Vice-Chancellor, Ladies and Gentlemen,

It's a privilege to speak here today; and daunting to join the eminent roll-call of Romanes Lecturers. Of the scientists among them, the most formidably eloquent was the immunologist Peter Medawar. His lecture was entitled 'Science and Literature', and this is how he began:

> "I hope I shall not be thought ungracious if I say at the outset that nothing on earth would have induced me to attend the kind of lecture you may think I am going to give."

Unlike Medawar, I won't deride your judgement in coming today. Indeed I'm relieved that you weren't put off by my title, 'The Limits of Science'. It's unalluring and vague. But its ambiguity was deliberate. I want to address different kinds of 'limits'. First, I'll scan some horizons in my own field of astronomy. Then broadening into other fields, I'll ask if there are intrinsic limits to our scientific grasp – phenomena within the remit of science that nonetheless transcend human understanding. And then, thirdly, I'll speak, simply as an anxious member of the human race, about more practical concerns: the threats and opportunities science presents, and the limitations on how it's applied that are set by politics, prudence or ethics. Mindful that this is an audience with diverse interests, I shall offer variety rather than detail – and that you'll be charitable enough to regard it as a smorgasbord rather than a dog's breakfast.

But let's start with a flashback. One can't stand here in the Sheldonian without recalling Christopher Wren – not only an architect, but Oxford's Savilian Professor of Astronomy; and, with Boyle, Wilkins and Hooke, one of the founders of the Royal Society. Indeed Wren lectured at the Royal Society's very first meeting. From the 1660s onward, the Society's Fellows met regularly. They peered through newly-invented telescopes and microscopes; they dissected weird animals; they heard travellers' tales. They experimented with airpumps, explosions, and poisons. And some meetings were more gruesome. Samuel Pepys recorded one in his diary where he witnessed a blood transfusion from a sheep to a man – who, amazingly, survived. Pepys conversed with him after the operation and found him 'cracked a little in his head, though he speaks very reasonably and very well' (and noted that the victim was a Cambridge graduate!).

But, as well as being 'ingenious and curious', these men were immersed in the practical agenda of their era – improving navigation, exploring the New World, and rebuilding London after the Great Fire. They were inspired by Francis Bacon – they were, in his phrases, 'merchants of light', but committed also to 'the relief of man's estate'.

Our horizons have hugely expanded since the 17th century. No new continents remain to be discovered; our Earth no longer offers an open frontier, but seems constricted and crowded – a speck in the immensities of space. But today's scientists have the same motives and enthusiasms as Wren's contemporaries – the curiosity to probe nature's laws; the delight in ingenious devices. (Though I have to report that health and safety regulations make Royal Society meetings duller now than they were then.)

And the fruits of science have transformed our lives. We're ever more dependent on elaborate technology – but also more vulnerable to its failures and misuse. Many are anxious that genetics, brain science and artificial intelligence may 'run away' too fast to be handled wisely. That's why, as I'll argue later, today's scientists should also emulate their 17th century forbears in engaging with society and public affairs.

By the way, I'm using the word 'science' in a broad sense, to encompass technology and engineering – this is not just to save words, but because they're symbiotically linked. 'Problem solving' motivates us all – whether one is an astronomer probing the remote cosmos, or an engineer facing a down-to-earth design conundrum. The latter is at least as challenging – a point neatly made by an old cartoon showing two beavers looking up at a hydroelectric dam. One beaver says 'I didn't actually build it, but it's based on my idea'. The Swedish engineer who invented the zip fastener made a greater intellectual leap than most 'pure' academics achieve.

Scientists can't now be polymaths. Research is professionalised, arcane and technical. There's consequently a communication barrier – between different specialisms, and with the wider public. Darwin (who was of course a friend of George Romanes) was the last great scientist whose discoveries could be fully presented in accessible prose – indeed in fine literature. I believe nonetheless that the essence of today's science can be conveyed, without undue distortion, in a form sufficiently free of technicalities to be accessible to all.

