Fully Automated Luxury … Dancing? (A futuristic conspiracy theory in the making)

Vic Grout, *Professor of Computing Futures*, Wrexham Glyndŵr University

> *[Note/Disclaimer: Some of the discussion in this piece is shockingly brief. A limit of 10,000 words was planned and (just) adhered to.]*

We've encountered Michael Moorcock's masterpiece, Dancers at the End of Time, before on these pages: both as an example of sci-fi doing what it does best (providing a blank canvas for a bigger discussion) and the problems futurologists have with not seeing key disruptive technology (the Internet, in Moorcock's case). But, for this post, an entirely different question to ponder: who exactly ARE 'The Dancers'?

Because answering that puzzle (there aren't *that* many clues to go on in the novel itself and obviously it *is* only a story) takes us to considering problems in (apparently) entirely different fields: *environment, politics, economics, etc.* (which is the important point really, of course) and may lead us to a view of the future quite at odds with current thinking right across the political spectrum. Specifically, what's usually wrong with long-term 'futuristic' political and economic prophesising? Particularly the very well-intentioned left-wing stuff. What's the one thing that everything from *Karl Marx*'s *Das Capital* to *Aaron Bastani*'s *Fully Automated Luxury Communism* appear to take for granted? (Spoiler alert: in simple terms it's the belief that just because a political/economic system's crap, it will naturally yield to something better – but we'll come to that.)

OK, 'Dancers' is only fiction. And yet it's such a great framework for this discussion that it's worth a quick recap of the premise behind the story. In fact, like most good sci-fi, the premise is almost all that matters. It creates the special scenario that's needed for what follows (as only sci-fi can); it sets the scene for the wider discussion the author wants to have. And, considered on that level, Moorcock's world and Bastani's manifesto are remarkably similar … *except for the numbers*. Both visions are futuristic in the extreme. Both describe an effectively *'anything's possible'* world. Both include technology/intelligence beyond the (existing) human. Both talk of drawing resources from elsewhere in the universe to benefit Earthlings. But while Bastani outlines how future communism could sustain nine billion people (and we're not contesting that it *could*), Moorcock's Dancers appear to number at most a few hundred. Why is that? And, again, *who are they?*

Moorcock describes an *"Earth so far in the future that the remaining remnant of the human race has, by virtue of their access to the accumulated reality-altering technology of all preceding eras, at last ceased to take itself seriously."* The scientists have done their work and are now dead, long gone and forgotten. Energy has been, and still is, drawn from remote regions of space using technology established millennia ago. The last of the humans have (through 'power rings' they wear) immediate access to almost unlimited control they have no understanding of whatsoever. A twist of a finger and anything can be created – or destroyed

– instantly at a whim, human bodies can be altered, regenerated, the past can effectively be undone. No-one knows how it works anymore – or needs to. In fact, there's no education: no-one needs to *know how* to do anything – because they just have the *power* to *do* it.

[As an aside, extending the comments made at the beginning:

1. For Moorcock, this was – as much as anything – a thought experiment. If you have unlimited power (literally) at your fingertips, then how do you behave? But if everyone else also has that same power, how do you – and they – behave? It becomes a moral discussion, not a technological one. We're not terribly interested in the answer here because this piece is more about the journey than the destination (although it's recommended reading) but this is what sci-fi's so good at, and it does point clearly to the fact that answers to technological questions often lie beyond the technology itself.
2. For all that though, Moorcock didn't foresee the Internet – despite it already existing behind closed doors in 1972, when he started writing the trilogy. If the Dancers want to see what's going on on the other side of the world, they can make a supersonic flying train and be there in a few minutes but Moorcock (apparently) failed to predict the emergence of ubiquitous, instantaneous 24/7 global electronic communication – (or did he?) This is always a worry for futurists: that 'big thing' that's currently not even in view.]

But, back to the main thread, why are there so few Dancers? And why might this fictitious handful be more realistic than Moorcock, or Marx, realised? Where has Bastani got it factually right … but potentially wrong?

OK, let's leave the fiction and turn to the actual science, and the technology, and the economics, and the politics. In fact, let's remind ourselves of this …

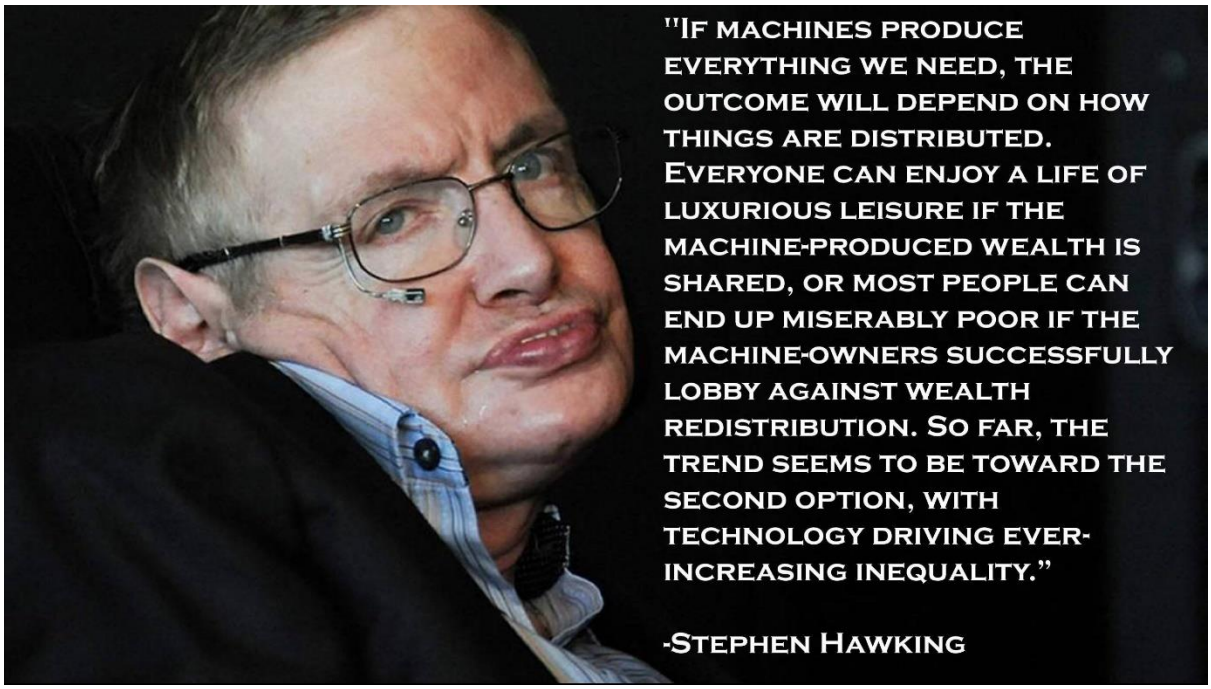

Remember how Hawking, just a few years before he died, was utterly vilified by the economic smartarses? Well, we've got a few more years under the belt now to see how that's turning out. We're beginning to see even more clearly how unprepared capitalism is for

climate change and embryonic automation, for example, and how's it's been unable to put people before profit in the current Covid-19 pandemic. The merest glimmer at the end of the lockdown tunnel and the airports open again, for example. If the virus was (you choose) God's/nature's way of giving us a warning, we've clearly not taken it. It's looking more and more as if economists telling scientists they're unrealistic about technological fallout is like computer gamers telling real soldiers they don't know anything about war. (Let's face it, some people are better than others at playing Monopoly but, when there's a real problem outside, sane people stop playing the game and go and deal with it.)

