

# The Twenty-First Century: A Bumpy Ride

**Professor Lord Rees**

---

*Martin Rees, Lord Rees of Ludlow OM FRS, is the UK's Astronomer Royal. He is based at Cambridge University where he is a Fellow (and former Master) of Trinity College. He is a former President of the Royal Society and a member of many foreign academies. His research interests include space exploration, high energy astrophysics, and cosmology. He is co-founder of the Centre for the Study of Existential Risk at Cambridge University (CSER), and has served on many bodies connected with education, space research, arms control, and international collaboration in science. In addition to his research publications he has written many general articles and ten books, including most recently On the Future: Prospects for Humanity (paperback version due in October 2021).*

*Professor Lord Rees introduces ideas presented to the American Philosophical Society on the risks posed by climate change and technology. He also outlines how governments can mitigate them.*

---

## Introduction

COVID-19 should not have struck us so unawares: similar viruses, SARS and MERS, had emerged within the last 20 years, and global pandemics had been widely discussed. So why were even rich countries so unprepared? It's because politicians and the public have a local focus. They downplay the long-term and the global. They ignore Nate Silver's maxim: 'The unfamiliar is not the same as the improbable.'

Indeed, we're in denial about a whole raft of newly emergent threats to our interconnected world, that could be devastating. Pandemics and massive cyberattacks, for instance, are immediately destructive. Their probability may seem low, but they could happen at any time. The worst of them could be so devastating that one occurrence would be too many. And their probability and potential severity is increasing. Indeed, I fear we are guaranteed a bumpy ride through this century. COVID-19 must be a wake-up call, reminding us—and our governments—that we're vulnerable.

Humans are now so numerous, and have now such a heavy collective 'footprint', that they can transform, or even ravage, the entire biosphere. The world is growing, and a more demanding population puts the natural environment under strain. Our collective actions could trigger dangerous climate change and mass extinctions if 'tipping points' are crossed—outcomes that would bequeath a depleted and impoverished world to future generations. We're familiar with these threats, but fail to prioritise countermeasures, because their worst impact stretches beyond the time horizon of political and investment decisions. It's like the proverbial boiling frog—contented in a warming tank until it's too late to save itself. We have endured a 'plague year', and it remains unclear when, or indeed if, the world will revert to anything close to its 'old normal'.

The 'global spasm' that we have collectively experienced—a spasm that is, at the time of this writing, far from over—shows clearly that the ability to make wise decisions based on science has a direct impact on survival—not just personally, but collectively. Because our entire world is so interconnected, a catastrophe in any region can cascade globally, making our society vulnerable to breakdowns. But well-directed, internationally deployed science and technology can offer salvation.

The potentials of biotech and the cyberworld are exhilarating—but they're frightening too. We are already, individually and collectively, so greatly empowered by rapidly changing technology that we can—by design, or as unintended consequences—engender global changes that will resonate for centuries.

## Climate and environment

There are some things we can confidently predict. For instance, there's firm evidence for climate change. Even within the next 20 years, regional shifts in climatic patterns, and more extreme weather, will aggravate pressures on food and water, and enhance migration pressure. Moreover, under 'business as usual' scenarios we can't rule out, later in the century, really catastrophic global warming, and tipping points triggering long-term trends like the melting of Greenland's ice sheet. But even those who accept these statements have diverse views on the best policy response. These divergences stem from differences in economics and ethics—in particular, in how much obligation we should feel towards future generations.

The Danish campaigner Bjørn Lomborg has bogeyman status among environmentalists—somewhat unfairly, as he doesn't contest the science. But his 'Copenhagen Consensus' of economists

downplays the priority of addressing climate change in comparison with shorter-term efforts to help the world's poor. That's because he applies a 'standard' discount rate—and in effect writes off what happens beyond 2050. But if you care about those who'll live into the twenty-second century and beyond, then, as economists like Lord Stern and Professor Martin Weitzman argue, it is worth paying an insurance premium now, to protect those generations against the worst-case longer-term scenarios.<sup>1</sup>

So, even those who agree that there's a significant risk of climate catastrophe a century hence, will differ in how urgently they advocate action today. Their assessment will depend on expectations of future growth, and optimism about technological fixes. But, above all, it depends on an ethical issue—in optimising people's life-chances, should we discriminate on grounds of date of birth?

