

Viewpoint: Existential ethics — Pondering the extinction of the global human population

It's an ominous sign of the times that human extinction is an increasingly common topic of discussion. If you search for 'human extinction' in the Google Ngram Viewer, which charts the frequency of words in Google's vast corpora of digitised books, you'll [see](#) that it's rarely mentioned before the 1930s. This changes slightly after the Second World War – the beginning of the Atomic Age – and then there's a sudden spike in the 1980s, when Cold War tensions rose, followed by a decline as the Cold War came to an end. Since the 2000s, though, the term's frequency has risen sharply, perhaps exponentially.

This is no doubt due to growing awareness of the climate crisis, as well as the various dangers posed by emerging technologies, from gene editing to artificial intelligence, the obstacle course of existential hazards before us seems to be expanding. Indeed many scholars have argued that, to [quote](#) Noam Chomsky, the overall risk of extinction this century is 'unprecedented in the history of *Homo sapiens*'. Similarly, Stephen Hawking declared in 2016 that 'we are at the most dangerous moment in the development of humanity'. Meanwhile, the Doomsday Clock, maintained by the venerable *Bulletin of the Atomic Scientists*, is currently showing that it's a mere 90 seconds before midnight, or doom – the closest it's ever been since this clock was created in 1947. The direness of our situation is also widely acknowledged by the general public, with one [survey](#) reporting that a whopping 24 per cent of people in the United States, the United Kingdom, Canada and Australia 'rated the risk of humans being wiped out' within the next 100 years 'at 50 per cent or greater'.

Follow the latest news and policy debates on sustainable agriculture, biomedicine, and other 'disruptive' innovations. Subscribe to our newsletter.

[SIGN UP](#)

But so *what* if we're wiped out? What does it matter if *Homo sapiens* no longer exists? The astonishing fact is that, despite acquiring the ability to annihilate ourselves back in the 1950s, when thermonuclear weapons were invented, very few philosophers in the West have paid much attention to the ethics of human extinction. Would our species dying out be bad, or would it in some way be good – or just neutral? Would it be morally wrong, or perhaps morally right, to cause or allow our extinction to occur? What arguments could support a 'yes' or 'no' answer?

These are just some of the questions that I [place](#) within a field called 'existential ethics', which, as noted, has been largely ignored by the philosophical community. This is a real shame for several reasons: first, even if you *don't* think our extinction is likely this century, reflecting on the questions above can provide clarity to a wide range of philosophical issues. The fact is that existential ethics touches upon some of the most fundamental questions about value, meaning, ethics and existence, which makes meditating on why our species might – or might not – be worth saving a very useful exercise. Second, if you *do* agree with Chomsky, Hawking and the *Bulletin of the Atomic Scientists* that our extinction is more probable now than in centuries past, shouldn't we want to know whether, and why, tumbling into the eternal grave would be right or wrong, good or bad, better or worse? Here we are, inches away from the precipice, tempting the

same fate that swallowed up the dinosaurs and the dodo, with hardly anyone thinking seriously about the ethical implications of this possibility. Surely this is a situation that we should not only rectify, but do so with a degree of moral urgency – or so I would argue.

This points to another question: why exactly has existential ethics been so neglected? Why has it languished in relative obscurity while so many other fields – machine ethics, business ethics, animal ethics, bioethics, and so on – have become thriving areas of research over the past several decades? One explanation is that philosophers have, in general, failed to appreciate just how rich and complicated the topic is. For example, the question ‘Would human extinction be bad?’ looks simple and straightforward, yet it conceals a treasure trove of fascinating complexity. Consider that ‘human’ and ‘extinction’ can be defined in many different, equally legitimate ways.