Incidentally, scientists habitually bemoan the meagre public grasp of their subject. But I think they protest too much. On the contrary, it's surprising and gratifying that there's wide interest in topics as far from everyday concerns as dinosaurs, the beginning of the universe, or alien life. Some people, admittedly, can't distinguish a proton from a protein; but just as many are ignorant of their nation's history, and can't find Korea or Syria on a map – and that's equally sad.

It's surely a <u>cultural</u> deprivation not to appreciate the panorama offered by modern cosmology, DNA, and Darwinian evolution: the chain of emergent complexity leading from some still-mysterious beginning to atoms, stars, and planets – and, on our planet, to a biosphere containing creatures with brains able to ponder the wonder of it all. This common understanding should transcend all national differences – and all faiths too.

Science is a uniquely global culture. Its findings are 'objective': they can be evaluated by criteria that don't depend on how they were motivated and arrived at. This universality is specially compelling in astronomy. All humans, throughout history, have gazed up at the same 'vault of heaven', but interpreted it in diverse ways. And it's led them to ponder some still-open questions: Was there a beginning? Is space infinite? Does alien life exist?

I'd like to devote the next ten minutes to the crescendo of recent discoveries in astronomy. These are indeed remarkable. They are owed primarily to more powerful telescopes in all wavebands, better sensors for faint radiation, and space technology. Armchair theorists like myself deserve little credit.

We've learnt the extraordinary history of atoms – a history that links us to the stars more intimately than the ancients envisaged. We are ashes of long-dead stars – 'nuclear waste' from the fusion power that makes stars shine. Each of us contains carbon, oxygen and iron atoms forged from pristine hydrogen in thousands of ancient stars from all over the Milky Way.

And we've learnt, since the 1990s, something that makes the night sky far more interesting. Many stars – perhaps even most of them – are orbited by retinues of planets, just as the Earth, Mars and Jupiter circle around our own star, the Sun. These planetary systems display a surprising variety. Planets as big as Jupiter are orbiting so close to their star that their 'year' lasts only a few days. Some are on very eccentric orbits. One planet is orbiting a binary star; it would have two 'suns' in its sky.

These planets aren't yet actually seen: their sizes and orbits are inferred from very tiny effects on the motion and apparent brightness of their parent stars. But we'd really like to image them directly. This is harder. To realise just how hard, suppose an alien astronomer with a powerful telescope was viewing the Earth from (say) 30 light years away – the distance of a nearby star. Our planet would seem, in Carl Sagan's phrase, a 'pale blue dot', very close to a star (our Sun) that outshines it by many billions: a firefly next to a searchlight.

But by careful measurements, the aliens could infer quite a bit about us. They'd see a different shade of blue depending on whether the Pacific Ocean or the Eurasian land mass was facing them. They could infer that our Earth had continents and oceans, the length of the 'day', the seasons, and the climate.

In 20 years we'll have telescopes in space or on the ground that can make just such measurements of Earth-like planets orbiting Sun-like stars.

Do any of these planets harbour life? Could some even be inhabited by beings that we might recognise as intelligent? Even in the 17th century, John Wilkins and others speculated about a 'plurality of inhabited worlds' – and we're as baffled as they were. Indeed, how life began here on Earth is still a mystery.

It could have been inevitable in the conditions prevailing on the young Earth; or it could have been a fluke, like shuffling a deck of cards into perfect order. So we can't yet say *a priori* say whether alien life is likely or unlikely – nor decide where it's most promising to search for it.

Moreover, even if simple life were widespread, it would be a separate question whether it is likely to evolve into anything intelligent or complex – still less into the more restricted category we might recognize as such.

Perhaps 'ET' doesn't exist. Earth's intricate biosphere may be unique. That may disappoint the searchers. But it would have its upside: it would entitle us to be less cosmically modest. Our planet, though tiny, could then be uniquely important – perhaps even a 'seed' from which life could spread through the entire Galaxy.

On the other hand, one day some astronomer might discover a signal that's clearly artificial – or even some artifact. I wouldn't hold my breath for success, but it's surely well worth gambling modest resources on such searches. Even a very boring signal – a list of prime numbers, for instance, or the digits of pi - would carry the momentous message that concepts of logic and physics (if not consciousness) aren't limited to the hardware in human skulls.

