Vice-Chancellor, Lillian Martin, ladies and gentlemen

It's an honour – but a daunting one – to be speaking here today. I'm privileged to have known James Martin, and, as Ian Goldin noted, to have been involved in the School since its inception.

The title I gave for this lecture was "Will technology's dark side forestall humanity's bright future?"

When a title ends with a question mark the answer is usually no. I'm not optimistic enough to give straight 'no'. My answer will be 'probably not' – I think that's what Jim would have said too.

Jim was a software pioneer and visionary – a hyper-geek, but also a humanist.

His 2006 book 'The Meaning of the 21st Century' covers a vast canvas – and reflects the mix of optimism and anxiety that I'll try to reflect today. He was an evangelist for new technology – but anxious that it's advancing so fast that we may not properly cope with it. And he was distressed by the immense gulf between the way the world could be and the way it actually is – especially by the plight of destitute nations.

Why is the 21st century special? Our planet has existed for 45 million centuries, but this the first when one species – ours – can determine the biosphere's fate. New technologies are transforming our lives and society – they're even changing human beings themselves. And they're inducing novel vulnerabilities. Ecosystems are being stressed because there are more of us (world population is higher) and we're all more demanding of resources. We're deep into what some call the Anthropocene.

And we've had one lucky escape already. At any time in the Cold War era, the superpowers could have stumbled towards nuclear Armageddon through muddle and miscalculation. Robert McNamara, US defence secretary at the time of the Cuba crisis, said after he retired 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.

After World War II some ex-Los Alamos physicists started a journal called the Bulletin of Atomic Scientists. The 'logo' on its cover is a clock, the closeness of whose hands to midnight indicates the Editorial Board's judgement on how precarious the world situation is.

When the Cold War ended, the clock was put back to 17 minutes to midnight. There was thereafter far less chance of tens of thousands of bombs devastating our civilisation.

But the clock's been creeping forward. That's partly because there's a rising risk of some nuclear weapon going off in a localised conflict, and because arms control seems to have stalled. And because of threats stemming from humanity's ever-heavier collective 'footprint'. Taking all these into account, the clock was last week moved forward to just three minutes to midnight – the closest for 40 years.

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 first, some comments on more insidious threats.

Jim Martin worried about population growth. This topic seems currently under-discussed. That is maybe because doom-laden forecasts in the 1970s, by the Club of Rome, Paul Erlich and others, have proved off the mark, and because it's deemed by some a taboo subject -- tainted by association with eugenics in the 1920s and 30s, with Indian policies under Indira Gandhi, and more recently with China's hard-line one-child policy.

Fifty years ago, world population was about 3 billion. It now exceeds 7 billion. But the growth is slowing. Indeed the number of births per year, worldwide, peaked a few years ago and is going down.

Nonetheless it's forecast to rise to around 9 billion by 2050. That's partly because most people in the developing world are young. They are yet to have children, and they will live longer. The age histogram in the developing world will become more like that in Europe.

But it's also because the demographic transition hasn't reached parts of India and Sub-Saharan Africa.

To quote some numbers: a hundred years ago, Ethiopia's population was 5 million. It is now about 80 million; it will almost double by 2050; and Sudan and Uganda will more than double by mid-century.

And it's naïve to specify the world's 'carrying capacity' or an 'optimum population'. That's because we can't confidently conceive what people's lifestyles, diet, travel patterns and energy needs will be. We can, though, predict continuing urbanization – 70 percent of people in cities by 2050. Even by 2030 Lagos, San Paulo and Delhi will have populations above 30 million. To prevent megacities becoming turbulent dystopias will surely be a major challenge. And indeed, as Mo Ibrahim (a member of the School's advisory board) insistently emphasises – improved governance is crucial.

The world couldn't sustain even its present population if everyone lived like present-day Americans. But my layman's impression from reading the work of experts, particularly those here in Oxford, is that the prospect of 9 billion people by 2050 isn't in itself frightening. Improved agriculture – low-till, water-conserving, and perhaps involving GM crops – together with better engineering to reduce waste, improve irrigation, and so forth, could sustainably feed that by mid-century. And advances in healthcare, IT and clean energy offer grounds for being techno-optimists.

Trends beyond 2050 will depend on what people now in their teens and 20s decide about the number and spacing of their children. Enhanced education and empowerment of women -- surely a benign priority in itself -- could reduce fertility rates where they're now highest. And that's surely welcome, because, whatever the actual limits are, the higher world 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.

