, the vast majority of this job destruction took place during economic downturns. One influential study suggests that, since the mid-1980s, 88 per cent of these middle-skill job losses took place within a year of a recession.[7]

What's more, our particular downturn is no ordinary recession. The pandemic also creates new and unique reasons to worry about the threat of automation. Most obviously, COVID-19 strengthens the incentive to replace human beings with machines. A machine, after all, will not pass the virus on to co-workers or customers; it will not fall ill and need to take time off from work; it will not need to isolate to protect its peers.

So far, this incentive has been kept in check to a certain extent by government interventions. The UK government, for instance, was at one point paying up to 80 per cent of the wages of 9.6 million workers – more than a third of all UK employees – to protect them from unemployment.[8] But many governments have not taken this route. And when those interventions that do exist are relaxed – as they inevitably will be – the incentive to automate will grow even stronger. For businesses looking to boost productivity during the downturn or cut labour costs as revenue falls, replacing workers with machines for particular activities might seem increasingly attractive. At the start of the pandemic, for instance, a survey of global business executives by the consulting firm EY found that 41 per cent were investing in accelerating automation.[9]

Finally, the pandemic may have softened some of the cultural resistance that accompanies the use of new technologies in the workplace. The barriers to automation, after all, are not simply technological ('is it possible to automate a task?'), economic ('is it profitable to automate

a task?'), or regulatory ('is it permitted to automate a task?'). They are also cultural: what is automated depends, in part, on whether people find it palatable to do something with a machine. And to the extent that any of us – whether as business owners, employers, employees or consumers – might have had a bias against new technologies before the pandemic, this crisis is likely to have weakened it. One poll, for instance, suggests that all age groups in Britain now 'feel more positively' towards technology; another, that a third of Brits have become 'more confident in using technology' as well.[10] Out of necessity, we have been forced to use technology in ways that would have simply seemed unimaginable a few months ago – and it has largely been a success. And so, any particular act of automation in the future is now likely to feel far less of an unprecedented leap.

Take medicine, for example. Before the pandemic began, about 80 per cent of doctors' appointments in England and Wales were conducted face-to-face; now, that proportion has fallen to only about 7 per cent.[11] It is hard to believe that virtual appointments will stop once the pandemic is over; on the other hand, it is easy to imagine that other parts of medicine – diagnostics, for instance – could also start being done differently through technology, and perhaps without involving doctors at all. Or take the law. In many jurisdictions, bricks-and-mortar courtrooms were shuttered, and practically overnight court became an online service rather than a physical place. As with medicine, not only is it possible to see how such a virtual setup might now become the norm in certain corners of the criminal justice system, but bolder technological proposals – for instance, that some cases, such as low-value civil disputes, might be settled without any human deliberation at all – seem far less radical than they would have only a few months ago.

LOWER-PAID WORKERS AT RISK

At the moment, to be sure, technology mostly appears to be keeping people in work rather than pushing them out. Many have been able to use technology to work remotely, to an extent that would have seemed unimaginable until recently: in the US and UK, as the crisis

began, about two-thirds of those still working did so.[12] However, not everyone can work from home, and those who can tend to be in better-paid, white-collar roles. A US survey found that while 71 per cent of people earning more than $180,000 could work remotely during the pandemic, only 41 percent of those earning less than $24,000 could. Another study reported that while 62 per cent of workers with a bachelor's degree or more could perform their jobs from home, only 9 per cent of those who did not graduate from high school could do so.[13] Remote work is simply not an option for many blue-collar workers, such as those who work in restaurants, shops and warehouses.