We may learn this century whether biological evolution is unique to the 'pale blue dot' in the cosmos that is our home, or whether Darwin's writ runs through a wider universe that teems with life – even with intelligence. But even in the latter case, such intelligence could be qualitatively different from our own – assemblages of superintelligent 'social insects', or computers. And, of course, there may be a lot more out there than we could ever detect. Absence of evidence wouldn't be evidence of absence.

Ludwig Wittgenstein famously said "If a lion could speak, we couldn't understand him". So, even if we detected aliens would the 'culture gap' be unbridgeable? Not necessarily. Any signal would of course take decades in transit – there will never be scope for snappy repartee. But if they had developed advanced technology, they would share with us an understanding not just of mathematics but of physics, and astronomy. They'd gaze out, if they had eyes, at the same cosmos – they'd trace their origins back, as we now can, to a still mysterious beginning nearly 14 billion years ago.

So what can we confidently say about how our universe has evolved? A range of interlocking evidence allows us to trace cosmic history back to an era when everything was squeezed hotter and denser than the centre of a star. Such inferences are as evidence-based as anything a geologist might tell you about the history of our Earth – we observe 'fossils' of those early eras, and can confidently infer how hot and dense things were even just a nano-second after the big bang, when every particle carried as much energy as can be generated by the huge Large Hadron Collider (LHC) in Geneva, and the entire observable universe was squeezed into dimensions no larger than the Solar System.

But, as always in science, each advance brings into focus some new questions that couldn't previously have even been posed – for instance "Why is the universe expanding the way it is?" and "how did it acquire its observed mix of particles and radiation?". The answers to these lie in the brief instants when our universe was hugely more compressed still, and conditions far hotter and denser than we can simulate in the lab. We consequently lose any firm foothold in experiment and get beyond any consensual understanding

One of my favourite magazine covers showed a sphere, with the caption: 'the universe when it was a trillionth of a trillionth of a trillionth of a second old – actual size'. According to a popular conjecture, the entire volume we can see with our telescopes 'inflated' from a hyper-dense blob no bigger than a tennis ball. How can we firm up such an idea?

The twin pillars of 20th century physics are Einstein's theory of gravity (general relativity) and quantum theory. But these haven't yet been meshed together or unified. In most contexts, this is no impediment because their domains of relevance do not overlap. Astronomers can ignore quantum fuzziness when calculating the motions of planets and stars. Conversely, quantum chemists can safely ignore gravitational attraction between individual atoms because this force is about 40 powers of ten feebler than the electrical forces between them.

But during the ultra-compressed earliest instants after the 'big bang', quantum fluctuations could, as it were, 'shake' the entire universe. To address the overwhelming question of what banged and why it banged therefore requires a synthesis or completion of these two great 20th century theories.

Einstein's theory is incomplete because it treats space and time as smooth and continuous. We know that no material can be chopped into arbitrarily small pieces: eventually, you get down to discrete atoms. Likewise, it's thought that space has a grainy and 'atomic' structure – but on a scale a trillion trillion times smaller than atoms. But what is the nature of this structure? According to the front-running idea – superstring theory – every apparent point in our 3-dimensional space, if hugely magnified, may actually embody an intricate structure: a tightly wound origami in six extra dimensions. So space would actually have 10 dimensions. We're unaware of the extra ones because they're wrapped up tightly. By analogy, a long hose-pipe may look like a line (with just one dimension) when viewed from a distance, but from closer up we realise that it extends in three dimensions.

A theory that unified cosmos and quantum, if achieved, would complete a unification programme that started with Newton, and continued through Faraday and Maxwell, and their successors. It might even realise the Pythagorean vision of reducing all nature's complexities to geometry. It would firm up our understanding if the ultra-early universe. It would also incidentally, elucidate the discovery (recognised by this year's Nobel Physics Prize) that a mysterious force is latent even in empty space, which pushes galaxies away from each other at an accelerating rate.