If our collective activities push too hard against so-called 'planetary boundaries', the resultant 'ecological shock' could irreversibly impoverish our biosphere. Extinction rates are rising – to quote Bob May, 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, preserving the richness of our biosphere has value in its own right, over and above what it means to us humans To quote the great ecologist E O Wilson 'mass extinction is the sin that future generations will least forgive us for'.

Pressures on food and water, and consequent migration and international disputes are all on the Martin School agenda. And these will be aggravated by climate change.

Nick Stern addressed this last month from this podium, with far more authority than I can profess – and it's certainly not underdiscussed.

One thing isn't controversial. The concentration of CO2 in the air is rising. This rise is mainly due to the burning of fossil fuels. And it causes warming, superimposed on all the other trends that make climate vary.

But what's less well understood is how much the warming due to CO2 itself is amplified by associated changes in water vapour and clouds. The fifth IPCC report presents a spread of projections reflecting this uncertainty.

There are two messages:

1. **There will be some severe regional disruptions to weather patterns within the next 20-30 years, But**
2. **Under 'business as usual' scenarios we can't rule out, later in the century, further positive feedbacks triggering a warming of more than 4 degrees., triggering long-term trends like the melting of Greenland's icecap.**

But even among those who accept both these statements, there are divergent views on policy. These stem less from differences about the science than from differences in economics and ethics -- in particular, in how much obligation we should feel towards

future generations.

Economists who apply a standard discount rate (as, for instance, Bjorn Lomberg's Copenhagen Consensus does) are in effect writing off what happens beyond 2050 -- so unsurprisingly they downplay the priority of addressing climate change in comparison with shorter-term efforts to help the world's poor.

But if you care about those who'll live into the 22st century and beyond, then, as Stern, Weizman and others argue, you may deem it worth paying an insurance premium now, to protect those future generations against the worst-case scenarios.

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

(As a parenthesis, I'd note that there's one policy context when an essentially zero discount rate is applied – radioactive waste disposal, where the depositories are required to prevent leakage for 10000 years – somewhat ironic when we can't plan the rest of energy policy even 30 years ahead)

[Consider this analogy… Suppose astronomers had tracked an asteroid, and calculated that it would hit the Earth in 2080, 65 years from now – not with certainty, but with (say) 10 percent probability. Would we relax, saying that it's a problem that can be set on one side for 50 years – people will then be richer, and it may turn out then that it's going to miss the Earth anyway? I don't think we would. There would surely be a consensus that we should start straight away and do our damnedest to find ways to deflect it, or mitigate its effects.]

What will actually happen on the climate front? Even if the Paris conference is deemed a success, CO2 concentration will rise at an accelerating rate throughout the next 20 years. But by then, we'll know with far more confidence -- from a longer timebase of data, and from better modelling -- just how strong the feedback from water vapour and clouds actually is. If the so-called 'climate sensitivity' is low, we'll relax. But if it's large, and climate consequently seems on an irreversible 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 either a massive investment in carbon capture and storage, or else by geoengineering.

It seems feasible to throw enough material into the stratosphere to cool the world's climate -- indeed what is scary is that this might be within the resources of a single nation, or even a single corporation. 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 a political nightmare: not all nations would want to adjust

thethermostat the same way. **Very elaborate climatic modelling would be needed in order to calculate the regional impacts of an artificial intervention. (The only beneficiaries would be lawyers. They'd have a bonanza if nations could litigate over bad weather!).**

Steve Rayner and colleagues here are at the forefront of these debates. I think it's 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.

Nonetheless there are grounds for optimism. Three measures that could mitigate climate change seem politically realistic.

First, all countries could improve energy-efficiency, insulate buildings better, and so forth—and thereby actually save money.

Second, we could target cuts to methane, black carbon and CFC emissions. These are subsidiary contributors to warming. But unlike CO2 they cause local pollution too – in Chinese cities, for instance – so there's a stronger incentive to reduce them.

[Of course the unique difficulty of motivating efforts to cut CO2 is that that the impact of any action not only lies decades ahead, but is globally diffused and brings no special benefit to the actor.]

But -- third and most important -- nations should hugely expand R and D into all forms of low-carbon energy generation (renewables, 4th generation nuclear, fusion, and the rest). And into other technologies where parallel progress is crucial – especially storage (batteries, compressed air, pumped storage, flywheels, etc) and smart grids.

The faster research in this area advances, the sooner will 'clean energy' systems become affordable to developing countries. This will need more generating capacity, and will otherwise feel pressured to build coal-fired power stations.

At the moment, only 2 percent of publicly funded R and D is devoted to these challenges. Why shouldn't the percentage be comparable to spending on medical or defence research?

