

The Case for Longtermism

Aug. 5, 2022 New York Times Opinion

By William MacAskill

A professor of philosophy at Oxford University and the author of “What We Owe the Future,” from which this essay has been adapted.

Imagine living the life of every human being who has ever existed — in order of birth.

Your first life begins about 300,000 years ago in Africa. After living that life and dying, you travel back in time to be reincarnated as the second-ever person, born slightly later than the first, then the third-ever person, and so on.

[One hundred billion \(or so\)](#) lives later, you are the youngest person alive today. Your life has lasted somewhere in the ballpark of four trillion years. You have spent approximately 10 percent of it as a hunter-gatherer and 60 percent as a farmer, a full 20 percent raising children, and over 1 percent suffering from malaria or smallpox. You spent 1.5 billion years having sex and 250 million giving birth.

That’s your life so far — from the birth of *Homo sapiens* until the present.

But now imagine that you live all future lives, too. Your life, we hope, would be just beginning. Even if humanity lasts only as long as the typical mammal species (about one million years), and even if the world population falls to a tenth of its current size, 99.5 percent of your life would still be ahead of you. On the scale of a typical human life, you in the present would be just a few months old. The future is big.

I offer this thought experiment because morality, at its core, is about putting ourselves in others’ shoes and treating their interests as we do our own. When we do this at the full scale of human history, the future — where almost everyone lives and where almost all potential for joy and misery lies — comes to the fore.

If you knew you were going to live all these future lives, what would you hope we do in the present? How much carbon dioxide would you want us to emit into the atmosphere? How careful would you want us to be with new technologies that could destroy, or permanently derail, your future? How much attention would you want us to give to the impact of today’s actions on the long term?

These are some of the questions that motivate *longtermism*: the idea that positively influencing the long-term future is a key moral priority of our time.

Longtermism is about taking seriously just how big the future could be and how high the stakes are in shaping it. If humanity survives to even a fraction of its potential life span, then, strange as it may seem, we are the ancients: we live at the very beginning of history, in its most distant past. What we do now will affect untold numbers of future people. We need to act wisely.

It took me a long time to come around to longtermism. Over the past 12 years, I've been an advocate of *effective altruism* — the use of evidence and reason to help others as much as possible. In 2009, I co-founded [an organization](#) that has raised hundreds of millions of dollars to help pay for bed nets to protect families against malaria and medicine to cure children of intestinal worms, among other causes. These activities had a tangible impact. By contrast, the thought of trying to improve the lives of unknown future people initially left me cold.

But some simple ideas exerted a persistent force on my mind: Future people count. There could be a lot of them. And we can make their lives better. To help others as much as possible, we must think about the long-term impact of our actions.

The idea that future people count is common sense. Suppose that I drop a glass bottle while hiking. If I don't clean it up, a child might cut herself on the shards. Does it matter when the child will cut herself — a week, or a decade, or a century from now? No. Harm is harm, whenever it occurs.

Future people, after all, are people. They will exist. They will have hopes and joys and pains and regrets, just like the rest of us. They just don't exist *yet*.

But society tends to neglect the future in favor of the present. Future people are utterly disenfranchised. They can't vote or lobby or run for public office, so politicians have scant incentive to think about them. They can't tweet, or write articles, or march in the streets. They are the true silent majority. And though we can't give political power to future people, we can at least give them fair consideration. **We can renounce the tyranny of the present over the future and act as trustees for all of humanity, helping to create a flourishing world for the generations to come.**

We face massive problems today, and the world is full of needless suffering, but in some ways we have made remarkable progress over the past few hundred years. Three hundred years ago, the average life expectancy was [less than 40 years](#); today, [it's over 70](#). More than [80 percent](#) of the world lived in extreme poverty; now just around [10 percent](#) does. Women were, for the most part, unable to attend universities, and the feminist movement didn't exist. No one lived in a democracy; now [over half](#) the world does. We've come a long way.

We have the power to encourage these positive trends. And we can reverse course on negative trends, too, like increases in carbon emissions and factory farming. We have the potential to build a world where everyone lives like the happiest people in the most well-off countries today.

And we could do better still — far better. Much of the progress we've made since 1700 would have been very difficult for people back then to anticipate. And that's with only a three-century gap. Humanity could, theoretically, last for millions of centuries on Earth alone. If we anchor our sense of humanity's potential to a fixed-up version of our present world, we risk dramatically underestimating just how good life in the future could be.

When I first began thinking about longtermism, my biggest reservation was practical. Even if future generations matter, what can we *actually do* to benefit them? **But as I learned more about the history-shaping events that could occur in the near future, I realized that we might soon be approaching a critical juncture in the human story.** Technological development is creating new threats and opportunities, putting the lives of future people on the line. Whether we get a future that's beautiful and just, or flawed and dystopian, or whether civilization ends and we get no future at all — that depends, in significant part, on what we do today.

Some of the ways we affect the long-term future are familiar. We drive. We fly. We emit greenhouse gases that can remain in the atmosphere and impact the environment for hundreds of thousands of years.

But reducing fossil fuel use is not the only way to improve the long term. Other challenges are at least as important, and often radically more neglected.