That the world will get warmer is a confident prediction. And with similar confidence we expect that it will get more crowded during this century. 50 years ago, world population was about 3.5 billion. It's now about 7.7 billion. The growth has been mainly in Asia and Africa. The number of births per year, worldwide, peaked a few years ago and is going down. Nonetheless, world population is forecast to rise to around nine billion by 2050. That's partly because most people in the developing world are young. They are yet to have children, and they will live longer. The age histogram in the developing world will become more like it is in Europe. By mid-century, Africa will have five times Europe's population. Lagos and other megacities could have populations of around 40 million.

Population growth seems under-discussed. That's partly, perhaps, because doom-laden forecasts in the late 1960s—for instance, by the Club of Rome and by Paul Ehrlich—proved off the mark. Also, some deem population growth to be a taboo subject—tainted by association with eugenics in the 1920s and 1930s, with Indian policies under Indira Gandhi, and more recently with China's hard-line one-child policy. As it's turned out, food production and resource extraction have kept pace with the rising population. Famines still occur, but they're due to conflict or maldistribution, not overall scarcity.

To feed nine billion in 2050 will require further-improved agriculture—low-till, water-conserving, and GM crops. It may also require dietary innovations—converting insects, highly nutritious and rich in proteins, into palatable food; and making artificial meat. To quote Gandhi—enough for everyone's need but not for everyone's greed.

Demographics beyond 2050 are uncertain. It's not even clear whether there'll be a continuing global rise, or a fall. Urbanisation, declining infant mortality, and women's education trigger the transition towards lower birthrates—but there could be countervailing cultural influences.

If, for whatever reason, families in Africa remain large, then according to the UN that continent's population could double again by 2100, to four billion, thereby raising the global population to 11 billion. Nigeria alone would by then have as big a population as Europe and North America combined.

Optimists may note that each extra mouth brings two hands and a

<sup>1</sup> I'd note that there's one policy context when an essentially zero discount rate is applied: radioactive waste disposal, where the depositories are required to prevent leakage for at least 10,000 years. This is somewhat ironic, since we can't plan the rest of energy policy even 30 years ahead.

brain. But the potential geopolitical stresses of runaway population growth are deeply worrying. As compared to the fatalism of earlier generations, those in poor countries now know, via the Internet etc, what they're missing. And migration is easier. Moreover, the advent of robots, and 'reshoring' of manufacturing, mean that still-poor countries won't be able to grow their economies by offering cheap skilled labour, as the Asian tiger economies did. It's a portent for disaffection and instability—multiple megaversions of the tragic loads of boat people crossing the Mediterranean today. Wealthy nations, especially those in Europe, should urgently promote prosperity in Africa, and not just for altruistic reasons.

And another thing: if humanity's collective impact on land use and climate is too deep, the resultant 'ecological shock' could cause mass extinctions. We'd be destroying the book of life before we'd read it. Already, there's more biomass in chickens and turkeys than in all the world's wild birds. And the biomass in humans, cows, and domestic animals is 20 times that in wild mammals.

Biodiversity is a crucial component of human wellbeing. We're clearly harmed if fish stocks dwindle to extinction. There are plants in the rain forest whose gene pool might be useful to us. And insects are crucial for the food chain and fertilisation. But for many environmentalists, preserving the richness of our biosphere has value in its own right, over and above what it means to us humans. To quote the great ecologist EO Wilson, 'mass extinction is the sin that future generations will least forgive us for'.

