Utilitas (2026), 1–15
doi:10.1017/S0953820826100351

CAMBRIDGE UNIVERSITY PRESS

ARTICLES

Longtermism, Technology, and Well-Being

Alexander Guerrero

Philosophy, Rutgers University New Brunswick, USA
Email: alex.guerrero@rutgers.edu

(Received 27 September 2024; revised 9 June 2025; accepted 1 April 2026)

Abstract

Longtermists such as William MacAskill maintain that if we are concerned about human well-being and the longterm future, we should embrace continued technological advancement and development, and that, indeed, one of the most significant threats to longterm human well-being and the well-being of other sentient creatures is technological stagnation – remaining at roughly the present technological level. This paper argues that we should be much more skeptical regarding claims that link continued technological development with positive contributions to well-being. We should take seriously that even if technological development has contributed significantly to improvements in well-being in the past, that we may reach or have already reached a point at which future technological development no longer does so, and that, in that event, thinking about longterm ethical concerns actually recommends developing the capability to slow down, stop, or even roll back (some) technological development.

Keywords: Longtermism; well-being; technology; William MacAskill

I

William MacAskill's *What We Owe the Future* (2022) tells a story about human history, civilizational, technological, and moral progress, where we have been, where we are now, and where we might go, depending on our choices now. I want to suggest that the story he tells has an inadequate theory of well-being at its center, that the problems with this theory infect his set of longterm concerns and responses, and, in particular, that a capitalist, consumerist, technologist ideology leads him to embrace claims about well-being and progress and the role of technology in both that should be questioned and rejected. Doing so recommends a very different path forward for the longtermist than the one that he (often) seems to favor. The focal difference here ends up concerning technology and its contribution to well-being, but I will work my way toward that.

What We Owe the Future is a remarkable, expansive manifesto. It draws on the work of many researchers and co-authors and area experts to situate a philosophical perspective in broader historical context, and to consider the implications of that philosophical perspective. That perspective consists of several components:

© The Author(s), 2026. Published by Cambridge University Press. This is an Open Access article, distributed under the terms of the Creative Commons Attribution licence (https://creativecommons.org/licenses/by/4.0/), which permits unrestricted re-use, distribution and reproduction, provided the original article is properly cited.

https://doi.org/10.1017/S0953820826100351 Published online by Cambridge University Press

> *No discounting*: "Impartially considered, future people should count for no less, morally, than the present generation."
> *Huge number of people*: "There may be a huge number of future people."
> *Wide life quality range*: "Life, for them, could be extraordinarily good or inordinately bad."
> *Causal effect*: "we really can make a difference to the world they inhabit."[1]

There are questions that we might ask about each of these, and questions about their cumulative implications.

One set of questions we might ask concerns whether these four claims together imply (or recommend) *strong longtermism*, the view (defended by Greaves and MacAskill 2025) that impact on the far future is the *most* important feature of our actions today. Strong longtermism might seem to follow straightforwardly from the sheer number of people, the no discounting, and the suggestion that we can make a significant difference in the quality of life that those trillions of people will have. Those who accept the above four claims might be drawn to a view we could call *swamping*, where the interests of all the trillions and trillions (or octillions and octillions) of possible future people simply swamp the interests of those billions of us currently living.

The practical implications of accepting strong longtermism are debatable. MacAskill and other longtermists maintain that there will be substantial *overlap* between what strong longtermism will recommend by way of action and what weak longtermism or even some much more modest neartermism will recommend: worrying about nuclear war, climate change, pandemic response, protecting democratic institutions, and responding to concerns about the social and other effects of developing technology like AI and potentially AGI ("artificial general intelligence"). Call this the *overlap* thesis.

The overlap thesis is convenient; it would be very hard to sell an ethical theory that says that what we should be doing now is, for example, completely ignoring climate change, pandemic response, protecting democratic institutions, and worrying about nuclear war in favor of working exclusively on building the technology to allow interstellar travel and space colonization.

