Some Thoughts on 2050 and Beyond[1]

MARTIN REES
Former Master, Trinity College at the University of Cambridge
Astronomer Royal

The flames that engulfed Notre Dame Cathedral in April 2019 threatened one of the "peaks" of our Western heritage—and gave us cause to marvel at the ambition and vision that its construction entailed. Its builders knew of nothing beyond Europe; they thought there might be an apocalypse within a millennium. But despite these constricted horizons, in both time and space, despite the deprivation and harshness of their lives, despite their primitive technology and meager resources, they built this immense and glorious building—extending the frontiers of what was possible. Those who conceived the cathedral knew they wouldn't live to see it finished. Their legacy still elevates our spirits, nearly a millennium later.

Unlike our forebears, we know a great deal about our world—and indeed about what lies beyond. Many phenomena still make us fearful, but the advance of science spares us from irrational dread. We know that we are stewards of a precious "pale blue dot" in a vast cosmos—a planet with a future measured in billions of years—whose fate depends on humanity's collective actions this century. But all too often our focus is short-term and parochial. Planning horizons span decades at most; they're not matched to the global challenges that face us in the 21st century.

There's an explanation for this seeming paradox. Medieval people's lives played out against a "backdrop" that changed slowly; they thought their children and grandchildren would lead similar lives and share their faith and culture. But, unlike our remote forebears, we expect the lives and priorities of new generations to be unpredictably different from ours. That's why we, and our governments, react with torpor to the compelling concerns about the future: we retreat into inaction because we're not confident enough of any scenario to commit to it.

Humans are now so numerous and have such a heavy collective "footprint" that they have the ability to transform, or even ravage, the entire biosphere. The world's growing and more demanding population puts the natural environment under strain; our collective actions

1 Read 27 April 2019.

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.

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 warming and tipping points triggering long-term trends like the melting of Greenland's ice cap. But even those who accept these statements have diverse views on the policy response. These divergences stem from differences in economics and ethics—in particular, in how much obligation we should feel toward future generations.

The Danish campaigner Bjorn Lomberg has boogeyman 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 22nd century and beyond, as economists like Nicholas Stern and Martin Weitzman argued, then you deem it worth paying an insurance premium now, to protect those generations against the worst-case longer-term scenarios.

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 optimizing people's life chances, should we discriminate on grounds of date of birth?[2]

2 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—somewhat ironic when we can't plan the rest of energy policy even 30 years ahead.

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. Fifty years ago, the 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; fertility in most countries is now going down. Nonetheless, world population is forecast to rise to around 9 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. And because the demographic transition hasn't yet occurred in much of Africa and India, by mid-century, Africa will have five times Europe's population—Lagos and other megacities could have populations around 40 million.

Population growth seems under-discussed. That's partly, perhaps, because doom-laden forecasts in the late 1960s—by, for instance, the Club of Rome and Paul Ehrlich—proved off the mark. Also, some deem it 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 turned out, food production has kept pace with rising population; famines still occur, but they're due to conflict or maldistribution, not overall scarcity.

To feed 9 billion people in 2050 will require further-improved agriculture—low-till, water-conserving, and genetically modified crops—and maybe dietary innovations, such as 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 rise in global population or a fall. Urbanization, declining infant mortality, and women's education trigger the transition toward lower birth rates, but there could be countervailing cultural influences.

If, for whatever reason, families in Africa remain large, then according to the United Nations that continent's population could double again by 2100, to 4 billion, thereby raising the global population to 11 billion. Nigeria alone would have as big a population as Europe and North America combined.

Optimists may note that each extra mouth brings two hands and a 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 labor, as the Asian Tiger states did. It's a portent for disaffection and instability—multiple mega-versions of the tragic cargoes 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 pushes too hard, the resultant "ecological shock" could trigger mass extinctions—we'd be destroying the book of life before we've 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 rainforest whose gene pool might be useful to us. And insects are crucial for the food chain and fertilization. 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 E. O. 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 and 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, India, for instance, can leapfrog to it. The health of the poor is jeopardized 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 of us are anxious that innovation is proceeding so fast that we may not properly cope with its downsides and failures 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 analogs of what happened in 2008 to the financial system. Air travel can spread a pandemic worldwide within days. And social media can spread panic and rumor, and psychic and economic contagion, literally at the speed of light.[3]

Biotech offers huge prospects for enhancing health and food 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 U.S. federal government ceased funding these "gain of function" experiments.

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" programs 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 restrain some kinds of scientific experiments. But I'd worry that whatever regulations are imposed can't be enforced worldwide—any more than the drug laws can, or the tax laws. Whatever can be done will be done by someone, somewhere.