This is an arena where public, private and philanthropic efforts need to mesh together. And it would be hard to think of a more inspiring challenge for young engineers than devising clean energy systems for the world – to replace fossil fuel globally, and to bring local benefits to the billion whose health is jeopardized by smoky stoves burning wood or dung.

All renewables have their niches, but an attractive scenario for Europe would be large-scale solar energy, coupled with a transcontinental DC smart grid network (north-south to transmit power from the sunny south southern Spain or even Morocco) , and east-west to smooth over peak demand in different time-zones)--- with efficient storage as well,

For the rest of this talk I'll focus on a different topic –the promise, and the dark side, of novel technologies that change society and empower individuals – and I'll venture some speculations about the far future.

We live in a 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 networks are highly resilient, their 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 literally at the speed of light.

It's imperative to guard against the downsides of such an interconnected world. Plainly this requires international collaboration. (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 magnitude of the societal breakdown from pandemics is far higher than in earlier centuries. English villages in the 14th century continued to function even when the black death almost halved their populations. In contrast, our social framework would be vulnerable to breakdown as soon as hospitals overflowed and health services were overwhelmed– which could occur when the fatality rate was still a fraction of one percent. And the human cost would be worst in the megacities of the developing world.]

Advances in microbiology -- diagnostics, vaccines and antibiotics -- offer prospects of containing pandemics. But the same research has controversial downsides. For instance, in 2012 researchers showed that it was surprisingly easy to make a virus both more virulent and transmissible. Last October, the US federal government decided to cease funding these so-called 'gain of function' experiments.

The new CRISPR techniques for gene-editing have been in the news because of ethical concerns raised by Chinese experiments on human embryos.

Malign or foolhardy individuals have far more leverage than in the past. It is hard to make a clandestine H-bomb. In contrast, biotech involves small-scale dual use equipment. Millions will one day have the capability to misuse it, just as they can misuse cybertech today. Indeed, biohacking is burgeoning even as a hobby and competitive game.

Back in the early days of DNA research, a group of biologists met in Asilomar, California and agreed guidelines on what experiments should and shouldn't be done. This seemed an encouraging precedent, and there's a call for similar regulation of the new bio-techniques. But today, 40 years later, the research community is far more broadly international, and

more influenced by commercial pressures. Whatever regulations are imposed, on prudential or ethical grounds, can't be enforced worldwide – any more than the drug laws can --- or the tax laws. Whatever can be done will be done by someone, somewhere.

We know all too well that technical expertise doesn't guarantee balanced rationality. The global village will have its village idiots and they'll have global range. The rising empowerment of tech-savvy groups (or even individuals), by bio as well as cyber technology will pose an intractable challenge to governments and aggravate the tension between freedom, privacy and security.

James Martin envisaged high tech monitoring via wearable technology, microwave and neutron beams, and suchlike. These might not be as vexatious as current security checks, but of course that doesn't make us relaxed about their intrusiveness.

This reminds us of another issue. The order in which futuristic technologies develop can be crucial. Monitoring techniques, vaccines, and so forth should be prioritized above the technologies that render them necessary.

These bio-concerns are relatively near-term – within 10 or 15 years. What about 2050 and beyond?

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 transformative advances that may now seem science fiction.

On the bio front, the great physicist Freeman Dyson conjectures 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.

And what about another transformative technology: robotics and artificial intelligence (AI)?

It's 20 years since IBM's 'Deep Blue' beat Kasparov, the world chess champion. More recently, another IBM computer won a TV gameshow – not the mindless kind featuring bubble-headed celebs, but one called 'Jeopardy' that required wide knowledge, and crossword-clue style questions.

Computers use 'brute force' methods. They learn to identify dogs, cats and human faces by 'crunching' through millions of images – not the way babies learn. They learn to translate by reading millions of pages of (for example) multilingual EU documents (they never get bored!).

There's been exciting advances in what's called generalized machine learning – Deep Mind (a small London company that Google recently bought) created a machine that can figure out the rules of old Atari games without being told, and then play them better than

humans.

Advances in sensors and motor-skills have been slower. Robots are still clumsy in moving pieces on a real chessboard. They can't tie your shoelaces or cut your toenails. But sensor technology, speech recognition, information searches and so forth are advancing apace.

They won't just take over manual work (indeed plumbing and gardening will be among the hardest jobs to automate), but routine legal work (conveyancing and suchlike) and medical diagnostics and operations.

Can robots cope with emergencies? For instance, if an obstruction suddenly appears on a crowded highway, can Google's driverless car discriminate whether it's a paper bag, a dog or a child? The likely answer is that its judgement will never be perfect, but will be better than the average driver – machine errors will occur, but not as often as human error. But when accidents occur, they will create a legal minefield. Who should be held responsible – the 'driver', the owner, or the designer?

