

# How Close Is Humanity to the Edge?

 [newyorker.com/culture/annals-of-inquiry/how-close-is-humanity-to-the-edge](https://www.newyorker.com/culture/annals-of-inquiry/how-close-is-humanity-to-the-edge)

Corinne Purtill

In mid-January, Toby Ord, a philosopher and senior research fellow at Oxford University, was reviewing the final proofs for his first book, “The Precipice: Existential Risk and the Future of Humanity.” Ord works in the university’s Future of Humanity Institute, which specializes in considering our collective fate. He had noticed that a few of his colleagues—those who worked on “bio-risk”—were tracking a new virus in Asia. Occasionally, they e-mailed around projections, which Ord found intriguing, in a hypothetical way. Among other subjects, “The Precipice” deals with the risk posed to our species by pandemics both natural and engineered. He wondered if the coronavirus might make his book more topical.

In February, the U.S. leg of Ord’s book tour, which was scheduled for the spring and was to include stops at Stanford, M.I.T., and Princeton, was cancelled. “The Precipice” was published in the United Kingdom on March 5th; two weeks later, Ord was sheltering in place at home. His wife, Bernadette Young, an infectious-disease specialist at John Radcliffe Hospital, in Oxford, began working overtime, while he cared for their daughter, Rose, who was then five. “I’d already known that, during a crisis, the unthinkable can quickly become the inevitable,” Ord told me, earlier this year. “But, despite having this intellectual knowledge, it was still quite something to see such a thing unfold before my eyes.”

For someone with Ord’s interests, living through a pandemic is an opportunity to contemplate alternate histories. What might have happened in a world in which *COVID-19* didn’t exist, or was handled differently? What if the virus had been more deadly? Ord’s book reckons with these divergences on a grand scale, considering both the grim futures that await us if existential threats to humanity aren’t addressed and the far more promising outcomes that become possible if they are. Ord has given the name “the precipice” to our current phase of history, which, he writes, began at 11:29 A.M. Coordinated Universal Time, on July 16, 1945—the moment of the Trinity test, when the first nuclear bomb was detonated. It will end, he suggests, with either a shared global effort to insure humanity’s continued survival or the extinction of our species.

[Read \*The New Yorker’s\* complete news coverage and analysis of the coronavirus pandemic.](#)

Ord places the risk of our extinction during the twenty-first century at one in six—the odds of an unlucky shot in Russian roulette. Should we manage to avoid a tumble off the precipice, he thinks, it will be our era’s defining achievement. The book catalogues many possible catastrophes. There are the natural risks we’ve always lived with, such as

asteroids, super-volcanic eruptions, and stellar explosions. “None of them keep me awake at night,” Ord writes. Then there are the large-scale threats we have created for ourselves: nuclear war, climate change, pandemics (which are made more likely by our way of life), and other novel methods of man-made destruction still to come. Ord is most concerned about two possibilities: empowered artificial intelligence unaligned with human values (he gives it a one-in-ten chance of ending humanity within the next hundred years) and engineered pandemics (he thinks they have a one-in-thirty chance of bringing down the curtain). The pandemic we are currently experiencing is the sort of event that Ord describes as a “warning shot”—a smaller-scale catastrophe that, though frightening, tragic, and disruptive, might also spur attempts to prevent disasters of greater magnitude in the future.

Unlike doomsday preppers who seem, on some level, to relish the idea of social breakdown, Ord believes in humanity’s potential for greatness. Ord was born in Melbourne, in 1979, and his interest in existential risk grew out of an early focus on ethics, health, and poverty. With his Oxford colleague Will MacAskill, he helped create a school of thought known as “effective altruism”: it holds that we should use reason and evidence to determine how, as individuals, we can best use our resources for the greater good. In 2009, Ord founded Giving What We Can, a society whose members pledge to donate at least ten per cent of their incomes, in the course of their careers, to charity. By his own accounting, Ord, who is now forty-one, has given away twenty-seven per cent of his lifetime earnings. He has pledged to donate the proceeds from “The Precipice” to organizations working on the problems that the book outlines.

