

I'll start with a flashback to 1902. In that year the young HG Wells gave a celebrated lecture at the Royal Institution in London. He spoke mainly in visionary mode. "Humanity", he proclaimed, "has come some way, and the distance we have travelled gives us some earnest of the way we have to go. All the past is but the beginning of a beginning; all that the human mind has accomplished is but the dream before the awakening." His rather purple prose still resonates more than a hundred years later -- he realised that we humans aren't the culmination of emergent life.

But Wells wasn't an optimist. He also highlighted the risk of global disaster: "It is impossible to show why certain things should not utterly destroy and end the human story .. and make all our 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".

I quote Wells because he reflects the mix of optimism and anxiety – and of speculation and science – which I'll try to offer in this lecture. Were he writing today he would have been elated by our expanded vision of life and the cosmos -- but he'd have been even more anxious about the perils we might face. The stakes are indeed getting higher: new science offers huge opportunities, but its consequences could jeopardise our survival. Many are concerned that it is 'running away' so fast that neither politicians nor the lay public can assimilate or cope with it.

My own expertise is in astronomy and space technology. So you may guess that I'm kept awake at night by worry about asteroid impacts. Not so. Indeed this is one of the few threats that we can quantify. Every ten million years or so, a body, a few kilometers across will hit the Earth, causing global catastrophe – there's a few chances in a million that this is how we'll die. But there are larger numbers of smaller asteroids that would cause regional or local devastation. A body (say) 300 metres across, if it fell into the Atlantic, would produce huge tsunamis that would devastate the East Coast of the US, as well as much of Europe. And till smaller impacts are more frequent. One in Siberia in 1908 released energy equivalent to 5 megatons..

Can we be forewarned of these impacts? The answer is yes.. There are plans to survey the million potential earth-crossing asteroids bigger than 50 metres and track their orbits precisely enough to predict possible impacts. With forewarning of an impact. action could be taken to evacuate the most vulnerable areas. Even better news is that during this century we could develop the technology to protect us. A 'nudge', imparted a few years before the threatened impact, would only need to change an asteroid's velocity by a millimeter per second in order to deflect its path away from the Earth.

If you calculate an insurance premium in the usual way, by multiplying probability by consequences, it turns out to be worth spending a billion dollars a year to reduce asteroid risk

Other natural threats – earthquakes and volcanoes – are less predictable. But there's one reassuring thing about them, as there is about asteroids: the annual risk they pose isn't getting bigger. It's the same for us as it was for the Neanderthals – or indeed for the dinosaurs.

HUMAN-INDUCED THREATS

In contrast, the hazards that are the focus of this talk are those that humans themselves engender – and they now loom far larger. And in discussing them I'm straying far from my 'comfort zone' of expertise. So I comment as a 'citizen scientist', and as a worried member of the human race. I'll skate over a range of topics, in the hope of being controversial enough to provoke discussion.

Ten years ago I wrote a book which I entitled 'Our Final Century?' My publisher deleted the question-mark. The American publishers changed the title to 'Our Final Hour'. (Americans seek instant gratification – and the converse).

My theme was this. Our Earth is 45 million centuries old. But this century is the first when one species – ours – can determine the biosphere's fate. I didn't think we'd wipe ourselves out. But I did think we'd be lucky to avoid devastating setbacks. That's because of unsustainable anthropogenic stresses to ecosystems because there are more of us (world population is higher) and we're all more demanding of resources. And – most important of all – because we're empowered by new technology, which exposes us to novel vulnerabilities.

And we've had one lucky escape already.

