London Review of Books

The Robots Are Coming

John Lanchester

In 1996, in response to the 1992 Russo-American moratorium on nuclear testing, the US government started a programme called the Accelerated Strategic Computing Initiative. The suspension of testing had created a need to be able to run complex computer simulations of how old weapons were ageing, for safety reasons, and also – it's a dangerous world out there! – to design new weapons without breaching the terms of the moratorium. To do that, ASCI needed more computing power than could be delivered by any existing machine. Its response was to commission a computer called ASCI Red, designed to be the first supercomputer to process more than one teraflop. A 'flop' is a floating point operation, i.e. a calculation involving numbers which include decimal points (these are computationally much more demanding than calculations involving binary ones and zeros). A teraflop is a trillion such calculations per second. Once Red was up and running at full speed, by 1997, it really was a specimen. Its power was such that it could process 1.8 teraflops. That's 18 followed by 11 zeros. Red continued to be the most powerful supercomputer in the world until about the end of 2000.

I was playing on Red only yesterday – I wasn't really, but I did have a go on a machine that can process 1.8 teraflops. This Red equivalent is called the PS3: it was launched by Sony in 2005 and went on sale in 2006. Red was only a little smaller than a tennis court, used as much electricity as eight hundred houses, and cost $55 million. The PS3 fits underneath a television, runs off a normal power socket, and you can buy one for under two hundred quid. Within a decade, a computer able to process 1.8 teraflops went from being something that could only be made by the world's richest government for purposes at the furthest reaches of

computational possibility, to something a teenager could reasonably expect to find under the Christmas tree.

The force at work here is a principle known as Moore's law. This isn't really a law at all, but rather the extrapolation of an observation made by Gordon Moore, one of the founders of the computer chip company Intel. By 1965, Moore had noticed that silicon chips had for a number of years been getting more powerful, in relation to their price, at a remarkably consistent rate. He published a paper predicting that they would go on doing so 'for at least ten years'. That might sound mild, but it was, as Erik Brynjolfsson and Andrew McAfee point out in their fascinating book, *The Second Machine Age*, actually a very bold statement, since it implied that by 1975, computer chips would be five hundred times more powerful for the same price. 'Integrated circuits,' Moore said, would 'lead to such wonders as home computers – or at least terminals connected to a central computer – automatic controls for automobiles and personal portable communications equipment'. Right on all three. If anything he was too cautious. Moore's law, now usually stated as the principle that computer chips double in power or halve in price every 18 months, has continued operating for half a century. This is what has led to such now ordinary miracles as the jump from the Red to the PS3. There has never been an invention in human history which has improved at such speed over such a long period.

In parallel with computing power having grown exponentially and become vastly cheaper, humans have also got better at programming. A high-profile piece of evidence to that effect came in 2011, with the public victory of a project called Watson, run by IBM. The idea behind Watson was to build a computer that could understand ordinary language well enough to win a popular TV quiz show called *Jeopardy!*, playing against not just ordinary contestants, but record-holding champions.[1] This would be, as Brynjolfsson and McAfee say, 'a stern test of a computer's pattern matching and complex communication abilities', much more demanding than another IBM project, the chess computer Deep Blue, which won a match against the world champion Gary Kasparov in 1997. Chess is vulnerable to brute computational force; I have a program on my smartphone which can easily beat the best chess player in the world. General knowledge-based quizzes, particularly ones like *Jeopardy!* with a colloquial and allusive component, are much less easily solved by sheer computing power.

The outcome is already a locus classicus in the study of computing, robotics and futurism, and is discussed at length in both John Kelly and Steve Hamm's *Smart Machines* and Tyler Cowen's *Average Is Over*.[2] Watson won, easily. Its performance wasn't perfect: it thought Toronto was in the US, and when asked about a word with the double meaning 'stylish elegance, or students who all graduated in the same year', it answered 'chic', not 'class'. Still,

its cash score at the end of the two-day contest was more than three times higher than that of the best of its human opponents. 'Quiz-show contestant may be the first job made redundant by Watson,' one of the vanquished men said, 'but I'm sure it won't be the last.'