Because of this acceleration, incidentally, the galaxies we now see will eventually disappear over (as it were) a 'horizon' into a domain that's not observable even in principle. If you're on a ship, you expect the ocean to extend fr beyond your horizon. Likewise, the observable domain that astronomers call 'the universe' could be a tiny part of the aftermath of our big bang. That's an inference that most cosmologists would take seriously.

But there are strong, albeit controversial, grounds for conjecturing further. "Our" big bang may not be the only one. For instance other space-times could exist 'alongside' ours. Imagine ants crawling around on a large sheet of paper (their two-dimensional 'universe'). They would be unaware of a similar sheet that's parallel to it. Likewise, there could be another entire universe (with 3-dimensional space, like ours) less than a millimetre away from us, but we'd be oblivious to it if that millimetre were measured in a fourth spatial dimension, while we are imprisoned in just three.

So a further Copernican demotion may loom ahead – not only are we in just one planetary system among billions, in one galaxy among billions, but we perhaps live in the aftermath of one big bang among many. A 'health warning' is in order here. Although there's compelling evidence for a hot dense beginning, the so called 'multiverse' concept is one where we're still groping for the truth – where, in the fashion of ancient cartographers, we must still inscribe 'here be dragons'.

Support or refutation for these speculations must await firmer links between the theories of the very large (the cosmos) and the very small (the quantum). But such insights won't have credibility unless buttressed by experiments or observations. And this requires huge and expensive equipment – telescopes or particle accelerators. The LHC at CERN in Geneva is the world's most elaborate scientific instrument. It has generated enthusiastic razzmatazz; but some query such a large investment in a seemingly recondite science. I'd respond by noting that it's costing the UK about 2 percent of our overall budget for academic science. This doesn't seem a disproportionate allocation to a field so challenging and fundamental (and where the UK has a specially strong record and can aspire to more than its pro-rata share of the action). But what <u>is</u> distinctive about this branch of science is that its practitioners world-wide have chosen to pool resources, and make a 20-year commitment to a single shared facility. Telescopes are now international facilities too. These global collaborations to probe nature's fundamental mysteries – and push technology to its limits – are surely something our society can take pride in.

So let me now move on from my own rather atypical science, and venture some more general comments.

Odd though it may seem, some of the best-understood phenomena are far away in the cosmos. Back in the 17th century, Newton could describe the 'clockwork of the heavens'; eclipses could be both understood and predicted. (Indeed even in Babylonian times, regularities were discerned and some prediction was possible even in ignorance of the underlying mechanism.) But few other things are so predictable, even when we understand them. For instance, it's hard to forecast, even a day before, whether those who go to view an eclipse will encounter clouds or clear skies. And our grasp of some familiar matters that interest us all – diet and child care for instance – is still so meagre that 'expert' advice changes from year to year.

But that's because astronomy's far simpler than the biological and human sciences. The smallest insect, with its layer upon layer of intricate structure, is far more complex than a star, where intense heat and compression by gravity precludes complex chemistry.

Incidentally, more powerful computers had been as crucial to many sciences as advances in instrumentation. That's especially so in fields where we can't do real experiments but can only observe. In the 'virtual world' inside a computer astronomers can crash another planet into the Earth to see if that's how our Moon might have formed; meteorologists can simulate the atmosphere (though 'chaos theory' sets fundamental limits to how well we can ever predict the weather); brain scientists can simulate how neurons interact. Just as video games get more elaborate as their consoles get more powerful, so, as computer-power grows, these 'virtual' experiments become more realistic and useful.

Perhaps I can inject a bit of advice to any undergraduates here who are thinking of embarking on research -- I hope some are, and they're not all going into finance. You may worry that the easy things have been done, and you'll have to tackle problems that defeated your predecessors. But you don't have to be cleverer than them. You'll have access to instruments and computer power that they never had; you can explore realms that they could never envisage. Choose a subject where things are happening fast – where the experience of the older generation is at a heavy discount. And remember there's a range of research styles. Some topics can be tackled by one person working alone; at the other extreme, some demand quasi-industrial teamwork; but most involve collaboration and debate in a small research group. And some people aspire to write a pioneering paper opening up a subject: others gain more satisfaction from writing a definitive monograph tidying up and codifying a topic. You must pick a topic to suit your personality.