At any time in the Cold War era -- when armament levels escalated beyond all reason --- the superpowers could have stumbled towards armageddon through muddle and miscalculation. During the Cuba crisis I and my fellow-students participated anxiously in vigils and demonstrations. But we would have been even more scared had we then realised just how close we were to catastrophe. Kennedy was later quoted as having saying as one stage that the odds were 'between one in three and evens'. And only when he was long retired did Robert McNamara, state frankly that "[w]e came within a hairbreadth of nuclear war without realizing it. It's no credit to us that we escaped – Khrushchev and Kennedy were lucky as well as wise." Be that as it may, we were surely at far greater hazard from nuclear catastrophe than from anything nature could do. Indeed the annual risk of thermonuclear destruction during the Cold War was about 10000 times higher than from asteroid impact.

It is now conventionally asserted that nuclear deterrence worked. In a sense, it did. But that doesn't mean it was a wise policy. If you play Russian roulette with one or two bullets in the barrel, you are more likely to survive than not, but the stakes would need to be astonishingly high -- or the value you place on your life inordinately low -- for this to seem a wise gamble. But we were dragooned into just such a gamble.

throughout the Cold War era. It would be interesting to know what level of risk other leaders thought they were exposing us to, and what odds most European citizens would have accepted, if they'd been asked to give informed consent. For my part, I would not have chosen to risk a one in three -- or even one in six -- chance of a disaster that would have killed hundreds of millions and shattered the historic fabric of all our cities, even if the alternative were certain Soviet dominance of Western Europe. And of course the devastating consequences of thermonuclear war would have spread far beyond the countries that faced a direct threat especially if a nuclear winter were triggered.

The threat of global annihilation involving tens of thousands of H-bombs is thankfully in abeyance; there is, though, currently more risk that smaller nuclear arsenals might be used in a regional context, or even by terrorists. But we can't rule out, later in the century, a geopolitical realignment leading to a standoff between new superpowers. So a new generation may face its own "Cuba" - and one that could be handled less well or less luckily than the 1962 crisis was.

Nuclear weapons are based on 20th century science. I'll return later in my talk to the 21st century sciences -- bio, cyber, and AI - and what they might portend.

But before that let's focus on the potential devastation that could be wrought by human-induced environmental degradation and climate change. These threats are long-term and insidious. They stem from humanity's ever-heavier collective 'footprint', which is stressing our finite planet's ecology beyond sustainable limits..

There's nothing new about these concerns. Doom-laden predictions of environmental catastrophe famously came in the 1970s from the Club of Rome, and from Paul Erlich and other groups. These proved wide off the mark. Unsurprisingly, such memories engender scepticism about the worst-case environmental and climatic projections. But the hazards have merely have been postponed - the pressures are now far higher.

For one thing, the world is more crowded. Fifty years ago, world population was below 3 billion. It now exceeds 7 billion.. And by 2050 it's projected to be between 8.5 and 10 billion, the growth being mainly in Africa and India. We must surely hope for a demographic transition in those countries whose populations are still rising fast, because the higher the post-2050 population becomes, the greater will be all pressures on resources [especially if the developing world, narrows its gap with the developed world in its per capita consumption].

Humans already appropriate around 40 percent of the world's biomass, and that fraction is growing. The resultant 'ecological shock' could irreversibly impoverish our biosphere. Extinction rates are rising -- we're destroying the book of life before we've read it. 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. But for many environmentalists these 'instrumental'

- and anthropocentric - arguments aren't the only compelling ones. For them there are further ethical issues: preserving the richness of our biosphere has value in its own right, over and above what it means to us humans.

Pressures on food supplies, and on the entire biosphere will be aggravated by climate change. And climate change exemplifies the tension between the science, the public and the politicians. One thing isn't controversial. The atmospheric CO₂ concentration is rising - and this is mainly due to the burning of fossil fuels. Straightforward physics tells us that this build-up will induce a long-term warming trend, superimposed on all the other complicated effects that make climate fluctuate. So far so good.

