

# William MacAskill: ‘There are 80 trillion people yet to come. They need us to start protecting them’

*Andrew Anthony*

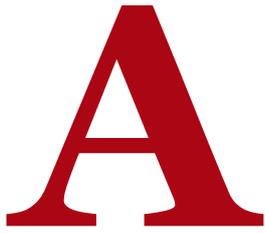


For the many, not the few: William MacAskill. Photograph: Matt Crockett

**The moral philosopher gives most of his earnings to charity, says we need more not less economic growth to tackle global heating, and in a striking new book argues that it’s our duty to ensure the wellbeing of our distant descendants**

Sun 21 Aug 2022 13.00 BST

- 
- 
-



Although most cultures, particularly in the west, provide a great many commemorations of distant ancestors – statues, portraits, buildings – we are much less willing to consider our far-off descendants. We might invoke grandchildren, at a push great-grandchildren, but after that, it all becomes a bit vague and, well, unimaginable.

And while we look with awe and fascination at the Egyptian pyramids, built 5,000 years ago, we seem incapable of thinking, or even contemplating, 5,000 years in the future. That lies in the realm of science fiction, which is tantamount to fantasy. But the chances are, barring a global catastrophe, humanity will still be very much around in 5,000 years, and going by the average existence of mammal species, should still be thriving in 500,000 years. If we play our cards right, we could even be here in 5m or 500m years, which means that there may be thousands or even millions times more human beings to come than have already existed.

All those numbers seem incalculably abstract but, according to the moral philosopher **William MacAskill**, they should command our attention. He is a proponent of what's known as longtermism – the view that the deep future is something we have to address now. How long we last as a species and what kind of state of wellbeing we achieve, says MacAskill, may have a lot to do with what decisions we make and actions we take at the moment and in the foreseeable future.

That, in a nutshell, is the thesis of his new book, *What We Owe the Future: A Million-Year View*. The Dutch historian and writer Rutger Bregman calls the book's publication “a monumental event”, while the US neuroscientist Sam Harris says that “no living philosopher has had a greater impact” upon his ethics.

We tend to think of moral philosophers as whiskery sages, but MacAskill is a youthful 35 and a disarmingly informal character in person, or rather on a Zoom call from San Francisco, where he is promoting the book.

Associate professor at Lincoln College, Oxford, he is president of the **Centre for Effective Altruism**, a body he co-founded to bring data-based analysis to the business of charity, thus making donations as effective as possible. He also co-founded with fellow moral philosopher Toby Ord **Giving What We Can**, an organisation whose members pledge to give at least 10% of their earnings to effective charities (MacAskill himself gives the majority of his earnings), and is also president of **80,000 Hours**, a non-profit group that advises on what careers have the most positive social impact.

“It was in the course of reasoning on the basis of Effective Altruism,” he says, “that led me to think about issues that impact not just on the present generation but also the longterm future too.”

MacAskill grew up in a comfortable middle-class Glasgow family and attended private school. He’d always had an altruistic side. As a 15-year-old, prompted by the knowledge of how many people were dying in the Aids crisis, he decided he wanted to become wealthy and give half his money away. He did voluntary work for a disabled scout group, but it wasn’t until he got to Cambridge, where he studied philosophy and played saxophone in a funk band, that his moral outlook took on a more intellectual form. Reading **Peter Singer’s *Famine, Affluence, and Morality*** propelled him into a lifetime of not just philosophical but practical commitment.

So it’s clear that MacAskill is not one of those armchair thinkers who talks the talk but doesn’t follow up on his ideas. Although he wouldn’t describe himself as a utilitarian, he is concerned with the maximum good, and firmly believes in minimising human suffering and maximising human wellbeing. He also believes in minimising the suffering of animals, something humanity almost certainly increased. But for humans to flourish they first have to be alive, and his argument is that the more humans there are who live, and the happier lives they lead, the better.

No one can say with great accuracy how many humans have lived, but one recent estimate by the US **Population Reference Bureau** says that about 120

billion Homo sapiens have so far been born. MacAskill says that if we assume that our population continues at its current size and we last as long as typical mammals, that would mean “there would be 80 trillion people yet to come; future people would outnumber us 10,000 to one”.

The moral argument is that, by sheer weight of numbers, our descendants’ needs should loom large in our deliberations. With alarming signs that the **climate crisis** is already upon us, MacAskill states the obvious and urgent need for decarbonisation. This problem, however, is not the main focus of his book. Rather he employs the climate crisis as a proof of longtermism: “We all contribute to a problem that literally has effects for hundreds of thousands of years,” he says.

Second, he explains: “Climate change is much less neglected than other concerns like pandemic prevention, nuclear war and AI safety: it is already widely agreed to be among the world’s most important problems and there are large social movements dedicated to solving climate change.”

The debate around the issue also provides a model of how to deal with uncertainty.