Watson's achievement is a sign of how much progress has been made in machine learning, the process by which computer algorithms self-improve at tasks involving analysis and prediction. The techniques involved are primarily statistical: through trial and error the machine learns which answer has the highest probability of being correct. That sounds rough and ready, but because, as per Moore's law, computers have become so astonishingly powerful, the loops of trial and error can take place at great speed, and the machine can quickly improve out of all recognition. The process can be seen at work in Google's translation software. Translate was a page on Google into which you could type text and see it rendered into a short list of other languages. When the software first launched, in 2006, it was an impressive joke: impressive because it existed at all, but a joke because the translations were wildly inaccurate and syntactically garbled. If you gave up on Google Translate at that point, you have missed many changes. The latest version of Translate comes in the form of a smartphone app, into which you can not only type but also speak text, and not just read the answer but also have it spoken aloud. The app can scan text using the phone's camera, and translate that too. For a language you know, and especially with text of any length, Translate is still somewhere between poor and embarrassing – though handy nonetheless, if you momentarily can't remember what the German is for 'collateralised debt obligation' or 'haemorrhoid'. For a language you don't know, it can be invaluable; and it's worth taking a moment to reflect on the marvel that you can install on your phone a device which will translate Malay into Igbo, or Hungarian into Japanese, or indeed anything into anything, for free.

Google Translate hasn't got better because roomfuls of impecunious polymaths have been spending man-years copying out and cross-referencing vocabulary lists. Its improvement is a triumph of machine learning. The software matches texts in parallel languages, so that its learning is a process of finding which text is statistically most likely to match the text in another language. Translate has hoovered up gigantic quantities of parallel texts into its database. A particularly fertile source of these useful things, apparently, is the European Union's set of official publications, which are translated into all Community languages. There was a point a few years ago when the software, after improving for a bit, stopped doing so, as the harvesting of parallel texts began to gather in texts which had already been translated by Translate. I don't know how, but they must have fixed that problem, because it's been getting better again. You could argue that this isn't really 'learning' at all, and indeed it probably isn't in any human sense. The process is analogous, though, in terms of the outcome, if that outcome is defined as getting better at a specific task.

Put all this together, and we can start to see why many people think a big shift is about to come in the impact of computing and technology on our daily lives. Computers have got dramatically more powerful and become so cheap that they are effectively ubiquitous. So have the sensors they use to monitor the physical world. The software they run has improved dramatically too. We are, Brynjolfsson and McAfee argue, on the verge of a new industrial revolution, one which will have as much impact on the world as the first one. Whole categories of work will be transformed by the power of computing, and in particular by the impact of robots.

*

For many years the problem with robots has been that computers are very good at things we find difficult but very bad at things we find easy. They are brilliant at chess but terrible at the cognitive skills we take for granted, one of the most important being something scientists call SLAM, for 'simultaneous localisation and mapping': the ability to look at a space and see it and know how to move through it, all simultaneously, and with good recall. That, and other skills essential to advanced robotics, is something computers are useless at. A robot chess player can thrash the best chess player in the world, but can't (or couldn't) match the motor and perceptual skills of a one-year-old baby. A famous demonstration of the principle came in 2006, when scientists at Honda staged a public unveiling of their amazing new healthcare robot, the Asimo. Asimo is short (4'3") and white with a black facemask and a metal backpack. It resembles an unusually small astronaut. In the video Asimo advances towards a staircase and starts climbing while turning his face towards the audience as if to say, à la Bender from *Futurama*, 'check out my shiny metal ass'. He goes up two steps and then falls over. Tittering ensues. It is evident that a new day in robotics has not yet dawned.

That, though, was nine years ago, and Moore's law and machine learning have been at work. The new generation of robots are not ridiculous. Take a look online at the latest generation of Kiva robots employed by Amazon in the 'fulfilment centres' where it makes up and dispatches its parcels. (Though pause first to enjoy the full resonance of 'fulfilment centres'.) The robots are low, slow, accessorised in a friendly orange. They can lift three thousand pounds at a time and carry an entire stack of shelves in one go. Directed wirelessly along preprogrammed paths, they swivel and dance around each other with surprising elegance, then pick up their packages according to the instructions printed on automatically scanned barcodes. They are not alarming, but they are inexorable, and they aren't going away: the labour being done by these robots is work that will never again be done by people. It looks like the future predicted by Wassily Leontief, a Nobel laureate in economics, who said in 1983 that 'the role of humans as the most important factor of production is bound to diminish in the same way that the role of horses in agricultural production was first diminished and then eliminated by the

introduction of tractors.'