> [However, another thing we need to clarify before we go much further is to do with conspiracy theories – and conspiracies. In simple terms, there *aren't* any; not decent ones anyway. OK, yes, there are always motives – sometimes hidden – to what people do; particularly those in power. (A few months previous to the time of writing, the leader of the UK opposition had jumped on the first flimsiest excuse he could find to sack a rival.) But, broadly, those guys *did* go to the moon, the Earth *is* round, and Covid-19 (probably) wasn't created deliberately – certainly not aided by 5G masts or suchlike. Even the loose idea of capitalism itself as a massive global conspiracy – with key players pulling the ropes, manipulating and deliberately dividing the populous, starting wars as needed, etc., probably isn't that accurate – or at least *complete*. The point is, it doesn't need to be. Capitalism simply seems to be a *stable* system. We *all* contribute to that stability, from the few at the top making those big decisions down to the masses by and large going along with it because we're too busy fighting amongst ourselves. It's stable because, taken globally, it's proven itself able to restore its equilibrium many times over the years despite some serious forces trying to push us away from it. (Two quick points though: (1) being stable doesn't make it *right*, and (2) we're not looking to the *future* just yet.) So we're not big on existing conspiracy theories per se: instead, we're going to build a new one – *one that hasn't happened yet.*]

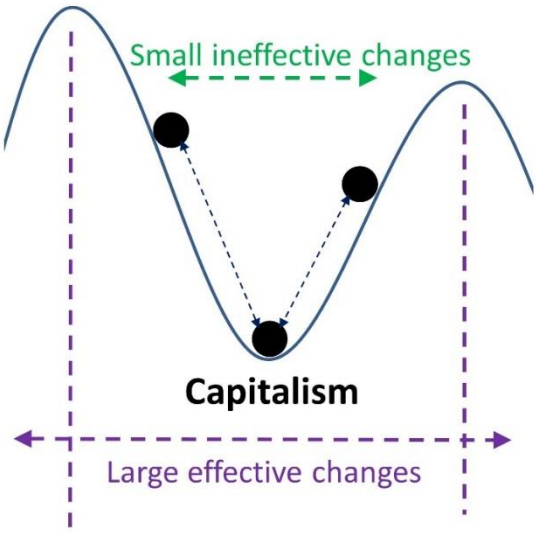

Point(s) of Stability/Instability of a Political/Economic System

Hawking didn't really say a lot about future technology and its social impact but what he did say he was right about. Bastani says a lot more and is mostly right. He's certainly right about technology changing our world beyond recognition over the next few years. But it isn't just 'automation' we need to consider. There are several technological drivers working together right now to make our future lives very different indeed. Robotics & automation, yes, but also artificial intelligence, the *Internet of Things ('Smart' devices)*, *Big Data analytics* and 'always on' ('ubiquitous') instantaneous global computing power and connectivity. Each working alone will have a huge effect on us. Predicting what's likely to be accelerating progress in any one of them in isolation is tricky. Visioning the world of the future with them all working in unison is very hard indeed. A connected, autonomous, automated, global network of intelligent devices, constantly evolving from increasing data awareness is almost inconceivable … and yet that's exactly what's on its way! But it doesn't even stop there because we then need to look beyond the technology itself to its impact on *us*. That's both the most important yet hardest bit of all – it might be impossible right now. However, we really should be trying.

Instead, though, we seem to be consoling ourselves with either or both of two arguments, for which there's absolutely no hard evidence whatsoever:

1. "This is just another industrial revolution. Like before, the boring jobs will be taken over by the machinery and newer, better ones will be created" and/or
2. "Have faith in the human spirit. It got us through world wars, etc. It'll get us through this"

Neither platitude takes into account the scale and severity of the threats facing us over the next five to ten years. But predicting *exactly what's* going to happen is hugely difficult because there are so many dimensions to have to combine. Where even to start?

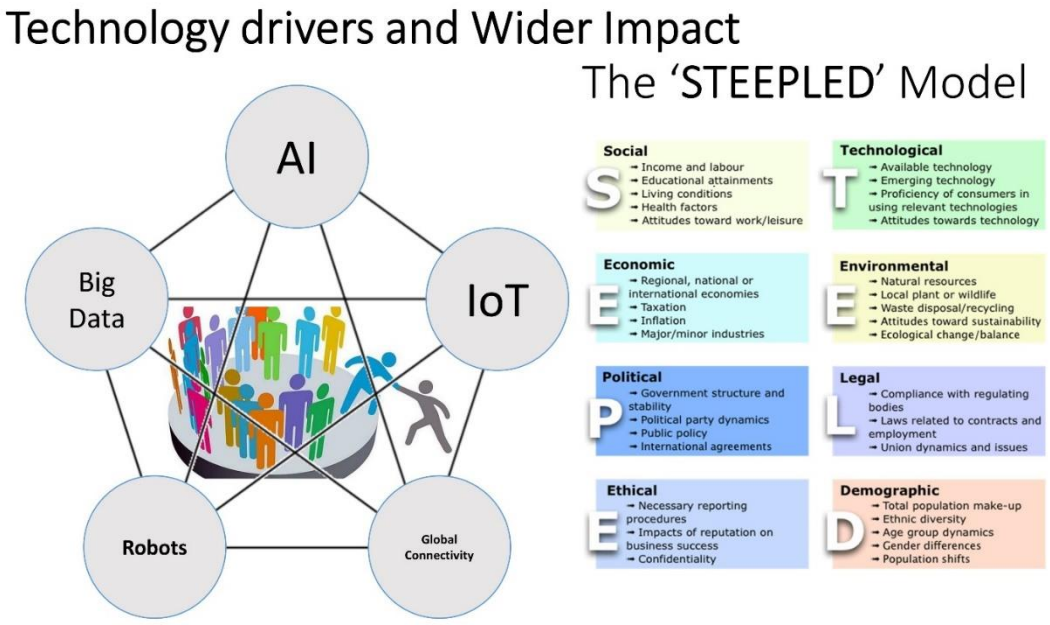

[Another quick aside here though: we can happily discuss AI – even *AGI* – in this post so far as it combines with all the other drivers to (maybe) create Bastani's perfect world, but we're largely going to steer away from the actual *Technological*

Singularity. That may or may not happen sometime and both the arguments for and against, and what it would mean were it to happen, need to be a different debate. Yes, it's critically important (or it might be) but it's also very uncertain on so many levels and we'll lose the plot entirely if we chase that dragon's tail. We'll consider the concept of machines building the next generation of machines, right enough, in the next sections – that's almost trivial – but that, in itself, isn't the *singularity*.]

So yes, as Bastani says, technology is going to change the world and, as Hawking says, that could be for better or worse. And, in case it's not obvious:

> *(a) which it turns out to be has absolutely nothing to do with the technology itself,* and
>
> *(b) the answer isn't pre-ordained by either scientific or political theory in isolation: yes, there's 'science' but there are also 'choices'.*

Rather it depends on the economic/political framework we place the technology in. For example, in terms of (a), whether massive levels of automation free up people for more leisure time and better lives or condemn them to hardship and misery isn't going to be decided by the technology. Unemployment right now (apart from for a privileged few) means hardship. If unemployment in future is going to mean something different, then that requires a political, social and economic sea change way beyond that that the technology can deliver. And then there's that ever-present logical risk here, (b), that, just because something *could* happen, that would be *good*, perhaps *better* than what we have now, we simply assume that it *will*. That's a critical fallacy.

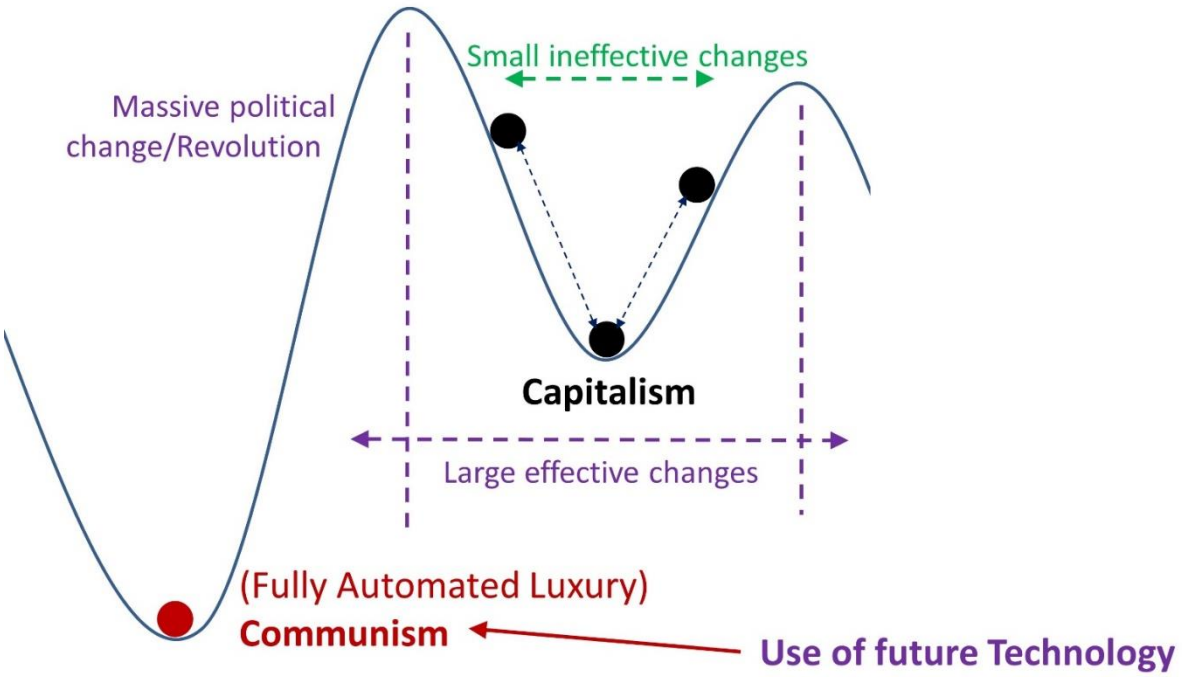

Because there's something else entirely that could happen and we'll get to it … soon …

OK, no-one's suggesting this is simple so let's break away from a single thread and look at this from a number of key perspectives – and points of view. The 'STEEPLED' model in the diagram above is a good place to start (although there's always a lot of overlap). Essentially,

‘STEEPLED analysis’ (an extension over the years from PEST analysis) allows us to consider difficult or complex issues from specific points of view.

