

## Chapter 4: Driverless Democracy

### What Happens to Citizens When AI Takes All the Work?

---

*Science fiction is fast becoming science fact, as rapidly improving artificial intelligence starts to impact our economy. However, rather than speculation about a 'jobless future', we should be worrying about growing inequality and whether the coming tech revolution will wipe out the middle class.*

LIKE MANY SILICON VALLEY start-ups, Starsky Robotics was founded by two twenty-somethings who regard sleep as optional. Any successful new tech firm needs someone who understands technology and someone who understands business, and these two distinct skills are rarely found in one individual. At Starsky, Kartik Tiwari is the 'tech guy' who specialises in artificial intelligence, and Stefan Seltz-Axmacher is the 'serial entrepreneur' guy, who mostly specialises in starting-up start-ups. Neither, you'll note, specialise in trucking, which was the domain of their company. This did not seem to bother investors, since this 11-person business, with ambitions to revolutionise the entire trucking industry by building a self-driving fleet, managed to raise millions of dollars of funding from venture capitalists.

'Everyone thought I was mad,' 27-year-old Stefan told me when I visited Starsky's Florida headquarters, a large rented property in a gated community, a few months ago. These days however, like many other industries, trucking is being disrupted by data science, artificial intelligence and venture capital.

Stefan had agreed to let me drive in Starsky's newest and shiniest truck with their resident driver Tony Hughes, a diminutive and friendly man with 20 years' experience, who is perhaps better described as part-driver, part-machine supervisor. Tony is in his fifties, with a high school diploma in

general studies from Shawnee Mission Northwest (Kansas) and a 'solid track record of achieving efficient, cost-effective transportation operations', but now finds himself training the machines that might eventually put him out of a job. He's spent months driving the Starsky truck up and down certain routes, over and over, so the software can collect data on how he does it. From this experience, it 'learns' how he behaves and how to mimic him. The law requires a human in the cab, and anyway the software is still in 'dev mode', so Tony supervises – there needs to be someone responsible in the cab, in case something goes wrong. Neither Stefan nor Kartik, both wannabe truck magnates, have any idea about how to actually drive one.

Figuring they might as well make some money rather than just transport an empty truck in circles around the country, Starsky vehicles sometimes make actual deliveries while they're training the software. So I jumped aboard 'Rosebud' (everyone names their trucks, apparently) in Orlando, Florida along with Tony, Stefan, Kartik and 5,000 pounds of empty milk cartons. Our destination was a depot in Deerfield Beach, around 200 miles south. Underneath Tony's feet and behind his oversized wheel were work-in-progress wires, pumps, shiny levers and many cogs. They were connected to computers in the back of the cab, which were under Kartik's command. The software controlled the pedals and steering wheel, which constantly adjust to real-time data collected by mounted radars and computer vision sensors that covered the vehicle: position, speed, road markings, other cars' positions and speed and so on.

We left the narrow residential roads, and joined Freeway 95.

Tony turned to Stefan: 'I could kick the system on if you guys are ready.'

'Rosebud on,' Stefan shouted into his walkie-talkie to other crew members, who were following in a car.

Tony flicked a little blue switch, and we went from 'manual' to 'auto'. The truck made a tiny lurch, as if to unshackle itself from Tony, who stretched back and crossed his legs, relaxed. 'I trust Rosebud completely,' he told me. 'I've trained her for months. She can almost drive as well as me now.'

It was, of course, terrifying and exhilarating to thunder along in a 40-tonne HGV without a driver. As the first major bend in the interstate approached, I did what I do when I'm on a flight with turbulence, and looked at the crew. Tony seemed calm. Rosebud took the corner smoothly, of course. I slowly relaxed, and learned that nothing really happens on the long, straight, wide American freeways anyway. 'They're ideal roads for

machines,' Stefan told me as we chugged along. Around 70 per cent of our journey to Deerfield Beach was operated entirely by machine. I soon became a little bored, as novelty gave way to tedium. This felt appropriate; some engineers worry about the risk of 'vigilance decrement', as humans lose practice and become less able to deal with emergencies. My tedium in turn gave way to the dawning realisation that what I'd thought was sci-fi was in fact fast becoming sci-fact.\*

In 2004, respected AI-researchers from the Massachusetts Institute of Technology (MIT) concluded that autonomous vehicles were a pipedream, driving being a skill that required too much human intuition and motor skills.<sup>1</sup> But we should never underestimate the speed at which digital technology can advance. Millions of dollars of investment are now pouring in from Uber, Google, Tesla, Mercedes, Volvo, Starsky and others. Several countries (including the UK) are encouraging 'real world' testing, and the British Chancellor expects that driverless cars will be on the road by 2021. The way it's going, regulations and insurance issues are more likely to slow this down than the technology.