Most people intuitively equate ‘human’ with our species, *Homo sapiens*, yet scholars sometimes use the word to mean ‘*Homo sapiens* and whatever descendants we might have’. On the latter definition, *Homo sapiens* could disappear completely and forever without human extinction having occurred. Indeed, one way to ‘go extinct’ would be to evolve into a new posthuman species, something that will inevitably occur over the next million years if only because of Darwinian evolution. Would this be bad? Or we might ‘go extinct’ by replacing ourselves with, say, a population of intelligent machines. Some would see this as a dystopian outcome, though others have recently [argued](#) that we should want it to happen. In his [book](#) *Mind Children* (1988), the computer scientist Hans Moravec, for instance, not only views this kind of extinction as desirable, but hopes to actively bring it about. He thus holds that the extinction of *Homo sapiens* would constitute a great tragedy – *unless* it were to coincide with the creation of machinic replacements, in which case it would be very good.

Would it be bad for a child to die prematurely? What about an elderly person, or someone middle-aged?

On my count, there are at least six distinct types of extinction that are relevant to existential ethics, although for our purposes we can focus on what I call the ‘prototypical conception’, whereby *Homo sapiens* disappears entirely and forever without leaving behind any successors. In other words, our extinction marks a complete and final end to the human story, which is why I label it ‘final’ human extinction.

So, this is one bundle of complexity hidden behind what appears to be a simple question: would human extinction be bad? Whenever someone asks this, the first thing you should do is reply: what do you mean by ‘human’? And which type of ‘extinction’ are you talking about? Once you’re clear on these issues, there’s a second complication to navigate. Consider the following, which isn’t a trick question: would it be bad for a child to die prematurely? What about an elderly person, or someone in the middle years of life? Would it be bad if their deaths were preceded by lots of physical suffering, anxiety or fear? My guess is that you’ve answered ‘Yes, obviously!’ to these questions. If so, then you’ll have to believe that human extinction *would be very bad* if caused by a worldwide catastrophe. Indeed, since an extinction-causing catastrophe would literally kill *everyone* on Earth, it would be the absolute worst disaster possible. There is no disaster with a higher body count. This is so widely accepted – only the most sadistic, ghoulish person would reject it – that I call it the ‘default view’. Do we all agree, then, that extinction would be bad?

Is there nothing left to say?

Hardly! But to see what's left, it's crucial to distinguish between two aspects of extinction: first, there's the process or event of Going Extinct and, second, there's the state or condition of Being Extinct. You could draw a rough analogy with individual death: many friends tell me that they aren't afraid of being dead – after all, they won't be around any longer to worry about the world, to suffer, or even to experience FOMO (fear of missing out). They are, however, terrified of the pain that dying might entail. They're scared about what might lead to death, but not about their subsequent nonexistence. Not everyone will agree with this, of course – some find the thought of having perished quite frightening, as do I. The point is that there are two features of extinction that one might identify as *sources* of its badness: maybe extinction is bad because of how Going Extinct unfolds, or perhaps there's something about Being Extinct that *also* contributes some badness. Almost everyone – except sadists and ghouls – will agree that, if Going Extinct involves a catastrophe, then our extinction would be very bad in this sense. Indeed, this is more or less trivially true, given that the definition of 'catastrophe' is, to quote the Merriam-Webster dictionary, 'a momentous tragic event ranging from extreme misfortune to utter overthrow or ruin'. Yet there is an extraordinary array of positions that go well beyond this simple, widely accepted idea.

For example, imagine a scenario in which every person around the world decides not to have children. Over the next 100 years or so, the human population gradually dwindles to zero, and our species stops existing – not because of a violent catastrophe, but because of the freely chosen actions of everyone on the planet. Would *this* be bad? Would there be something bad about our extinction even if Going Extinct were completely voluntary, didn't cut anyone's life short, and didn't introduce any additional suffering?

Some philosophers would vociferously respond: 'Yes, because there's something bad about Being Extinct totally independent of how Going Extinct takes place.' In defending this position, such philosophers would point to some further loss associated with Being Extinct, some opportunity cost arising from our species no longer existing. What might these further losses, or opportunity costs, be? In his [book](#) *Reasons and Persons* (1984), the moral theorist Derek Parfit proposed two answers: on the one hand, Being Extinct would prevent human happiness from existing in the future and, given that Earth will remain habitable for another ~1 billion years, the total amount of future happiness could be enormous. All of this happiness would be lost if humanity were to cease existing, even if the cause were voluntary and harmless. That would be bad. On the other hand, the future could witness extraordinary developments in science, the arts and even morality and, since Being Extinct would preclude such developments, it would – once again – be bad, independent of how Going Extinct unfolds. This led Parfit to declare that the elimination of humanity, however it happens, 'would be by far the greatest of all conceivable crimes', an idea first articulated, using almost the exact same words, by the utilitarian Henry Sidgwick in *The Methods of Ethics* (1874).