But what's less well understood is how big the effect is. Doubling of CO₂ in itself just causes 1.2 degrees warming. But the effect can be amplified by associated changes in water vapour and clouds. We don't know how important these feedback processes are. The recent fifth report from the IPCC presents a spread of projections. But some things are clear. In particular, if annual CO₂ emissions continue to rise soon we risk triggering drastic climate change—leading to the devastating scenarios later in this century portrayed in the recent book by Naomi Oreskes and Erik Conway [and the initiation of irreversible melting of the Greenland and Antarctic ice, which would eventually raise sea level by many metres]

Many still hope that we can segue towards a low-carbon future without trauma and disaster. But politician won't gain much resonance by advocating a bare bones approach that entails unwelcome lifestyle changes -- especially if the benefits are far away and decades into the future. There are however three politically realistic measures that should be pushed. First, all countries could promote measures that actually save money - better energy-efficiency, insulating buildings better, and so forth. Second, efforts could focus on reduction of pollutants, methane and black carbon. These are minor contributors to global warming. But their reduction would (unlike that of CO₂) have more manifest local side-benefits - especially in Asia. And third, there should be a step change in R and D into clean energy - who shouldn't it be on a scale comparable to medical research?.

The climate debate has been marred by too much blurring between the science, the politics and the commercial interests. Those who don't like the implications of the IPCC projections have rubbished the science rather than calling for better science. But even if the science were clear-cut, there is wide scope for debate on the policy response. Those who apply a standard discount rate (as, for instance, Bjorn Lomborg's Copenhagen Consensus recommendations do) are in effect writing off what happens beyond 2050. There is indeed little risk of catastrophe within that time-horizon, so unsurprisingly they downplay the priority of addressing climate change. But if you apply a lower discount rate -and in effect, don't discriminate on grounds of date of birth, and care about those who'll live into the 22nd century and beyond, -- then you may deem it worth making an investment now, to protect those future

generations against the worst-case scenario – and to prevent triggering really long-term changes like the melting of Greenland’s ice. Marty Weitzman here has been the most distinguished analyst of these issues..

So what will actually happen on the climate front? My pessimistic guess is that political efforts to decarbonise energy production won’t gain traction,, and that the CO2 concentration in the atmosphere will rise at an accelerating rate throughout the next 20 years. But by then, we’ll know with far more confidence -- perhaps from advanced computer modelling, but also from how much global temperatures have actually risen by then – just how strongly the feedback from water vapour and clouds amplifies the effect of CO2 itself in creating a 'greenhouse effect'. If the effect is strong, and the world’s climate consequently seems on a trajectory into dangerous territory, there may then be a pressure for 'panic measures'. These would have to involve a 'plan B' -- being fatalistic about continuing dependence on fossil fuels, but combating its effects by some form of geoengineering.

The ‘greenhouse warming’ could be counteracted by (for instance) putting reflecting aerosols in the upper atmosphere, or even vast sunshades in space. It seems feasible to throw enough material into the stratosphere to change the world’s climate -- indeed what is scary is that this might be within the resources of a single nation, or perhaps even a single corporation.. The political problems of such geoengineering may be overwhelming. There could be unintended side- effects. Moreover, the warming would return with a vengeance if the countermeasures were ever discontinued; and other consequences of rising CO2 (especially the deleterious effects of ocean acidification) would be unchecked.

Geoengineering would be an utter political nightmare: not all nations would want to adjust the thermostat the same way. Very elaborate climatic modelling would be needed in order to calculate the regional impacts of any artificial intervention. (It would be a bonanza for lawyers if an individual or a nation could be blamed for bad weather!). Dan Schrag , who’ll be commenting later, is an expert on this topic. But as a non-expert I’d think it prudent to explore geoengineering techniques enough to clarify which options make sense, and perhaps damp down undue optimism about a technical 'quick fix' of our climate.

So we’re deep into what Paul Creutzen dubbed the ‘anthropocene’. We’re under long-term threat from anthropogenic global changes to climate and biodiversity – due to rising population, all more demanding of food, energy and other resources. All these issues are widely discussed. What’s depressing is the inaction – for politicians the immediate trumps the long-term; the parochial trumps the global. We need to ask whether nations need to give up more sovereignty to new organizations along the lines of IAEA, WHO, etc.