## A world without work?

Driverless vehicles are just one application of an artificial intelligence revolution that is sweeping through the global economy. The leaps forward in AI – which were first gradual and then sudden – are behind the first genuine mass panic of the twenty-first century: that we are entering a world in which robots will take every job currently done by humans, putting us all out of work. Media outlets seem to enjoy writing scary headlines about this.\*

There are a couple of widely held misconceptions about AI that need to be cleared up. Despite the Hollywood movies and the breathless headlines, no machines are remotely close to reaching a human level of intelligence, which we can define as 'performing as well as humans in a series of different domains' (often known as 'general AI'). Although divided, most experts don't think this level of intelligence will be possible for another 50 to 100 years – but to be honest, no one really has a clue. And whether machines will ever achieve consciousness is altogether another question entirely, and probably one best left to philosophers rather than roboticists.

The obsession with marching machines, *Terminator 2*'s Skynet and humanoids unhelpfully distracts from the real action, which is 'domain-specific' AI and often uses a technique known as 'machine learning'. A

human feeds an algorithm with data and teaches it what each input means. From that it can detect patterns, from which it can mimic a particular human behaviour or perform a very specific task, whether driving up the interstate, predicting the weather, giving credit scores, reading license plates and so on.

Machine learning has been around for many years and is already embedded in many aspects of our economies, whether suggesting a next Amazon purchase or a new Facebook friend. ML relies on data to learn and, because we now produce so much of it, it has been able to grow quickly. The falling prices of computing have helped bring about a powerful self-reinforcing feedback loop. More data fed in makes ML better, which allows it to make more sense of new data, which makes it better still, and so on. More sophisticated ML is being developed all the time. The latest involves teaching machines to solve problems for themselves rather than just feeding them examples, by setting out rules and leaving them to get on with it. This is sometimes called ‘deep learning’, which attempts to mimic the activity that occurs in layers of brain neurons through a neural network to spot patterns or images very effectively.<sup>2</sup>

To understand how this is different and potentially more powerful than classic ML, consider the ancient Chinese game Go. Machines have been beating humans at chess for years, but Go is more difficult for machines because of the sheer number of possible moves: in the course of a match, there are more possible combinations than there are atoms in the universe. A few years ago, DeepMind, a Google-owned AI firm, built software to play the game, called AlphaGo. It was trained the ‘classic’ ML way, using thousands of human games; for example, being taught that in position x humans played move y; and in position a, humans played move b, and so on. From that starting point AlphaGo played itself billions of times to improve its knowledge of the game. In 2016, to the surprise of many experts, AlphaGo decisively beat the world’s best Go player, Lee Sedol. This stunning result was quickly surpassed when, in late 2017, Deep Mind released AlphaGo Zero, a software that was given no human examples at all and was taught the rules of how to win by using a deep learning technique with no prior examples. It started off dreadfully bad but improved slightly with each game, and within 40 days of constant self-play it had become so strong that it thrashed the original AlphaGo 100–0. Go is now firmly in the category of ‘games that humans will never win against machines again’.

Most people in Silicon Valley agree that machine learning is the next big thing, although some are more optimistic than others. Tesla and SpaceX

boss Elon Musk recently said that AI is like ‘summoning the demon’, while others have compared its significance to the ‘scientific method, on steroids’, the invention of penicillin and even electricity. Andrew Ng, former chief scientist at Baidu, reckons that there isn’t a single industry that won’t shortly be ‘transformed’.