I classify this as the 'further-loss view' for the obvious reason that it hinges on the idea of further losses arising from our nonexistence. Many contemporary philosophers take Parfit's side, including advocates of an ethical framework called '[longtermism](#)' recently promoted by the Oxford philosopher William MacAskill in his [book](#) *What We Owe the Future* (2022). In fact, not only do longtermists see Being Extinct as one source of extinction's badness, but most would argue that Being Extinct is the worst aspect of our extinction by a long shot. This is to say, even if Going Extinct were to involve horrendous amounts of suffering, pain, anguish and death, these harms would be utterly dwarfed by the further loss of all future

happiness and progress. Here's how the philosophers Nick Beckstead, Peter Singer and Matt Wage (the first of whom [laid](#) the foundations for longtermism) express the idea in their [article](#) 'Preventing Human Extinction' (2013):

One very bad thing about human extinction would be that billions of people would likely die painful deaths. But in our view, this is, by far, not the worst thing about human extinction. The worst thing about human extinction is that there would be no future generations.

Other philosophers – myself included – reject this further-loss view. We argue that there can't be anything bad about Being Extinct because there wouldn't be anyone around to experience this badness. And if there isn't anyone around to suffer the loss of future happiness and progress, then Being Extinct doesn't actually harm anyone. To quote Jonathan Schell's magisterial [book](#) *The Fate of the Earth* (1982): 'although extinction might appear to be the largest misfortune that mankind could ever suffer, it doesn't seem to happen *to anybody*.' The reason is that 'we, the living, will not suffer it; we will be dead. Nor will the unborn shed any tears over their lost chance to exist; to do so they would have to exist already.' Similarly, the philosopher Elizabeth Finneron-Burns [asks](#): 'If there is no form of intelligent life in the future, who would there be to lament its loss?'

If Going Extinct involves a catastrophe, then obviously extinction would be bad

I call this the 'equivalence view', because it claims that the badness or wrongness of extinction comes down *entirely* to the badness or wrongness of Going Extinct. This means that, if there isn't anything bad or wrong about Going Extinct, then there isn't anything bad or wrong about extinction – full stop. Applying this to the scenario mentioned above, since there's nothing bad or wrong about people not having children, there wouldn't be anything bad or wrong about our extinction if caused by everyone choosing to be childless. The answer one gives to 'Would our extinction be bad?' is *equivalent* to the answer one gives to 'Would this or that way of Going Extinct be bad?'

There are some important similarities and differences between further-loss and equivalence views. Both accept the default view, of course, although they diverge in a crucial way about whether the default view says everything there is to say about extinction. Philosophers like myself believe that the default view is the whole story about why extinction would be bad or wrong. If Going Extinct involves a catastrophe, then obviously extinction would be bad, since an extinction-causing catastrophe 'would entail maximum mortality, likely preceded by unprecedented human suffering,' to [quote](#) the philosopher Karin Kuhlemann, who (along with Finneron-Burns and myself) accepts the equivalence view. But if Going Extinct doesn't cause lots of suffering and death, if it happens through some peaceful, voluntary means, then extinction wouldn't be bad or wrong. Those who accept a further-loss view will strenuously disagree: they claim that, if Going Extinct entails suffering and death, this would be one source of extinction's badness. But it wouldn't be the only source, nor even the biggest source: the subsequent state of Being Extinct would also be very bad, because of all the attendant further losses.