Large categories of work, especially work that is mechanically precise and repetitive, have already been automated; technologists are working on the other categories, too. Brynjolfsson and McAfee:

> Rodney Brooks, who co-founded iRobot, noticed something else about modern, highly automated factory floors: people are scarce, but they're not absent. And a lot of the work they do is repetitive and mindless. On a line that fills up jelly jars, for example, machines squirt a precise amount of jelly into each jar, screw on the top, and stick on the label, but a person places the empty jars on the conveyor belt to start the process. Why hasn't this step been automated? Because in this case the jars are delivered to the line 12 at a time in cardboard boxes that don't hold them firmly in place. This imprecision presents no problem to a person (who simply sees the jars in the box, grabs them, and puts them on the conveyor belt), but traditional industrial automation has great difficulty with jelly jars that don't show up in exactly the same place every time.

It's that problem, and others like it, that many observers think robots are beginning to solve. This isn't just a First World issue. The Taiwanese company Foxconn is the world's largest manufacturer of consumer electronics. If you're reading this on an electronic gadget, there is a good chance that it was made in one of Foxconn's factories, since the firm makes iPhones, iPads, iPods, Kindles, Dell parts, and phones for Nokia and Motorola and Microsoft. It employs about 1.2 million people around the world, many of them in China. At least that's how many it currently employs, but the company's founder, Terry Gou, has spoken of an ambition to buy and deploy a million robots in the company's factories. This is nowhere near happening at the moment, but the very fact that the plan has been outlined makes the point: it isn't only jobs in the rich part of the world that are at risk from robots. The kind of work done in most factories, and anywhere else that requires repetitive manual labour, is going, going, and about to be gone.

And it's not just manual labour. Consider this report from the Associated Press:

> CUPERTINO, Calif. (AP) Apple Inc. (AAPL) on Tuesday reported fiscal first-quarter net income of $18.02 billion.

> The Cupertino, California-based company said it had profit of $3.06 per share.

> The results surpassed Wall Street expectations ... The maker of iPhones, iPads and other products posted revenue of $74.6 billion in the period, also exceeding

Street forecasts. Analysts expected $67.38 billion.

For the current quarter ending in March, Apple said it expects revenue in the range of $52 billion to $55 billion. Analysts surveyed by Zacks had expected revenue of $53.65 billion.

Apple shares have declined 1 per cent since the beginning of the year, while the Standard & Poor's 500 index has declined slightly more than 1 per cent. In the final minutes of trading on Tuesday, shares hit $109.14, an increase of 39 per cent in the last 12 months.

We'll be returning to the content of that news story in a moment. For now, the fact to concentrate on is that it wasn't written by a human being. This has been a joke or riff for so long – such and such 'reads like it was written by a computer' – that it's difficult to get one's head around the fact that computer-generated news has become a reality. A company called Automated Insights owns the software which wrote that AP story. Automated Insights specialises in generating automatic reports on company earnings: it takes the raw data and turns them into a news piece. The prose is not Updikean, but it's better than E.L. James, and it gets the job done, since that job is very narrowly defined: to tell readers what Apple's results are. The thing is, though, that quite a few traditionally white-collar jobs are in essence just as mechanical and formulaic as writing a news story about a company earnings report. We are used to the thought that the kind of work done by assembly-line workers in a factory will be automated. We're less used to the thought that the kinds of work done by clerks, or lawyers, or financial analysts, or journalists, or librarians, can be automated. The fact is that it can be, and will be, and in many cases already is. Tyler Cowen's *Average Is Over* points towards a future in which all the rewards are likely to be captured by people at the top of the income distribution, especially those who become most adept at working with smart machines.

So what's going to happen now? Your preferred answer depends on your view of history, though it also depends on whether you think the lessons of history are useful in economics. The authors of these books are interested in history, but plenty of economists aren't; a hostility to history is, to an outsider, a peculiarly strong bias in the field. It's connected, I suspect, to an ambition to be considered a science. If economics is a science, the lessons of history are 'in the equations' – they are already incorporated in the mathematical models. I don't think it's glib to say that a reluctance to learn from history is one of the reasons economics is so bad at predicting the future.