AIs are starting to outperform humans in a small-but-quietly-growing number of narrow tasks. Over the last year alone inroads have been made into things such as driving, bricklaying, fruit-picking, burger-flipping, banking, trading and automated stock-taking. Legal software firms are developing statistical prediction algorithms that can analyse past cases and recommend trial strategies. Tools to analyse CVs are now routinely used by companies to help them filter out obviously unsuitable candidates (for example, those who do not have the correct A Levels). Using complex data models, software can now make predictions on investment strategies. According to the consulting firm McKinsey & Company, which examined 2,000 work activities in 800 occupations, 45 per cent of tasks people are paid to do at the moment could be done by *currently proven technologies*. Similarly, the Bank of England recently suggested that up to 15 million British jobs might be unnecessary within a generation.

I don’t take these predictions all that seriously. Many of these applications are still young. Every new technological revolution unleashes similar speculation, and it is often wide of the mark. Even our wisest heads get it wrong – back in the 1930s John Maynard Keynes believed that the UK was witnessing ‘technological unemployment’, as the ability of machines to take over jobs outpaced the economy’s ability to generate new ones. We’ve had tech-led disruption before, and we have usually found new (and often better) jobs. After all, machines tend to drive up productivity, which in turn stimulates more investment and demand.<sup>3</sup> A recent analysis of the American workforce between 1982 and 2012 found that employment *grew* in several areas where computers were used (gaming, graphic design and programming).<sup>4</sup> And in many instances, productivity gains driven by technology won’t mean fewer jobs, but rather improvements in current ones. When AI techniques transform medical diagnosis – within the next few years – it won’t mean fewer doctors, but better patient care because our busy doctors won’t need to spend hours staring at scans. We are terrible at predicting what the jobs and industries of the future might be. Millions of people are already employed in roles that didn’t exist 20 years ago: web developers, app designers, Uber drivers, lifestyle coaches and a thousand other things besides.

Gunning down Interstate 95 thinking about these issues while Rosebud took care of the driving, it struck me that the real challenge might not be jobs, but inequality. Driverless trucks will create plenty of great new employment opportunities.\* Starsky's employees are all highly educated, enthusiastic young people doing very good jobs: robotics specialists, engineers, machine learning experts. They have created these jobs, and good for them. If they succeed, however, they will also remove the need for other jobs. Not all drivers will vanish, of course (Stefan says some will sit in an office and control several trucks at once, taking over from the truck while they manoeuvre in and out of busy depots or around complicated roundabouts). Plus, in the US at least, there is still a shortage of truckers. But over the coming years, it's likely that we won't need anywhere near the number of drivers that we currently do.<sup>5</sup>

For Tony, trucking has been his profession since the mid-1990s. It is a difficult and lonely job. But for many Americans without a college degree, it is one of the best-paid jobs available – especially in the poorer states such as Iowa or North Dakota.<sup>6</sup> Three per cent of Americans are employed in a driving job of one sort or another.<sup>7</sup> Will the drivers who lose out take any of the snazzy new jobs created by Stefan and Kartik? A handful might; some, like Tony, are already training the software. Perhaps others will re-skill, and claw their way up to the winners' table. I was told repeatedly in the magical Silicon Valley bubble where everything is possible that unemployed truckers in their fifties could retrain as web-developers and machine learning specialists – a convenient self-delusion that no one really believes. It is far more likely that many truck drivers, without the necessary skills, will drift off to more precarious, piecemeal, low-paid work – perhaps becoming taxi drivers (assuming they still exist) or Amazon warehouse operators or Mechanical Turk labourers who are paid an hourly rate to train software or fill in surveys. Perhaps they could clean the machines that clean the machines that repair the driverless trucks that they once occupied.\*

## Routine and non

We should be reasonably confident that AI will result in forward leaps in productivity and overall wealth. The big question is how the spoils of that wealth will be shared out. Just because machines can beat other machines that beat us at board games, it does not follow that they can undertake every useful task that humans currently perform. Many of the 'thinking' things humans do, such as extremely complicated calculations, can be easily

replicated by machines. They can undertake routine or predictable tasks at a speed and rate of accuracy far beyond anything we can manage. Conversely, a lot of things we do unthinkingly, like picking up a deck of dropped cards or tying shoelaces, are far harder for machines. They are significantly worse than humans at dealing with unpredictable situations, especially ones requiring sensorimotor skills. This is sometimes known as Moravec's Paradox, after the roboticist Hans Moravec who realised that high-level reasoning often requires little computational power, but low-level sensorimotor skills need a lot.\* Not to mention, of course, that humans still win hands-down on breadth and range of intelligence.