There's another thought experiment that helps to foreground the major disagreement between these positions. Imagine two worlds, A and B. Let's say that world A contains 11 billion people and world B

contains 10 billion. Now, a terrible disaster rocks both worlds, killing exactly 10 billion in each. There are two questions we can ask about what happens here. The first is pretty straightforward: how many events occur in world A versus world B? Most will agree that, on the highest level of abstraction, one event happens in world A – the loss of 10 billion people in a sudden disaster – while *two* events happen in world B – the loss of 10 billion people *plus* the extinction of humanity, since the total population was 10 billion.

That's the first question. The second is whether this extra event in world B – the extinction of humanity – is morally relevant. Does it matter? Does it somehow make the disaster of world B worse than the disaster of world A? If a homicidal maniac named Joe causes both disasters, does he do something extra-wrong in world B?

phil torres

Image not found or type unknown

Credit: Phil Torres

If you accept a further-loss view, then you're going to say: 'Absolutely, the catastrophe in world B is much worse, and hence Joe did something extra-wrong in B compared with A.' But if you accept the equivalence view, you'll say: 'No, the badness or wrongness of these catastrophes is identical. The fact that the catastrophe causes our extinction in world B is irrelevant, because Being Extinct is not itself a

source of badness.’ The intriguing implication of this answer is that, according to equivalence views, human extinction does not present a unique moral problem. There is nothing special about extinction itself: it doesn’t introduce a distinct moral conundrum; the badness or wrongness of extinction is wholly reducible to how it happens, beyond which there is nothing left to say. Since Being Extinct is not bad, the question of whether our extinction is bad depends entirely on the details of Going Extinct. Advocates of the further-loss view will of course say, in response, that this is deeply misguided: extinction does pose a unique moral problem. Why? Precisely because it would entail some further losses – things of value, perhaps immense value, that wouldn’t be lost forever if our species were to continue existing, such as future happiness and progress (though there are many other further losses that one could point to, which we won’t discuss here). For them, assessing the badness or wrongness of extinction requires looking at both the details of Going Extinct *and* the various losses that Being Extinct would involve.

We’re starting to see the lay of the land at this point, the different positions that one could take within existential ethics. But there’s another option that we haven’t yet touched upon: whereas further-loss views say that Being Extinct is bad, and equivalence views assert that Being Extinct is not bad, you might also think, for some reason or other, that Being Extinct would be *less bad*, or maybe even positively *good*. This points to a third family of views that I call ‘pro-extinctionist views’, which a surprising number of philosophers have accepted. Right away, it’s important to be clear about this position: virtually all pro-extinctionists accept the default view. They would agree that a catastrophic end to humanity would be horrible and tragic, and that we should try to avoid this. No sane person would want billions to die. However, pro-extinctionists would add that the outcome of Being Extinct would nonetheless be better than Being Extant – ie, continuing to exist. Why? There are many possible answers. One is that, by no longer existing, we would prevent future human suffering from existing too, which would be less bad, or more good, than our current situation. There are several ways of thinking about this.

The first concerns the possibility that the future contains huge amounts of human suffering. According to the cosmologist [Carl Sagan](#) – who defended a further-loss view in his writings – if humanity survives for another 10 million years on Earth, there [could](#) come to exist some 500 trillion people. This number is much larger if we colonise space: the philosopher [Nick Bostrom](#) in 2003 [put](#) the figure at 10^{23} (that’s a 1 followed by 23 zeros) biological humans per century within the Virgo Supercluster, a large agglomeration of galaxies that includes our own Milky Way, although the number rises to 10^{38} if we include digital people living in virtual-reality computer simulations. Subsequent number-crunching in Bostrom’s [book](#) *Superintelligence* (2014) found that some 10^{58} digital people could exist within the Universe as a whole. What motivated these calculations from Sagan and others was the further-loss view that the nonexistence of all these future people, and hence all the value they could have created, is a major opportunity cost of Being Extinct. However, one could flip this around and argue that, even if these future people have lives that are overall worthwhile, the suffering they will unavoidably experience during the course of these lives could still add up to be, in absolute terms, very large. However, if we go extinct, none of this suffering will exist.