STEEPLED

- **S**ocial
- **T**echnological
- **E**conomic
- **E**nvironmental
- **P**olitical
- **L**egal
- **E**thical
- **D**emographic

‘STEEPLED’ isn’t a magic wand or a mathematically complete system that guarantees correct answers or deductions. It’s naturally imperfect (because it’s done by *us*) and there’s no particular process to follow but it does, at least, give us *some chance* of considering the matter from all relevant angles, without missing out anything obvious or important. So it’s a useful framework for this complex discussion; it’s not anything more than that though.

This will go all over the place now but we’ll pull it together eventually. Here goes …

SOCIAL

This is a difficult one to start with (if we insist on considering them in order like this) because it so obviously cuts across all the others. It’s equally hard to summarise in a few words but, to take a pre-emptive stab at it … well, let’s face it, things are getting *nasty*, aren’t they? And, along with increasing unpleasantness, there’s increased divergence. (A few weeks before the time of writing, a 12-year old boy had sent racist messages to a black footballer, who retweeted them. The outrage across the political spectrum immediately polarised from ‘left-wing’ howls of recrimination to ‘right-wing’ accusations that the footballer had somehow ‘got a 12-year old arrested for nothing’.) We can (sometimes all too) easily dismiss opposing views as invalid but we can’t stop them existing. If anything, the pace of this descent into multifaceted, dysfunctional tribalism is increasing. *‘Cancel culture’*, for example became a meaningless term within weeks of finding its way into widespread use.

There are various theories as to why this social breakdown is happening – and accelerating – beyond the scope of this piece. Many blame social media. Others suggest that social media may be merely facilitating the collapse that’s taking place anyway: the pre-ripples of the catastrophe that’s on its way as disruptive technological evolution accelerates, natural resources become scarce and *survival* on Earth takes a more literal turn. (We’ll get to that.) True enough, we’re not going to suddenly wake one day to find the robots have taken over or there’s no energy, food, water (although this *does* happen already in some parts of the world), clothing, housing, etc. It’s a more gradual process and social tensions will naturally ramp up in advance and, implicitly, subconsciously or otherwise, in anticipation – perhaps preparation. Is this what we’re seeing already? The embryonic tensions of a global society beginning to fray at the edges? Perhaps, but that’s too tangential for this discussion. (However, the general principle that the changes over the next five to ten years – although massive – will still be incremental, rather than cataclysmic, is a sound one; and central to everything that follows.)

Perhaps also we were just never designed/evolved to have the level of connectivity that social media and suchlike, through the global network, give us. In days gone by, if someone woke up with a sore head, they could take it out on a few family members or neighbours at most. Historically, periods of uncertainty, upheaval, even global war, have often followed significant advances in widespread communication media. Today, a hateful thought can be 'articulated' around the planet in seconds. All we can say with some certainty is that, more than through any other period in history – even including wartime, individual humans now seem able and willing to wish harm – even death – on particular others they've never met, on the other side of the world. OK, these are just a few close-ups of a much bigger picture but it's apparent that the whole scene is not socially healthy – and it's getting worse.

The rest we'll consider as we go through the other sections …

TECHNOLOGICAL

The usual purpose of the 'STEEPLED' model in technological forecasting is to avoid overlooking the wider picture *beyond* the technology, that aspect being assumed to have already had a fair hearing. The 'tech angle' can thus be short. However, it's front and centre here precisely for the effect it will have on all those other aspects of our lives and there's a lot we haven't said yet or won't say otherwise. So, here goes …

Moorcock's choice of introduction for *Dancers* is interesting: '… *the human race … access to the accumulated reality-altering technology …'.* It could well just be flowery words but it does happen to point quite effectively to a vision of where future technology could take us, namely that …

> *'… at some point in the future, pretty much anything … will be possible; at least in simulation and without much difficulty if we're prepared to accept limited quality in the early stages.'*
>
> … this being borne out by past experience of technology adoption with early gaps and glitches. (Back to the principle of incrementality.)

But what does that technological future look like more precisely? *Is* 'anything' going to be possible?

Well, perhaps not literally, but fairly close in two key senses:

1. As it has throughout history, technology will make *more and more* things possible: new, better, faster, easier, more fun, etc. We've already noted this as an accelerating process. Now, the combination of automation, AI, big data, IoT smart devices and seamless global integration will speed it up even more. But, at the same time …
2. Whatever it is, *we'll* have to do less and less: the robots (significantly driven by all the other emerging tech.) will do it for us.

Now, if this all sounds a bit simplistic, … *it is*. But the real reality check isn't that some of this stuff won't happen; it's simply that it won't happen overnight – all at once. We'll get to this future by smooth (although potentially quite fast) transition. And this period of transition – as much as the destination – warrants a bit more attention than we've been giving it.

Firstly, the process has clearly already begun. We're seeing smart devices talking to each other, AI and big data combining to design new solutions, automation taking over production roles and faster and faster networks connecting it all together. Baby steps: we've hardly started yet but we can at least partially see where it's going.

> [Just a quick (this time, *relevant*) aside on why AI, in particular, is so important: even more significant than automation and the rest of the tech. 'drivers' because, ultimately, it will drive all of them. AI gives us something technologically – no *philosophically* – different to anything we've ever had. For millennia, humans have solved problems by, er, … *knowing how to solve them*! (Obvious? Just wait …) The first computers didn't change that paradigm significantly. Computers could solve problems faster and more accurately but we'd still have to know *how* to do it (an *algorithm*) because we'd have to program the computer. Maybe we were justified in thinking it would always be that way? But AI changes all that. Rather than deterministic lines of code, directly mapped to the problem to be solved, AI programming consists of setting up abstract learning structures (*'neural networks'*, for example) , then throwing training data at them to be tuned – to eventually become very good at it. *We're* no longer solving the problems: the machines are – and we don't necessarily know how. (So for example, rather than tell the computer to recognise a tree by looking for branches and leaves (which is what we might do), the computer trains itself to recognise a tree by analysing lots of trees and figuring out it needs to look for … well, that's it: *we wouldn't necessarily know!*) This is a critical moment in our history. For the first time, we're now able to solve problems we don't know how to solve! And to build machines that can do it. That's huge.]

But let's take the obvious example of automation but push it forward a few years. The trend towards replacing people with machinery started a long time ago and will continue: not just in scale but scope too. Increased levels of miniaturisation and the adaptive flexibility that AI/big data can bring will eventually bring many current 'human' roles within reach of machinery – or at least, technology; *most* of them in fact: *yes seriously* …

Because there's a common misunderstanding here … Go to https://willrobotstakemyjob.com/ and try something like 'production worker'. Apparently, there's a 90+% chance that you'll lose your job to automation and that's probably not unreasonable. (It might be too low!) But now try 'teacher' and see what happens. Does that 1% or so make you feel comfortable (if you're in that line of work)? It shouldn't because it's not just automation we should be considering but *all* the tech. drivers, AI, big data, IoT, Internet, automation, working *together*.