One historically informed view of the present moment says that the new industrial revolution has already happened. Computers are not a new invention, yet their impact on economic

growth has been slow to manifest itself. Bob Solow, another Nobel laureate quoted by Brynjolfsson and McAfee, observed as long ago as 1987 that 'we see the computer age everywhere, except in the productivity statistics.' The most thorough and considered version of this argument is in the work of Robert Gordon, an American economist who in 2012 published a provocative and compelling paper called 'Is US Economic Growth Over?' in which he contrasted the impact of computing and information technology with the effect of the second industrial revolution, between 1875 and 1900, which brought electric lightbulbs and the electric power station, the internal combustion engine, the telephone, radio, recorded music and cinema.[3] As he points out in a *Wall Street Journal* op-ed, it also introduced 'running water and indoor plumbing, the greatest event in the history of female liberation, as women were freed from carrying literally tons of water each year'. (A non-economist might be tempted to ask why it was the women were carrying the water in the first place.) Gordon's view is that we coasted on the aftermaths and sequelae of these inventions until about 1970, when

> the computer revolution took over and allowed the economy to remain on our historic path of 2 per cent annual growth. Computers replaced human labour and thus contributed to productivity, but the bulk of these benefits came early in the Electronics Era. In the 1960s, mainframe computers churned out bank statements and telephone bills, reducing clerical labour. In the 1970s, memory typewriters replaced repetitive retyping by armies of legal clerks. In the 1980s, PCs with word-wrap were introduced, as were ATMs that replaced bank tellers and barcode scanning that replaced retail workers.

These were real and important changes, and got rid of a lot of drudgery. What happened subsequently, though, with the impact of Moore's law and miniaturisation, was a little different:

> The iPod replaced the CD Walkman; the smartphone replaced the garden-variety 'dumb' cellphone with functions that in part replaced desktop and laptop computers; and the iPad provided further competition with traditional personal computers. These innovations were enthusiastically adopted, but they provided new opportunities for consumption on the job and in leisure hours rather than a continuation of the historical tradition of replacing human labour with machines.

In other words, most of the real productivity benefits of the computing revolution happened a few decades ago. We have more and cooler devices, but what these gadgets do, for the most part, is entertain and distract us. They do nothing to aid productivity, and may even diminish

it. The lightbulb changed the world; Facebook is just a way of letting people click 'like' on photos of cats that resemble Colonel Gaddafi. On this view, Moore's law has mainly led to an explosion of digital activity of a not very consequential type. Real change would involve something like a ten or hundredfold increase in the potential of batteries; but that requires progress in chemistry, which is a lot harder than cramming more circuits into a silicon chip.

Gordon's analysis is in line with a longstanding vein of thinking concerning 'technological unemployment'. The term was coined by Keynes to describe 'our discovery of means of economising the use of labour outrunning the pace at which we can find new uses for labour'. It's a form of progress which makes jobs go away through the sheer speed of its impact. A foundational axiom of economics is that economic processes are based on human wants; and since human wants are infinite, the process of supplying them is also infinite. The economy won't come to a halt until human wants do. Since that will never happen, there will be enough work for everyone, for ever, except during occasional recessions and depressions and crises.

The reassuring lesson from history seems to confirm that. Although it might be possible in theory for some new invention to come along and make a category of work disappear so quickly that there is no alternative work to replace it, in practice that hasn't happened. Innovation takes away some jobs and replaces them with others. Just to be clear: the disappearance of work happens to individuals, not to entire economies. A job lost in one place is replaced by a new job, which may be somewhere else. In 1810, agriculture employed 90 per cent of the American workforce. A hundred years later, the figure was about 30 per cent (today it's less than 2 per cent). That might sound like a recipe for chaotic disruption and endemic unemployment, but the US economy managed the transition fine, thanks in large part to the effect of the technologies mentioned by Robert Gordon (plus the railways). So, by extension and analogy, maybe we don't need to fear technological unemployment this time either.

*

That isn't the view forcefully put by Brynjolfsson, McAfee and Cowen. In effect, the two sides disagree about where we are in relation to the impact of the IT revolution: Gordon thinks it's already happened, the other writers that it's about to. The point they all make – which, again, is an argument by extension and analogy – is that the big effects of the first industrial revolution took a long time to arrive. Watt improved the efficiency of the steam engine by 300 per cent between 1765 and 1776, but it took decades for that transformation to have its full effect on the economy – railroads didn't arrive in their full glory until the latter part of the 19th century. In Gordon's view, it took at least 150 years for the first industrial revolution to have its 'full range of effects'. He doesn't see a similar trajectory this time, but what if he's wrong? What if we are a few decades into the new industrial revolution, living in a

contemporary version of 1780, a few decades after the first successful steam engine (1712) but still some way before the first commercially successful steam-powered railway journey (1804)?