This has some important ramifications. It means that the jobs most at risk will be those involving the routine tasks that can be most easily done by machines. The safest, and most likely to be created, work will be 'non-routine'. The strange thing about our economy is that non-routine jobs tend to be either very well or very badly paid. A specialist in machine learning at Starsky Robotics or Google performs a non-routine job, since it involves a lot of intuition, creativity and independent thinking in unpredictable situations. So does a gardener, carer or Deliveroo cyclist. It's the jobs in the middle – what you might call 'routine cognitive' jobs – that will be most at risk. If you are a train operator, a mortgage adviser, a stock analyst, a paralegal, a credit analyst, a loan officer, a bookkeeper, a tax accountant or a radiologist, you might consider retraining.

If we lose these jobs, we will be heading towards what David Autor, a labour market economist at MIT, calls a 'barbell-shaped economy', a kind of extreme inequality. Anyone with the skills, diligence, money or fortune to work closely with technology – and especially AI – will likely see huge leaps in productivity and wages. There will be plenty more low-paid, insecure service jobs that won't disappear due to automation, too. But with millions of people competing to wait on, care for and serve food to the winners of the great tech revolution, wages for those jobs will be driven down further. The labour market in 2030: either a well-paid job at Facebook, or the opportunity to deliver food on a bicycle on minimum wage to these busy and important people. But good luck trying to get a stable non-cognitive job as a local news journalist, paralegal, truck driver or tax accountant.

I don't wish to put AI in the dock for crimes as yet uncommitted. But this kind of tech-fuelled inequality is a familiar pattern. 'There is no economic law that says that all workers, or even a majority of workers, will benefit from technological progress,' write McAfee and Brynjolfsson in their influential book *The Second Machine Age*, which argues persuasively that,

while other factors are of course at play – including globalisation – technological advance over the last 30 years has been the main factor behind growing economic inequality. Skilled workers, they explain, tend to benefit most from new technologies, while others fall further behind. In the US, productivity has been rising, shiny new buildings are being built and corporate profits are increasing, but average salaries are falling. In a similar fashion, productivity in the UK increased by 80 per cent between 1973 and 2011 (although it is still low by the standards of the Organisation for Economic Co-operation and Development) but the hourly compensation of the median worker went up by only 10 per cent in real terms. All over the world – including in socialist Sweden and Mittelstand Germany – top earners and top jobs have been doing just fine, while for a lot of people in the middle and bottom, earnings and wealth haven't increased at all in real terms since the 1970s.

There are other forms of inequality that no one is thinking about at play here too. As a general rule, technology empowers those who have either the money or the skills to take advantage of it. The more powerful the tech, the more powerful the tendency. So let's say that people start owning personal AI assistants, as I outlined in Chapter One. These AI assistants would help their owners identify the best prices, the cheapest holidays, the most sound legal advice, write the strongest CVs and so on. Those who can afford the best AI assistants will see their prospects rise further, while those without will lag further behind. What makes this sort of inequality especially troublesome is that, unlike wage levels or house values, which are collected by government departments or academics, this would be a form of disadvantage that would be extremely difficult to spot.

In addition to favouring more skilled workers, digital technology also increases the financial returns to capital owners over labour. Machines don't demand a share of the profits, which means any machine-driven productivity gains accrue to whoever owns them, and that's usually the wealthy. It is of note that the percentage share of GDP going to labour relative to capital has been falling in recent years. For much of the twentieth century, the ratio of national wealth in the US between labour and capital was 66/33. It is now 58/42. The great defence against these trends for much of the twentieth century were trade unions, which ensured that the spoils of corporate profit were spread around. The unions' slow decline has been disastrous for wealth equality and – in a cruel twist – new technology is likely to further militate against worker unionisation, by both making it harder for 'gig economy' workers to band together, and by giving bosses new ways to monitor and



control their workforce.\* At the most extreme end of this economic bifurcation, the world's richest eight men own more than the bottom half of the world's population – and four of them are the founders of technology companies.<sup>8</sup>