A concern for existential risk seemed, to Ord, to be the next logical expansion of a broadening moral circle. If we can learn to value the lives of people in other places and circumstances equally to our own, then we can do the same for people situated at a different moment in time. Those future people, whose quality of life and very existence will be intimately affected by our choices today, matter as much as we do; from the perspective of our species, they are us and we are them. Ord compares humanity’s current situation to adolescence, a treacherous period when strength and desire outpace wisdom and self-control, and when one’s future life seems remote and unreal. According to fossil records, the typical lifetime of a mammalian species is a million years. “If we think of one million in terms of a single, eighty-year life,” he writes, then today humanity would be “sixteen years old; just coming into our power; just old enough to get ourselves in serious trouble.”

Teen-agers have parents, or older siblings, who can offer advice. But we are the first species to accumulate the power to destroy itself through conscious action. Ideally, we’d possess the collective self-restraint to limit our access to deadly technologies until we acquire the maturity to use them safely. Ord advocates the establishment of international agreements that might create such restraints by defining collective goals and setting out rules and timetables for achieving them. (The Montreal Protocol, which led to the phasing out of

ozone-depleting chemicals, beginning in 1989, is one model.) Looking forward, he proposes a tax on certain high-tech businesses, which could be used to fund initiatives that guard against the dangers of new technologies. (At the moment, he writes, entrepreneurs profit from technological advancements while society shoulders the existential risk they present; by his calculations, in any given year, the world's investment in protective technologies is dwarfed by its spending on ice cream.) He urges the investors, governments, and universities who support new technologies to prioritize those that protect humanity over those that present new ways to destroy it—a concept that the philosopher Nick Bostrom, the founder and director of the Future of Humanity Institute, calls “differential technological development.” It's not that we're powerless against existential risks, Ord argues; we just have to muster the collective will to counter them.

Contemplating such well-intentioned ideas, a reader already horrified by our inadequate efforts to address climate change might feel a pit in her stomach. During the pandemic, Ord has been heartened by the coöperative work of the world's scientists, and by the resilience of individuals and communities: in Oxford, flyers from local Good Samaritans have come through his mail slot, offering assistance to neighbors in need; a street near his own devised a color-coded system of Post-its, placed in windows—green means all is fine, yellow means someone inside is ill and self-isolating, and red is a plea for help. The responses of national governments, on the other hand—the inability of so many individual nations to work effectively on behalf of their own citizens and the global community—has worried him. The conflict in the United States between the Trump Administration and its own scientific advisers has been a particularly unsettling example of our self-defeating tendencies. “In ecological terms, it is not a human that is remarkable, but humanity,” Ord writes, in his book. And yet it's easy for humanity's aims to be undermined by the short-term political calculations and selfish desires of individuals. We are a blip, Ord argues, in a chain of human life that stretches back two hundred thousand years, and that could in theory extend for hundreds of thousands, if not millions, more. One would have to be ruinously shortsighted or grotesquely narcissistic to risk obliterating all of that. And yet that is how we often seem to be.

Ord is not the first philosopher to propose that the end is nigh. In 1951, Bertrand Russell wrote that the twentieth century would end in either human extinction, violent social collapse, or a unified world government. Stephen Hawking warned, in 2014, that artificial intelligence could spell our doom. Not all thinkers are alarmed by what they foresee. In “Notes from an Apocalypse: A Personal Journey to the End of the World and Back,” published this year, the Irish journalist Mark O'Connell recounts wondering whether human extinction would really be so bad: “Why was it so unthinkable that we ourselves—not necessarily tomorrow or the next day, but eventually—follow the same well-beaten trail toward oblivion as the dodo, the black rhinoceros, the passenger pigeon . . . and all the countless other species whom we ourselves had driven from the face of the earth?”

For many people, the idea of all of humanity disappearing one day in the future, when no one we now know and love is around to see it, isn't as alarming as the prospect of individual human suffering in the present. A study conducted last year, by three experimental psychologists at Oxford, found that individuals considered total human extinction to be only slightly worse than a catastrophic event that wipes out eighty per cent of the population. Respondents' opposition to extinction rose when they were asked to consider the specific consequences of all of human culture being extinguished forever. Still, without that prompt, it was hard to grasp how much worse total obliteration would be compared with merely cataclysmic death.