So, yes OK, there's a pretty slim chance that, in 5-10 years' time, a two-legged metal robot will be standing where the teacher stood, waving their arms and writing on a board. But *that's* not how teachers get replaced. They'll get replaced by vast – and exponentially increasing – quantities of material on the Internet, initially written by humans, then collected through big data processes, curated by AI, made available anywhere through smart mobile devices and customised – by all of those things – to the precise requirements and learning needs of individual students. From a student perspective, it'll be fantastic: who *wouldn't* want it? But *that's* how teachers will lose their jobs. That's how *most* of us will lose our jobs. Moreover, many teachers, lecturers, etc. are currently being forced to produce the very material that's going to put them out of work! Not just teachers, of course: nearly all of us: nurses, doctors, engineers, drivers, distributors, builders, fitters, repairers, etc. – individually

or collectively – are contributing to our own redundancy in one form or another through increased (but incremental) reliance on the technology that will one day replace us. The scale of the automation revolution is being underestimated almost everywhere because very few people are considering the effects of all these disruptive technologies working *together*.

The ultimate destination then, of these integrated technologies: AI figuring out things we can't using big data captured through smart devices over global networks, then implemented through automation is that the machines will design and build themselves. In fact, individually, this has been happening for years. Automated 'built by robots' production lines have been with us for decades: that's covered. Also though, modern computers are designed by computers. The layout of electronic logic, for example, has been beyond human calculation for a long time now: chips, etc. are routinely optimised by software. Many other aspects of design today rely on computers. Put together the software and the hardware, the *design* and the *production*, complete the loop and you have machine *reproduction*. Add a few lines of AI-generated code to each generation and you have improvement – machine *evolution*.

> [In fact, although we're not going that far in this piece, this critical future point of machine reproduction/evolution is perhaps a better notion of the 'technological singularity' (TS) than vaguely connected notions of AGI. A crisper definition would be when (in simple terms) a 3D printer is able to make an identical 3D printer. Then, with a tweak or two, the offspring makes a better one, and so on. Then we'd be rather worried about how fast this evolution might accelerate. OK, it's not really quite as straightforward as that. The printer, aside from its knowledge/instructions, for example, would need materials. It would either need to be connected to some material supply network or have mobility to find its own. Both are a bit fanciful at present but perhaps not entirely so in the brave new world. A kind of Wikipedia for Robots to share/use collected knowledge is already a research thing.]

It's (hopefully) going to be up to us to what extent we're prepared to let robots engage in philosophy, write poetry or music, play sports, etc. But, in terms of the routine stuff: primary resource mining/acquisition, design, production, distribution, maintenance, troubleshooting, knowledge transfer, etc., there's not likely to be much they can't do. And they'll be better than humans: faster, stronger, more reliable and, critically in the current model … *cheaper*.

So what's that going to mean?

ECONOMIC

There's no space here for a lecture on capitalism; and we don't need it: we know the way it works. A very small number of people own and control the world – not just its resources, but the right to exist and occupy space within it too. Small portions of this ownership and control are delegated down in layers to those with appropriate aspirations, enthusiastically embracing the system. The process continues, with the numbers growing and return diminishing until it reaches most people, who have either nothing or (considering any normal calculations of global averages) next to nothing. These people spend their time, ruin their health, and waste their lives, processing materials or managing space owned or controlled by others higher in the hierarchy, only to be rewarded by being handed back (depending on their place in the hierarchy) just enough material or space to permit their survival. At the lowest levels, of course, they often can't even do that.

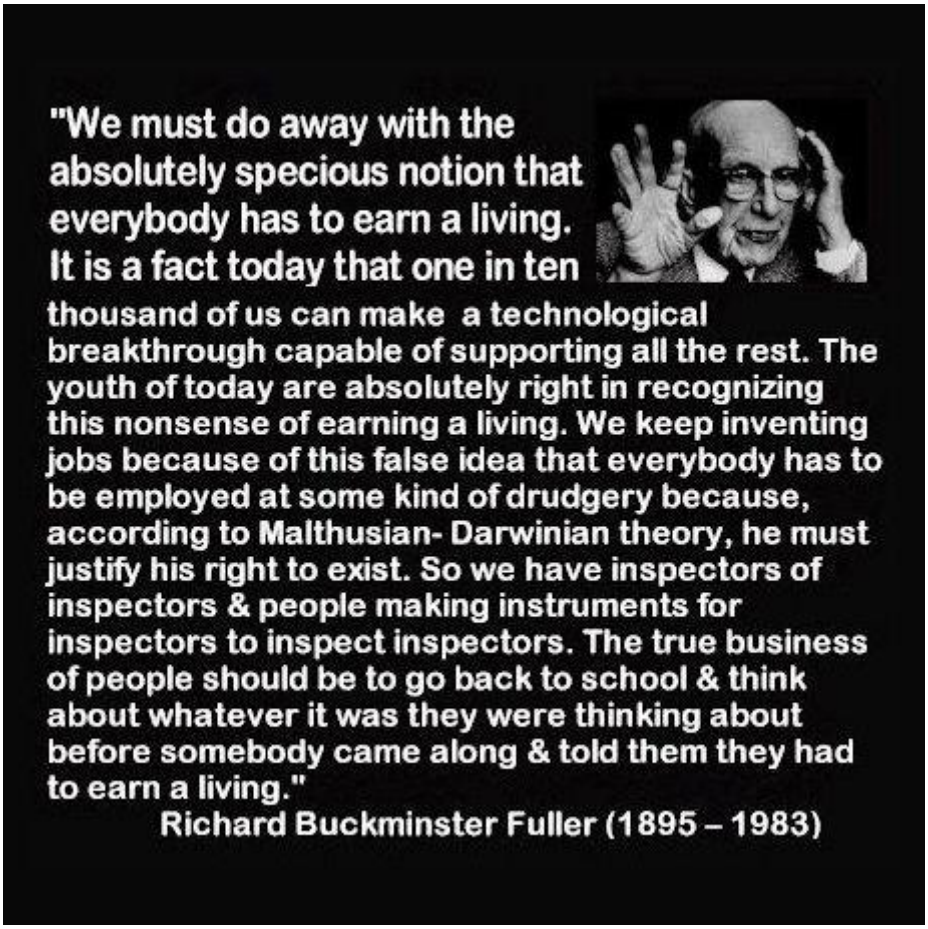

So, how will future technology (automation driven by AI, smart connectivity and data) change that? It's too easy to say 'unemployment'. 'Unemployment' only makes sense in relation to 'employment', which in turn, only has meaning in terms of 'work', assumed generally (in economic, not scientific terms) to mean 'something that has to be done'. *What* will have to be done in the future world? There's no simple answer to that because it depends on what that future world looks like and that depends on some big choices we're making (or not making) round about now. What *will* happen is anyone's guess. Maybe it makes more sense to ask, what work *should* be done?

Well, that's also complicated so let's flag it but park it for now and come back in the important conclusion. What we do know is that there will be more conventional work done by machines, so less by people. The positive spin on this takes us to the model of technocapitalism, which, in simple terms argues that, as *machines* do more of the work turning raw materials into goods, *people* occupy themselves more imaginatively turning *data* into *information*. Apart from the obvious contradiction that this has with everything we know about the combined AI/big data/smart devices/automated/connected future tech model, once again it just doesn't *scale* to the new world. Several billion people won't occupy themselves being creative with data: even if we are to enjoy the fruits of automation, there are only so many films/books/games we can watch/read/play – and there's no credible economic model that would support that anyway.

Fairly obviously, a lot of this will come down to who *owns* this future technology. If machines can do more of the work, then fewer people are needed. In *conventional* terms, that increases 'unemployment'. But what 'unemployment' means depends entirely on the

political/economic framework – not the technology that caused it. Again, Hawking's spot on. It *could* mean Bastani's future of cooperative, collective luxury (if the technology and its benefits are shared) or it *could* mean massively increased levels of global inequality, hardship and division (if the technology continues to be controlled by a small world elite). Not only that but, if those in control need fewer people working for them, they'll also need fewer people to 'sell' to to maintain their 'position' – but we'll revisit that towards the end of this piece as well. We'll also talk about *which way* it might go but, critically – and taking into account all the other factors, it can't stay the way is for much longer.