THREATS FROM NOVEL TECHNOLOGY

But for the rest of this talk I’ll address a different topic –our vulnerability to powerful

technologies – those we depend on today, and those that still seem futuristic, even science fiction. Unlike climate and environment these are still under-discussed.

Those of us with cushioned lives in the developed world fret too much about minor hazards: improbable air crashes, carcinogens in food, low radiation doses, and so forth. But we are less secure than we think. We (and our political masters) don't worry enough about scenarios that have thankfully not yet happened – events that could arise as unexpectedly as the 2008 financial crisis, but which could cause world-wide disruption, and deal shattering blows to our society.

We live in an interconnected world increasingly dependent on elaborate networks: electric-power grids, air traffic control, international finance, just-in-time delivery, globally-dispersed manufacturing, and so forth. Unless these globalised networks are highly resilient, their manifest benefits could be outweighed by catastrophic (albeit rare) breakdowns -- real-world analogues of what happened in 2008 to the financial system. Our cities would be paralysed without electricity. Supermarket shelves would be empty within days if supply chains were disrupted. Air travel can spread a pandemic worldwide within days. And social media can spread panic and rumour, and psychic and economic contagion, literally at the speed of light.

The issues impel us to plan internationally. (For instance, whether or not a pandemic gets global grip may hinge on how quickly a Vietnamese poultry farmer can report any strange sickness.). And, by the way, the risk that pandemics cause societal breakdown is far higher 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 breakdown as soon as hospitals overflowed and health services were overwhelmed– which would occur when the fatality rate was still a fraction of one percent. But the human cost would be worst in the shambolic but burgeoning megacities of the developing world

Advances in microbiology offer better prospects of containing such disasters. But the same research has downsides too. For instance, in 2012 researchers at Wisconsin, and also at Erasmus University in Holland, showed that it was surprisingly easy to make an influenza virus both virulent and transmissible. When they published they were pressured to redact some details. And the Wisconsin group has been experimenting on H1N1, the virus that let do the catastrophic 1918 epidemic. Last month the US government decided to cease funding and impose a moratorium on so-called 'gain of function' experiments. The concern here was partly that it would be aiding terrorists, but partly also that if such experiments weren't conducted everywhere to the very highest safety and containment standards, there would be a risk of bioerror.

It is hard to make a clandestine H-bomb. In contrast, millions will one day have the capability to misuse biotech, just as they can misuse cybertech today. In the 1970s, in the early days of recombinant DNA research, a group of biologists led by Paul Berg formulated the 'Asilomar Declaration', advocating a moratorium on certain

types of experiments, and setting up guidelines. In retrospect, this move was perhaps over-cautious, but it seemed an encouraging precedent. However, the research community is far larger, far more broadly international, and far more influenced by commercial pressures. Whatever regulations are imposed, on prudential or ethical grounds, they could never be enforced worldwide – any more than the drug laws can. Whatever can be done will be done by someone, somewhere.

In consequence, maybe the most intractable challenges to all governments will stem from the rising empowerment of tech-savvy groups (or even individuals), by bio or cyber technology that becomes potentially ever more devastating – to the extent that even one episode could be too many. This will aggravate the tension between freedom, privacy and security.

The consequences of releasing dangerous pathogens are so incalculable that bioterror isn't likely to be deployed by extremist groups with well-defined political aims. But such concerns would not give pause to an eco-fanatic, empowered by the bio-hacking expertise that may soon be routine, who believes that 'Gaia' is being threatened by the presence of a few billion too many humans. That's my worst nightmare (Most devastating would be a potentially fatal virus that was readily transmissible and had a long latency period).

The global village will have its village idiots and they'll have global range.