But we might still be worried about places of non-overlap. It might be, for example, that much of the relatively impoverished world is simply too far away from having the educational and technological infrastructure to allow most people in those places to contribute to addressing central concerns of the extremely long term like AGI, bio-engineered pandemics, and space travel. Just given the numbers involved, it would seem like ignoring the plight of the world's worst off billion people would be justified and even required to redirect resources to help the future's trillions and trillions of people. That is a harder sell. It puts great pressure on our confidence in the "no discounting" commitment. We – also, ask who the "we" is supposed to be – are going to watch a lot of people starve and suffer, while investing a ton in fancy technological enterprises. It also highlights that even if we accept the four parts of the longtermist perspective articulated above, there is a fifth, epistemic commitment required by the perspective for it to have any interesting practical implications:

> *Epistemic tractability ("cluefulness")*: we can, from our current vantage point, make intelligent, reasonable, rational, epistemically justified assessments about what we

[1]MacAskill (2022: 5).

https://doi.org/10.1017/S0953820826100351 Published online by Cambridge University Press

> should do now to make things better for those many trillions upon trillions of future people who might exist more than 100,000 years from now.

Presumably, any plausible ethical theory suggests that we should think about the consequences of our actions – and not just over the next day or week, but even over the next year, decade, or even many generations into the future.[2] The number 100,000 was chosen arbitrarily; it might be that the epistemic difficulty comes much sooner – perhaps even 100 years out, or 1000 years, or 10,000 years out. The important point is that wherever we set that number, the longtermist claim is that there will or might be trillions and trillions of more people later, depending on what we do and on what happens, and that the future interests of those people matter.

There are many reasons to question epistemic tractability, and many arguments we might make that put pressure on it. Some of these might point out how badly we have done at predicting or anticipating even very near-term events. Others might highlight how unhelpful and misguided the suggestions of those living even a few hundred years ago would have been for us now. MacAskill takes up these questions, albeit somewhat briefly, in the book.

For MacAskill, the main defense of epistemic tractability comes in the form of the extensive and exhaustive discussions in the book about what we should be concerned about: bad value lock-in, extinction from engineered pathogens, extinction from nuclear war, extinction from bad AGI, civilizational collapse, and, maybe most significantly for MacAskill, technological stagnation (as technological advance can help deal with all manner of problems, including many we can't now anticipate). The focus on these concerns takes up the heart of the book, and results in a list of "what to do now" recommendations at the end of the book. It is a "proof is in the pudding" effort, and MacAskill, along with the many researchers, have done a lot of work to make their case.

Some might directly address the question of epistemic tractability and raise worries about cluelessness of various forms. Others might challenge "no discounting" on various grounds.

I want to raise a more direct challenge to the specific recommendations of things to be worried about and what to do and focus on, taking on board a moral requirement to take seriously the interests of those who might exist in the very very long term. I will suggest that longtermist considerations actually recommend a very different path forward than the one that MacAskill puts in view for us. This indirectly provides some support for cluelessness – MacAskill says we should do X; I argue that we should do not-X and instead Y; and it is (perhaps) hard to see how we should decide between us. But it also might just provide support for a kind of hedging strategy, if a bifurcated or n-furcated path can be made to work. I will take up those questions toward the end.

[2]That is why longtermism is not thought to depend on a background commitment to any form of optimizing consequentialism (the view that in acting we ought to bring about the *best* consequences), nor a commitment to any form of exclusive consequentialism (the view that the *only* thing relevant to the moral assessment of actions are their consequences). An ethical view that completely rejected the relevance of consequences to moral assessment would be unaffected by longtermist considerations, but those ethical views are implausible. The sheer magnitude of the longterm is supposed to generate profound effects for all other ethical views.

https://doi.org/10.1017/S0953820826100351 Published online by Cambridge University Press

II

MacAskill in the book is officially uncommitted with respect to a theory of well-being – recounting, but not opting for any of, the familiar big three: preference-satisfaction views, hedonism, and objective list theories. That said, much of the discussion throughout the book concerning well-being – and improvements in well-being due to technology, which we will turn to soon – suggests an implicit commitment to either simple preference-satisfactionism or hedonism. This is concerning, given that there are powerful arguments against both views.