• • •

This is not a book about the economics of digital technology – there are plenty of those already – but about politics. While some degree of inequality is inevitable and necessary in a free market economy, too much is bad for democracy. It is well-documented that a healthy democracy depends on a vibrant, sizeable middle class. This sector of a population is democracy's backbone – the people that buy newspapers, join political parties, sponsor charities, vote and participate in community projects. Based on decades of research into the impact of inequality, we can predict with some confidence that a 'barbell economy' would result in, among other things: a shrinking tax base, growing levels of crime, depression, addiction and infant mortality, lower life expectancy and poorer health.<sup>9</sup>

High levels of inequality also wear away the fabric of society. The more unequal life gets, the less we spend time with people not like ourselves, and the less we trust each other. Ironically, trusting societies are more likely to be innovative and entrepreneurial, since they see the world as full of reliable people and good opportunities.<sup>10</sup> Most importantly of all, the middle classes are democracy's most fervent supporters. One of the reasons Marx's projected revolution didn't take place in the UK – somewhere he expected it to occur first – was because the working classes evolved into a broad, asset-owning middle class. Having something to protect and a stake in society, this group is repeatedly found in studies to value individual freedom, property rights and democratic accountability more than other groups.<sup>11</sup> The emergence of middle-class societies, especially in Europe and America increased the legitimacy of liberal democracy as a political system in both the nineteenth and twentieth centuries.<sup>12</sup>

To see what happens when tech-fuelled inequality takes off, there's no better place to start than the home of it all, Silicon Valley and its increasingly put-upon neighbour, San Francisco. There are two worlds in Silicon Valley, and they barely ever meet. There's the exciting start-up open-plan offices, with beanbags, table football, TED Talks and flip-flops, where the region's half a million tech workers can expect to earn on average

well over a hundred thousand dollars a year. (For the biggest companies, the median salary is higher still.) Mostly under 40, they want to live in nearby bustling San Francisco, since Silicon Valley can resemble *The Stepford Wives*. Each morning, thousands of tech workers hop on private, Wi-Fi-enabled coaches from one of the dozens of pick-up points in San Fran's increasingly gentrified streets, and head down Highway 101 into Menlo Park (for Facebook), Sunnyvale (for Yahoo) or Mountain View (for Google). It's impossible to ignore the buzz, the thrill, and the enterprise of the place.

Alongside it, though, is another world, inhabited by the people who are left behind in the mad rush towards progress: the ignored women in tech start-ups who complain about misogyny, the Uber drivers who can only afford to live 70 miles away and have to work on zero-hour contracts, the long-time residents who are turfed out so their landlords can rent out their homes on Airbnb. It's a place where minorities struggle on low-wage service jobs, serving the largely white affluent tech workers. The median house value in both San Francisco and Silicon Valley is now around a million dollars, and the average rent is over three thousand per month for a two-bedroom flat: beyond the reach of almost anyone but tech workers. (The average salary in San Francisco is \$46,000, and less if you don't work for Facebook, Google, etc.) In one of the richest cities in America, 15,000 registered homeless people – one of the highest numbers per capita in the country – struggle on, often with serious mental health and drug addiction problems. This has been a long-standing problem in San Francisco (and in much of California), though when I was there recently, locals told me that it has never been as bad as it is now. There are parts of this glorious and gleaming metropolis that reek of destitution, used needles, human waste and food banks, some of it literally in the shadows of the world's biggest and coolest companies. One morning I witnessed junkies openly shooting up on a busy pavement: it wasn't yet 9a.m. And, on the same street, techies wearing white earbuds entered the gleaming offices of a company that promises to let you 'belong anywhere'.

## **Epilogue: Universal Basic Income**

At some point all this creative destruction becomes bad even for the winners. No one wants to live in a world comprising a handful of trillionaires and hordes of unemployed or extremely poorly paid people – not even the trillionaires. A growing number of people are proposing a bold new idea to deal with this.

In 2017 I interviewed Sam Altman, the president of Y Combinator, the most important fund in Silicon Valley for tech start-ups. Thousands of businesses apply every year to access Y Combinator's funding and guidance, in exchange for a small slice of their company. Sam is a Princeton dropout and frequently wears a hoodie, yet when I met him, he was only 31 years old and already a multi-millionaire. He is often described as 'the man who invents the future'. The companies Y Combinator have funded include Airbnb and Starsky Robotics, and are now altogether valued at \$80 billion.