This particular inequality in the ability of workers to adapt to the pandemic through technology is symptomatic of a deeper problem. When the crisis began, it was said that COVID-19 would be a 'great leveller'. This disease, many proclaimed, would not discriminate according to a person's ethnicity or their wealth; all of us were equally at risk. We now know that this was a myth. To start with, the medical impact of the virus has been extremely unequal. In the UK, people from ethnic minority backgrounds made up 14 per cent of the population but 34 per cent of critically ill COVID-19 patients; in the US, black people were almost five times more likely to be hospitalized by, and more than twice as likely to die from, the virus than white people.[14] And the *economic* impact of the virus has been extremely unequal, too. The job losses, for instance, have been concentrated among the lower-paid workers: one study suggests that, in the US, workers in the bottom 20 per cent of earners were about four times more likely to lose their job at the start of the pandemic than those in the top 20 per cent of earners.[15]

These inequalities are striking in themselves, but they are also important for thinking about the looming threat of automation. It is likely that the pandemic has both increased that threat and made it clear that workers who are already economically disadvantaged will be hardest hit.

Over the last few decades, lower-paid workers have mostly been protected from automation. This is because their jobs very often involve personal interaction or manual labour, and until recently these tasks have proven tricky to automate. But the cruel irony of the last few months is that these workers have in fact been hardest hit by

the pandemic precisely *because* of those properties of their jobs: the virus spreads through personal interaction, and it flourishes in poorly ventilated indoor spaces such as factories and warehouses. As a result, many of these people have found themselves unable to work.

As the pandemic has increased the incentive to automate, therefore, it is likely that these hands-on workers are the most at risk: they cannot readily work in their traditional workplaces, nor can they retreat to a home office and do their job from there. It is no surprise that so many recently reported technological developments seem aimed directly at them: machines that stock shelves, prepare packages, greet customers, deliver goods, clean floors, take temperatures, and so on.

Does the prospect of an effective vaccine mean that this incentive to automate, however strong it may be at the moment, will eventually fizzle out? That may indeed be the case. But it is not clear that such a development – however magnificent it would be from a medical point of view – would bring the threat of automation back down. To begin with, the cultural shifts mentioned before may persist: if the pandemic has made us more welcoming towards technology, that new attitude will probably remain. More significantly, the pandemic has also transformed the fundamental rhythms of how many of us live our lives: we eat out less, shop online more, avoid travel if we can, stay away from theatres and cinemas and sporting events, work from home if possible, and so on. Even when the pandemic fades away and government restrictions are relaxed, these changes in habits and behaviours are unlikely to reverse completely.[16]

Those who say that the pandemic spells the 'end of the office', the 'death of the high street', or the 'collapse of the city centre' are probably overstating their case: though offices and shopping destinations were abandoned for a while, people are slowly starting to return.[17] Nevertheless, it is entirely plausible that such places will remain diminished versions of their former selves for quite some time – perhaps indefinitely. And if that is right, it does not bode well for workers who depend on these places: security guards, receptionists and cleaners in offices; waiters, sandwich makers and baristas in nearby streets; retailers, transport workers, hotel staff and entertainers in city centres, and so on. In this scenario, of course, the decline in demand for their work

may be due more to the effects of the pandemic rather than to technology as such. But when thinking about the threat of automation, it is crucial to consider these shifts, because the lower-paid, hands-on jobs are precisely the ones that provided people displaced by machines with work in the past – and their future is now in doubt.

In a sense, the pandemic has been a pilot scheme in how we ought to respond to a world with less work. This exercise has been unplanned and unwanted, but it has also proven to be informative and revelatory. I hope that in the months and years to come we are able to reflect on this vast social experiment, to understand what has worked in responding to this crisis and to be honest about where we have fallen short. At the moment, we are only temporary visitors in a world with less work. This pandemic, like all those before it, will eventually pass, and the problems that consume us today will fall away. But when the COVID-19 crisis recedes, the threat of automation may have only increased. And then the challenges which we have caught an unsettling glimpse of during the pandemic will re-emerge and start to trouble and test us once again.