LOOKING BEYOND 2050

These concerns are relatively near-term.. Trends beyond 2050 should make us even more anxious. I'll venture a word about these – but a tentative word, because scientists have a rotten record as forecasters. Ernest Rutherford, the greatest nuclear physicist of his time, said in the 1930s that nuclear energy was 'moonshine'. One of my predecessors as Astronomer Royal said, as late as the 1950s, that space travel was 'utter bilge'. My own crystal ball is very cloudy.

In the latter part of the 21st century the world will be warmer and more crowded – that's one of the few confident predictions.. But we can't predict how our lives might then have been changed by novel technologies. After all, the speedy societal transformation brought about by the smartphone, the internet and their ancillaries would have seemed magic even 20 years ago. So, looking several decades ahead we must keep our minds open, or at least ajar, to prospects that may now seem science fiction.

The physicist Freeman Dyson foresees a time when children will be able to design and create new organisms just as routinely as his generation played with chemistry sets. I'd guess that this is comfortably beyond the 'SF fringe', but were even part of this scenario to come about, our ecology (and even our species) surely would not long survive unscathed.

But what about another fast-advancing technology: robotics and machine intelligence? Even back in the 1990s IBM's 'Deep Blue' beat Kasparov, the world chess champion. More recently 'Watson' won a TV gameshow. Maybe a new-generation 'hyper computer' could achieve oracular powers that offered its controller dominance of international finance and strategy.

Advances in software and sensors have been slower than in number-crunching capacity. Robots still can't match the facility of a child in recognising and moving the pieces on a real chessboard. They can't tie your shoelaces or cut your toenails. But machine learning and sensor technology are advancing apace. If robots could observe and interpret their environment as adeptly as we do they would truly be perceived as intelligent beings, to which (or to whom) we can relate, at least in some respects, as we to other people. And their greater processing speed may give them an advantage over us.

But will robots remain docile rather than 'going rogue'? And what if a hyper-computer developed a mind of its own. If it could infiltrate the internet - and the internet of things - it could manipulate the rest of the world. It may have goals utterly orthogonal to human wishes - or even treat humans as an encumbrance.

Indeed, as early as the 1960s the British mathematician I J Good pointed out that a super-intelligent robot (were it sufficiently versatile) could be the last invention that humans need ever make. Once machines have surpassed human capabilities, they could themselves design and assemble a new generation of even more powerful ones.

Ray Kurzweil, now working at Google, is the leading evangelist for this so-called 'singularity'. He thinks that humans could transcend biology by merging with computers, maybe losing their individuality and evolving into a common consciousness. In old-style spiritualist parlance, they would 'go over to the other side'. But he's worried that it may not happen in his lifetime. So he wants his body frozen until this nirvana is reached. I was once interviewed by a group of 'cryonic' enthusiasts -- in California (where else!)-- called the 'society for the abolition of involuntary death'. They will freeze your body, so that when immortality's on offer you can be resurrected. I said I'd rather end my days in an English churchyard than a Californian refrigerator.. They derided me as a 'deathist'. (I was surprised to find that three Oxford professors were Cryonic enthusiasts. Two had paid the full whack; a third has taken the cut-price option of just having his head frozen).

In regard to all these speculations, we don't know where the boundary lies between what may happen, and what will remain science fiction -- just as we don't know whether to take seriously Freeman Dyson's vision of bio-hacking by children. There are widely divergent views. Some experts, for instance Stuart Russell at Berkeley, and Demis Hassabis of Deep Mind think that the AI field, like synthetic biotech, already needs guidelines for 'responsible innovation' But others, like Rodney Brooks,

think these concerns are 'misguided', and too far from realization to be worth worrying about. And the whole concept is philosophically contentious. – John Searle has an article in a recent NYRB dismissing the entire concept that a machine could have a mind of its own.

Be that as it may, it's likely that before 2100, our society and its economy will be transformed by autonomous robots, even though these may be 'idiot savants' rather than displaying full human capabilities.