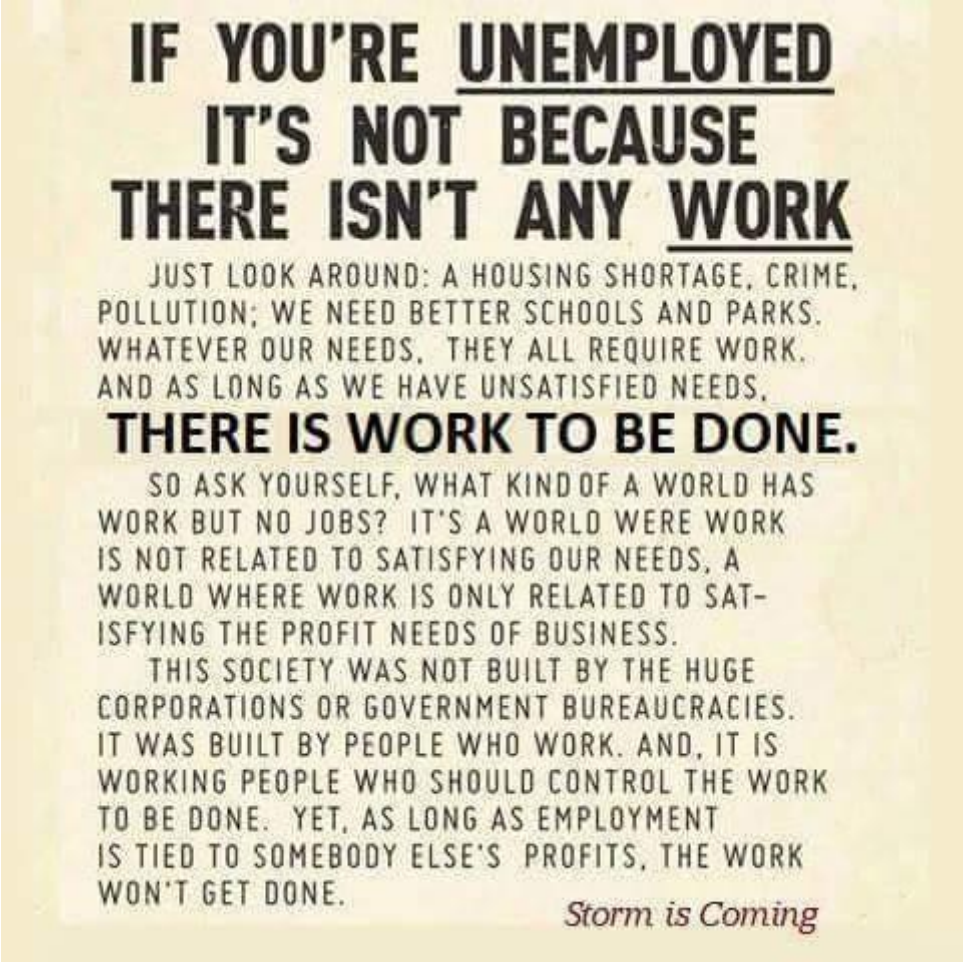

In simple terms, tweak as you might, the days of the conventional capitalist economic model are numbered – the numbers don't add up any more: even the capitalists (those far enough up the hierarchy) are realising that. Bastani proposes an entirely credible alternative. So is that what we'll get? Well, let's work further through the STEEPLED sections …

ENVIRONMENTAL

We have to start this section with the uncomfortable admission that it may already be too late. We may be locked into biological cycles we can't escape from. The plastic floating around in our oceans was thrown away years ago. Plastic use/disposal increased exponentially over the past few decades. What's it going to look like when all *that* ends up in there? Those massive ice islands that have broken away at the North Pole aren't going to get stuck back on again. They're *going* to drift away and melt: water levels *will* rise. Most of the trees *are* gone. Most of the ground and sea and air *is* poisoned. We can only *hope* that

there's a way back, even if we start right now; not that we're showing much sign of doing that, of course. It's going to take more than turning down a plastic straw in a pub or putting onions in a string bag; it should be challenging the big industries, and their exploitative processes, that really do the damage but there seems to be little social or political appetite for that at the moment. In fact, those who try are often vilified.

But this piece is essentially about *people* so let's try to stick to that. Essentially, what have we done to the planet so far, what's it done to us and, even more importantly, how that going to play out in both directions in the future? Let's try to answer that by considering explicitly if the planet's *overpopulated* (by humans)?

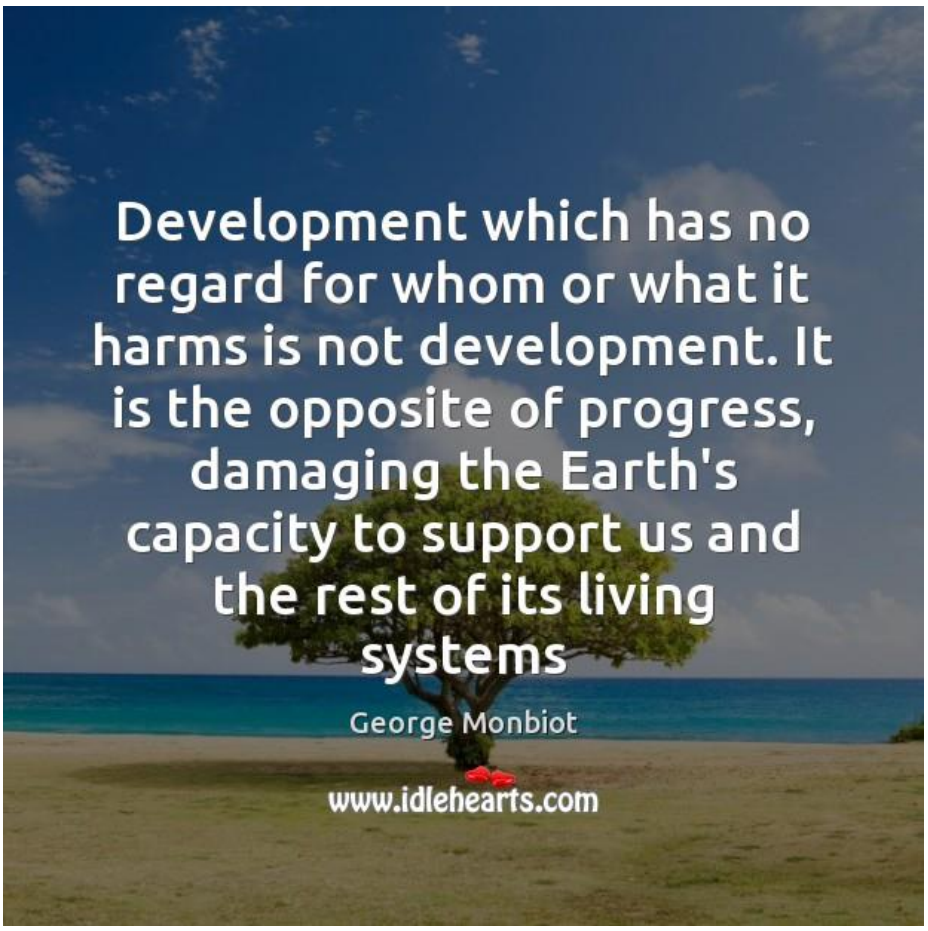

Well, in pure biological terms, yes probably. Various ecologists and anthropologists have different views on how you could calculate how many humans there *should* be on the planet but there's some consensus across orders of magnitude from 50 to 500 million. So it could be as little as the population of the UK, or a larger country, but distributed across the globe. That's an estimate of the number of humans that can live in some sort of natural balance with their environment(s) and other species. It's a population that wouldn't deplete natural resources, such as water, critically and would make only a sustainable impact on other species of animals and plants.

But, for good or bad, humans *aren't* in balance with the rest of the natural world and haven't been for some time. On the downside, the global population is roughly 100 times the 'spread out UK' figure and we know we *have* been doing damage. We've wiped out species, polluted and depleted resources to a critical level. On the upside, we're pretty clever and

(ought to be) capable of doing something about it. In this context, technology for example, can be the nail in the coffin or a saviour – depending on how its deployed (and by whom).

But it's not just about simple population: it's about *consumption* too. Each of us doesn't *need* to have to have the footprint we currently do. In the developed world, consumption per head is between three and seven times that in the developing world. (And, of course, *development* is a daft concept anyway: it's effectively measured in how much damage a country does to the Earth's survivability.) We don't *have* to import food from all over the world, drive that car, fly in that plane, etc. but we still do because we've grown up with it and it's sold to us as 'freedom'. But more on that in the next section.