The second consideration is based on the idea, which I find very plausible, that there are some kinds of suffering that no amount of happiness could ever counterbalance. Imagine, for example, that you need to undergo a painful surgical procedure that will leave you bedridden for a month. You go through with the

surgery and, after recovering, you live another 50 very happy years. The surgery and recovery process might have been dreadful – not something you'd want to relive – but you might still say that it was 'worth it'. In other words, the decades of happiness you experienced after surgery counterbalanced the pain you had to endure. But now consider some of the worst things that happen in the world: abuse of children; genocides, ethnic cleansings, massacres; people tortured in prisons; and so on. Ask yourself if there's any amount of happiness that can make these things 'worth it'. Is there any heap of goodness large enough to counterbalance such atrocities? If you think of particular historical horrors, you might find the question outright offensive: 'No, *of course* that genocide can't somehow be counterbalanced by lots of happiness experienced by other people elsewhere!' So, the argument goes, since continuing to exist carries the risk of similar atrocities in the future, it would be better if we didn't exist at all. The good things in life just aren't worth gambling with the bad things.

If you think that Being Extinct is better than Being Extant, how should we actually bring this about?

Considerations like these are why pro-extinctionists will argue that the scenario of world B above is actually preferable to the scenario of world A. Here's what they might say about this: 'Look, there's no question that the disaster in world B is terrible. It's absolutely wretched that 10 billion people died. I say that unequivocally because I, like everyone else, accept the default view! However, for precisely the *same reason* that I think this disaster is very bad, I also maintain that the second event in world B – the extinction of humanity – makes this scenario better than the scenario of A, as it would mean no more future suffering. At least world B's disaster has an upside – unlike in world A, where 10 billion die *and* future people will suffer.'

The major practical problem for pro-extinctionism is one of getting from here to there: if you think that Being Extinct is better than Being Extant, how should we actually bring this about? There are three main options on the menu: [antinatalism](#), whereby enough people around the world stop having children for humanity to die out; pro-mortalism, whereby enough people kill themselves for this to happen; and omnicide, whereby someone, or some group, takes it upon themselves to kill everyone on Earth. (The word 'omnicide' was defined in 1959 by the theatre critic Kenneth Tynan as 'the murder of everyone', although, curiously, a chemical company had earlier trademarked it as the name for one of its insecticides.)

A very small number of pro-extinctionists have advocated for omnicide – mostly fringe environmental extremists who see humanity as a 'cancer' on the biosphere that must be excised. This was defended in an article from the *Earth First! Journal* titled 'Eco-Kamikazes Wanted' (1989) and later [endorsed](#) by a group called the Gaia Liberation Front. But omnicide also appears to be an implication of 'negative utilitarianism', an ethical theory that, in its strongest form, asserts that the *only* thing that matters is the reduction of suffering. As the philosopher R N Smart [noted](#) in 1958, this means that one should become a 'benevolent world-exploder' who destroys humanity to eliminate all human suffering, which Smart described as patently 'wicked'. However, the Oxford philosopher Roger Crisp (who isn't a negative utilitarian) recently [contended](#) that if you were to discover a huge asteroid barrelling toward Earth, and if you could do something to redirect it, you should seriously consider letting it slam into our planet,

assuming it would kill everyone immediately. This would, in effect, be omnicide by inaction rather than action, although Crisp never says that you should definitely do this, only that you should think very hard about it, given that Being Extinct 'might' be 'good', a tentative conclusion based on the possibility that some suffering cannot be counterbalanced by any amount of happiness.

Other pro-extinctionists have advocated for both antinatalism and pro-mortalism. An example is the 19th-century German pessimist Philipp Mainländer, who argued that we should not just refrain from procreating, but never have sex in the first place – in other words, we should all remain virgins. He also endorsed suicide, and indeed shortly after receiving the first copies of Volume I of his magnum opus *The Philosophy of Redemption* (1876), he placed them on the floor, stood on top of them, stepped off, and hanged himself. He was only 34 years old, and wasn't the only person in his family to commit suicide: his older brother and sister did, too.