## Prospects for technology

It would be hard to think of a more inspiring challenge for young scientists and engineers than devising clean and economical energy systems—and sustainable, humane agriculture—for the entire world. Nations should accelerate R&D into all forms of low-carbon energy generation, and into other technologies where parallel progress is crucial—especially storage (batteries, compressed air, pumped storage, flywheels, etc) and smart grids. If carbon-free energy gets cheap enough, then India, for instance, can leapfrog to it. The health of the poor is jeopardised by smoky stoves burning wood or dung, and there would otherwise be pressure to build coal-fired power stations. Likewise, public health should be a global priority.

But we need wisely directed technology. Indeed, many are anxious that innovation is proceeding so fast that we may not properly cope with it—and that we'll have a bumpy ride through this century. We're ever more dependent on elaborate networks: electric power grids, air traffic control, international finance, just-in-time delivery, globally dispersed manufacturing, and so forth. Unless these networks are highly resilient, their manifest benefits could be outweighed by catastrophic (albeit rare) breakdowns that cascade globally—real-world analogues of what happened in 2008 to the financial system. Air travel can spread a pandemic worldwide within days.<sup>2</sup> And social media can spread panic and rumour, and psychic and economic contagion, literally at the speed of light.

Biotech offers huge prospects for enhancing health and food

<sup>2</sup> Pandemics also cause far more societal breakdown than in earlier centuries. English villages in the fourteenth century continued to function even when the black death halved their populations. In contrast, our societies would be vulnerable to serious unrest as soon as hospitals were overwhelmed— which would occur before the fatality rate was even one percent. (And there's likewise huge societal risk from cyberattacks on infrastructure, etc.)

production. But there are downsides, from both ethical and prudential perspectives. It offers, for instance, the ability to modify viruses. In 2012, experiments done in Wisconsin and in Holland showed that it was surprisingly easy to make the influenza virus more virulent and more transmissible. This seemed a portent, and in 2014 the US federal government ceased funding these ‘gain of function’ experiments. Similar manipulations can be carried out on coronaviruses. There is of course no suggestion that COVID-19 was malevolently engineered, though there is an ongoing debate about the possibility that it could have been an accidental release from the Wuhan Institute of Virology, where it is known that gain of function experiments were being done.

The new CRISPR-Cas9 technique for gene editing is hugely promising, but there are already ethical concerns—for instance, about Chinese experiments modifying embryos—and anxiety about possible runaway ecological consequences of ‘gene drive’ programmes to wipe out species as diverse as mosquitos or grey squirrels.

Governments will surely adopt a stringent and precautionary attitude to the applications of biotech—and even to the kinds of experiment that can be legally pursued. But I’d worry that whatever regulations are imposed can’t be enforced worldwide, any more than the drug laws or tax laws. Whatever can be done will be done by someone, somewhere.

An atomic bomb can’t be built without large-scale special-purpose facilities. But biotech involves small-scale dual-use equipment. Indeed, biohacking is burgeoning even as a hobby. The rising empowerment of tech-savvy groups (or even individuals), by bio-as well as cyber-technology, will pose an intractable challenge to governments and aggravate the tension between freedom, privacy, and security. The global village will have its village idiots, and they’ll have global range. These concerns are relatively near-term—within ten or 15 years.

By mid-century we might expect two things: a better understanding of the combinations of genes that determine key characteristics of humans and animals; and the ability to synthesise genomes that match these features. If it becomes possible to ‘play God on a kitchen table’, our ecology (and even our species) may not long survive unscathed.

And what about another transformative technology: robotics and artificial intelligence (AI)? DeepMind’s ‘AlphaGo Zero’ computer famously achieved world championship level in the games of Go and chess in just a few hours—it was given just the rules, and learnt by playing against itself over and over again. Its processing speed allowed it to complete several games every second.

Already AI can cope better than humans with complex fast-changing networks—traffic flow, or electric grids. It could let the Chinese gather and process all the information needed to run an efficient planned economy that Marx could only dream of. And in science, its capacity to explore millions of options could allow it to discover recipes for better drugs, or a material that conducts electricity with zero resistance at room temperature. Computers learn to identify dogs, cats, and human faces by ‘crunching’ through millions of images—not the way babies learn. They learn to translate by reading millions of pages of multilingual text—EU documents, for instance (their boredom threshold is infinite!).