There is a term for this outlook: "scope neglect," the cognitive bias that makes it harder to understand the full scale of problems the larger those problems get. It's the struggle, as Ord puts it, to care ten times as much about something that's ten times more important than an alternative. "One of the aspects in which I'm an outlier is that I take scale really seriously, and always have," Ord told me. "You can see that all through the book, really, including taking the scale of the universe seriously."

As "Precipice" closes, Ord zooms out to the cosmos and, against the backdrop of its unfathomable vastness, asks us to grasp the scale of what we risk losing if the human story ends prematurely. He writes that, just as our early forebears, huddled around some Paleolithic fire, couldn't have imagined the creative and sensory experiences available to us today, we, too, are ill-equipped to conceive of what is possible for those who will follow us. Humanity's potential is worth preserving, he argues, not because we are so great now but because of the possibility, however small, that we are a bridge to something far greater. "How strange it would be if this single species of ape, equipped by evolution with this limited set of sensory and cognitive capacities, after only a few thousand years of civilization, ended up anywhere near the maximum possible quality of life," he writes. "I think that we have barely begun the ascent."

These days, Ord contemplates our possible futures from home, a Victorian terrace house, built in 1875, on the eastern side of Oxford. His daughter's alarm sounds at 7:30 A.M.; he wakes at 7:23, having determined that he needs exactly seven minutes to dress and pull his thoughts together. In the ground-floor kitchen, a bowl of cold-brew coffee steeps on the counter, beneath a cardboard cover. A table is cluttered with the ephemera of family life in lockdown: a fractal that Ord drew for Rose on a scrap of paper, to illustrate an impromptu home lesson on division; a card she made for her mother after a stint of night shifts ("Dear Mum, welcome back to normal"). This summer, Rose celebrated her sixth birthday at a local park. Ord staged an elaborate treasure hunt, involving riddles, maps, chocolates, and the many foreign coins for which he has no use during the pandemic.

Like the majority of schools in the U.K., Rose's has resumed in-person classes. Once breakfast is done, Ord walks her the few blocks to the entrance, then returns home to answer e-mails, write, and read through the day's news online. There is no backup

generator in the house, no stockpile of canned food; it is not practical or even advisable, Ord believes, for individuals to devote themselves to preparing for grand catastrophes. High-level risks demand high-level coordination. As individuals, we can do our part by giving to causes that support humanity's survival, and by having the conversations that coalesce, over time, into collective action. Much depends on what is imaginable. Ord points out that the possibility of nuclear apocalypse took time to feel real, growing easier to grasp as the Cold War progressed. Today, as forests burn and cities flood, the consequences of climate change are becoming more vivid. One of Ord's goals is to help us imagine better, sooner. "At the moment, it seems that we have to wait about forty years after a threat is scientifically recognized before culture gets to it. And I think that's too slow," he told me, over Skype. Experience is a powerful teacher—countries that previously battled SARS and MERS have had comparatively better responses to the pandemic—and yet warning shots can only alarm us for so long. The challenge is to adopt a new frame of mind in which distant threats aren't confused with impossible ones.

Ord proposed "a simple piece of mathematics" that, for him, "has helped make it all visceral": even a once-in-a-century catastrophe has a five-per-cent chance of happening during any five-year Parliamentary term in the U.K., and a four-per-cent chance of occurring during any single Presidential term in the U.S. This risk calculus may seem commonsensical to us now, because we are living through such an event. Still, he said, "we'll forget it!" He laughed. "We have five years or so where we are really protected from things like this, because we're conscious of it," he went on. "Beyond that, it relies on us to create institutions." It's for this reason, especially, that being consumed with dread serves no purpose. Fear must be motivating, or it's pointless. Safeguarding the future requires believing in one.

---

## More on the Coronavirus

---

- To protect American lives and revive the economy, Donald Trump and Jared Kushner should listen to Anthony Fauci rather than trash him.
- We should look to students to conceive of appropriate school-reopening plans. It is not too late to ask what they really want.
- A pregnant pediatrician on what children need during the crisis.
- Trump is helping tycoons who have donated to his reelection campaign exploit the pandemic to maximize profits.
- Meet the high-finance mogul in charge of our economic recovery.
- The coronavirus is likely to reshape architecture. What kinds of space are we willing to live and work in now?