And another thing: in choosing your problem, don't head straight for the most important one. You should multiply the importance of the problem by the probability that you'll solve it, and maximise that product. So don't all swarm into the kind of theoretical physics I mentioned earlier – the unification of cosmos and quantum – even though it's plainly one of the highest intellectual summits we aspire to reach.

A unified theory is sometimes, incidentally, called the 'theory of everything'. But that phrase is hubristic and misleading. Such a theory is irrelevant to 99 percent of scientists. Problems in biology, and in environmental and human sciences, remain unsolved because it's hard to elucidate their complexities – not because we don't understand subatomic physics well enough.

The sciences are sometimes likened to different levels of a tall building – particle physics on the ground floor, then the rest of physics, then chemistry, and so forth: all the way up to psychology (and the economists in the penthouse). There is a corresponding hierarchy of complexity – atoms, molecules, cells, organisms, and so forth.

But the analogy with a building is poor. The 'higher level' sciences dealing with complex systems aren't imperiled by an insecure base, as a building is. Each science has its own distinct concepts and explanations.

For instance, mathematicians trying to understand why flows go turbulent, or why waves break, don't care that water is H2O. They treat the fluid as a continuum.

An albatross returns predictably to its nest after wandering ten thousand miles in the southern oceans. But this is not the same kind of prediction as astronomers make of celestial orbits and eclipses.

Everything, however, complicated – breaking waves, migrating birds, and tropical forests – is made of atoms and obeys the basic equations of quantum physics. But even if those equations could be solved for immense aggregates of atoms, they wouldn't offer the enlightenment that scientists seek. Reductionism is true in a sense. But it's seldom true in a <u>useful</u> sense. Each science has its autonomous concepts and laws. To take another example, the best explanation of what's happening on a computer screen is in terms of software; not a 'bottom up' description of the forces on all the electrons.

The path towards a consensual understanding is often winding – with many blind alleys being explored before reaching it. Sometimes a maverick is vindicated. We all enjoy seeing this happen – but such instances are rarer than is commonly supposed, and perhaps rarer than would be inferred from reading the popular press. And sometimes a prior consensus is overturned – though Thomas Kuhn's famous book on 'Scientific Revolutions' perhaps exaggerates how often this happens. The Copernican cosmology, overthrowing the concept of a geocentric cosmos, would qualify as a genuine revolution, as would quantum theory. But most advances transcend and generalise the concepts that went before, rather than contradicting them. For instance, Einstein didn't overthrow Newton. He transcended Newton, giving us a new perspective offering broader scope and deeper insights.

(By the way, if neutrinos really travelled faster than light, it might trigger a real revolution – which would be wonderful news. But extraordinary claims demand extraordinary evidence, which is why most experts are deeply sceptical of the recent much-hyped claim. It's properly the subject of a wide debate – a debate that's constructive because everyone plays by the same evidence-based rules. In contrast, I've never found it fruitful to debate astrologers nor creationists. Many crucial aspects of nature are still perplexing. But we should strive hard for a better understanding, and not let a craving for quick answers drive us towards the illusory comfort and reassurance that pseudo-sciences offer.)

Most scientists of my vintage would find (as I do) that the issues being debated in their student days have been settled; we're now addressing questions that couldn't then have been posed. And our successors will address questions that we can't yet even formulate: Donald Rumsfeld's famous 'unknown unknowns' (what a pity, incidentally, that Rumsfeld didn't stick to philosophy!).

But I want to pose the question: is science really an unending quest? Or will we, perhaps far down the line, encounter limits – hit the buffers?

This could happen for two reasons. Obviously, some topics get cleaned up and codified – atomic physics, for instance – and researchers then move on towards new horizons.

But maybe we should be open-minded about the obverse possibility – that we hit the buffers because our brains don't have enough conceptual grasp. Einstein averred that "The most incomprehensible thing about the universe is that it is comprehensible". He was right to be astonished. Our minds, evolved to cope with the life of our remote ancestors on the African savannah. It's amazing these minds can comprehend so much of the counterintuitive microworld of atoms, and phenomena billions of lightyears away.