Taking this heading purely in isolation, whatever views there may be on other matters, it's a simple scientific fact that we can't continue the way we are for much longer. (And that's a pattern that's beginning to develop here right across our STEEPLED model.) Nature will put a stop to it before long. Either we change our ways drastically or we will lose the means of survival. It's a stark choice but it's one we have to face or suffer the default option.

Again, Bastani offers a pretty credible technological (and the rest) model for how such an (ecologically) overpopulated human world can sustain a balance. But can we ever get there? Unfortunately, there will be even bigger forces to contend with …

POLITICAL

Bastani's talking communism, of course, and that will instantly put people off – a certain type of person anyway, and sadly there are plenty enough of those. In fact, if we were to reluctantly applaud capitalism for its two most astonishing achievements, they would be these:

1. From birth, capitalism conditions us (or, at least, enough of us) to doubt and scorn any other political/economic framework – *even when* it's clearly about to kill us all and alternatives could save us. [Again, we could shout 'conspiracy' here: politicians being controlled by big business and the likes of us by the media; but it could equally be just part of its essential stability as a system.]
2. Capitalism's final victory may be that, even as it *does* destroy us, it will divide – not unite – us and have us all blaming each other for the catastrophe. [However the end actually comes, a global war over resources, say, we'll still blame whoever's top of our programmed hate list. It somehow won't be capitalism's fault. It's sometimes hard to avoid an ironic smile when you realise that.]

Yes, capitalism thrives on forcing us, not just to compete for status and survival, but to actually despise anyone – or other group – that threatens us – or is even *perceived* to threaten us. A race with the communication skills, intelligence, etc. to cooperate and share, is programmed from birth to do the exact opposite.

Criticise capitalism in the slightest degree and there will be those who immediately resort to reminding you that millions died under various communist regimes. Aside from glossing over the various situations in which that happened, there's never any mention of the tens of millions who die unnecessarily *every year* under global capitalism; deaths that would be avoided with fairer distribution of resources. It's not Stalin and Mao that should be at the top of those 'death lists': it's Murdoch, Branson, Bezos, etc. But because, that's considered

'normal', no-one really notices. Unfortunately, *'Yeah but Stalin'* seems to be all that's needed nowadays to slap down any talk of any alternative social/political/economic direction that might save us!

Particularly in the developed world, where maybe things aren't quite so desperate yet, and even more particularly the West, we're largely controlled by governments steering a right-of-centre course between the hard-line, uncompromisingly capitalist (but misleadingly labelled 'neo-liberal') requirements of big business and the more fanciful, but often even further to the right, populism of billionaire-owned media. We're sold on 'freedom' and any suggestion that we might have to mend our ways becomes an 'attack on freedom'. But it's a very odd notion of freedom: it means being allowed to drive a big car, take foreign holidays, if you can afford it, and choose from 100 different types of shampoo; but it doesn't guarantee that everyone else has enough to eat or somewhere safe to live.

Even so, the ease with which popular opinion has shifted from, *'it's OK: capitalism's got this covered'* to *'I'd rather die than suffer communism*' is genuinely astonishing and may say more about where the human race has got to right now than anything does? Is this an actual malaise of the human race, evidence of the power of the controlling forces of business and media or just another indication of the current stability of the system? Whatever it is, the notion of us 'sleepwalking' into oblivion seems horribly accurate. It's not admirable or stoic to be resigned to a fate when we could actually be doing something about it.

LEGAL

Remember, there's nothing deterministic about the STEEPLED framework; it's just a way of trying to cover all the angles. Sometimes the different sections can be considered independently; sometimes there are huge areas of overlap between them. Legal aspects of the discussion look at first as if they should be considered separately (after all the 'legislature', 'executive', 'judiciary' democratic pillars tell us so) but we've already seen that there are greater powers at work in politics and economics so it's no surprise to find this driven by other factors. In fact, to a large extent, we've already touched on what needs to be said here.

In terms of what it does now, and what we might hope it would do in a better future, we have to recognise that the primary purpose of the law is to uphold the current political/economic system. Yes, that's obvious on a superficial level but the extent to which assumptions are carried through in this way merits a moment's thought. How does the law work in a capitalist setting and how might it work in a socialist one? In fact, the law is much *harder* to implement under capitalism.

- In a capitalist system, inequality is not just permitted but *encouraged*. [To take a deliberately simple example, it's perfectly OK for one person to be able to afford five TVs, and someone else none, or one.] Much of human life is spent trying to achieve an advantage over others through the mechanism of money [which buys the TVs] and much of this is through the acquisition – legal or otherwise – of goods and resources – *deals*. Some deals are legitimate whereas other activities are simple theft. But there's a large swamp in between where it's either not so clear or hard to prove that a crime has been permitted. Our current legal system spends much of its time – and employs huge numbers of people, in determining which side of an arbitrary line in the sand a particular action falls. And, from time to time, the line *moves* – as suits the system.

- In a socialist system, in which everyone was allowed three TVs, the simple fact of *having* four would be the 'crime': there'd be no need to discuss the finer detail of how they got them! It's easier to measure the *result* than examine the *process*. It's just possible that's why so many lawyers, often proud 'liberals', can't quite bring themselves to move further to the left. Perhaps they subconsciously recognise that full-on socialism would need fewer of them?

OK, yes, this is horribly simplified (legal discussions and word counts don't marry well) but the take-home point is that – the 'obvious' stuff/basic human ethics aside – there's nothing morally 'right' about any legal system, other than in the context of the political system it supports.

Let's take another daft – largely fictitious, but credible – illustration to hammer home this quite important point … Local councils maintain roads and mend potholes to prevent damage to cars. But with, tighter budgets and more cars on the roads, they can't keep up. A motorist wins a test case and sues the council for expensive repairs. Increasing numbers of motorists do the same. The council's finances can't cope. The 'system' is threatened. There would be various relatively sensible solutions to the problem ranging from (economic) increasing council budgets to (environmental) encouraging drivers out of their cars but what would really happen, of course, is that the law would be changed to prevent further claims. What was a 'right' law becomes a 'wrong' law to support the system.

> [Actually, it's really not that daft an example … We have an underfunded National Health Service. It's understaffed and under pressure. Mistakes happen. An entire private legal industry has built up around suing dedicated professionals for negligence because there's profit in it. Billionaires also sue it for loss of contract opportunities. As a result, it loses more money it already hasn't got.]

But, specifically, in terms of legislation relating to emerging and future technology, this brings us quickly to …

ETHICAL

So what chance for *ethics* in all of this? Probably not much, frankly! We're leaving it rather late here to introduce two other principles that have served as axioms in other discussions. But they're still essential to understanding how technology develops, why, and for/by whom:

- Technology – existing and emerging – is generally used and abused in about equal measure, and will probably continue to be, although it's often a matter of personal opinion as to what's actually good or bad. Just as one man's terrorist is another's freedom-fighter, one woman's pleasure is another's sin. [Consider arguments for/against encryption – and bypassing encryption – over the Internet, for example.]
- Appropriate legislation always lags some way behind the changes brought about by emerging technology and there's really no credible history of short-term social, moral or ethical objections being effective in restricting long-term technological advance and deployment. [Look at stem cell research and genetic modification, for example.]

This is odd, in a way, because there's no shortage of 'ethical codes' in technological circles. All professional/learned societies have an ethical code of conduct and, in principle,

no research is funded – at least publicly – without appropriate 'ethical screening'. The problem it tuns out isn't so much with the ethical code: it's with the practical implementation of that code. Ethical screening can run through a checklist of a dozen or so key headings (humans, children, animals, environment, privacy, environment, etc.) but, to pass the screening, research proposers by-and-large have to demonstrate that they're *aware* of their ethical obligations and have the *capability* to discharge them rather than much in the way of monitoring that they actually *do*.

And the real problem lies with private research, of course. Ultimately, technology is developed because someone, somewhere profits by it. Yes, there is a veneer of 'for the common good' in the way its presented (just as public services to – *usually* – stop people dying in the street present an acceptable veneer within a capitalist government) but, when it really matters, ethics rarely really get in way of development if the returns are high enough.

Then there's the problem of it often being difficult to *see* the long-term effects of current development. Even if we always could though, it probably wouldn't have made any difference. We have some safeguards in the west but a blind spot for what goes on elsewhere. Where we get our clothes, domestic goods, technology, etc. from nowadays often simply transfers our ethical/environmental impact to another part of the world where we can't see it. Seriously, if we *had* had a clearer view 50 or 100 years ago of the damage we were beginning to do (a view shared by the majority, that is, not just a few 'cranks'), *would* it have changed our ways? Someone's profit would still have come first; someone's profit *did* come first. Maybe it always will.

