

'Longtermism'—why the million-year philosophy can't be ignored

December 19 2022, by Katie Steele

World population




Chart: The Conversation • Source: Our World In Data • Created with Datawrapper

Credit: The Conversation

In 2017, the Scottish philosopher William MacAskill [coined](#) the name "longtermism" to describe the idea "that positively affecting the long-run future is a key moral priority of our time." The label took off among like-minded philosophers and members of the "effective altruism" movement, which sets out to use evidence and reason to determine how individuals can best help the world.

This year, the notion has leapt from philosophical discussions to headlines. In August, MacAskill published a [book](#) on his ideas, accompanied by a barrage of media coverage and endorsements from the likes of [Elon Musk](#). November saw more media attention as a company set up by [Sam Bankman-Fried](#), a prominent financial backer of the movement, collapsed in spectacular fashion.

Critics say longtermism relies on [making impossible predictions](#) about the future, gets caught up in speculation about [robot apocalypses and asteroid strikes](#), depends on wrongheaded moral views, and ultimately fails to give present needs the attention they deserve.

But it would be a mistake to simply dismiss longtermism. It raises thorny philosophical problems—and even if we disagree with some of the answers, we can't ignore the questions.

Why all the fuss?

It's hardly novel to note that modern society has a huge impact on the prospects of future generations. Environmentalists and peace activists have been making this point for a long time—and emphasizing the importance of wielding our power responsibly.

In particular, "[intergenerational justice](#)" has become a familiar phrase, most often with reference to climate change.

Seen in this light, longtermism may look like simple common sense. So why the buzz and rapid uptake of this term? Does the novelty lie simply in bold speculation about the future of technology—such as [biotechnology and artificial intelligence](#)—and its implications for humanity's future?

For example, MacAskill acknowledges we are not doing enough about the threat of climate change, but points out other potential future sources of human misery or extinction that could be even worse. What about a tyrannical regime enabled by AI from which there is no escape? Or an engineered biological pathogen that wipes out the human species?

These are conceivable scenarios, but there is a real danger in getting carried away with sci-fi thrills. To the extent that longtermism chases headlines through rash predictions about unfamiliar future threats, the movement is wide open for criticism.

Moreover, the predictions that really matter are about whether and how we can *change* the probability of any given future threat. What sort of actions would best protect humankind?

Longtermism, like effective altruism more broadly, has been [criticized](#) for a bias towards philanthropic direct action—targeted, outcome-oriented projects—to save humanity from specific ills. It is quite plausible that less direct strategies, such as building solidarity and strengthening shared institutions, would be better ways to equip the world to respond to future challenges, however surprising they turn out to be.

Optimizing the future

There are in any case interesting and probing insights to be found in longtermism. Its novelty arguably lies not in the way it might guide our

particular choices, but in how it provokes us to reckon with the reasoning *behind* our choices.

A core principle of effective altruism is that, regardless of how large an effort we make towards promoting the "general good"—or benefiting others from an impartial point of view—we should try to optimize: we should try to do as much good as possible with our effort. By this test, most of us may be less altruistic than we thought.

For example, say you volunteer for a local charity supporting [homeless people](#), and you think you are doing this for the "general good." If you would better achieve that end, however, by joining a different campaign, you are either making a strategic mistake or else your motivations are more nuanced. *For better or worse*, perhaps you are less impartial, and more committed to special relationships with particular local people, than you thought.

In this context, impartiality means regarding all people's well-being as equally worthy of promotion. Effective altruism was initially preoccupied with what this demands in the spatial sense: equal concern for people's well-being wherever they are in the world.

Longtermism extends this thinking to what impartiality demands in the temporal sense: equal concern for people's well-being wherever they are *in time*. If we care about the well-being of unborn people in the distant future, we can't outright dismiss potential far-off threats to humanity—especially since there may be truly staggering numbers of future people.

How should we think about future generations and risky ethical choices?

An explicit focus on the well-being of future people unearths difficult questions that tend to get glossed over in traditional discussions of altruism and intergenerational justice.

For instance: is a world history containing more lives of positive well-being, all else being equal, better? If the answer is yes, it clearly raises the stakes of preventing human extinction.

A number of philosophers [insist the answer is no](#)—more positive lives is not better. Some suggest that, once we realize this, we see that longtermism is overblown or else uninteresting.

But the implications of this moral stance are less simple and intuitive than its proponents might wish. And premature human extinction is not the only concern of longtermism.

Speculation about the future also provokes reflection on how an altruist should respond to uncertainty.

For instance, is doing something with a 1% chance of helping a trillion people in the future better than doing something that is certain to help a billion people today? (The "expectation value" of the number of people helped by the speculative action is 1% of a trillion, or 10 billion—so it might outweigh the billion people to be helped today.)

For many people, this may seem like gambling with people's lives—and not a great idea. But what about gambles with more favorable odds, and which involve only contemporaneous people?

There are important philosophical questions here about apt risk aversion when lives are at stake. And, going back a step, there are philosophical questions about the authority of any prediction: how certain can we be about whether a possible catastrophe will eventuate, given various

actions we might take?

Making philosophy everybody's business

As we have seen, longtermist reasoning can lead to counter-intuitive places. Some critics respond by eschewing rational choice and "optimization" altogether. But where would that leave us?

The wiser response is to reflect on the combination of moral and empirical assumptions underpinning how we see a given choice. And to consider how changes to these assumptions would change the optimal choice.

Philosophers are used to dealing in extreme hypothetical scenarios. Our reactions to these can illuminate commitments that are ordinarily obscured.

The longtermism movement makes this kind of philosophical reflection everybody's business, by tabling extreme [future](#) threats as real possibilities.

But there remains a big jump between what is *possible* (and provokes clearer thinking) and what is in the end *pertinent* to our actual choices. Even whether we should further investigate any such jump is a complex, partly empirical question.

Humanity already faces many threats that we understand quite well, like [climate change](#) and massive loss of biodiversity. And, in responding to those threats, time is not on our side.

This article is republished from [The Conversation](#) under a Creative Commons license. Read the [original article](#).

Provided by The Conversation

Citation: 'Longtermism'—why the million-year philosophy can't be ignored (2022, December 19)
retrieved 18 July 2024 from <https://phys.org/news/2022-12-longtermismwhy-million-year-philosophy.html>

This document is subject to copyright. Apart from any fair dealing for the purpose of private study or research, no part may be reproduced without the written permission. The content is provided for information purposes only.