Nonetheless – and here I'm sticking my neck out – maybe some aspects of reality are intrinsically beyond us, in that their comprehension would require some post-human intellect – just as quantum theory was beyond the first primates. In his provocative recent book 'The Beginning of Infinity', David Deutsch contests this view by claiming that any physical process is in principle computable. Maybe, but this isn't the same as being conceptually graspable

But rather than get enmeshed in philosophy, let's home in towards the here and now. I'm sometimes asked: 'Do astronomers, because of their avocation, think about life differently?' Well, from a life spent amongst them I can confirm that contemplation of huge expanses of space and time doesn't make them any more serene in everyday life. But I'd highlight one difference – we're more mindful than most of the immense future that lies ahead.

The stupendous timespans of the evolutionary past are now part of common culture – albeit still not in some creationist circles. But most educated people, though fully aware that our biosphere is the outcome of several billion years of Darwinian evolution, still somehow think that humans are the culmination of the evolutionary tree. That hardly seems credible to astronomers. They're mindful that our Sun formed 4.5 billion years ago, but also that it's got 6 billion years more before the fuel runs out. It then flares up, engulfing the inner planets. 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'.

Any creatures witnessing the Sun's demise won't be human – they'll be as different from us as we are from a bug. Post human evolution, here on Earth or far beyond, could be as prolonged as the Darwinian evolution that has led to us – and even more wonderful. Darwin himself realised that 'no living species will preserve its unaltered likeness for a distant futurity'.

Moreover, evolution could in future proceed far more rapidly than in the past, driven by technology rather than natural selection. Indeed, if one day there are communities living in space, they'd surely wish to use the resources of genetics to adapt their offspring to an alien environment – and we earthlings would surely wish them good luck, whatever ethical constraints we'd want to impose on such techniques here on Earth. The post-human era could begin within a few centuries. Whether the really long-range future lies with organic post-humans or with intelligent machines is a matter for debate.

We're all familiar with the pictures of our planet from space – Earth's delicate biosphere contrasting with the sterile moonscape where the astronauts left their footprints. We've had these images for 40 years; they're iconic for environmentalists. But let me offer a cosmic vignette.

Suppose hypothetical aliens had been watching the Earth for its entire history, what would they have seen? Over nearly all that immense time, 4.5 billion years, its surface would have altered very gradually. The continents drifted; the ice cover waxed and waned; successive species emerged, evolved and became extinct.

But in just a tiny sliver of the Earth's history – the last one millionth part, a few thousand years – the patterns of vegetation altered much faster than before.

This signalled the start of agriculture. The pace of change accelerated as human populations rose.

Then came even faster changes. Within fifty years – little more than one hundredth of a millionth of the Earth's age, the carbon dioxide in the atmosphere began to rise anomalously fast. The planet became an intense emitter of radio waves (the total output from all TV, cellphone, and radar transmissions.) And something else unprecedented happened: small projectiles launched from the planet's surface escaped the biosphere completely. Some were propelled into orbits around the Earth; some journeyed to the Moon and planets.

If they understood astrophysics, the aliens could confidently predict that the biosphere would face doom in a few billion years when the Sun flares up and dies. But could they have predicted this unprecedented runaway fever – less than half way through the Earth's life? And what might they see if they watched for another century? Will this spasm be followed by silence? Will sustainable stability ensue? And will more projectiles leave the Earth to establish oases of life elsewhere?

Earth's lifespan is more than 100 million centuries. But this century is special. It's the first in our planet's history where one species – ours – has Earth's future in its hands, and could jeopardise life's immense potential. We've entered a geological era called the anthropocene. And this leads to my final theme today: hopes and fears for the coming decades – the scope and limits of science-driven changes.

The anthropocene began with the advent of thermonuclear weapons. The threat of <u>global</u> nuclear annihilation involving tens of thousands of bombs has been in abeyance since the Cold War ended. But later in the century, a global political realignment leading to a standoff between new superpowers, that could be handled less well or less luckily than the Cuba crisis was. And in the meantime there is more risk than ever that smaller nuclear arsenals are used in a regional context or even by terrorists.