The implications for our society are already ambivalent. If there is

a ‘bug’ in the software of an AI system, it is not always possible to track it down. This is likely to create public concern if the system’s ‘decisions’ have potentially grave consequences for individuals. If we are sentenced to a term in prison, recommended for surgery, or even given a poor credit rating, we would expect the reasons to be accessible to us, and contestable by us. If such decisions were delegated to an algorithm, we would be entitled to feel uneasy, even if presented with compelling evidence that, on average, the machines make better decisions than the humans they have usurped.

AI systems will become more intrusive and pervasive. Records of all our movements, our health, and our financial transactions, will be in the ‘cloud’, managed by a multinational quasi-monopoly. The data may be used for benign reasons (for instance, for medical research, or to warn us of incipient health risks), but its availability to internet companies is already shifting the balance of power from governments to globe-spanning conglomerates.

There will be other privacy concerns. Are you happy if a random stranger sitting near you in a restaurant or on public transportation can, via facial recognition, identify you and invade your privacy? Or if fake videos of you become so convincing that visual evidence can no longer be trusted? Or if a machine knows enough about you to compose emails that seem to come from you? The ‘arms race’ between cybercriminals and those trying to defend against them will become still more expensive and vexatious when drones, driverless cars, etc proliferate.

Many experts think that AI, like synthetic biotech, already needs guidelines for ‘responsible innovation’. But others, like the roboticist Rodney Brooks (creator of the Baxter robot and the Roomba vacuum cleaner), think that for many decades artificial intelligence will be less of a concern than real stupidity. And machines are still clumsy compared to children in sensing and interacting with the real world.

The incipient shifts in the nature of work have been addressed in several excellent books by economists and social scientists. Clearly, machines will take over much of the work of manufacturing and retail distribution. They can supplement, if not replace, many white-collar jobs: routine legal work, accountancy, computer coding, medical diagnostics, and even surgery. Many ‘professionals’ will find their hard-earned skills in less demand. In contrast, some skilled service-sector jobs—plumbing and gardening, for instance—require non-routine interactions with the external world and will be among the hardest jobs to automate.

The digital revolution generates enormous wealth for innovators and global companies, but preserving a healthy society will surely require redistribution of that wealth. There is talk of using it to provide a universal income. It is better when all who are capable of doing so can perform socially useful work rather than receive a handout.

Indeed, to create a humane society, governments will need to vastly enhance the number and status of those who care for the old, the young, and the sick. There are currently far too few, and they’re poorly paid, inadequately esteemed, and insecure in their positions. Such work is more fulfilling than a job in a call centre or Amazon warehouse. I can foresee this benign redeployment happening in Scandinavia, though there might be ideological barriers in some other nations. We surely hope, when old, to be cared for by someone with real, not synthetic, empathy. We want young children to be told stories by real people who can share and understand their emotions. It is likely that society will be transformed by autonomous robots,

even though the jury is out on whether they will be *idiots savants* or display superhuman capabilities.

If robots become less clumsy in interacting with the world, would they truly be perceived as intelligent beings? Would we then have obligations towards them? Should we feel guilty if they are underemployed or bored?

Ray Kurzweil, author of *The Age of Spiritual Machines*, even foresees that humans could transcend biology by merging with computers. In old-style spiritualist parlance, they would 'go over to the other side'. We then confront the classic philosophical problem of personal identity. If your brain were downloaded into a machine, in what sense would it still be 'you'? Or are the input into our sense organs, and our physical interactions with the real external world, so essential to our being that this transition would be not only abhorrent but also impossible? These are ancient conundrums for philosophers, but practical ethicists may soon need to address them.

Not even Kurzweil thinks this will happen in his lifetime, so he wants his body frozen until immortality's on offer, and he can be resurrected into some posthuman world.<sup>3</sup> But of course research on ageing is being seriously prioritised. Some think it's a 'disease' that can be cured. Dramatic life extension would plainly have huge ramifications, for society and population projections.