Daniel Susskind
London
30 September 2020

NOTES

1 For the US, see Dominic Rushe, 'US Economy Suffers Worst Quarter Since the Second World War as GDP Shrinks by 32.9%', *Guardian*, 30 July 2020; for the UK, monthly GDP fell back to where it was in July 2002 – see the ONS data in Delphine Strauss, 'UK Economy Suffers Worst Slump in Europe in Second Quarter', *Financial Times*, 12 August 2020.

2 <https://twitter.com/AndrewYang/status/1238095725721944065> (accessed 30 September 2020).

3 For the US, see BBC News, 'Coronavirus: US to Borrow Record $3tn as Spending Soars', 4 May 2020; for UK, see Chris Giles, 'UK Public Finances Continue on Path to Record Peacetime Deficit', *Financial Times*, 25 September 2020.

4 Richard Henderson, 'Big Tech Presents a Problem for Investors as Well as Congress', *Financial Times*, 1 August 2020.

5 BBC News, 'Apple More Valuable than the Entire FTSE 100', 2 September 2020.

6 For the UK, see Office for National Statistics, 'Coronavirus and Depression in Adults, Great Britain: June 2020', 18 August 2020; for the US, see Paige Winfield Cunningham, 'The Health 202: Texts to Federal Government Mental Health Hotline Up Roughly 1,000 percent', *Washington Post*, 4 May 2020.

7 Nir Jaimovich and Henry Siu, 'Job Polarization and Jobless Recoveries', *Review of Economics and Statistics*, 102:1 (2020), p. 129–47. Whether this holds in other countries is less clear. See Joel Bilt, 'Automation and Reallocation: Will COVID-19 Usher in the Future of Work?', *Canadian Public Policy, Project Muse* (2020) for evidence it holds in Canada. Though also see Georg Graetz and Guy Michaels, 'Is Modern Technology Responsible for Jobless Recoveries?', *American Economic Review*, 107:5 (2017) p. 168–73 for evidence it holds less strongly elsewhere.

8 Brigid Francis-Devine, Andrew Powell, and Niamh Foley, 'Coronavirus: Impact on the Labour Market', *House of Commons Library Briefing Paper* 8898, 12 August 2020. There were 28.02 million employees in the UK for April to June 2020.

9 EY, 'Capital Strategies being Rewritten as C-Suite grapples with immediate impact of new reality', 30 March 2020, published online at https://www.ey.com/en_gl/news/ (accessed 29 September 2020).

10 Ellen Daniel, 'Brits More Positive about Technology following Covid-19 Pandemic', Verdict, 18 June 2020. Published online at https://www.verdict.co.uk/covid-19-technology-vodafone/ (accessed 29 September 2020); Vinous Ali, 'Survey Results: Lockdown and Changing Attitudes Towards Tech', TechUK, 17 July 2020, published online at https://www.techuk.org/ (accessed 29 September 2020).

11 Paul Lynch and Daniel Wainwright, 'Coronavirus: How GPs Have Stopped Seeing Most Patients in Person', BBC News, 11 April 2020.

12 'Sixty-two percent of employed Americans currently say they have worked from home during the crisis', see Megan Brenan, 'U.S. Workers Discovering Affinity for Remote Work', Gallup, 3 April 2020; 'of those still working, 61 % were working remotely all the time', see CIPD, 'Impact of COVID-19 on Working Lives: Findings from Our April 2020 Survey', 3 September 2020, published online at https://www.cipd.co.uk/ (accessed 29 September 2020).

13 Olga Khazan, 'How the Coronavirus Could Create a New Working Class', *Atlantic*, 15 April 2020; Rakesh Kochhar and Jeffrey Passel, 'Telework May Save U.S. Jobs in COVID-19 Downturn, Especially Among College Graduates', Pew Research Centre, 6 May 2020.

14 For the UK, see Suzie Bailey and Michael West, 'Ethnic Minority Deaths and COVID-19: What are We To Do?', The Kings Fund, 30 April 2020; for the US, see CDC, 'COVID-19 Hospitalization and Death by Race/Ethnicity', 2020, https://www.cdc.gov/coronavirus/2019-ncov/. This data is from the 18 August update.