DEMOGRAPHIC

Leaving aside the trend within countries for population shift from rural to urban settings, the current headline demographic issue is *migration between countries*. (Although, in most cases, 'refugee' would be a more appropriate term than 'migrant'.) The first point to make is an obvious one: refugees are what Western countries should *expect* on their doorstep when part of their 'wealth' comes from sending weaponry to politically unstable parts of the world. Instead, however, they're exploited by 'divide and conquer' politics to divert attention from financial cronyism, underfunded public services, major corporations not paying tax, etc.

But a second point is that 'sensible' migration is a *good* thing – so long as we define 'sensible' *sensibly*, hopefully *ethically*. The Earth is not homogeneous in its ability to support human life: some parts are better than others. (Yes, technology can help with this but perhaps never entirely.) OK, it's really not the way it's happening at the moment but there *would* be considerable merit in finding a better, smoother global distribution. Unfortunately, that will require global cooperation, not the competition between countries we see now. There are credible arguments for both small, completely autonomous local anarchies *or* full-on global government but the system we have now where large nations exist in a state of perpetual tension – but still without internal or external balance – is the worst of all options.

Of course, one thing research, technology, developments in healthcare, etc. *have* achieved is a general (though not uniform) global increase in life-expectancy. We should be:

1. Rejoicing in this, and
2. Using technology in order to support elderly people – and those with disabilities – to live longer, happier, fulfilling, independent lives.

Instead, of course, it all comes down to money once more. The falling 'carer to cared' ratio is immediately translated (nationally) into a financial burden on the taxpayer and (corporately) into reduced profits for businesses in that line of work. What should be celebrated as a rare (effective) social achievement of the human race is seen as a potential problem because of the ridiculous system we seem addicted to.

PUTTING IT ALL TOGETHER

OK, yes, so that did indeed go all over the place; because we were trying to look at the question of what *is*, and what *could be*, from as many angles as possible. The STEEPLED model sort of helps but, if some of it didn't seem relevant … well maybe it wasn't … but we were trying to get that elusive 'uncluttered but all-encompassing' view that effective futurism requires. It's not easy: most people only think in one direction at a time and that's not enough with the scale of the change just about to hit us. We could have continued discussing each of these threads almost indefinitely; but we have to draw a line somewhere and we've, hopefully, covered the essential detail. So where has it got us? Well, let's recap/summarise those STEEPLED sections:

> SOCIAL: It's getting nasty out there: social divisions are appearing everywhere. Conflicts are beginning; they'll get worse. (Not 'before they get better', they'll just continue to get worse.) Something has to happen; *something* will.
>
> TECHNOLOGICAL: Technology requires fewer people to get things done and will yield huge benefits to those who own/control it; but who will that be? That question is critical.
>
> ECONOMIC: As long as an elite retain control of the technology, they need fewer people to work for them and fewer people to exploit/control ('sell to') to maintain their position.
>
> ENVIRONMENTAL: There are vastly too many people on the planet (under the current economic system). Either we find a fairer/better way or there *will* be massive depopulation (*something*, 'natural' or otherwise, will force this).
>
> POLITICAL: There are clear alternatives to what we're doing now but there's scant evidence that we're prepared to accept them (and there are those who would clearly not want us to).
>
> LEGAL: The law essentially resists change (social or otherwise) unless that change is palatable to those already in power. Beneficiaries of a system change the law in anticipation of further benefits with greater ease than its victims ever achieve redress.
>
> ETHICAL: If something *can* be developed, it will be, so long as there's a profit (or other benefit) in it for someone. (To put it bluntly, we don't *really* care about ethics when it matters to capitalism … or something even worse.)
>
> DEMOGRAPHIC: A healthier, ageing population should be something a civilised world celebrates but, without social change, it just exacerbates pretty much all of the above. Similarly, a 'spreading-out' of the world's population should be welcomed at a global level but creates tension in a non-cooperative, competitive world.

Doesn't look great, does it?

PUTTING IT ALL TOGETHER FROM A PARTICULAR PERSPECTIVE

So the final assembly of our 'build-your-own-conspiracy-theory' comes from combining these previous STEEPLED sections and considering it all from a particular point of view …

Just imagine now you're one of the world elite.

> [By *'elite'* here, we're not talking millionaires or billionaires, 'rich' or even 'super-rich'. We're not talking about managers or bosses. Not factory owners or business owners or company directors nor the right-wing hacks that do their bidding. We're not even talking about the people that own the companies that really own these people. We're talking about the people that own the people that manage the organisations that control the global infrastructures that run the companies that have swallowed up the businesses that the rest of us work for. We're talking about the top few hundred people in the world – a thousand or so at most.]

How does this all look from your perspective?

Well, at the moment, things are chugging along just fine, pretty much as outlined in the previous sections. Although working conditions have improved in some parts of the world over the past few decades and that's maybe cost you a Caribbean island or two, the basic political/economic framework – particularly in relation to 'work' and 'employment' – hasn't changed much since The *Ragged Trousered Philanthropists*. The 'he works for her but she works for him but then he, etc., etc., … and everyone competes and cheats' structure probably has a few more layers than it used to but it still leads up to the same place … *you*.

You sit atop a vast structure made strong by tension. At all levels below you – from individuals to whole nations, people and populations both vie with each other for resources and finance each other's survival through the economics of work and employment, and competition and warfare. Critically, you essentially own the resources – or the raw materials. Below you is an exchange and distribution network in which, through many layers, people work to turn your raw materials into products that they themselves need. You hand out arbitrary tokens of value (*money*) to people to work to produce things, that you still own, that they give you back the value tokens in exchange for, so they can continue to live, so they can continue to work, to live, to work, to … and on it goes. It's your own personal perpetual energy machine. Very few challenge the structure effectively because the tensions of accountability, blame and competition make it hard to do. Most don't challenge it at all because they're conditioned from birth to accept it as the norm. The exceptions that 'escape' the structure often achieve little except cutting themselves off from the survival resources the system drips down. You don't have to actively manage this structure: it manages itself. *It's stable*. To a great extent, the people who most suffer from it do the most to ensure its stability.

But what are the *real* benefits of this system to you? It's not exactly right to say *'wealth'* because that's only a concept in terms of the arbitrary value system you've created. What it really means is 'first-in-line', privileged, almost unlimited access to everything the world has to offer (the arbitrary value system merely quantifies and normalises this) and the highest levels of freedom that can be achieved. Without expending any effort at all, you can do *more*

things, do the *best* things, and – critically – enjoy the highest quality of life that current technology (and eventually, emerging and future technology) can deliver. You own the technology, of course: because the rest of the world takes your raw materials, works to build it for you, then returns it back to you. And this quality of life extends to *survivability* too: your position gives you the best levels of safety, security, health care, medical attention if something goes wrong, etc.. It's essentially *your* planet. Aside from the inconvenience of having to share it with billions of plebeians, you're already on your way to being one of Moorcock's dancers – and your mates will be the others. Very speculative at this stage maybe but, for example, if transhumanist technology was ever to look like delivering anything close to immortality, you'd be the first in line. More realistically, if technology ever looked like it was going to force disruption upon the world order, you're still in pole position to steer the change in the direction you'd like it to go. It's much harder for anyone else to do that because – once again – you own the technology that will do the changing.

Because technology *is indeed* threatening to change the world order. The AI/automation/big data/smart systems/global connectivity technology revolution is going to mean that more work will be done by intelligent machinery and less by humans. But that, in turn, will place an even greater demand on the world's dwindling natural resources whilst those bloody plebs are increasing in number all the time.

So how does it really look to you?

Well, assuming *you* continue to own and control this emerging technology, you're going to need fewer people to work for you. You're going to be paying smaller wages for turning your materials into products because you don't pay machines (in the same way). That means (still in terms of this arbitrary value system) that you need fewer people to sell the goods back to, which in turn means you need to produce less, so there's less to work to be done. What work *is* left to be done can be – more and more – done by technology, and back through the loop it goes, the whole model shrinking as it cycles. The destination (certainly not an overnight journey) is a much smaller amount of technology working purely for *you*. There's no longer any need to turn raw materials into products (which is a good job because those will have run out anyway) so you don't need anyone to sell them to. All you need is *your* technology (and a few human technicians for repairs, etc. in the early years) looking after *you* (and eventually your descendants, of course). Even that arbitrary value system becomes redundant eventually because it was all a front in the first place. Instead, on the back of all the technological development *other people* have done for you, *you've* now become *self-sufficient*. You finally are a *Dancer*!