[Books like "The Second Machine Age" have addressed the economic and social disruption that will ensue when robots replace not just factory workers, but white-collar workers as well (even lawyers are under threat!).]

A short digression:

One context where robots surely have a future is in space. In the second part of this century the whole solar system will be explored by flotillas of miniaturized robots. And, on a larger scale, robotic fabricators may build vast lightweight structures floating in space (solar energy collectors, for instance), perhaps mining raw materials from asteroids.

These robotic advances will erode the practical case for human spaceflight. Nonetheless, I hope people will follow the robots, though it will be as risk-seeking adventurers rather than for practical goals. The most promising developments are spearheaded by private companies. For instance SpaceX, led by Elon Musk, who also makes Tesla electric cars, has launched unmanned payloads and docked with the Space Station. He hopes soon to offer orbital flights to paying customers. Wealthy adventurers are already signing up for a week-long trip round the far side of the Moon – voyaging further from Earth than anyone has been before (but avoiding the greater challenge of a Moon landing and blast-off). I'm told they've sold a ticket for the second flight but not for the first flight. We should surely cheer on these private enterprise efforts in space – they can tolerate higher risks than a western government could impose on publicly-funded civilians, and thereby cut costs.

By 2100, groups of pioneers may have established 'bases' independent from the Earth – on Mars, or maybe on asteroids. Whatever ethical constraints we impose here on the ground, we should surely wish these adventurers good luck in using all the resources of genetic and cyborg technology to adapt themselves and their progeny to alien environments. This might be the first step towards divergence into a new species: the beginning of the post-human era. And it would also ensure that advanced life would survive, even if the worst conceivable catastrophe befell our planet.

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. It's a dangerous delusion to think that space offers an escape from Earth's problems.

And here on Earth we may indeed have a bumpy ride through this century. The scenarios I've described – environmental degradation, extreme climate change, or

unintended consequences of advanced technology -- could trigger serious, even catastrophic, setbacks to our civilization. But they wouldn't wipe us all out. They're extreme, but strictly speaking not 'existential'.

TRULY EXISTENTIAL RISKS?

Are there conceivable events that could snuff out all life? Promethian concerns of this kind were raised by scientists working on the atomic bomb project during the Second World War. Could we be absolutely sure that a nuclear explosion wouldn't ignite all the world's atmosphere or oceans? Before the Trinity bomb test in New Mexico, Hans Bethe and two colleagues addressed this issue; they convinced themselves that there was a large safety factor. And luckily they were right. We now know for certain that a single nuclear weapons, devastating though it is, can't trigger a nuclear chain reaction that would utterly destroy the Earth or its atmosphere.

But what about even more extreme experiments? Physicists were (in my view quite rightly) pressured to address the speculative 'existential risks' that could be triggered by powerful accelerators in Brookhaven and Geneva that generate unprecedented concentrations of energy. Could physicists unwittingly convert the entire Earth into particles called 'strangelets' - or, even worse, trigger a 'phase transition' that would shatter the fabric of space itself? Fortunately, reassurance could be offered: indeed I was one of those who pointed out that cosmic rays of much higher energies collide so frequently in the Galaxy, but haven't ripped space apart. And they have penetrated white dwarf and neutron stars without triggering their conversion into 'strangelets'.

But physicists should surely be circumspect and precautionary about carrying out experiments that generate conditions with no precedent even in the cosmos - just as biologists should avoid release of potentially-devastating genetically-modified pathogens.

So how risk-averse should we be? Some would argue that odds of 10 million to one against an existential disaster would be good enough, because that is below the chance that, within the next year, an asteroid large enough to cause global devastation will hit the Earth. (This is like arguing that the extra carcinogenic effects of artificial radiation is acceptable if it doesn't so much as double the risk from natural radiation.) But to some, this limit may not seem stringent enough. If there were a threat to the entire Earth, the public might properly demand assurance that the probability is below one in a billion -- even one in a trillion -- before sanctioning such an experiment.