15 Tomaz Cajner, Leland Crane, Ryan Decker, *et al.*, 'The U.S. Labor Market During the Beginning of the Pandemic Recession', Becker Friedman Institute Working Paper No. 2020-58 (2020).

16 See, for instance, McKinsey & Company, 'Survey: US Consumer Sentiment During the Coronavirus Crisis', 28 August 2020.

17 See David Autor and Elisabeth Reynolds, 'The Nature of Work After the COVID Crisis: Too Few Low-wage Jobs', The Hamilton Project, Essay 2020-14, July 2020. The authors make a similar argument to the one that follows.

Introduction

The 'Great Manure Crisis' of the 1890s should have come as no surprise.[1] For some time, in big cities like London and New York, the most popular forms of transport had relied upon horses – hundreds of thousands of them – to heave cabs, carts, wagons, wains, and a variety of other vehicles through the streets. As locomotives, horses were not particularly efficient: they had to take a break to rest and recover every few miles, which partly explains why quite so many were needed.[2] Operating a basic carriage, for example, required at least three animals: two working in rotation to pull it along, plus one in reserve in case of a breakdown. The horse-drawn tram, the transit mode of choice for New Yorkers, relied on a team of eight, which took turns dragging it on a set of specially laid tracks. And in London, thousands of horse-drawn double-decker buses, modestly sized versions of today's red ones, demanded about a dozen animals apiece for the task.[3]

With these horses came manure – and lots of it. A healthy horse produces somewhere between fifteen and thirty pounds of manure a day, almost the weight of a two-year-old child.[4] One enthusiastic health officer working in Rochester, New York, calculated that the horses in his city alone produced enough in a year to cover an acre of land to a height of 175 feet, almost as high as the Leaning Tower of Pisa.[5] Apocryphally, people at the time extrapolated from these calculations to an inescapably manure-filled future: a New York commentator who predicted that piles would soon reach the height of third-storey windows, a London reporter who imagined that by the middle of the twentieth century the streets would be buried under nine feet of the stuff.[6] Nor was the crisis simply about manure.

Thousands of putrefying dead horses littered the roads, many deliberately left to decay to a size that made for easier disposal. In 1880 alone, about 15,000 horse carcases were removed from New York City.[7]

It is said that policymakers did not know what to do.[8] They couldn't simply ban horses from the streets: the animals were far too important. In 1872, when the so-called Horse Plague hit the United States, with horses struck down by one of the worst outbreaks of equine flu in recorded history, large parts of the country's economy came to a halt.[9] Some even blame the epidemic for that year's Great Fire of Boston; seven hundred buildings burned to the ground, they claim, because there were not enough horses to pull firefighting equipment to the scene.[10] But the twist in the tale is that, in the end, policymakers didn't need to worry. In the 1870s, the first internal combustion engine was built. In the 1880s, it was installed in the first automobile. And only a few decades later, Henry Ford brought cars to the mass market with his famous Model T. By 1912, New York had more cars than horses. Five years after that, the last horse-drawn tram was decommissioned in the city.[11] The Great Manure Crisis was over.

The 'Parable of Horseshit', as Elizabeth Kolbert called it in the *New Yorker*, has been told many times over the years.[12] In most versions of the story, the decline of horses is cast in an optimistic light, as a tale of technological triumph, a reassuring reminder that it is important to remain open-minded even when you find yourself knee-deep in a foul, seemingly intractable problem. But for Wassily Leontief, the Russian-American economist who won the Nobel Prize in 1973, the same events suggested a more unsettling conclusion. What he saw instead was how a new technology, the combustion engine, had taken a creature that, for millennia, had played a central role in economic life – not only in cities but on farms and fields – and, in only a matter of decades, had banished it to the sidelines. In a set of articles written in the early 1980s, Leontief made one of the most infamous claims in modern economic thought. What technological progress had done to horses, he said, it would eventually do to human beings as well: drive us out of work. What cars and tractors were to them, he thought, computers and robots would be to us.[13]

