Foreword

In 1903, H.G. Wells gave a lecture at the Royal Institution in London, highlighting the risk of global disaster: 'It is impossible', proclaimed the young Wells, 'to show why certain things should not utterly destroy and end the human race and story; why night should not presently come down and make all our dreams and efforts vain. ... something from space, or pestilence, or some great disease of the atmosphere, some trailing cometary poison, some great emanation of vapour from the interior of the earth, or new animals to prey on us, or some drug or wrecking madness in the mind of man.' Wells' pessimism deepened in his later years; he lived long enough to learn about Hiroshima and Nagasaki and died in 1946.

In that year, some physicists at Chicago started a journal called the *Bulletin of Atomic Scientists*, aimed at promoting arms control. The 'logo' on the Bulletin's cover is a clock, the closeness of whose hands to midnight indicates the editor's judgement on how precarious the world situation is. Every few years the minute hand is shifted, either forwards or backwards.

Throughout the decades of the Cold War, the entire Western World was at great hazard. The superpowers could have stumbled towards Armageddon through muddle and miscalculation. We are not very rational in assessing relative risk. In some contexts, we are absurdly risk-averse. We fret about statistically tiny risks; carcinogens in food, a one-in-a-million chance of being killed in train crashes, and so forth. But most of us were 'in denial' about the far greater risk of death in a nuclear catastrophe.

In 1989, the Bulletin's clock was put back to seventeen minutes to midnight. There is now far less chance of tens of thousands of bombs devastating our civilization. But there is a growing risk of a few going off in a localized conflict. We are confronted by proliferation of nuclear weapons among more nations – and perhaps even the risk of their use by terrorist groups.

Moreover, the threat of global nuclear catastrophe could be merely in temporary abeyance. During the last century the Soviet Union rose and fell; there were two world wars. In the next hundred years, geopolitical realignments could be just as drastic, leading to a nuclear stand-off between new superpowers, which might be handled less adeptly (or less luckily) than the Cuba crisis, and the other tense moments of the Cold War era. The nuclear

threat will always be with us – it is based on fundamental (and public) scientific ideas that date from the 1930s.

Despite the hazards, there are, today, some genuine grounds for being a techno-optimist. For most people in most nations, there has never been a better time to be alive. The innovations that will drive economic advance – information technology, biotechnology and nanotechnology – can boost the developing as well as the developed world. Twenty-first century technologies could offer lifestyles that are environmentally benign – involving lower demands on energy or resources than those demanded by what we consider a good life today. And we could readily raise the funds – were there the political will – to lift the world's two billion most-deprived people from their extreme poverty.

But, along with these hopes, twenty-first century technology will confront us with new global threats – stemming from bio-, cyber- and environmental-science, as well as from physics – that could be as grave as the bomb. The Bulletin's clock is now closer to midnight again. These threats may not trigger sudden worldwide catastrophe – the doomsday clock is not such a good metaphor – but they are, in aggregate, disquieting and challenging. The tensions between benign and damaging spin-offs from new technologies, and the threats posed by the Promethean power science, are disquietingly real. Wells' pessimism might even have deepened further were he writing today.

One type of threat comes from humanity's collective actions; we are eroding natural resources, changing the climate, ravaging the biosphere and driving many species to extinction.

Climate change looms as the twenty-first century's number-one environmental challenge. The most vulnerable people – for instance, in Africa or Bangladesh – are the least able to adapt. Because of the burning of fossil fuels, the CO_2 concentration in the atmosphere is already higher than it has ever been in the last half million years – and it is rising ever faster. The higher CO_2 rises, the greater the warming – and, more important still, the greater will be the chance of triggering something grave and irreversible: rising sea levels due to the melting of Greenland's icecap and so forth. The global warming induced by the fossil fuels we burn this century could lead to sea level rises that continue for a millennium or more.

The science of climate change is intricate. But it is simple compared to the economic and political challenge of responding to it. The market failure that leads to global warming poses a unique challenge for two reasons. First, unlike the consequences of more familiar kinds of pollution, the effect is diffuse: the CO_2 emissions from the UK have no more effect here than they do in Australia, and vice versa. That means that any credible framework for mitigation has to be broadly international. Second, the main downsides are not immediate but lie a century or more in the future: inter-generational justice comes into play; how do we rate the rights and interests of future generations compared to our own?

The solution requires coordinated action by all major nations. It also requires far-sightedness – altruism towards our descendants. History will judge us harshly if we discount too heavily what might happen when our grandchildren grow old. It is deeply worrying that there is no satisfactory fix yet on the horizon that will allow the world to break away from dependence on coal and oil – or else to capture the CO_2 that power stations emit. To quote Al Gore, 'We must not leap from denial to despair. We can do something and we must.'

The prognosis is indeed uncertain, but what should weigh most heavily and motivate policy-makers most strongly – is the 'worst case' end of the range of predictions: a 'runaway' process that would render much of the Earth uninhabitable.

Our global society confronts other 'threats without enemies', apart from (although linked with) climate change. High among them is the threat to biological diversity. There have been five great extinctions in the geological past. Humans are now causing a sixth. The extinction rate is one thousand times higher than normal and is increasing. We are destroying the book of life before we have read it. There are probably upwards of ten million species, most not even recorded – mainly insects, plants and bacteria.