Against simple preference-satisfaction views, we might worry about adaptive preferences and bad preferences, born of pernicious ideologies.[3] Deeply consumerist and capitalist societies will have many institutions and mechanisms working to inculcate preferences in people, regardless of whether satisfaction of those preferences improves their well-being or contributes to their good. People might end up with strong preferences to have access to and own the latest technology: iPhone 17; Instagram, Snapchat, TikTok, and whatever is next; to have ChatGPT write their essays; to spend hours each day bingeing Netflix reality TV shows; to have larger and more powerful cars; to have food and other items delivered at the press of a button; and so on. A plausible if somewhat controversial claim: satisfying those preferences is not good (or at least not always good) for those people. More refined and idealized preference-satisfaction views are available; I raise here only the red flag of concern about the somewhat crude versions that seem to permeate much of the discussion, albeit non-explicitly.

Against hedonism, we might worry about things along the way to experience machines and Brave New World style drugs like soma.

There are also related concerns about the methods by which some attempt to measure individual and population well-being and happiness, relying heavily on various kinds of surveys and self-reports. There are many concerns about these methods; one central concern is that we should expect them to be inaccurate if neither how "satisfied" or "happy" people feel (or report that they feel) are good measures of well-being or the good.

We need a plausible theory of well-being and the good in order to evaluate claims about progress and the contribution that technology makes to our good. That is missing in *What We Owe the Future*, and I worry that in its place we can discern either a misguided view about well-being or claims that are only plausible on a misguided view of the good.

Throughout the book, there are many discussions of "civilization," "standard of living," "progress," "improvement," "advancement," and "growth" in the telling of the story about human history, as well as in the documenting of concerns about value lock-in, extinction, collapse, and stagnation. These concepts are normative concepts, and should be treated as such, but MacAskill often seems to credulously embrace a normative understanding of these informed by economists, capitalists, and a certain kind of techno-optimist. For example, MacAskill writes that "even huge crises have failed to knock global civilization off course" and the support for this is that "world GDP has only shrunk in a single year a handful of times, and it has always completely rebounded

[3]There is a large literature on so-called "adaptive preferences." See, e.g., Sen (1985). Serene Khader provides an account of adaptive preferences on which they are "(1) preferences inconsistent with basic flourishing (2) that are formed under conditions nonconducive to basic flourishing and (3) that we believe people might be persuaded to transform upon normative scrutiny of their preferences and exposure to conditions more conducive to flourishing." Khader (2011: 42).

https://doi.org/10.1017/S0953820826100351 Published online by Cambridge University Press

within a couple of years."[4] Is it plausible that we can assess the course of global civilization in terms of how world GDP is doing? There are many other similar examples that one could point to.

There are long discussions of "progress" or "advancement" in our human situation (or the situation of at least some or many human beings). In some instances, these focus on moral change: women being allowed to attend university and acquiring a host of other rights, the abolition of slavery, the rise of electoral democracy. But there is also an extended focus on technological change and its contribution to well-being. This is most apparent in Chapter Seven, which focuses on one of MacAskill's main concerns: technological stagnation.

III

MacAskill's concern about technological "stagnation" is a concern about "a near halt to growth and a plateauing of technological advancement," with even a period of stagnation serving to "increase the risks of extinction and permanent collapse."[5] He draws significantly on ideas offered by the economist Tyler Cowen concerning the need to keep pushing for technological "advancement" and "growth" to avoid stagnation.

A central villain in this story is stagnating or even declining "total factor productivity," a measure of a society's ability to "get more output from the same inputs."[6] What is "output" on this story? Is it a thoughtful measure of well-being? No. Certainly not directly. It might be things like how many calories of food can be produced by an acre of farmland. We see an improvement in "output" here through the use of extensive pesticides, fertilizers, and "modern farming techniques." These economic measures always require substantial flattening and simplification. The nutritional and aesthetic quality of the calories, the joy and pride involved in the work of producing them, the externalities from longer term environmental damage – these will not be included in this measure of "productivity." It is troubling, then, that MacAskill says things like "[t]otal factor productivity measures this ability [to get more output from the same inputs] and *represents technological advancement*."[7] This is not an attractive philosophical theory of technological advancement, at least not without substantial refinement and qualification.