The big social and economic question is this: Will robotics be like earlier disruptive technologies – the car, for instance –and create as many jobs as it destroys? Or is it really different this time?

The fine work of Carl Frey and Michael Osborne addresses this key question.

These innovations could generate huge wealth for an elite. It's not just lefties but people like Martin Wolf of the FT who argue the need for massive redistribution to ensure that everyone had at least a 'living wage'. And to create and upgrade public-service jobs where the human element is crucial and is now undervalued – carers for young and old, custodians, gardeners in public parks and so on.

But let's look further ahead.

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.

Such machines pervade popular culture —in movies like Her, Transcendence and Ex Machina.

In his scary and scholarly book 'Superintelligence', Nick Bostrom speculates what cold happen if a machine developed a mind of its own. **Would it stay docile, or 'go rogue'?** 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.**

Some of the serious AI pundits think the field already needs guidelines – just as biotech does. But others regard these concerns as premature – and worry less about artificial

intelligence than about natural stupidity.

Be that as it may, it's likely that during this century our society and its economy will be transformed by autonomous robots, even though the jury's out on whether they'll be 'idiot savants' or display superhuman capabilities.

[There's disagreement about the route towards human-level intelligence. Some think we should emulate nature, and reverse-engineer the human brain. Others say that's as misguided as designing a flying machine by copying how birds flap their wings. (And philosophers debate whether "consciousness" is special to the wet, organic brains of humans, apes and dogs — so that robots, even if their intellects seem superhuman, will still lack self-awareness or inner life).]

Ray Kurzweil, now working at Google, argues that once machines have surpassed human capabilities, they could themselves design and assemble a new generation of even more powerful ones – an intelligence explosion. 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'.

Kurzweil is the most prominent proponent of this so-called 'singularity'. 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 or your brain downloaded. If you can't afford the full whack there's a cut price option of having just your head frozen,

I told them I'd rather end my days in an English churchyard than a Californian refrigerator. They derided me as a 'deathist'. (But I later read (in the Daily Mail – so who could doubt it?!) that Nick Bostrom and two of his colleagues are Cryonic enthusiasts who've signed up).

And that reminds me that research on ageing is being seriously prioritized. Will the benefits be incremental? Or is ageing a 'disease' that can be cured? Dramatic life-extension would plainly be a real wild card in population projections, with huge social ramifications. But it may happen, along with human enhancement in other forms.

Technology brings with it great hopes, but also great fears. The Martin School is uniquely well placed to assess which scenarios 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.

We mustn't forget an important maxim: the unfamiliar is not the same as the improbable.

And now a digression into my special interest – space. This is where robots surely have a future. During this century the whole solar system will be explored by flotillas of miniaturized probes – far more advanced than those on Rosetta and which surveyed Pluto

– which were built 15 years ago. And giant robotic fabricators may build vast lightweight structures floating in space (solar energy collectors, gossamer-thin radio reflectors, for instance) -- mining raw materials from the Moon or asteroids.

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. These private enterprise efforts can tolerate higher risks than a western government could impose on publicly-funded civilians, and thereby cut costs compared to NASA or ESA. And they should be promoted as adventure or extreme sports -- the phrase 'space tourism' should be avoided. It lulls people into unrealistic confidence.

By 2100 courageous pioneers in the mould of Ranolf Fiennes (the adventurer who dragged a sledge to the south pole in Antarctic winter) or the man who broke the sound barrier in free fall from a high-altitude balloon may have established 'bases' independent from the Earth – on Mars, or maybe on asteroids. Musk himself (aged 44) says he wants to die on Mars – but not on impact.

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 I've argued that 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 society. But they wouldn't wipe us all out. They're extreme, but strictly speaking not 'existential'.

So are there conceivable events that could snuff out all life?

Physicists were (in my view quite rightly) pressured to address the speculative 'existential risks' that could be triggered by CERN's LHC accelerator, which generated unprecedented concentrations of energy. Could it 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 calculated that cosmic rays of much higher energies collide frequently in the Galaxy, but haven’t ripped space apart.

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

So how risk-averse should we be? 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. If a US 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 uneasy 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 an unlimited and unenvisioned source of energy?” I’d again offer fairly monstrous 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 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 Oxford 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 trillions of future people – and indeed an open-ended post-human future.

This immense future, incidentally, is something that astronomers are specially aware of.

The stupendous timespans of the evolutionary past are now part of common culture (outside ‘fundamentalist’ circles, at any rate). But most people still tend to regard humans as 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. To quote Woody Allen, eternity is very long, especially towards the end.