It's a good question. What if that's where we are, and – to use the shorthand phrase relished by economists and futurists – 'robots are going to eat all the jobs'? A thorough, considered and disconcerting study of that possibility was undertaken by two Oxford economists, Carl Benedikt Frey and Michael Osborne, in a paper from 2013 called 'The Future of Employment: How Susceptible Are Jobs to Computerisation?'[4] They came up with some new mathematical and statistical techniques to calculate the likely impact of technological change on a sweeping range of 702 occupations, from podiatrists to tour guides, animal trainers to personal finance advisers and floor sanders. It ranks them, from 1 (you'll be fine) to 702 (best start shining up the CV). In case you're wondering, here are the top five occupations:

1. Recreational Therapists
2. First-Line Supervisors of Mechanics, Installers and Repairers
3. Emergency Management Directors
4. Mental Health and Substance Abuse Social Workers
5. Audiologists

And here are the bottom five:

698. Insurance Underwriters
699. Mathematical Technicians
700. Sewers, Hand
701. Title Examiners, Abstractors and Searchers
702. Telemarketers

The theme is clear: human-to-human interaction and judgment is in demand, routine tasks are not. Some of the judgments seem odd: is it really the case that choreographers come in at 13, ahead of physicians and surgeons at 15, and a long way ahead of, say, anthropologists and archaeologists at 39, not to mention writers at 123 and editors at 140? Nonetheless, the paper's methodology is sober and it makes clear just how far-ranging the impact of technological change is in white as well as blue-collar work. Software is having a big impact on the legal profession, for instance: the work of scanning documents for forensic mentions and matches is dramatically cheaper when done by machines. (I've seen specialist legal search software at work, and it's really something.) There are similar things happening in financial services and medical record-keeping. The news that *Average Is Over* is, for most of

us, bad news, since most of us, by definition, are average.

Frey and Osborne's conclusion is stark. In the next two decades, 47 per cent of employment is 'in the high-risk category', meaning it is 'potentially automatable'. Interestingly, though not especially cheeringly, it is mainly less well-paid workers who are most at risk. Recent decades have seen a polarisation in the job market, with increased employment at the top and bottom of the pay distribution, and a squeeze on middle incomes. 'Rather than reducing the demand for middle-income occupations, which has been the pattern over the past decades, our model predicts that computerisation will mainly substitute for low-skill and low-wage jobs in the near future. By contrast, high-skill and high-wage occupations are the least susceptible to computer capital.' So the poor will be hurt, the middle will do slightly better than it has been doing, and the rich – surprise! – will be fine.

Note that in this future world, productivity will go up sharply. Productivity is the amount produced per worker per hour. It is the single most important number in determining whether a country is getting richer or poorer. GDP gets more attention, but is often misleading, since other things being equal, GDP goes up when the population goes up: you can have rising GDP and falling living standards if the population is growing. Productivity is a more accurate measure of trends in living standards – or at least, it used to be. In recent decades, however, productivity has become disconnected from pay. The typical worker's income in the US has barely gone up since 1979, and has actually fallen since 1999, while her productivity has gone up in a nice straightish line. The amount of work done per worker has gone up, but pay hasn't. This means that the proceeds of increased profitability are accruing to capital rather than to labour. The culprit is not clear, but Brynjolfsson and McAfee argue, persuasively, that the force to blame is increased automation.

That is a worrying trend. Imagine an economy in which the 0.1 per cent own the machines, the rest of the 1 per cent manage their operation, and the 99 per cent either do the remaining scraps of unautomatable work, or are unemployed. That is the world implied by developments in productivity and automation. It is Pikettyworld, in which capital is increasingly triumphant over labour. We get a glimpse of it in those quarterly numbers from Apple, about which my robot colleague wrote so evocatively. Apple's quarter was the most profitable of any company in history: $74.6 billion in turnover, and $18 billion in profit. Tim Cook, the boss of Apple, said that these numbers are 'hard to comprehend'. He's right: it's hard to process the fact that the company sold 34,000 iPhones every hour for three months. Bravo – though we should think about the trends implied in those figures. For the sake of argument, say that Apple's achievement is annualised, so their whole year is as much of an improvement on the one before as that quarter was. That would give them $88.9 billion in profits. In 1960, the most profitable company in the world's biggest economy was General