Chief among these is the development of advanced artificial intelligence. According to [leading economic models](#), advanced A.I. could greatly accelerate economic growth and technological progress. But equipped with A.I.-enabled capabilities, bad political actors could potentially increase and entrench their power. Our future could be a perpetual totalitarian dystopia.

Or we could [lose control](#) over the A.I. systems we've created. Once artificial intelligence far exceeds human intelligence, we could find ourselves with as little power over our future as chimpanzees have over theirs. Civilization could be governed by the A.I.'s aims and goals, which could be utterly alien and valueless from our perspective.

And we may not even make it to the development of advanced A.I. We still live under the shadow of 9,000 nuclear warheads, each far more powerful than the bombs dropped on Hiroshima and Nagasaki. [Some experts](#) put the chances of a third world war by 2070 at over 20 percent. An all-out nuclear war could cause the collapse of civilization, and we might never recover.

Advances in biotechnology could create weapons of even greater destructive power. Engineered viruses could be much more deadly than natural diseases because they could, in theory, be modified to have dangerous new properties: the lethality of Ebola and the contagiousness of measles. In the worst-case scenario, the release of an

engineered bioweapon could kill billions, possibly beyond the point where humanity could recover. Our future would be permanently destroyed.

These are daunting challenges. In his book “The Precipice,” my colleague Toby Ord puts the probability of an existential catastrophe in the next century at one in six — roughly equivalent to playing Russian roulette. This is an unacceptable level of risk.

We aren’t helpless in the face of these challenges. Longtermism can inspire concrete actions, here and now. Constant monitoring of wastewater could ensure that we respond to any new viruses as soon as they arise. Developing and distributing advanced personal protective equipment would protect essential workers. Forms of far-ultraviolet lighting can safely [sterilize](#) a room; if proven safe and widely installed, this could prevent airborne pandemics while eliminating all respiratory diseases along the way.

On A.I. risk, too, there is much to do. We need the brightest technical minds to figure out what goes on under the hood of increasingly inscrutable A.I. systems, and to ensure they are helpful, harmless and honest. We need scholars and policymakers to design new governance systems to ensure that A.I. is developed for the benefit of all humankind. And we need courageous leaders to prevent new arms races and catastrophic wars between the great powers.

If we are careful and farsighted, we have the power to help build a better future for our great-grandchildren, and their great-grandchildren in turn — down through hundreds of generations. **But positive change is not inevitable. It’s the result of long, hard work by thinkers and activists. No outside force will prevent civilization from stumbling into dystopia or oblivion. It’s on us.**

Does longtermism imply that we must sacrifice the present on the altar of posterity? No. Just as caring more about our children doesn’t mean ignoring the interests of strangers, caring more about our contemporaries doesn’t mean ignoring the interests of our descendants.

In fact, as I’ve learned more about longtermism, **I’ve realized that there is remarkable overlap between the best ways we can promote the common good for people living right now and for our posterity.**

Every year [millions](#) of people, disproportionately in poor countries, die prematurely because fossil fuel burning pollutes the air with particulates that cause lung cancer, heart disease and respiratory infections. Moving off carbon is a win-win for both the near and the long term. The same holds for preventing pandemics, controlling artificial intelligence and decreasing the risk of nuclear war.

The idea that we could affect the long-term future, and that there could be so much at stake, might just seem too wild to be true. This is how things initially seemed to me. But **I think this wildness comes not from the moral premises**

that underlie longtermism but from the fact that we live at such an unusual time.

Our era is undergoing an [unprecedented amount of change](#). Currently, the world economy doubles in size about every [19 years](#). But before the Industrial Revolution, it took hundreds of years for the world economy to double; and for hundreds of thousands of years before that, growth rates were close to zero. What's more, the current rate of growth cannot continue forever; within just 10,000 years, there would be a trillion civilizations' worth of economic output for every reachable atom.

All this indicates that we are living through a unique and precarious chapter in humanity's story. Out of the hundreds of thousands of years in humanity's past — and the potentially billions of years in its future — we find ourselves living *now*, at a time of extraordinary change.

A time marked by thousands of nuclear warheads standing ready to fire. A time when we are rapidly burning fossil fuels, producing pollution that might last hundreds of thousands of years. A time when we can see catastrophes on the horizon — from engineered viruses to A.I.-enabled totalitarianism — and can act to prevent them.

To be alive at such a time is both an exceptional opportunity and a profound responsibility: We can be pivotal in steering the future onto a better trajectory. There's no better time for a movement to stand up, not just for our generation or even our children's generation, but for all the generations yet to come.

William MacAskill is a professor of philosophy at Oxford University and the author of the forthcoming book "What We Owe the Future," from which this essay has been adapted.

The Times is committed to publishing [a diversity of letters](#) to the editor. We'd like to hear what you think about this or any of our articles. Here are some [tips](#). And here's our email: letters@nytimes.com.

Follow The New York Times Opinion section on [Facebook](#), [Twitter \(@NYTopinion\)](#) and [Instagram](#).

A version of this article appears in print on Aug. 7, 2022, Section SR, Page 5 of the New York edition with the headline: The Case for Longtermism. [Order Reprints](#) | [Today's Paper](#) | [Subscribe](#)