**He argues that humanity could soon enter a phase in which values, good or bad, will become ‘locked-in’**

“Climate change sceptics often point to our uncertainty as a reason for inaction...,” he writes. “But this is a terrible argument... Crucially, the uncertainty around climate change is not symmetric: greater uncertainty should prompt more concern about worst-case outcomes and this shift is not offset by a higher chance of best-case outcomes, because the worst-case outcomes are worse than the best-case outcomes are good.”

To deal with the uncertainty that is inherent in longterm thinking, MacAskill believes we can make greater use of a method of probability assessment called “expected value theory”. It’s a means of prioritising or assigning values to outcomes in situations with multiple possibilities. Professional gamblers use it but MacAskill says its application could help guide us through the complex contingencies ahead.

He gives the example of how the Intergovernmental Panel on Climate Change estimates that in a medium-low-emissions scenario there will be 2.5C (36.5F) of warming by the end of the century. “But this is uncertain,” he writes. “There is a one in 10 chance that we get 2 degrees or less. But that should not reassure us, because there is also a one in 10 chance that we get more than 3.5 degrees. Less than 2 degrees would be something of a relief compared with the best-guess estimate, but more than 3.5 degrees would be much worse. The uncertainty gives us *more* reason to worry, not less.”

The good news is that with the threat of an engineered pandemic, which he says is rapidly increasing, he believes there are specific steps that can be taken to avoid a breakout.

“One partial solution I’m excited about is called far ultraviolet C radiation,” he says. “We know that ultraviolet light sterilises the surfaces it hits, but most ultraviolet light harms humans as well. However, there’s a narrow-spectrum far UVC specific type that seems to be safe for humans while still having sterilising properties.”

The cost for a far UVC lightbulb at the moment is about \$1,000 (£820) per bulb. But he suggests that with research and development and philanthropic funding, it could come down to \$10 or even \$1 and could then be made part of building codes. He runs through the scenario with a breezy kind of optimism, but one founded on science-based pragmatism.

Less tractable is the threat that **artificial intelligence** (AI) may lead to a dystopian outcome. He argues that humanity is currently in a phase of history in which our values are plastic – they can still be shaped and changed – but that we could soon enter a phase in which values, good or bad, will become “locked-in”.

Illustration by Lehel Kovács.

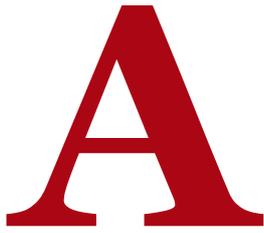
“Imagine,” he says, “if the Nazis had won the second world war, held on to power for a few hundred years, established a world government and then got to the point of developing AGI [artificial general intelligence], then you could have a future that would be guided and controlled by Nazi ideology for ever.”

The critical point in this analogy is not so much the Nazis, who represent humanity's potential for doing ill, but AGI. Put simply, AGI is the technological state in which an intelligent machine can learn and enact any intellectual task that a human being can do. From that point, the potential to control ideas and social development is almost limitless, which brings into focus the possibility of an unending dystopia.

For MacAskill, the time to address that possibility is now, because later may well prove too late. No one can be certain that AGI is achievable and, if it is, when it will be achieved, but most scientists working in the field think that it will happen. A significant minority believe that it's probable within the next 50 years and some think it may take only 20. MacAskill himself estimates a 10% chance of AGI in the next 50 years. Short of stopping research across the planet – and how could that be enforced? – what can be done?

“This is tough,” he acknowledges. “It’s not nearly as clearcut as preventing pandemics. But I think there are some things we can do. For example, the idea of slowing down some areas of AI research. AI will be hugely beneficial, but you can get a lot of the gains without going all the way. Do we need to have AI systems that are engaged in longterm planning? Do we need to have AI systems that are enormously multimodal and able to do very many different things rather than just narrow tasks?”

While these choices are not yet directly upon us, he says that there are steps to be taken immediately. One area that demands more support is the field of interpretability research in AI. “At the moment,” he explains, “we’ve got these black boxes with input data, this enormously complex model, and then these outputs and we don’t really know what’s going on under the hood. There are enormous challenges but I can see a tractable path towards making progress on this.”



At the outset of his book, MacAskill presents a metaphor of a risky expedition into uncharted terrain. Just like early explorers, we don't know exactly what threats await us, but "we can scout out the landscape ahead of us, ensure the expedition is well resourced and well coordinated, and, despite uncertainty, guard against those threats we are aware of".

In making his case for the journey ahead, MacAskill dismisses some of the ideas that are held dear by many who are concerned about the future, particularly those looking at things from an environmental perspective. It's not uncommon in green circles to hear arguments against economic growth, against consumption and, indeed, even against bringing further children into the world.

MacAskill disagrees with all these positions. In the longterm the kind of growth we've seen in the past century or so – above 2% – is unsustainable. If it continued at a rate of 2% for the next 10,000 years, he writes, "we would produce 100tn trillion trillion trillion trillion trillion times as much output as we do now".

That means there would have to be 10m tn times as much output as our current world produces for every atom that we could in principle access. "Though of course we can't be certain," he writes with wry understatement, "this just doesn't seem possible."