Motors. In today's money, GM made $7.6 billion that year. It also employed 600,000 people.
Today's most profitable company employs 92,600. So where 600,000 workers would once generate $7.6 billion in profit, now 92,600 generate $89.9 billion, an improvement in profitability per worker of 76.65 times. Remember, this is pure profit for the company's owners, after all workers have been paid. Capital isn't just winning against labour: there's no contest. If it were a boxing match, the referee would stop the fight.

*

Under the current political and economic dispensation, automation is certain to reinforce these trends. Consider the driverless car being developed by Google. This is both miraculous, in that to an amazing extent it already works, and severely limited, in that there are many routine aspects of driving that it can't manage – it can't, for instance, overtake, or 'merge' with flowing traffic, which is no small issue in the land of the freeway. But imagine for a moment that all the outstanding technical issues are solved, and the fully driverless car is a reality. It would be astonishing, especially when/if it were combined with clean energy sources. Your car would take your children on the school run while they scramble to finish their homework, then come home to take you to work while you do your email, then drive off and self-park somewhere, then pick you up at the end of the day and take you to dinner, then drive you home while you sleep off that last regrettable tequila slammer, and all – thanks to self-co-ordinating networks of traffic information from other driverless cars – cleanly and frictionlessly. It's not clear that the car would even need to be 'your' car: it would just have to be a vehicle that you could summon whenever you needed it. This isn't just an urbanist vision, since there is a lot of immobility and isolation in the countryside that would be immeasurably helped by the driverless car.

The catch: all the money would be going to Google. An entire economy of drivers would disappear. The UK has 231,000 licensed cabs and minicabs alone – and there are far, far more people whose work is driving, and more still for whom driving is not their whole job, but a big part of what they are paid to do. I suspect we're talking about a total well into the millions of jobs. They would all disappear or, just as bad, be effectively demonetised. Say you're paid for a 40-hour week, half of which is driving and the other half loading and unloading goods, filling out delivery forms etc. The driving half just became worthless. Your employer isn't going to pay you the same amount for the other twenty hours' labour as she was paying you for forty, since for twenty of those hours all you're doing is sitting there while your car does the work. That's assuming the other part of your job doesn't get automated too. The world of driverless cars would be amazing, but it would also be a world in which the people who owned the cars, or who managed them, would be doing a lot better than the people who didn't. It would look like the world today, only more so.

This world would likely be one which suffered from severe deflation. If jobs are disappearing, then there is less and less money in most people's pockets, and when that happens, prices fall. This isn't exactly the kind of deflation we are starting to have today in wide swathes of the developed world; that's more to do with the oil price falling at the same time as economies stagnate and consumers lose confidence. But the different deflations could easily overlap. Larry Page, founder and CEO of Google, is sanguine about that, as he recently said in an interview reported in the *Financial Times*:

> He sees another boon in the effect that technology will have on the prices of many everyday goods and services. A massive deflation is coming: 'Even if there's going to be a disruption in people's jobs, in the short term that's likely to be made up by the decreasing cost of things we need, which I think is really important and not being talked about.'

> New technologies will make businesses not 10 per cent, but ten times more efficient, he says. Provided that flows through into lower prices: 'I think the things you want to live a comfortable life could get much, much, much cheaper.'

> Collapsing house prices could be another part of this equation. Even more than technology, he puts this down to policy changes needed to make land more readily available for construction. Rather than exceeding $1m, there's no reason why the median home in Palo Alto, in the heart of Silicon Valley, shouldn't cost $50,000, he says.

> For many, the thought of upheavals like this in their personal economics might seem pie in the sky – not to mention highly disturbing. The prospect of millions of jobs being rendered obsolete, private-home values collapsing and the prices of everyday goods going into a deflationary spiral hardly sounds like a recipe for nirvana. But in a capitalist system, he suggests, the elimination of inefficiency through technology has to be pursued to its logical conclusion.

These chilling views aren't unusual in Silicon Valley and the upper reaches of the overlord class. The tone is inevitabilist, deterministic and triumphalist. There's no point feeling sad about it, this is just what's going to happen. Yes, robots will eat the jobs – all the little people jobs, anyway.