Now, I'm not an across-the-board anti-capitalist or anti-technologist. It is plausible that capitalism and some modern technology has proven to be a powerful force for improving well-being, contributing to our good, even if we use a more attractive theory of well-being. Regarding capitalism, this is a complicated position, given the history of colonialism, slavery, and economies of unjust extraction which have been so intertwined with capitalism. I am confident that capitalism is better than autocratic, centrally controlled economic systems, while I think that it is less clear that capitalism is or would be better than something that might be called democratic socialism with substantial protections for private property and free markets. It is possible that thoughtful, well-regulated capitalism might be a force for good. But it is also plausible that largely unregulated, unthoughtful capitalism is unlikely to be a force for good, particularly as economic inequality alters and strains democratic politics. Indeed, it may be that

[4]MacAskill (2022: 125).
[5]MacAskill (2022: 144).
[6]MacAskill (2022: 145).
[7]MacAskill (2022: 145 (emphasis added)).

https://doi.org/10.1017/S0953820826100351 Published online by Cambridge University Press

capitalism has not recently been contributing positively to well-being in the places in which it has been in place for a long time.[8]

With respect to technological development, MacAskill is generally uncritical about how, when, and why technology improves well-being or contributes to our good. No explicit theory about this is defended, other than the above suggestion about technological advancement being tied to total factor productivity. But his commitments become evident in his discussion both about the supposed benefits that attach to technology, and the need to "keep climbing" up the technological ladder.

Borrowing a "thought experiment" from the economic historian Robert Gordon, MacAskill encourages us to imagine what our life would be like in 1870, then 1920, then 1970, then 2020. In each case, we are told a little about what to think about our situation might be like at that time. (The "we" here is something like middle class people in the world's wealthiest countries.) It is not an attempt to give a robust characterization of our life or what our overall well-being might have been like at those times. It focuses attention on a few things about our situation. On his telling, it is a story of great leaps in improved well-being as a result of technology from 1870 to 1920, and again from 1920 to 1970, but with no corresponding leap in well-being improvement from 1970 to 2020. The suggestion is that we are in danger of (or might actually be experiencing) technological stagnation or, if we aren't yet, we should be wary that we might be, soon. The details are worth attending to, as they reveal what MacAskill seems to think matters about the way in which technology might contribute to our good.

Thinking about life in 1870, MacAskill has us focus on "grueling physical labor," the "unrelenting toil" of life as a housewife (as evidenced in part by the need "to carry water 8 to 10 times a day"), the need to "rely on horses for transport."[9] We are told that "[m]ostly your life is one of isolation: the telephone doesn't yet exist, and the postal service doesn't reach your farm . . . modern forms of leisure are unknown. The tallest building in New York City is a church steeple."[10]

Note that we are not told anything about whether people are happy or satisfied, whether they have close knit friends and families, whether they feel connected to others and to the earth, whether they feel that their lives have meaning, value, or purpose. We are told that because people don't have phones or postal service that their lives are ones of isolation, but that seems an unwarranted inference.

There are, of course, aspects of life in 1870 that would have been much worse. Some people certainly had to do too much physical labor. There is no question that much suffering was caused by the lack of medical knowledge and medical technology. (MacAskill reports that life expectancy at birth was 39 years. We should be somewhat wary of being misled by life expectancy at birth statistics; at least it should be of interest also what life expectancy is after childhood. For much of human history, surviving childhood has involved significant challenges, such that life expectancy conditional upon surviving childhood would be much longer.[11]) Still, it is worth considering and scrutinizing what about people's lives was made better by technology, and in what ways. A somewhat unstated theme for MacAskill is that increases in convenience and "leisure" time straightforwardly contribute to our good. This strikes me as under-supported and tendentious.

[8]For an excellent discussion of concerns about capitalism and longtermism, see Venkatesh (2025).
[9]MacAskill (2022: 146).
[10]MacAskill (2022: 146).
[11]For relevant discussion and empirical research, see Roser (2023).

https://doi.org/10.1017/S0953820826100351 Published online by Cambridge University Press