Today, the world is gripped again by Leontief's fear. In the United States, 30 per cent of workers now believe their jobs are likely to be replaced by robots and computers in their lifetime. In the UK, the same proportion think it could happen in the next twenty years.[14] And in this book, I want to explain why we have to take these sorts of fears seriously – not always their substance, as we shall see, but certainly their spirit. Will there be enough work for everyone to do in the twenty-first century? This is one of the great questions of our time. In the pages that follow, I will argue that the answer is 'no' and explain why the threat of 'technological unemployment' is now real. I will describe the different problems this will create for us – both now and in the future – and, most importantly, set out how we might respond.

It was John Maynard Keynes, the great British economist, who popularized the term 'technological unemployment' almost fifty years before Leontief wrote down his worries, capturing in a pithy pairing of words the idea that new technologies might push people out of work. In what follows, I will draw on many of the economic arguments that have been developed since Keynes to try to gain a better look back at what happened in the past, and a clearer glimpse of what lies ahead. But I will also seek to go well beyond the narrow intellectual terrain inhabited by most economists working in this field. The future of work raises exciting and troubling questions that often have little to do with economics: questions about the nature of intelligence, about inequality and why it matters, about the political power of large technology companies, about what it means to live a meaningful life, about how we might live together in a world that looks very different from the one in which we have grown up. In my view, any story about the future of work that fails to engage with these questions as well is incomplete.

NOT A BIG BANG, BUT A GRADUAL WITHERING

An important starting point for thinking about the future of work is the fact that, in the past, many others have worried in similar ways

about what lies ahead – and been very wrong. Today is not the first time that automation anxiety has spread, nor did it first appear in the 1930s with Keynes. In fact, ever since modern economic growth began, centuries ago, people have periodically suffered from bouts of intense panic about being replaced by machines. Yet those fears, time and again, have turned out to be misplaced. Despite a relentless flow of technological advances over the years, there has always been enough demand for the work of human beings to avoid the emergence of large pools of permanently displaced people.

And so, in the first part of the book, I begin with this history, investigating why those who worried about being replaced by machines turned out repeatedly to be so wrong, and exploring how economists have changed their minds over time about the impact of technology on work. Then I turn to the history of artificial intelligence (AI) – a technology that has captured our collective imagination over the last few years, and which is largely responsible for the renewed sense of unease that many now feel about the future. AI research, in fact, began many decades ago, with an initial burst of enthusiasm and excitement, but that was followed by a slump into a long, deep winter when little progress was made. In recent years, though, there has been a rebirth, an intellectual and practical revolution that caught flat-footed many economists, computer scientists, and others who had tried to predict which activities machines could never do.

In the second part of the book, building on this history, and trying to sidestep the intellectual mistakes that others have made before, I explain how technological unemployment is likely to unfold in the twenty-first century. In a recent survey, leading computer scientists made the claim that there is a 50 per cent chance that machines will outperform human beings at 'every task' within forty-five years.[15] But the argument I make does not rely on dramatic predictions like this turning out to be true. In fact, I find it hard to believe that they will. Even at the century's end, tasks are likely to remain that are either hard to automate, unprofitable to automate, or possible and profitable to automate but which we will still prefer people to do. And despite the fears reflected in those polls of American and British workers, I also find it difficult to imagine that many of today's jobs will vanish completely in years to come (to say nothing about new

types of jobs that await in the future). Much of that work, I expect, will turn out to involve some tasks that lie beyond the reach of even the most capable machines.

The story I tell is a different one. Machines will not do everything in the future, but they will do *more*. And as they slowly, but relentlessly, take on more and more tasks, human beings will be forced to retreat to an ever-shrinking set of activities. It is unlikely that every person will be able to do what remains to be done; and there is no reason to imagine there will be enough demand for it to employ all those who are indeed able to do it.