But can we meaningfully give such assurances? We may offer these odds against the Sun not rising tomorrow, or against a fair die giving 100 sixes in a row; that's because we're confident that we understand these things. But if our understanding is shaky - as it plainly is at the frontiers of physics -- we can't really assign a probability, nor confidently assert that something is stupendously unlikely. It's surely presumptuous to place extreme confidence in any theories about what happens

when atoms are smashed together with unprecedented energy. If a congressional committee asked: 'Are you really claiming that there's less than a one in a billion chance that you're wrong?' I'd feel uncomfortable saying yes.

But on the other hand, if a congressman went on to ask: "Could such an experiment disclose a transformative discovery that -- for instance -- provided a new source of energy for the world?" I'd again offer high odds against it. The issue is then the relative likelihood of these two unlikely events -- one hugely beneficial, the other catastrophic. Innovation is often hazardous, but if we don't take risks we may forgo disproportionate benefits. Undiluted application of the 'precautionary principle' has a manifest downside. There is 'the hidden cost of saying no'.

And, by the way, the priority that we should assign to avoiding truly existential disasters depends on an ethical question posed by (for instance) the philosopher Derek Parfit, which is this. Consider two scenarios: scenario A wipes out 90 percent of humanity; scenario B wipes out 100 percent. How much worse is B than A? Some would say 10 percent worse: the body count is 10 percent higher. But others would say B was incomparably worse, because human extinction forecloses the existence of billions, even trillions, of future people -- and indeed an open-ended post-human future.

Especially if you accept the latter viewpoint, you'll agree that existential catastrophes deserve more attention. That's why some of us in (the other) Cambridge -- both natural and social scientists -- have inaugurated a research programme (the Centre for the Study of Existential Risks) to address these 'existential' risks, as well as the wider class of extreme risks I've discussed. We need to deploy the best scientific expertise to assess which alleged risks are pure science fiction, and which could conceivably become real; to consider how to enhance resilience against the more credible ones; and to warn against technological developments that could run out of control. . . And there are similar efforts elsewhere: at Oxford in the UK here at MIT and in other places.

Moreover, we shouldn't be complacent that all such probabilities are miniscule. We've no grounds for assuming that human-induced threats worse than those on our current risk register are improbable: they are newly emergent, so we have a limited timebase for exposure to them and can't be sanguine that we would survive them for long -- nor about the ability of governments to cope if disaster strikes.. Indeed we have zero grounds for confidence that we can survive the worst that future technologies could bring in their wake.

Technology brings with it great hopes, but also great fears. We mustn't forget an important maxim: the unfamiliar is not the same as the improbable.

Another digression:

I'm often asked: is there a special perspective that astronomers can offer to science and

philosophy? Having worked among them for many years, I have to tell you that contemplation of vast expanses of space and time doesn't make astronomers serene and relaxed. They fret about everyday hassles as much as anyone. But they do have one special perspective -- an awareness of an immense future.

The stupendous timespans of the evolutionary past are now part of common culture (outside 'fundamentalist' circles, at any rate). But most people still somehow think we humans are the culmination of the evolutionary tree. That hardly seems credible to an astronomer. Our Sun formed 4.5 billion years ago, but it's got 6 billion more before the fuel runs out. And the expanding universe will continue -- perhaps for ever -- destined to become ever colder, ever emptier. To quote Woody Allen, eternity is very long, especially towards the end.

Posthuman evolution -- here on Earth and far beyond -- could be as prolonged as the Darwinian evolution that's led to us -- and even more wonderful. Any creatures witnessing the Sun's demise 6 billion years hence won't be human -- they'll be as different from us as we are from a bug. Indeed evolution will be even faster than in the past -- on a technological not a natural selection timescale.

Even in this 'concentrated' timeline -- extending billions of years into the future, as well as into the past -- this century may be a defining moment where humans could jeopardise life's immense potential. That's why the avoidance of complete extinction has special resonance for an astronomer.