Most pro-extinctionists, though, have held that the only morally acceptable route to extinction is antinatalism – refusing to have children. The best-known advocate of this position today is [David Benatar](#), who argues in his [book](#) *Better Never to Have Been* (2006) that our collective nonexistence would be positively good, since it would mean the absence of suffering, and the absence of suffering is good. On the flip side, he notes that even though Being Extinct would entail the loss of future happiness, this wouldn't be bad because there'd be no one around to suffer such a loss. Hence, Being Extinct corresponds to a good (no suffering) and not-bad (no happiness) situation, which contrasts with our current state, Being Extant, which involves the presence of both happiness (good) and suffering (bad). He concludes that since a good/not-bad situation is obviously better than a good/bad situation, we should strive to bring about our extinction – by remaining childless.

My own view is a complicated mix of these considerations, which push and pull in diametrically opposite directions. The first thing I would emphasise is that our minds are totally ill-equipped to grasp just how awful Going Extinct would be if caused by a catastrophe. In a fascinating [paper](#) from 1962, the German philosopher [Günther Anders](#) declared that, with the invention of nuclear weapons, we became 'inverted Utopians'. Whereas 'ordinary Utopians are unable to actually produce what they are able to visualise, we are unable to visualise what we are actually producing,' ie, the possibility of self-annihilation. That is to say, there's a yawning chasm – he called it the 'Promethean gap' – between our capacity to destroy ourselves and our ability to feel, comprehend and imagine the true enormity of catastrophic extinction. This dovetails with a cognitive-emotional phenomenon called 'psychic numbing', which the psychologist Paul Slovic [describes](#) as the

inability to appreciate losses of life as they become larger. The importance of saving one life is great when it is the first, or only, life saved, but diminishes marginally as the total number of lives saved increases. Thus, psychologically, the importance of saving one life is diminished against the background of a larger threat – we will likely not 'feel' much different, nor value the difference, between saving 87 lives and saving 88, if these prospects are presented to us separately.

Now imagine that the number isn't 87 or 88 deaths but *8 billion* – the total human population on Earth

today. As a *Washington Post* article from 1947 quotes Joseph Stalin as saying, 'if only one man dies of hunger, that is a tragedy. If millions die, that's only statistics,' which is often truncated to: 'A single death is a tragedy, a million deaths are a statistic.' The point is that an extinction-causing catastrophe would be horrendous to a degree that our puny minds cannot even begin to grasp, intellectually or emotionally, although simply understanding this fact can help us better calibrate assessments of its badness, by compensating for this deficiency. In my opinion, the terribleness of such catastrophes – events with the highest possible body count – is more than enough reason to prioritise, as a species, projects aimed at reducing such risk. The default view about Going Extinct is even more profound, and its implications even more compelling, than one might initially believe.

But I also think that Being Extinct would be regrettable for many reasons. As Mary Shelley wrote in her novel *The Last Man* (1826) – one of the very first books to address the core questions of existential ethics – without humanity there would be no more poetry, philosophy, painting, music, theatre, laughter, knowledge and science, and that would be very sad. I feel the pull of this sentiment, and find myself especially moved by the fact that our extinction would bring the transgenerational enterprise of scientific understanding to a screeching halt. It would be a monumental shame if humanity had popped into existence, looked around at the Universe in puzzlement and awe, pondered the Leibnizian question of why there is something rather than nothing, and then vanished into the oblivion before knowing the answer. Maybe the answer is unknowable, but even discovering this fact could provide a degree of intellectual satisfaction and psychological closure, an 'ah-ha' moment that relieves and vindicates one's prior frustration.

There would be no more love, but there would also be no more heartbreak

This is a version of what's called the 'argument from unfinished business' and, while many people aren't persuaded by it, I am. However, I don't see it as a specifically *moral* position, but would instead classify it as a *non-moral* further-loss view. This matters because we typically see moral claims as having much more force than non-moral ones. There's a big difference between saying 'You shouldn't eat chocolate ice-cream because vanilla is better' and 'You shouldn't drown kittens in your bathtub for fun.' The first expresses a mere aesthetic preference, and hence carries much less weight than the second, which expresses a moral proposition. So, the unfinished business argument that I accept isn't very weighty. Other considerations – especially moral considerations – could easily override this personal preference of mine.