The timescale for human-level AI may be decades, or it may be centuries. Be that as it may,

it's but an instant compared to the future horizons, and indeed far shorter than timescales of the Darwinian selection that led to humanity's emergence.

I think it's likely that the machines will gain dominance on Earth. This is because there are chemical and metabolic limits to the size and processing power of 'wet' organic brains. Maybe we're close to these already. But no such limits constrain silicon based computers (still less, perhaps, quantum computers): for these, the potential for further development over the next billion years could be as dramatic as the evolution from pre-Cambrian organisms to humans. So, by any definition of 'thinking', the amount and intensity that's done by organic human-type brains will be utterly swamped by the future cerebrations of AI.

Moreover, the Earth's biosphere isn't the optimal environment for advanced AI – interplanetary and interstellar space may be the preferred arena where robotic fabricators will have the grandest scope for construction, and where non-biological 'brains' may develop powers than humans can't even imagine

But we humans shouldn't feel too humbled. We could be of special cosmic significance for jump-starting the transition to silicon-based (and potentially immortal) entities, spreading their influence far beyond the Earth, and far transcending our limitations.

So, even in this 'concertinered' 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.

That's the rationale for the Future of Humanity Institute, the element of the Martin School that addresses 'existential' risks on the science fiction fringe.

But enough of these speculations. In closing, let's focus back closer to here and now.

I'd argue that there's no <u>scientific</u> impediment to achieving a sustainable world, where all enjoy a lifestyle better than those in the 'west' do today. We live under the shadow of new risks – but these can be minimized by a culture of 'responsible innovation', especially in fields like biotech, advanced AI and geoengineering. And the thrust of the world's technological effort needs redirection. We can be technological optimists. Were that all, the answer to the question in my title would be a firm 'no'.

But the intractable politics and sociology -- the gap between potentialities and what actually happens -- engenders pessimism. There are many problems. The emergent threat from globally-empowered mavericks is growing. But the pre-eminent concern is the institutional failure to plan long-term, and to plan globally. For politicians, the local trumps the global, the immediate trumps the long term.

But almost all the issues I've addressed this evening have to be tackled, monitored and regulated internationally. This can happen via agreements of the kind being sought in Paris on climate – or via beefed-up global agencies resembling the IAEA or the WHO.

The Martin School's research and expertise in all these areas is a huge resource for decision-makers. But we want its findings to resonate far beyond academia. So how can its impact be maximized?

Obviously, links forged with politicians and senior officials can help – and links with NGOs and the private sector too. Our report 'Now for the Long Term' got significant traction. But scientists who've served as government advisors have often had frustratingly little influence. Experts by themselves can't generate political will.

Politicians are, however, influenced by their inbox, and by the press. So academics can sometimes achieve more as 'outsiders' and activists – promoting their message via widely-read books, via campaigning groups, via blogging and journalism, or through political activity. If their voices are echoes and amplified by a wide public, and by the media, long-term global causes will rise on the political agenda.

And sometimes, incidentally, the great religious faiths can be our allies. The Pope's recent Encyclical on environmental and climate was hugely welcome. The Catholic Church transcends normal political divides – there's no gainsaying its global reach, nor its durability and long-term vision, nor its focus on the world's poor. This Pope's message resonates in Latin America, Africa, and East Asia – even perhaps in the American Republican Party.

Universities span national and political divides too. Younger people, who may live into the 22nd century, have wider horizons. And they care: about the environment; and about the poor.

Student involvement in, for instance, 'effective altruism' campaigns is burgeoning, especially here in Oxford. William McAskill's book 'Doing Good Better' is a compelling manifesto. It reminds us that urgent and meaningful improvements to people's lives can be achieved by redeployment of existing resources and knowledge towards the developing or destitute nations – something that Jim Martin emphasized in his own lectures.

But Jim's gaze focused on the future.

it would surely be shameful if we persisted in short term policies that denied future generations a fair inheritance and left them with a more depleted and more hazardous world. Wise choices will require the effective efforts of natural scientists, environmentalists, social scientists and humanists ---all guided by the knowledge that 21st century science can offer – and inspired by values that science alone can't provide.

"Space-ship Earth" is hurtling through the void. 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 risk.

In Jim's own words:

“If we understand this century …. , our future will be magnificent. It we get it wrong we may be at the start of a new Dark Age.”

These concerns were his motivation for setting up the Martin School. It’s had a hugely productive first decade. Let’s hope it indeed fulfils Jim’s vision and helps ensure a bright future for generations to come.

Thank you.