It's certainly credible that human beings—in their mentality and their physique—may become malleable through genetic and cybernetic technologies. Moreover, this future evolution—a kind of secular 'intelligent design'—would take only centuries, in contrast to the thousands of centuries needed for Darwinian evolution. This is a game changer. When we admire the literature and artefacts that have survived from antiquity, we feel an affinity, across a time gulf of thousands of years, with those ancient artists and their civilisations. But we can have zero confidence that the dominant intelligences a few centuries hence will have any emotional resonance with us, even though they may have an algorithmic understanding of how we behaved.

## Prospects in space

And now I turn briefly to another technology: space. This is where robots surely have a future, and where I'd argue that these changes will happen fastest and should worry us less.

We depend every day on space for satnav, environmental monitoring, communication, and so forth. These are in large part now commercially funded, though projects with a focus on scientific research and planetary exploration are bankrolled by national or international agencies.

During this century the whole solar system will be explored by swarms of miniature probes, far more advanced than the probes that have beamed back pictures of Saturn's moons, of Pluto, and beyond—20,000 times further away than the Moon. Think back to the computers and phones of the 1990s, when these probes were designed, and realise how much better we can do today. The next

<sup>3</sup> I was surprised to find that three academics back in England had gone in for these 'cryonics'. Two had paid the full whack; the third had taken the cut-price option of wanting just his head frozen. I was glad they were from Oxford, not from my university. For my part, I'd rather end my days in an English churchyard than an American refrigerator.

step will be deployment in space of robotic fabricators, which can build large structures under zero gravity—for instance, solar energy collectors, or giant telescopes with huge gossamer-thin mirrors

What about manned spaceflight? The practical case gets weaker with each advance in robots and miniaturisation. Were I an American, I would only support NASA's *unmanned* programme. And I certainly wouldn't support a manned programme done by the European Space Agency. I would argue that private ventures like Elon Musk's SpaceX or Jeff Bezos' Blue Origin—bringing a Silicon Valley culture into a domain long dominated by NASA and a few aerospace conglomerates—should 'front' all manned missions. They can impose higher risks than can a Western country on publicly funded civilian astronauts, and thereby slash costs. There would still be many volunteers—some willing to accept the risk of 'one-way tickets'—driven by the same motives as early explorers, mountaineers, and the like.

By 2100, courageous thrill-seekers may have established 'bases' independent from the Earth—on Mars, or maybe on asteroids. Elon Musk says he wants to die on Mars (though not on impact). But don't ever expect mass emigration from Earth. Nowhere in our solar system offers an environment as clement as even the Antarctic or the top of Everest. Here I disagree with Musk and my late colleague Stephen Hawking. It's a dangerous delusion to think that space offers an escape from Earth's problems. Dealing with climate change on Earth is a doddle compared to terraforming Mars. There's no 'planet B' for ordinary risk-averse people.

But those pioneer adventurers who escape the Earth could be cosmically important. This is why. They'll be ill-adapted to their new environment, and beyond the clutches of our terrestrial regulators. They will use all the resources of genetics and cybernetics to adapt. They will change faster and could within a few centuries become a new species. Moreover, if they make the transition to fully inorganic intelligences, they won't need an atmosphere. They may prefer zero-G. They'll also be nearly immortal. So it's in deep space—not on Earth, nor even on Mars—that non-biological 'brains' may develop powers that humans can't even imagine.

This raises the question that astronomers are asked most often: is there life out there already? Or is a sterile cosmos awaiting our progeny? We know too little about how life began on Earth to lay confident odds. We don't know what triggered the transition from complex molecules to entities that can metabolise and reproduce. Moreover, even if simple life is common, it is not clear whether it's likely to evolve into anything intelligent or complex.

Maybe we'll one day find ET. On the other hand, Earth's intricate biosphere could be unique. But it's important that this wouldn't render life a cosmic sideshow. That's because there's abundant time ahead for posthuman life seeded from Earth to pervade the Galaxy. We're the outcome of four billion years of Darwinian evolution, but the Sun is less than half way through its life. And the universe may continue for ever. To quote Woody Allen, eternity is very long, especially towards the end.