This leads directly to the question of whether Being Extinct might be, in some morally relevant way, better than Being Extant. Here I find myself sympathetic with the sentiments behind pro-extinctionism. Although Being Extinct would mean no more happy experiences in the future, no more poetry, painting, music and laughter, it would also guarantee the absence of the worst atrocities imaginable – child abuse, genocides, and the like. There would be no more love, but there would also be no more heartbreak, and I suspect many would agree, upon reflection, that heartbreak can hurt worse than love feels good. It's also entirely possible that scientific and technological advancements make unspeakable new forms of suffering feasible. Imagine a world in which radical life-extension technologies enable totalitarian states to keep people alive in torture chambers for indefinitely long periods of time – perhaps hundreds or thousands of

years. Is our continued existence worth *risking* such agony and anguish? Those who answer 'yes' put themselves in the awkward position of saying that torture, child abuse, genocide and so on are somehow 'worth it' for the good things that might come to exist alongside them.

Considerations arising from our impact on the natural world, and the way we treat our fellow creatures on Earth, also support the pro-extinctionist view. Who can deny that humanity has been a force of great evil by obliterating ecosystems, razing forests, poisoning wildlife, polluting the oceans, hunting species to extinction and tormenting domesticated animals in factory farms? Without humanity, there would be no more humanity-caused evils, and surely that would be very good.

So where does this leave us? I'm inclined to agree with the philosopher Todd May, who argued in *The New York Times* in 2018 that human extinction would be a mixed bag. I reject the further-loss views of Parfit and the longtermists, and accept the equivalence view about the badness of extinction. But I'm also sympathetic with aspects of pro-extinctionism: all things considered, it's hard to avoid the conclusion that Being Extinct might, on balance, be positive – even though I'd be saddened if the business of revealing the arcana of the cosmos were left forever unfinished. (The sadness here, though, is not really of the moral kind: it's the same sort of sadness I'd experience if my favourite sports team were to lose the championship.)

That said, the horrors of Going Extinct in a global catastrophe are so enormous that we, as psychically numb inverted Utopians, should do everything in our power to reduce the likelihood of this happening. On my view, the only morally permissible route from Being Extant to Being Extinct would be voluntary antinatalism, yet as many antinatalists themselves have noted – such as Benatar – the probability of everyone around the planet choosing not to have children is approximately zero. The result is a rather unfortunate predicament in which those who agree with me are left anticipatorily mourning all the suffering and sorrow, terrors and torments that await humanity on the road ahead, while simultaneously working to ensure our continued survival, since by far the most probable ways of dying out would involve horrific disasters with the highest body count possible. The upshot of this position is that, since there's nothing uniquely bad about extinction, there's no justification for spending disproportionately large amounts of money on mitigating extinction-causing catastrophes compared with what have been called 'lesser' catastrophes, as the longtermists would have us do, given their further-loss views. However, the bigger the catastrophe, the worse the harm, and for *this reason alone* extinction-causing catastrophes should be of particular concern.

My aim here isn't to settle these issues, and indeed our discussion has hardly scratched the surface of existential ethics. Rather, my more modest hope is to provide a bit of philosophical clarity to an immensely rich and surprisingly complicated subject. In a very important sense, virtually everyone agrees that human extinction would be very bad. But beyond this default view, there's a great deal of disagreement. Perhaps there are other insights and perspectives that have not yet been discovered. And maybe, if humanity survives long enough, future philosophers will discover them.

Émile P Torres is a PhD candidate in philosophy at Leibniz Universität Hannover in Germany. Their writing has appeared in *Philosophy Now*, *Nautilus*, Motherboard and the *Bulletin of the Atomic Scientists*, among others. Follow Émile on Twitter [@xriskology](#)

A version of this article was originally posted at [Aeon](#) and has been reposted here with permission. Any reposting should credit the original author and provide links to both the GLP and the original article. Find Aeon on Twitter [@aeonmag](#)