However, he doesn't believe now is the time to slow growth because, he argues, we are not yet at a technological stage where that's possible without potentially calamitous effects. To illustrate the point, he uses the example of where we were 100 years ago. If we stopped growth then, we would have been left with two possibilities: either to return to the grinding privations of agricultural life or to burn through fossil fuels, leading to a climate catastrophe.

He believes both technological development and economic growth are needed to avoid threats of climate crisis, bioterrorism and much else besides. The other point he makes is that stopping growth would in any case be pointless unless all 193 countries signed up to it.

“Suppose that you managed to persuade 192 to stop growth, but one country continues to grow. Compound growth means that before too long that one country is the entire world economy and in the long run you’ve really not done anything.”

Similarly, rather than cut back on consumption, he argues, it’s much more effective to donate to causes that are dealing with the problems created by consumption.

“We need to get more fine-grained,” he says. “OK, some technologies have net bad effects, some have net good effects. And we can really push on the ones that are more beneficial.”

At the core of the book is the question of human values. These are obviously not set in stone, because we need only look at history to see how they have radically changed over time. A key example that MacAskill returns to is slavery and its abolition. At various periods and across most cultures slavery has existed and been deemed natural, or at least acceptable.

Arguments were made against it, but none so compelling that it ceased to be. What changed things was the Atlantic slave trade and the Enlightenment. The contradictions between the principles of universalism and owning and mistreating fellow humans became increasingly untenable, at least in a moral and intellectual sense.

But despite the claims that have been made, says MacAskill, there was no economic imperative to end slavery. Sugar plantations were not mechanised for many years after the end of slavery and sugar consumption continued to grow when slavery stopped. Rather than think of abolition as inevitable, he argues, we should acknowledge the part played by those who made the moral case. It seems so obvious to us now that it’s hard to imagine anyone could have opposed it. But powerful forces did. Moral progress is contingent, MacAskill emphasises, not inexorable. It just seems like that once progress has been achieved.

Going against that moral step forward, however, is the argument that rejects humanism, post-Enlightenment values and the whole liberal discourse as merely being the soft power of colonialism, a western imposition that runs roughshod over other cultures and values. MacAskill has little time for such relativist arguments.

“I think to conflate colonialism and [these aspects of liberalism] is a huge mistake,” he says. “Colonialism was absolutely horrific, one of the great abominations of history. But if you have this idea that all moral perspectives command equal respect, then are slave-owning societies and extreme patriarchal societies – the most common societies throughout history – OK, it’s just their way of being and we shouldn’t tell them they’re wrong? No.”

That said, MacAskill believes the west has a great deal to learn from other cultures. In the first instance, by appreciating just how well off most people are in developed countries by comparison with those who are not. He says he currently gives away everything he earns above post-tax £26,000 per year, which still places him in the top 5% in terms of wealth across the globe. Although he doesn’t have yet children, if he did, he would allow an extra £5,000 for each one, were he sharing the financial burden with another parent, and £10,000 if a single parent.

“I honestly think people in rich countries could be giving radically more than they’re currently giving at very little cost to their own wellbeing, while at the same time doing an enormous amount of good.”

There are also philosophical wisdoms to be gained from Indigenous populations.

“The Iroquois in their oral constitution advocated concern for future generations,” he says. “It seems like this was true for many Indigenous philosophies in Africa as well.”

He speculates that there are logical cultural reasons why this was the case. In hunter-gatherer societies technological change took place very slowly. So learning something from your ancestors 1,000 years ago or handing down those wisdoms to your descendants 1,000 years hence was viable, because it was likely to still be relevant.

But in societies undergoing rapid change, we feel more disconnected from the distant future because we struggle to conceive what it will be like.

“We’ve gone too far in that direction,” he says, suggesting that we have developed the conceptual tools to navigate our way through the great unknown to come and should make use of them. “We can now use expected value theory to hedge against uncertainty.”

One problem is that often when a society, or more specifically a regime, speaks of the longterm future, it’s to establish an epic stage to bolster their claims on governance. It’s what the Nazis attempted to do with the “thousand-year Reich”, just as the first Qin emperor spoke of an empire lasting 10,000 generations. In fact the Qin empire lasted for 15 years, three years longer than the Nazis.

What MacAskill is arguing for, though, is humility in the face of the astonishing expanse of time that humanity could fill. But that shouldn’t lead to complacency or paralysis. Standard economic forecasts for the next 100 years predict more of the same, with growth of approximately 2% a year. MacAskill says that we should take into account the possibility of a catastrophe that wipes out 50% of the population, or that there might a significant increase in growth.

We don’t know what’s going to happen, but we should put a lot more time and effort into preparing for different outcomes. We owe that to ourselves, says MacAskill, but we also owe it to the teeming billions yet to come.

- *What We Owe the Future: A Million-Year View* by William MacAskill is published by Oneworld (£20). To support the *Guardian* and *Observer* order your copy at [guardianbookshop.com](https://guardianbookshop.com). Delivery charges may apply