But even in this 'concertina'ed timeline, extending billions of years into the future as well as into the past, we're living in a special century. The century when humans could jump-start the transition to entities that far transcend our limitations, and eventually spread their influence far beyond the Earth. Or—to take a darker view—the century where our follies could foreclose the immense future potential and leave an anarchic and depleted planet.



## On our future, this century

Zooming back closer to the here and now, one can offer some tentative hopes, fears, and recipes.

Technologies offer huge promise. But our society is brittle, interconnected, and vulnerable. We fret unduly about small risks—air crashes, carcinogens in food, low radiation doses, etc. But we're in denial about some newly emergent threats that could be globally devastating. Some of these are environmental—the pressures of a growing and more demanding population. Others are the potential downsides of novel technologies.

And, of course, most of the challenges are global. Coping with potential shortage of food, water, and resources—and transitioning to low-carbon energy—can't be achieved by each nation separately. Nor can regulation of potentially threatening innovations. Indeed, a key issue is whether nations need to give up more sovereignty to new organisations along the lines of the International Atomic Energy Agency, the World Health Organization, etc.

Scientists have an obligation to promote beneficial applications of their work and warn against the downsides. Universities and academies need to assess which scary scenarios—eco-threats, or risks from misapplied technology—can be dismissed as science fiction, and how best to avoid the hazards that cannot be so dismissed.

The trouble is that even the best politicians focus mainly on the urgent and parochial. They do not focus on long-term global issues, or on averting possible catastrophes that haven't yet happened, unless such policies feature sufficiently prominently in the press and in their inboxes that they are confident they won't lose votes by endorsing them.

So concerned scientists must enhance their leverage—by involvement with NGOs, via blogging and journalism, and by enlisting charismatic individuals and the media to amplify their voices. Here are two recent instances:

The Papal encyclical *Laudato Si'* had a worldwide influence in the lead-up to the Paris climate conference in 2015. There's no gainsaying the Catholic Church's global reach, long-term vision, and concern for the world's poor.

And I doubt that we in the UK would have legislated against non-biodegradable plastic waste had it not been for the BBC's *Blue Planet II* television programmes fronted by our secular pope, David Attenborough. The images of albatrosses returning to their nests and regurgitating plastic debris are as iconic as the polar bear on the melting ice floe was in the climate debate.

It's encouraging to witness more activists among the young, who can hope to live to the end of the century. Their vocal commitment is welcome. It gives grounds for hope.

I'll end with a flashback, right back to the Middle Ages. For medieval people, the entire cosmology—from creation to apocalypse—spanned only a few thousand years. They were bewildered and helpless in the face of floods and pestilences, and prone to irrational dread. Large parts of the Earth were *terra incognita*.

But they built cathedrals, constructed with primitive technology by masons who knew they wouldn't live to see them finished—vast and glorious buildings, that still inspire us centuries later.

In contrast, our horizons in space and time are now vastly extended, as are our resources and knowledge. But we don't plan centuries ahead. This seems a paradox. But there is a reason. Medieval lives played out against a 'backdrop' that changed little from one generation to the next. They were confident that they'd have grandchildren who would appreciate the finished cathedral. But for us, unlike for them, the next century will be drastically different from the present. We can't foresee it, so it's harder to plan for it. There is now a huge disjunction between the ever-shortening timescales of social and technological change and the billion-year time spans of biology, geology, and cosmology.

'Spaceship Earth' is hurtling through the void. Its passengers are anxious and fractious. Their life-support system is vulnerable to disruption and breakdowns. But there is too little planning—too little horizon-scanning. This 'pale blue dot' in the cosmos is a special place. It may be a unique place. And we're its stewards at a specially crucial era. That's an important message for us all, whether or not we're astronomers.

We need to think globally, we need to think rationally, we need to think long-term. We need to be 'good ancestors', empowered by twenty-first-century technology but guided by values that science alone can't provide.