But devastation could arise insidiously rather than suddenly, through unsustainable pressure on energy supplies, food, water and other natural resources. Indeed these pressures are the prime 'threats without enemies' that confront us. The higher the population becomes, the more serious they will become – especially if the developing world, where most of the growth will be, narrows its gap with the developed world in its per capita consumption.

By 2050 the world's population is projected to reach 9 billion. It's now 7 billion – there's actually a year or two uncertainty in when this milestone is passed, but the UN is officially marking it this week. Whether the rising trend continues <u>beyond</u> 2050 will depend on what people now in their teens and 20s decide about the number and spacing of their children. Hundreds of millions of women are denied such a choice. Enhancing the life-chances of the world's poorest people – by providing clean water, primary education and other basics – should be a humanitarian imperative, and a readily achievable one. But it seems also a precondition for achieving, especially in Africa, the demographic transition that's already occurred elsewhere.

Humankind's collective 'footprint' is growing fast; we now appropriate around 40 percent of the world's biomass. This 'ecological shock' could irreversibly degrade our environment – and this trend is aggravated by climate change. Biodiversity is often proclaimed as a crucial component of human wellbeing and economic growth. It manifestly is: 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, preserving the richness of our biosphere has value in its own right, over and above what it means to us humans.

Despite these concerns, modern engineering and agriculture could provide food and energy for 9 billion by mid-century and avert irreversible degradation. And other advances, especially in healthcare and information technology, offer grounds for being techno-optimists. But we can't be so optimistic about nations achieving the cooperation that's needed if these benefits are to be shared by the developing world.

Moreover there's a downside: the same technologies that promise so much are opening up new vulnerabilities. For instance, global society depends on elaborate networks – electricity grids, air traffic control, international finance, just-in-time delivery and so forth. Unless these are highly resilient, their manifest benefits could be outweighed by catastrophic (albeit rare) breakdowns cascading through the system. And the threat is terror as well as error; concern about cyber-attack, by criminals or by hostile nations, is rising sharply. Synthetic biology, likewise, offers huge potential for medicine and agriculture -- but it could facilitate bioterror.

And we're kidding ourselves if we think that those with technical expertise will all be balanced and rational: expertise can be allied with fanaticism – not just the traditional fundamentalism that we're so mindful of today, but that exemplified by some 'new age' cults, extreme eco-freaks, violent animal rights campaigners and the like. And there will be individual weirdoes with the mindset of those who now unleash computer viruses. The global village will have its village idiots – and their idiocies can have global range. The huge empowerment of individuals or small groups by fast-developing technologies presents novel hazards.

Incidentally, there's a mismatch between public perception of very different risks and their actual seriousness. We fret unduly about carcinogens in food and low level radiation. But we are in denial about 'high-consequence' events, natural or man-made, that may be improbable, but where even one occurrence could be too many.

The technologies I've mentioned are those that already exist. But what about transformational new ones that could emerge later this century? Scientific forecasters have a dismal record. One of my predecessors as Astronomer Royal said space travel was "utter bilge". Few in the mid-20th century envisaged the transformative impact of the silicon chip or the double helix. So, looking 50 years ahead we must keep our minds open, or at least ajar, to what may now seem science fiction.

But one thing can be confidently predicted: the gulf between what science enables us to do, and what applications it's prudent or ethical (or economic) actually to pursue will get even wider than it already is.

For example, human nature and human character have changed little for millennia. Before long, however, new cognition-enhancing drugs, genetics, and 'cyborg' techniques may alter human beings themselves. That's something qualitatively new – and disquieting because it could portend more fundamental forms of inequality if these options were open only to a privileged few.

And we are living longer. Ongoing research into the genetics of ageing may explain why – indeed, a real 'wild card' in population projections is that future generations could achieve a really substantial enhancement in lifespan. This is still speculation – mainstream researchers are cautious about the prospect of improvements that are more than incremental. (And of course whether a longer lifespan is indeed an 'improvement' depends on whether it is the years of full activity or those of senile decrepitude that are prolonged.) But such caution hasn't stopped some Americans, worried that they'll die before this nirvana is reached, from bequeathing their bodies to be 'frozen', hoping that some future generations will resurrect them, or download their brains into a computer. For my part, I'd rather end my days in an English churchyard than a Californian refrigerator.