Aware of the potential turbulence that AI might unleash, Y Combinator recently started to fund a pilot in universal basic income. UBI, as it is commonly referred to, is an increasingly popular idea to deal with the possible rise of joblessness and tech-fuelled inequality. The basic concept is that governments should give everyone enough money to live on, with no strings attached. Several pilot schemes, including Oakland, California and Finland, are examining the idea (although it's too early to say how well they are working yet), and a number of serious thinkers and writers believe it is worth further investigation. In other words, UBI has become very fashionable. For some people on the political right, it is a way to keep capitalism ticking over in times of economic uncertainty. For some on the left, including a handful of radicals in the UK circling Labour leader Jeremy Corbyn, it represents a way to redistribute wealth more fairly. And for the utopians, it would allow people to do more meaningful things with their lives than monotonous labour.\*

Sam doesn't think anyone is ready for AI. 'We are going to need to have new distribution, new social safety nets,' he told me in his Y Combinator office. 'What happens if you just give people money to live on? . . . Say, "Here's enough money to have a house and eat and have fun".'

It's an interesting idea. There are an awful lot of jobs that people don't really want to do. If the things that people gain from work – economic means, structure, purpose – can be achieved in other ways, that's worth exploring. Proponents of UBI argue that this would be a 'basic' income – and not necessarily a replacement for work. Some people would work while others would choose not to, and could instead dedicate their time to other things.

However, I don't see how UBI would stop a tiny band of elites from becoming even wealthier than everyone else. And who would pay for UBI is not clear to me, especially at a time when the tech firms seem keen to pay as little tax as possible. (In fact, the issue of how to pay for UBI is a quite interesting problem. If you were to divide up all current US spending on

social welfare by capita, it would amount to only \$2,300 per person per year, which is clearly not enough. Advocates for UBI rely on assumptions about the falling future costs of essential goods. Sam Altman, for example, in a 2016 discussion on the subject, said that it would be affordable in the future because of huge increases in productivity and a reduction in the cost of necessities. I doubt this would be a strong enough basis to persuade most policy-makers in government). ‘It strains credulity,’ writes tech critic Nicholas Carr, ‘to imagine today’s technology moguls, with their libertarian leanings and impatience with government, agreeing to the kind of vast wealth distribution scheme that would be necessary.’<sup>13</sup> Certainly their behaviour to date does not bode well.

I asked Sam what felt to me like a very simple question: would people really be happy living in a society in which there are a small number of very rich people, and everyone else is given money to keep them occupied. What about dignity in work? What about growing inequality?

‘You have a very pessimistic view of the future,’ he replied. ‘I hope you’re wrong. I believe that someone doing mechanical labour is not the best fulfilment of their dreams and aspirations.’

Neither do I, of course, but that’s not really the point. ‘The thing that makes me nervous is that society will have to change dramatically, and that’s quite worrying,’ I said.

‘I believe society *will* have to change dramatically,’ he replied. ‘We’ve been through many of these changes before. Look, I understand that people have this spirit of “I’m going to hang on to the past at all costs [At this, Sam clenched his fists, looked up to the sky, and waved his arms] and I hate progress and I hate change”, and I get that from you.’

‘It’s not that!’ I cut back in. ‘It’s not hating progress. What if the progress that you are creating is not what other people want?’

‘There are 40 million people in the US that live in poverty,’ he said. ‘If technology can eliminate human suffering, we should do that; if technology can generate more wealth and we can figure out how to distribute it better, we should do that.’ There was no hint that tech has played some role in creating the problem that tech is now supposed to fix. I mentioned that journalists have to ask about the negative possibilities. That’s our job.

‘If you continue this thrust of “we should stop progress”, no one is going to take you seriously,’ he replied. ‘I think you can add an important voice, but I worry you are going in the wrong direction with this anti-progress angle.’

With that, the interview was over. Sam is considered one of Silicon Valley's most interesting and original thinkers, and there's no doubting his sharp mind. But I suspect on the current trajectory, the more likely scenario is a relentless drive toward a more unequal economy. The winners always convince themselves that they deserve their small fortunes and that they are helping society become more connected, even as they profit from tearing it apart. The dystopia we should fear is not robots with all the jobs, but a barbell-shaped economy where socially progressive tech millionaires live in gated communities well away from the masses who they either fear, patronise or detest. The feeling would, of course, be mutual.