And, in a sense, this *would* be the natural progression, it *would* happen, were it not for one thing: all those other bloody people breathing your air, drinking your water and getting stroppy – potentially confrontational – about starving to death as technology takes over and the planet is bled dry. *There's* your problem …

Because all of these individual threads have already started. Automation *is* happening, machine intelligence *is* developing, human population *is* increasing, resources *are* getting scarce, the planet *is* dying …

You can *see* that future utopia well enough but there's no guarantee you'll get there. Seven or eight billion plebs may have other ideas when they also see where it's going. They're going to prefer Bastani's model, aren't they? And they are many; you are few. When things

get so bad that even the dimmest realise what's going on, they will stop you, wreck *your* planet *trying* to stop you, or possibly even just wreck *your* planet by continually escalating their fighting amongst themselves. (*That* old distraction will work against you eventually.) One way or another, your perfect future world is going to get shunted off the rails by conflict. So what do you do?

Because, now, surely, looking to ahead to *your* ideal, fully-automated, luxury world, but with dwindling natural resources and several billion other humans, that are ultimately of no use to you whatsoever – but might be minded to try to stop you, *you have to come to a simple conclusion …*

> ***The world is already massively, and from your perspective, entirely unnecessarily, overpopulated.***
>
> ***And it's going to get more so.***
>
> ***And you might start thinking what to do about it?***

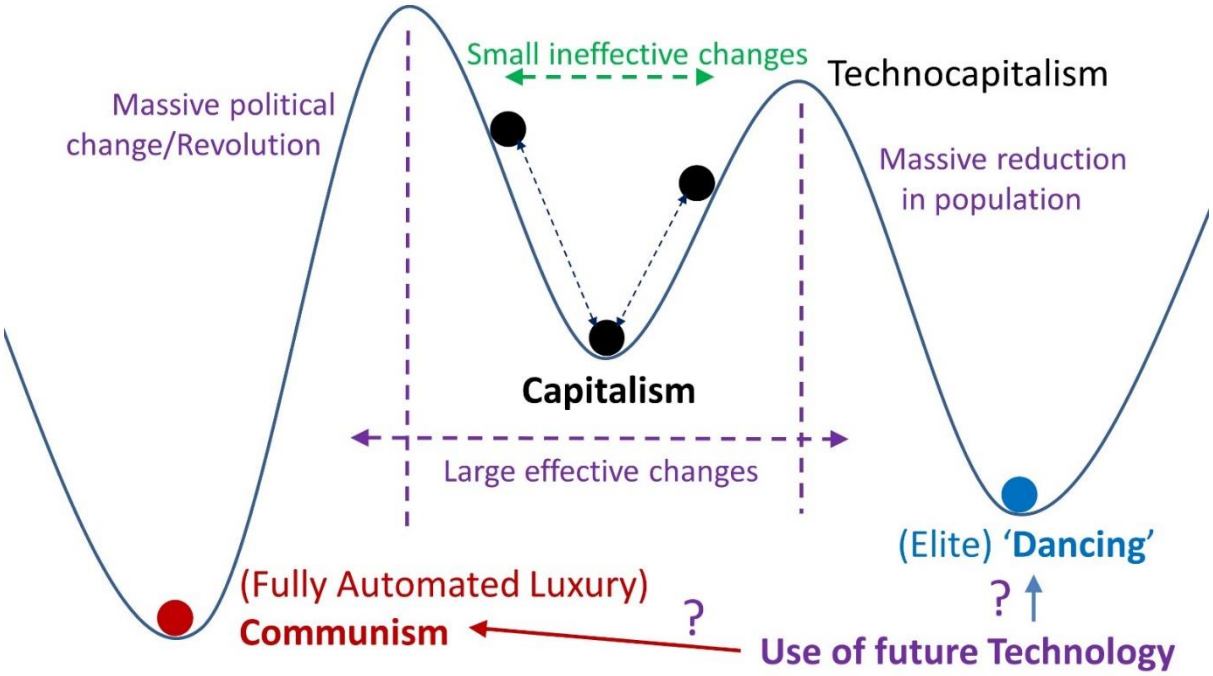

Now, we're going to stop short of suggesting explicitly, in the here and now, that *the world elite are just about to kill off a large part of the world's human population* but, now that we've scared ourselves into some sort of awareness, we have to realistically consider some certain, or very likely points:

- What we've accepted (disliked, maybe, but accepted nonetheless) as a stable human social/political/economic/legal framework is about to be made unstable by technology, population and resources.
- What new model we move towards will be determined by who can wield the most power/control over the next few years, which currently pitches conventional (potentially abstract) ownership of resources against (more concrete) human numbers.

- 'Natural' progression – perhaps even *peaceful* means of change – *should* take us towards Bastani's utopia: it's really the only thing that makes sense with the population we have. To *prevent* it – to send it some other way – will need definite action.
- There will those among the world elite who are aware of this *now* and having these conversations.

OK, yes, of course this is horribly simplified. And even depopulation conspiracy theories aren't entirely new: they go back at least to the 1970s. But we're *not* suggesting that this is something that's happened *already*. Nor are we simply suggesting that humans could be *replaced* by robots: it's more that they're (the vast majority of them anyway) not really *needed at all* in the automated future because there's no longer any need for significant levels of work to be done by anyone or anything. Even the Marxist model breaks eventually. Instead, we need to take a fresh look at what's happening in terms of today's technology, today's natural world and how much of it is going to (very clearly) develop over the next decade or so. *Something* has to happen – so what will it be? Even so, how can you sum up such a significant period of human history, with so many subtle variables, in ten thousand words? You can't: this should be a book but there's not time. And one final reminder: no, none of this will happen overnight. Hopefully, we're not going to wake up tomorrow with the water turned off and the elite suddenly with their own supply (more than they already have), for example. If it happens at all, it will be gradual and, to a large extent, they'll initially continue to use political means: good old 'divide and conquer', of course, while that still works. They'll try to get *us* to do the job for them. And, unsurprisingly, the game of the rich blaming the poor has already started. But, as George Monbiot points out in the article …

> *'At this year's World Economic Forum in Davos, the primatologist Dame Jane Goodall, who is a patron of the charity Population Matters, told the assembled pollutocrats, some of whom have ecological footprints thousands of times greater than the global average: "All these things we talk about wouldn't be a problem if there was the size of population that there was 500 years ago." I doubt that any of those who nodded and clapped were thinking, "yes, I urgently need to disappear".'*

So, yes, sadly, it may be necessary to take sides on this. If we don't, others will and they won't be on ours. The current Covid-19 crisis is a limited, but good example. It's not full-blown depopulation theory – and, even if it was, it wouldn't be touted as such – but it's striking how left-wing governments around the world are more likely to act quickly and forcefully to protect the wider population while the right-wing ones talk of 'herd immunity' and intervene on the side of the economy rather than the national health. (It's also a sub-plot to *Conscious*, of course.) If anything like this *does* happen, it's as likely to be spontaneous and opportunistic as specifically planned.

Or, we can do without the blatant conspiracy model entirely. Ultimately, and to come at this from a different – maybe more plausible – direction, if we look forward to what the world *could* be like in five, ten, fifty years, there are really only three broad possibilities:

1. We're still bumbling along pretty much as we are now. Considering the multiplicity of forced changes ahead, and the clear incompatibility of exponential trends with finite resources, *this really doesn't look possible*.

2. We've moved to a dramatically different, more equitable, more efficient system for sharing the world's resources. This would be great but *can we achieve this? Because – at the moment – we're not in control.*
3. A massive reduction in the human population of the planet, through conflict, famine, disease, etc., *which would suit the people who currently are in control.*

Just as there *has* to be an answer to the Fermi Paradox, for example, there *has* to be a *future* – with or without humans. We (or what's left of us, or even the planet without us) have to be doing *something* in a generation's time. So, roughly speaking, one of these three – or a variant – has to be true. Frankly on our current trajectory, the last one seems by far the most likely, leaving us merely to mull the question of *how much of it will just happen because it has to and how much of it might be helped on its way?*

'Time will tell,' as they say; but perhaps not in a good way?