In other words, if you picked up this book expecting an account of a dramatic technological big bang in the next few decades, after which lots of people suddenly wake up to find themselves without work, you will be disappointed. That scenario is not likely to happen: some work will almost certainly remain for quite some time to come. But, as time passes, that work is likely to sit beyond the reach of more and more people. And, as we move through the twenty-first century, the demand for the work of human beings is likely to wither away, gradually. Eventually, what is left will not be enough to provide everyone who wants it with traditional well-paid employment.

A useful way of thinking about what this means is to consider the impact that automation has already had on farming and manufacturing in many parts of the world. Farmers and factory workers are still needed: those jobs have not completely vanished. But the number of workers that are needed has fallen in both cases, sometimes precipitously – even though these sectors produce more output than ever before. There is, in short, no longer enough demand for the work of human beings in these corners of the economy to keep the same number of people in work. Of course, as we shall see, this comparison has its limits. But it is still helpful in highlighting what should actually be worrying us about the future: not a world without *any* work at all, as some predict, but a world without *enough* work for everyone to do.

There is a tendency to treat technological unemployment as a radical discontinuity from economic life today, to dismiss it as a fantastical idea plucked out of the ether by overly neurotic shock-haired economists. By exploring how technological unemployment

might actually happen, we will see why that attitude is a mistake. It is not a coincidence that, today, worries about economic inequality are intensifying at the exact same time that anxiety about automation is growing. These two problems – inequality and technological unemployment – are very closely related. Today, the labour market is the main way that we share out economic prosperity in society: most people's jobs are their main, if not their only, source of income. The vast inequalities we already see in the labour market, with some workers receiving far less for their efforts than others, show that this approach is already creaking. Technological unemployment is simply a more extreme version of that story, but one that ends with some workers receiving nothing at all.

In the final part of the book, I untangle the different problems created by a world with less work and describe what we should do about them. The first is the economic problem just mentioned: how to share prosperity in society when the traditional mechanism for doing so, paying people for the work that they do, is less effective than in the past. Then I turn to two issues that have little to do with economics at all. One is the rise of Big Tech, since, in the future, our lives are likely to become dominated by a small number of large technology companies. In the twentieth century, our main worry may have been the economic power of corporations: but in the twenty-first, that will be replaced by fears about their *political* power instead. The other issue is the challenge of finding meaning in life. It is often said that work is not simply a means to a wage but a source of direction: if that is right, then a world with less work may be a world with less purpose as well. These are the problems we will face, and each of them will demand a response.

A PERSONAL STORY

The stories and arguments in this book are, to some extent, personal ones. About a decade ago, I began to think about technology and work in a serious way. Well before this, however, it had been an informal interest, something I often mulled over. My father, Richard Susskind, had written his doctorate in the 1980s at Oxford University on artificial intelligence and law. During those years, he had squirrelled himself

away in a computing laboratory, trying to build machines that could solve legal problems. (In 1988, he went on to co-develop the world's first commercially available AI system in law.) In the decades that followed, his career built upon this work, so I grew up in a home where conundrums about technology were the sorts of things we chewed over in dinner-table conversation.

When I left home, I went to Oxford to study economics. And it was there, for the first time, that I was exposed to the way that economists tend to think about technology and work. It was enchanting. I was enthralled by the tightness of their prose, the precision of their models, the confidence of their claims. It seemed to me that they had found a way to strip away the disorienting messiness of real life and reveal the heart of the problems.