When we zoom forward to 1920, we are told that "[y]our standard of living is in the process of rapid and dramatic improvement."[12] There is no mention made of WWI, or child labor and the misery of factory work in this era of industrialization, nor is there any mention of the Great Depression on the horizon, let alone the horrors of WWII. More strikingly, there is no discussion of how happy people are, whether they see their lives as meaningful and valuable, whether they enjoy nature, whether they have close families and friendships. Instead, we hear about how electricity is becoming more prevalent and that it provides light that is "ten times brighter" than kerosene lamps and "a hundred times brighter" than candles; how we are beginning to use telephones, enabling "instant communication"; how "[m]ass produced cars are beginning to replace horses" and that "Skyscrapers are beginning to rise in New York City."[13] Many of these technological changes strike me as at best a mixed bag when it comes to contributing to our good or making our lives better. They certainly don't support any claim to "dramatic improvement" in our well-being or our good or even to our happiness. Again, medical improvements are different. The one example that seems to unequivocally support "improvement" is the way in which we would be "less likely to contract cholera or typhoid thanks to routine disinfection of drinking water."[14]

Zooming forward to 1970, we are told that again we see "an enormous difference in your life,"[15] where this is clearly meant to suggest an enormous positive difference. What makes for that difference? More human connection, engagement with art and literature, sense of meaning and purpose? No. Instead:

> Most households finally have an indoor flush toilet. You live in a spacious suburban home with a gas stove, a refrigerator, and central heating. Your household owns two cars, and if you want you can fly around the world on an aeroplane. You have a television, and on this TV you just watched a man land on the moon. . . . Your work is probably much less exhausting, and with a forty-hour workweek, vacations, and retirement, you have ample leisure time.[16]

Again, medical advances – penicillin, a polio vaccine, and other vaccines – make close to an unequivocal positive difference. But these other technological advances make a more complicated, equivocal contribution to our good. Convenience, entertainment, and leisure are doing well. But are *we* doing well? Life in cookie-cutter suburbs, with neighbors who are often strangers, everyone in their own home and their own cars, watching their own TVs, often living far from extended family, and having little sense of meaning or purpose – this is good compared to *some* lives we might imagine, and it is certainly good for capitalist firms looking to sell things to us. But is this what the good life consists in?

Finally, we arrive near the present, to 2020. Here we are told "comparatively speaking, this time your life is not all that different."[17] Why not? Well, "among your household appliances, the only difference is that you now have a microwave. Your television is bigger and higher definition, and you have a wider range of shows to watch."

[12]MacAskill (2022: 146).
[13]MacAskill (2022: 146–7).
[14]MacAskill (2022: 147).
[15]MacAskill (2022: 147).
[16]MacAskill (2022: 147).
[17]MacAskill (2022: 147).

https://doi.org/10.1017/S0953820826100351 Published online by Cambridge University Press

MacAskill notes that "[o]f course, there has been a revolution in information and communication technologies – you now have the computers and the internet, tablets and mobile phones," but "technological progress that meaningfully impacts your life has been confined nearly exclusively to those spheres."[18]

This strikes me as dramatically understating the changes wrought by computers and information technology. And it makes one wonder whether MacAskill has thought enough about the way in which technology affects social and political life. Almost every aspect of human existence is different now: how we socialize, how we date and find romantic partners, how we work, how we learn and what we learn about, how we engage in politics, how we interact with each other, who we can engage and communicate and collaborate with, and on and on. The scale and scope of our social and political life is radically different. The quality of our interactions with others has transformed from textured, personal, and in-person, to remote, flattened, and digital. We have much less interaction with the world outside, with many of us only rarely "touching grass" – let alone walking through an actual forest.[19] What it is like to be an adolescent or a young adult now is radically different than what it was in 1970.[20] It would be more different for all of us, but many of us do not entirely or fully inhabit the 21st technological world.

MacAskill is worried that our lives not being all that different (by his lights) is a sign of "growth slowdown" and potential looming "technological stagnation."

My worry about the present is that, in many places we have moved past the point of positive returns from many forms of technology. Many forms of modern technology do not make our lives better, even if they make them more convenient, more filled with a certain kind of "leisure" and entertainment. In places where capitalism is fully entrenched, and in which smartphones and advanced computing and networked technology (internet, social media) are ubiquitous and Amazon/Uber/Doordash convenience abounds, people are increasingly depressed, unhappy, alienated, socially isolated, friendless, angry, lonely, suicidal, inclined toward drug and alcohol abuse, disconnected from the natural world, living lives that seem (even from the inside) largely meaningless, joyless, purposeless.[21]

MacAskill sees it as morally urgent that we avoid technological stagnation because of the concerns that during such a period we might see an increased risk of collapse and extinction. But he also seems to think that it would be somehow more directly bad for people: "[c]learly, during the period of stagnation, people would be much poorer than they could have been if technological progress had continued."[22] But this seems to reveal a view on which technology just keeps making our lives better and better, contrary to the evidence. We see more of this picture early on, when he considers how things might be made better and better in as-yet unimaginable ways:

[18]MacAskill (2022: 147).