OBLIGATIONS OF SCIENTISTS

Finally, a few thoughts of special relevance to my hosts in STS. Sheila Jasinoff and others have discussed the obligations of scientists when their investigations have potential social, economic and ethical impact that concern all citizens. These issues are starkly relevant to the theme of this talk. So I'd like, before closing, to offer some thoughts -- though with diffidence in front of this audience. It's important to keep 'clear water' between science and policy. Risk assessment should be separate from risk management. Scientists should present policy options based on a consensus of expert opinion; but if they engage in advocacy they should recognise that on the economic, social and ethical aspects of any policy they speak as citizens and not as experts -- and will have a variety of views.

I'd highlight some fine exemplars from the past: for instance, the atomic scientists who developed the first nuclear weapons during World War II. Fate had assigned them a pivotal role in history. Many of them --- men such as Jo Rotblat, Hans Bethe, Rudolf Peierls and John Simpson (all of who I was privileged to know in their later years) -- returned with relief to peacetime academic pursuits. But the ivory tower wasn't, for them, a sanctuary. They continued not just as academics but as engaged citizens --- promoting efforts to control the power they had helped unleash, through national academies, the Pugwash movement, and other bodies..

They were the alchemists of their time, possessors of secret specialized knowledge. The technologies I've discussed today have implications just as momentous as

nuclear weapons. But in contrast to the 'atomic scientists', those engaged with the new challenges span almost all the sciences, are broadly international – and work in the commercial as well as public sector.

But they all have a responsibility. You would be a poor parent if you didn't care what happened to your children in adulthood, even though you may have little control over them. Likewise, scientists shouldn't be indifferent to the fruits of their ideas – their creations. They should try to foster benign spin-offs – commercial or otherwise. They should resist, so far as they can, dubious or threatening applications of their work, and alert politicians when appropriate. We need to foster a culture of 'responsible innovation', especially in fields like biotech, advanced AI and geoengineering.

But, more than that, choices on how technology is applied – what to prioritise, and what to regulate -- require wide public debate, and such debate must be informed and leveraged by 'scientific citizens' – who will have a range of political perspectives. They can do this via campaigning groups, via blogging and journalism, or through political activity. There is a role for national academies too.

A special obligation lies on those in academia or self-employed entrepreneurs -- they have more freedom to engage in public debate than those employed in government service or in industry. (Academics have a special privilege to influence students. Polls show, unsurprisingly, that younger people who expect to survive most of the century, are more engaged and anxious about long-term and global issues – we should respond to their concerns.)

More should be done to assess, and then minimize, the extreme risks I've addressed. But though we live under their shadow, there seems no scientific impediment to achieving a sustainable and secure world, where all enjoy a lifestyle better than those in the 'west' do today. We can be technological optimists, even though the balance of effort in technology needs redirection – and to be guided by values that science itself can't provide. But the intractable politics and sociology -- the gap between potentialities and what actually happens -- engenders pessimism. Politicians look to their own voters – and the next election. Stockholders expect a pay-off in the short run. We downplay what's happening even now in far-away countries. And we discount too heavily the problems we'll leave for new generations. Without a broader perspective – without realizing that we're all on this crowded world together – governments won't properly prioritise projects that are long-term in a political perspectives, even if a mere instant in the history of our planet.

"Space-ship Earth" is hurtling through space. Its passengers are anxious and fractious. Their life-support system is vulnerable to disruption and break-downs. But there is too little planning too little horizon-scanning, too little awareness of long-term risks.

There needs to be a serious research programme, involving natural and social

scientists, to compile a more complete register of these 'extreme risks', and to enhance resilience against the more credible ones. The stakes are so high that those involved in this effort will have earned their keep even if they reduce the probability of a catastrophe by one in the sixth decimal place.

But we mustn't leap from denial to despair. So, having started with HG Wells, I give the final word to another scientific sage the great 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."