There is a missing piece here. A great deal of modern economic discourse takes it as axiomatic that economic forces are the only ones that matter. This idea has bled into politics too, at least in the Western world: economic forces have been awarded the status of inexorable truths. The idea that a wave of economic change is so disruptive to the social order

that a society might rebel against it – that has, it seems, disappeared from the realms of the possible. But the disappearance of 47 per cent of jobs in two decades (as per Frey and Osborne) must be right on the edge of what a society can bear, not so much because of that 47 per cent, as because of the timeframe. Jobs do go away; it's happened many times. For jobs to go away with that speed, however, is a new thing, and the search for historical precedents, for examples from which we can learn, won't take us far. How would this speed of job disappearance, combined with extensive deflation, play out? The truth is nobody knows. In the absence of any template or precedent, the idea that the economic process will just roll ahead like a juggernaut, unopposed by any social or political counter-forces, is a stretch. The robots will only eat all the jobs if we decide to let them.

It's also worth noting what isn't being said about this robotified future. The scenario we're given – the one being made to feel inevitable – is of a hyper-capitalist dystopia. There's capital, doing better than ever; the robots, doing all the work; and the great mass of humanity, doing not much, but having fun playing with its gadgets. (Though if there's no work, there are going to be questions about who can afford to buy the gadgets.) There is a possible alternative, however, in which ownership and control of robots is disconnected from capital in its current form. The robots liberate most of humanity from work, and everybody benefits from the proceeds: we don't have to work in factories or go down mines or clean toilets or drive long-distance lorries, but we can choreograph and weave and garden and tell stories and invent things and set about creating a new universe of wants. This would be the world of unlimited wants described by economics, but with a distinction between the wants satisfied by humans and the work done by our machines. It seems to me that the only way that world would work is with alternative forms of ownership. The reason, the only reason, for thinking this better world is possible is that the dystopian future of capitalism-plus-robots may prove just too grim to be politically viable. This alternative future would be the kind of world dreamed of by William Morris, full of humans engaged in meaningful and sanely remunerated labour. Except with added robots. It says a lot about the current moment that as we stand facing a future which might resemble either a hyper-capitalist dystopia or a socialist paradise, the second option doesn't get a mention.

[1] The title comes from the fact that contestants can choose both the category of question they want to answer and the amount they want to risk in answering it. Get it right and they win the money, get it wrong and they lose – so there's a component of risk management added to the general knowledge. It is similar to but different from the quiz format of *Who Wants to Be a Millionaire*, in which the sums increased towards £1 million but with the contestant risking everything as the sums got higher. This isn't on TV any longer, a welcome casualty of the Great Recession: as times got harder, contestants realised it made much more sense to cash out their very useful £50,000 or £75,000 rather than risking the money just to

make better entertainment for the audience at home on their sofas.

[2] *Smart Machines: IBM's Watson and the Era of Cognitive Computing* by John Kelly and Steve Hamm (Columbia, 160 pp., £16, 2013, 978 0 231 16856 4).

[3] Gordon's paper is full of wonderful things: before the invention of the window screen, the national animal of the US was said to be the housefly; in 1885, the average housewife in North Carolina walked 148 miles a year while carrying 35 tons of water; one reason cars were such an economic boon is because looking after a horse for a year cost as much as buying one. And they call it the dismal science!

[4] The paper is focused on the US, but I'm taking it here as a proxy for First World employment in general.

Vol. 37 No. 5 · 5 March 2015 » John Lanchester » The Robots Are Coming
pages 3-8 | 6320 words

Letters
Vol. 37 No. 6 · 19 March 2015

'The robots will only eat all the jobs if we decide to let them,' John Lanchester writes (*LRB*, 5 March). This is surely mistaken. We might as well say: 'Devastating climate change will only happen if we decide to let it.' There seem to be psychological forces we cannot control that prevent collective action, even when we are aware of the problems that are about to clobber us (and of their solutions). The likely consequences of developments in artificial intelligence and robotics become clearer by the day, but the likelihood of governments doing anything to address them is close to zero. Why? Lanchester nails that: money. Just as preventing climate change means spending trillions today for no immediate profit, so machine intelligence promises more money (to a few) than Croesus could have dreamed of.

Robin Prior
San Diego, California