[19]In her extended discussion of the benefits of time spent in nature, Williams (2017) recounts that humans have spent 99.9% of our history living extensively in and engaging extensively with nature, but most Americans now spend less than 7% of our waking hours outdoors.

[20]For empirical discussion of these differences and their impact on well-being, see, for example, Haidt (2024), which documents the effects, often sharply negative, of smartphones and social media on childhood and adolescence.

[21]Case and Deaton (2020) document these effects with respect to the working class in the United States. There is a vast literature on the effects of, for example, smartphones on the quality of social interactions. See, e.g., Dwyer et al. (2018), Misra et al. (2014).

[22]MacAskill (2022: 159).

https://doi.org/10.1017/S0953820826100351 Published online by Cambridge University Press

> Much of the progress we've made since 1700 would have been very difficult for people back then to anticipate . . . 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.
>
> Consider the very best moments in your life – moments of joy, beauty, and energy, like falling in love, or achieving a lifelong goal, or having some creative insight. These moments provide proof of what is possible: we know that life can be at least as good as it is then. But they also show us a direction in which our lives can move, leading somewhere we have yet to go. If my best days can be hundreds of times better than my typically pleasant but humdrum life, then perhaps the best days of those in the future can be hundreds of times better again.[23]

I have a hard time knowing what to make of this suggestion. It suggests a bizarre conception of the good for creatures like us, where we could do better at just mainlining or neuro-stimulating the feeling of falling in love or achieving a lifetime goal, if given the right technology. And it's not clear to me that the way in which we've made progress since the 1700s would have been difficult to anticipate in terms of the *ends* that we've achieved. The specific *means*, yes. But the idea that everyone could have a better chance at living a long healthy life or wouldn't have to work quite as long or hard in the fields or would be able to communicate with people at some distance instantaneously, rather than only slowly through writing and horseback transportation – none of these seem all that difficult to imagine.

We should be suspicious of the ability of technology to radically improve human life indefinitely or infinitely. We should be wary that much technology – as in Aldous Huxley's *Brave New World* dystopia – will be used to entertain us, harnessing our psychological weaknesses and limits, getting us to buy and buy and consume and consume more and more, to be more and more obedient and passive participants in our own lives. We should be cognizant of who is creating the technology and for what ends. We should look critically at claims of big corporate capitalist embedded tech to be making the world a better place. We should take seriously that we might end up at a point in time when we should say: OK, we are doing well enough technologically; new tech is more likely to make us worse off than to make us better off. (Note that this is compatible with continuing to live in the modern world; it needn't imply moving to a cabin in the mountains. Note also that this is compatible with working to better distribute existing technology to everyone, so that, for example, people all over the world have access to advanced 21st century medical care.)

IV

MacAskill worries about this kind of view about the limits of technology, or the need to limit technology (even if doing so results in technological "stagnation"), because of worries (a) that stagnation brings various other harms and (b) that we need new technology to prevent extinction and to ensure sustainability over time. Neither concern seems well motivated.

Why is stagnation a concern on MacAskill's view?

[23]MacAskill (2022: 20–21).

https://doi.org/10.1017/S0953820826100351 Published online by Cambridge University Press

One suggestion he offers borrows an idea from economists like Benjamin Friedman which suggests that people are "more morally motivated in times of economic growth . . . they will be more supportive of generous, open, and tolerant social policies."[24] On this view, stagnation is a route to moral backsliding, intolerance, a morally vicious attitude of fighting for what one can.