And that's a nightmarish prospect. An atomic bomb can't be built without large-scale special-purpose facilities: regulation is feasible. 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 10 or 15 years. By mid-century we might expect two further developments: a better understanding of the combination of genes that determine key characteristics of humans and animals, and the ability to synthesis 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.

3 And, by the way, pandemics could cause far more societal breakdown today than in earlier centuries. English villages in the 14th 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 1 percent. (And there's likewise huge societal risk from cyber-attacks on infrastructure, etc.)

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 enable the Chinese to 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 humans' 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 cyber-criminals 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 robots and the Roomba vacuum cleaner) think that for many decades we'll be less concerned about artificial intelligence than about 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 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 that all who are capable of doing so should 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 center 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 "idiot 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 toward them? Should we feel guilty if they are underemployed or bored?

Ray Kurzweil, author of *The Age of Spiritual Machines*, 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 is the input into our sense organs, and

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.

But not even Kurzweil thinks this will happen in his lifetime, so he wants his body frozen until immortality is on offer and he can be resurrected into some post-human world.[4] But of course research on aging is being seriously prioritized. 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—their mentality and their physique—may become malleable through genetic and cyborg 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 artifacts that have survived from antiquity, we feel an affinity, across a time gulf of thousands of years, with those ancient artists and their civilizations. 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 their hegemony will happen fastest and should worry us less.

We depend every day on space for satnav, environmental monitoring, communication, and so forth. Much of this is 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 miniaturized probes—far more advanced than those which have already beamed back pictures of Saturn's moons, Pluto, and beyond (more than 10,000 times further away than the Moon). Think back to the computers and phones of the 1990s, when these probes were designed, and realize how much better we can do today. The next step will be deployment in space of robotic fabricators, which can build

4 I was surprised to find that three academics back in England had gone in for "cryonics." Two paid the full whack; the third has 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.

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 ever weaker with each advance in robots and miniaturization. Were I an American taxpayer I would only support NASA's *un*manned program, and I certainly wouldn't support a manned program done by ESA in Europe. I would argue that private-enterprise ventures like Elon Musk's SpaceX or Jeff Bezos's 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 take higher risks than a western country can impose on publicly funded civilian astronauts and thereby slash costs. There would still be many volunteers—some perhaps even accepting "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—but not on impact.) But don't ever expect mass emigration from Earth. Nowhere in our solar system offers an environment even as clement as 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; they'll be beyond the clutches of our terrestrial regulators. They will use all the resources of genetics and cyborg technology to adapt; they will (unlike us) have an incentive to change 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, and they'll be near-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.

But 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 metabolize 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 the latter wouldn't render life a cosmic

sideshow. That's because there's abundant time ahead for post-human life seeded from Earth to pervade the galaxy. We're the outcome of 4 billion years of Darwinian evolution. But the Sun is less than halfway through its life. And the universe may continue forever—to quote Woody Allen, eternity is very long, especially toward the end.

But even in this concertinaed 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

So finally, 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. A wise mantra is that "the unfamiliar is not the same as the improbable."

And, of course, most of the challenges are global. Coping with potential shortage of food, water, resources—and transitioning to low-carbon energy—can't be solved 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 organizations along the lines of the IAEA, WHO, etc.

Scientists have an obligation to promote beneficial applications of their work and warn against these downsides. Universities and academies need to assess which scary scenarios—ecothreats 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, not on long-term global issues, nor on averting possible catastrophes that haven't yet happened—unless such policies feature sufficiently prominently in the press and in their inboxes, so that they are confident they won't lose votes by endorsing them. Concerned scientists must therefore enhance their leverage—by involvement with NGOs, via blogging and journalism, and by enlisting

charismatic individuals and the media to amplify their voice. 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. The pope got a standing ovation at the UN: he has a billion followers in Latin America, Africa, and East Asia. 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 be legislating against non-degradable plastic waste had it not been for the BBC's *Blue Planet II* TV programs fronted by our secular pope, David Attenborough, especially the image of albatrosses returning to their nests and regurgitating plastic debris—an image 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 expect to live to the end of the century. Their vocal commitment is welcome; it gives grounds for hope.

"Spaceship Earth" is hurtling through the void. Its passengers are anxious and fractious. Their life-support system is vulnerable. But there is too little planning, too little horizon-scanning, too little long-term vision. We need to think globally, we need to think rationally, we need to think long-term—empowered by 21st-century technology but guided by values that science alone can't provide.

I give the last word to one of my scientific heroes, the eloquent immunologist Peter Medawar: "The bells that toll for mankind are . . . like the bells of Alpine cattle. They are attached to our own necks, and it must be our fault if they do not make a tuneful and melodious sound."