As time passed, my initial enchantment dulled. Eventually, it disappeared. After graduating, I joined the British government – first in the Prime Minister's Strategy Unit, then in the Policy Unit in 10 Downing Street. There, buoyed by technologically inclined colleagues, I started to think more carefully about the future of work and whether the government might have to help in some way. But when I turned for help to the economics I had learned as an undergraduate, it was far less insightful than I had hoped. Many economists, as a matter of principle, want to anchor the stories they tell in past evidence alone. As one eminent economist put it, 'Although we all enjoy science fiction, history books are usually a safer guide to the future.'[16] I was not convinced by this sort of view. What was unfolding in the economy before me looked radically different from experiences of what had come before. I found this very disconcerting.

And so, I left my role in British government and, after time spent studying in America, returned to academia to explore various questions about the future of work. I completed a doctorate in economics, challenging the way that economists had traditionally thought about technology and work, and tried to devise a new way to think about what was happening in the labour market. At the same time, I teamed up with my father to write *The Future of the Professions*, a book that explored the impact of technology on expert white-collar workers – lawyers, doctors, accountants, teachers, and others. When we began our research for that project a decade ago, there was a widespread

presumption that automation would only affect blue-collar workers. It was thought that professionals were somehow immune from change. We challenged that idea, describing how new technologies would allow us to solve some of the most important problems in society – providing access to justice, keeping people in good health, educating our children – without relying on traditional professionals as we had done in the past.[17]

Insights from both my academic research and our book on the professions will reappear in the pages that follow, sanded into better shape through subsequent experience and thinking. In short, then, this book captures my own personal journey, a decade spent thinking almost entirely about one particular issue – the future of work.

GOOD PROBLEMS TO HAVE

Although these opening words may suggest otherwise, this book is optimistic about the future. The reason is simple: in decades to come, technological progress is likely to solve the economic problem that has dominated humanity until now. If we think of the economy as a pie, as economists like to do, the traditional challenge has been to make that pie large enough for everyone to live on. At the turn of the first century AD, if the global economic pie had been divided into equal slices for everyone in the world, each person would have received just a few hundred of today's dollars per year. Most people lived around the poverty line. Roll forward a thousand years, and roughly the same would have been true. Some even claim that, as late as 1800, the average person was no more materially prosperous than her equivalent back in 100,000 BC.[18]

But over the last few hundred years, economic growth has soared, and this growth was driven by technological progress. Economic pies around the world have become much bigger. Today, global GDP per capita, the value of those equally sized individual slices, is already about $10,720 a year (an $80.7 trillion pie shared out among 7.53 billion people).[19] If economies continue to grow at 2 per cent per year, our children will be twice as rich as us. If we expect a measlier 1 per cent annual growth, then our grandchildren will be twice as well off

as we are today. We have, at least in principle, come very close to solving the problem that plagued our fellow human beings in the past. As the economist John Kenneth Galbraith so lyrically put it, 'man has escaped for the moment the poverty which was for so long his all-embracing fate.'[20]

Technological unemployment, in a strange way, will be a symptom of that success. In the twenty-first century, technological progress will solve one problem, the question of how to make the pie large enough for everyone to live on. But, as we have seen, it will replace it with three others: the problems of inequality, power, and purpose. There will be disagreement about how we should meet these challenges, about how we should share out economic prosperity, constrain the political power of Big Tech, and provide meaning in a world with less work. These problems will require us to engage with some of the most difficult questions we can ask – about what the state should and should not do, about the nature of our obligations to our fellow human beings, about what it means to live a meaningful life. But these are, in the final analysis, far more attractive difficulties to grapple with than the one that haunted our ancestors for centuries – how to create enough for everyone to live on in the first place.

Leontief once said that 'if horses could have joined the Democratic party and voted, what happened on farms might have been different.'[21] It is a playful phrase with a serious point. Horses did not have any control over their collective fate, but we do. I am not a technological determinist: I do not think the future must be a certain way. I agree with the philosopher Karl Popper, the enemy of those who believe that the iron rails of our fate have already been set down for us to trundle along, when he says that 'the future depends on ourselves, and we do not depend on any historical necessity.'[22] But I am also a technological realist: I do think that our discretion is constrained. In the twenty-first century, we will build systems and machines that are far more capable than those we have today. I don't believe we can escape that fact. These new technologies will continue to take on tasks that we thought only human beings would ever do. I do not believe we can avoid that, either. Our challenge, as I see it, is to take those unavoidable features of the future as given, and still build a world where all of us can flourish. That is what this book is about.