And what about robotics? Even back in the 1990s IBM's 'Deep Blue' beat Kasparov, the world chess champion. But robots can't yet recognise and move the pieces on a real chessboard as adeptly as a child can. Later this century, however, their more advanced successors may relate to their surroundings (and to people) as adeptly as we do. Moral questions then arise. We accept an obligation to ensure that other human beings, and indeed some animal species, can fulfil their 'natural potential. So what's our obligation towards sophisticated robots? Should we feel guilty about exploiting them? Should we fret if they are underemployed, frustrated, or bored? This may seem fanciful, but I mention it here because it's an issue that George Romanes himself would probably have addressed seriously: in his book "Animal Intelligence' he speculated even about the feelings of insects.

These are problems for the conjectural future. But already, possible applications of science confront us with hard choices. To take a few at random: Who should access the 'readout' of our personal genetic code? How will lengthening life-spans affect society? Should we build nuclear power stations – or wind farms – to keep the lights on? Should we plant GM crops? Should the law allow 'designer babies'? How can we best help Africa towards a more prosperous future?

None of these issues is purely scientific: they all involve ethics, economics and social policies as well. So a second reason why scientific literacy is important for everyone – apart from the cultural argument I mentioned earlier – is to ensure that public discussion of such issues can be broad, and that it can rise beyond tabloid slogans. In domains beyond their special expertise, scientists have no enhanced authority and will have a wide range of political and social perspectives But as "scientific citizens" they have a special obligation to engage – for instance by involvement with NGOs or campaigning groups, via blogging and journalism, or through political activity. You would be a poor parent if you didn't care about what happened to your children in adulthood, even though you may have little control over it. Likewise, scientists, whatever their expertise, shouldn't be indifferent to the fruits of their ideas. 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 the public and politicians to perceived dangers. But they shouldn't be bashful in proclaiming that despite the challenges there seems no scientific impediment to achieving a sustainable world, where all enjoy a lifestyle better than we in the west do today. Above all, they should urge greater priority for long-term global issues on the political agenda, where the urgent usually trumps the important.

I'll conclude with a personal perspective on this theme, triggered when I visit the greatest building near where I live – Ely Cathedral. Ely Cathedral overwhelms us today. But think of its impact 900 years ago – think of the vast enterprise its construction entailed. Most of its builders had never travelled more than 50 miles – the Fens were their world. Even the most educated knew of essentially nothing beyond Europe. They thought the world was a few thousand years old – and that it might not last another thousand.

But despite these constricted horizons in both time and space; despite the deprivation and harshness of their lives; despite their primitive technology and meagre resources; they built this cathedral – pushing the boundaries of what was possible. Those who conceived it knew they wouldn't live to see it finished. Their legacy still elevates our spirits, nearly a millennium later.

What a contrast to so much of our discourse today! Unlike our forebears, we know a great deal about our world – and indeed about what lies beyond.

Technologies that our ancestors couldn't have conceived enrich our lives and our understanding. 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.

But all too often the focus is parochial and short term. We downplay what's happening even now in impoverished far-away countries. And we discount too heavily the problems we'll leave for our grandchildren. In today's runaway world, we can't aspire to leave a monument lasting a thousand years, but it would surely be shameful if we denied future generations a fair inheritance.

To survive this century, we'll need the idealistic and effective efforts of natural scientists, environmentalists, social scientists and humanists. They must be guided by the best evidence, but inspired by values from beyond the limits of science.

I started by quoting Medawar; I'd like to end with him too. This quote is from his BBC Reith Lectures in 1959, but the message is more urgent today:

> "The bells which toll for mankind are – most of them, anyway – like the bells on Alpine cattle; they are attached to our own necks, and it must be <u>our</u> fault if they do not make a cheerful and harmonious sound."

Thank you,