Biodiversity is often proclaimed as a crucial component of human well-being. Manifestly it is: we are clearly harmed if fish stocks dwindle to extinction; there are plants in the rain forest whose gene pool might be useful to us. But for many of us these 'instrumental' – and anthropocentric – arguments are not the only compelling ones. Preserving the richness of our biosphere has value in its own right, over and above what it means to us humans.

But we face another novel set of vulnerabilities. These stem not from our collective impact but from the greater empowerment of individuals or small groups by twenty-first century technology.

The new techniques of synthetic biology could permit inexpensive synthesis of lethal biological weapons – on purpose, or even by mistake. Not even an organized network would be required: just a fanatic or a weirdo with the mindset of those who now design computer viruses – the mindset of an arsonist. Bio (and cyber) expertise will be accessible to millions. In our networked world, the impact of any runaway disaster could quickly become global.

Individuals will soon have far greater 'leverage' than present-day terrorists possess. Can our interconnected society be safeguarded against error or terror without having to sacrifice its diversity and individualism? This is a stark question, but I think it is a serious one.

We are kidding ourselves if we think that technical education leads to balanced rationality: it can be combined with fanaticism – not just the traditional fundamentalism that we are so mindful of today, but new age irrationalities too. There are disquieting portents – for instance, the Raelians (who claim to be cloning humans) and the Heavens Gate cult (who committed

collective suicide in hopes that a space-ship would take them to a 'higher sphere'). Such cults claim to be 'scientific' but have a precarious foothold in reality. And there are extreme eco-freaks who believe that the world would be better off if it were rid of humans. Can the global village cope with its village idiots – especially when even one could be too many?

These concerns are not remotely futuristic – we will surely confront them within next ten to twenty years. But what of the later decades of this century? It is hard to predict because some technologies could develop with runaway speed. Moreover, human character and physique themselves will soon be malleable, to an extent that is qualitatively new in our history. New drugs (and perhaps even implants into our brains) could change human character; the cyberworld has potential that is both exhilarating and frightening.

We cannot confidently guess lifestyles, attitudes, social structures or population sizes a century hence. Indeed, it is not even clear how much longer our descendants would remain distinctively 'human'. Darwin himself noted that 'not one living species will transmit its unaltered likeness to a distant futurity'. Our own species will surely change and diversify faster than any predecessor – via human-induced modifications (whether intelligently controlled or unintended), not by natural selection alone. The post-human era may be only centuries away. And what about Artificial Intelligence? Superintelligent machine could be the last invention that humans need ever make. We should keep our minds open, or at least ajar, to concepts that seem on the fringe of science fiction.

These thoughts might seem irrelevant to practical policy – something for speculative academics to discuss in our spare moments. I used to think this. But humans are now, individually and collectively, so greatly empowered by rapidly changing technology that we can – by design or as unintended consequences – engender irreversible global changes. It is surely irresponsible not to ponder what this could mean; and it is real political progress that the challenges stemming from new technologies are higher on the international agenda and that planners seriously address what might happen more than a century hence.

We cannot reap the benefits of science without accepting some risks – that has always been the case. Every new technology is risky in its pioneering stages. But there is now an important difference from the past. Most of the risks encountered in developing 'old' technology were localized: when, in the early days of steam, a boiler exploded, it was horrible, but there was an 'upper bound' to just how horrible. In our ever more interconnected world, however, there are new risks whose consequences could be global. Even a tiny probability of global catastrophe is deeply disquieting.

We cannot eliminate all threats to our civilization (even to the survival of our entire species). But it is surely incumbent on us to think the unthinkable and study how to apply twenty-first century technology optimally, while minimizing

the 'downsides'. If we apply to catastrophic risks the same prudent analysis that leads us to take everyday safety precautions, and sometimes to buy insurance – multiplying probability by consequences – we had surely conclude that some of the scenarios discussed in this book deserve more attention that they have received.

My background as a cosmologist, incidentally, offers an extra perspective – an extra motive for concern – with which I will briefly conclude.

The stupendous time spans of the evolutionary past are now part of common culture – except among some creationists and fundamentalists. But most educated people, even if they are fully aware that our emergence took billions of years, somehow think we humans are the culmination of the evolutionary tree. That is not so. Our Sun is less than half way through its life. It is slowly brightening, but Earth will remain habitable for another billion years. However, even in that cosmic time perspective – extending far into the future as well as into the past – the twenty-first century may be a defining moment. It is the first in our planet's history where one species – ours – has Earth's future in its hands and could jeopardise not only itself but also life's immense potential.

The decisions that we make, individually and collectively, will determine whether the outcomes of twenty-first century sciences are benign or devastating. We need to contend not only with threats to our environment but also with an entirely novel category of risks – with seemingly low probability, but with such colossal consequences that they merit far more attention than they have hitherto had. That is why we should welcome this fascinating and provocative book. The editors have brought together a distinguished set of authors with formidably wide-ranging expertise. The issues and arguments presented here should attract a wide readership – and deserve special attention from scientists, policy-makers and ethicists.

Martin J. Rees