PART ONE

The Context

1

A History of Misplaced Anxiety

Economic growth is a very recent phenomenon. In fact, for most of the 300,000 years that human beings have been around, economic life has been relatively stagnant. Our more distant ancestors simply hunted and gathered what little they needed to survive, and that was about it.[1] But over the last few hundred years, that economic inactivity came to an explosive end. The amount each person produced increased about thirteen-fold, and world output rocketed nearly 300-fold.[2] Imagine that the sum of human existence was an hour long:

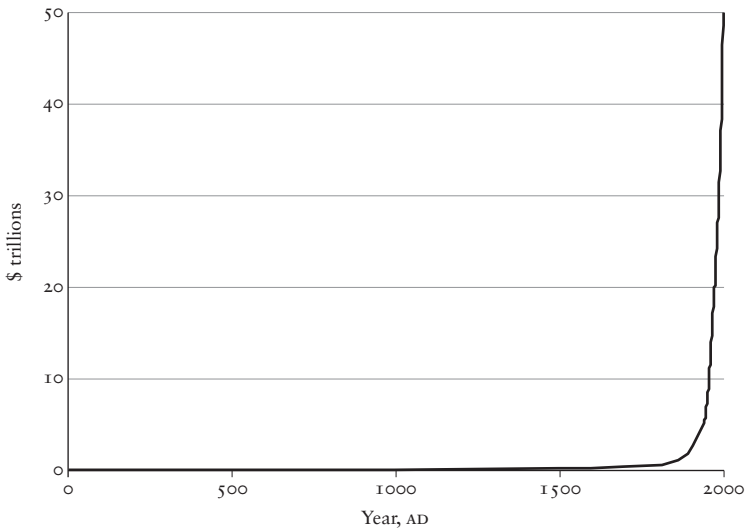

Figure 1.1: Global Output Since AD 1[3]

most of this action happened in the last half-second or so, in the literal blink of an eye.

Economists tend to agree with one another that this growth was propelled by sustained technological progress, though not on the reasons why it started just where and when it did – in Western Europe, towards the end of the eighteenth century.[4] One reason may be geographical: certain countries had bountiful resources, a hospitable climate, and easily traversable coastlines and rivers for trade. Another may be cultural: people in different communities, shaped by very different intellectual histories and religions, had different attitudes towards the scientific method, finance, hard work, and each other (the level of 'trust' in a society is said to be important). The most common explanation of all, though, is institutional: certain states protected property rights and enforced the rule of law in a way that encouraged risk-taking, hustle, and innovation, while others did not.

Whatever the particular reasons, it was Britain that led the economic charge, thundering ahead of others in the 1760s.[5] Over the following decades, new machines were invented and put to use that greatly improved the way that goods were produced. Some, like the steam engine, have become standard symbols of economic progress and technological ingenuity. And dramatic as the term 'revolution' may seem, it is probably still an understatement: the Industrial Revolution is one of the most significant moments in the history of humankind. Before this period, any economic growth had been limited, stuttering, and quickly fizzled out. Afterwards, it started to flow relatively bountifully and steadily. Today, we have become entirely dependent upon this economic fix. Think of the eruptions of anger and anxiety, the waves of frustration and despondency that crash through society each time economic growth stops or even slows. It is as if we can no longer live well without it.

The new technologies of the Industrial Revolution allowed manufacturers to operate more productively than ever before – in short, to make far more with far less.[6] And it is here, at the beginning of modern economic growth, that we can also detect the origins of 'automation anxiety'. People started to worry that using these machines to make more things would also mean less demand for their own work. From