But the evidence presented for this is not evidence that ties moral motivation to *economic growth* (let alone to technological development) but instead connects moral motivation to something that might also be present at times of economic growth: feelings of stability, safety, relatively good levels of welfare, feelings of optimism about the future, a feeling that all are doing well. In some cases, economic growth might correlate with or even cause these to obtain. But there is no reason to think that this is true in all or even most cases of economic growth, regardless of how the benefits of that growth or development are distributed, regardless of the substantive details of the technological changes that have been wrought. If we focus on technological change, rather than economic growth, the unstable connection with felt stability and optimism about the future becomes even more apparent. And we should note, also, the more direct technological attacks on other-regarding moral motivation. Consider the way in which media market segmentation, cable news, and algorithmically sorted social media and internet channeled news have affected the extent to which we are ensconced in echo chambers, led to see the other side as irrational at best, evil and immoral at worst. No serious evidence is presented to suggest that simply developing more technology causes us to be morally better or more morally motivated.

A distinct suggestion is made that "in a stagnant world, the economic reasons to engage in critical thinking and scientific inquiry would be much weaker. Instead, other values would be selected for, such as those favoring hierarchy and conformity, which have guided so many societies in the past."[25] But again, these claims are not supported by any serious evidence. And they presuppose that the only reasons we expect to see people engage in critical thinking or scientific inquiry are economic reasons. If anything, the technological developments we've seen over the past 50 years have only encouraged us to be less inclined toward critical thinking and scientific inquiry. We are now encouraged to sit back, watch TV in our "spacious suburban homes" while we ingest Fox News or MSNBC or "reality" television, numbing ourselves with various substances, passively accepting comfort and convenience over anything interesting or difficult, in Neil Postman's phrase, *Amusing Ourselves to Death*.

In a world with no or little new technology to distract us, we would still have our familiar and deeply human and humanistic reasons to engage in critical thinking, inquiry about the world, artistic and expressive and scientific pursuits, philosophical reflection to make sense of our situation, and all the wonder, joy, sorrow, and individuality that comes with being human and a part of a social and natural world. We could turn off the TV, brick our "smart" phones, and slowly start unplugging from the experience machine to which we've become attached through decades of consumerist and capitalist encouragement.

No argument or evidence is provided to think that technological "stagnation" will cause us to opt for the values of hierarchy and conformity. (Note that not having *new* technology is not the same as returning to the Middle Ages.) Indeed, we could think more about issues of equality and distribution, rather than just always trying to get the

[24]MacAskill (2022: 160).
[25]MacAskill (2022: 161).

https://doi.org/10.1017/S0953820826100351 Published online by Cambridge University Press

newest thing. And if anything both encourages conformity and encourages us to accept hierarchy, it is the story told by consumerist, capitalist culture, on which we all should just want the newest product, and we should accept that Jeff Bezos, Mark Zuckerberg, Elon Musk, and the rest should be able to buy and control everything, because of all the benefits to humanity they have brought. They do have economic reasons to build new technology that we will be more and more dependent upon, but that is very different than having reasons to engage in critical thinking and scientific inquiry that might contribute to the human good. MacAskill acknowledges that "this is all speculative" but one might wonder why the speculation goes in the direction that it does, given what evidence we have.

The final, most serious concern MacAskill has about technological stagnation is that it is a route to human extinction or civilizational collapse. As he puts it, "[a] different consideration is more clear-cut: a long period of stagnation could substantially increase the probability of extinction or civilizational collapse."[26] This is perhaps the most puzzling part of MacAskill's views about technology.

The idea seems to be something like this. Sometimes, we create technology that poses serious risks to the longterm. To address that, we must keep developing new technology that will counteract those risks. If we stop developing new technology while at a level of "unsustainable" technological advancement, we create a substantial risk of extinction or collapse. So, we must keep building new technology. The central case in his mind is the level of technological advancement in the 1920s, when we had technology that was emitting high levels of carbon dioxide and burning through the planet's fossil fuels. If we stayed there, we would end up with unavoidable extreme climate change, which would have made Earth uninhabitable and, given the current level of technology, we wouldn't have been able to do anything about that.

The puzzling part of this picture is that whatever we think of this point regarding some "unsustainable" level of technological advancement, if we are at a currently sustainable level of technological advancement – such that just keeping things where they are would pose no serious risk to the longterm future – then it seems completely fine to stay at that level. Perhaps there is an implicit premise in the background that we can't just stay at a sustainable level of technological advancement. But why not? Indeed, as MacAskill notes, when we press on to some new technological level, there are worries that we encounter new risks, and enter into new unsustainable epochs. Here is what he says about that, in suggesting the need to avoid technological stagnation:

> Our next level of technological advancement might be unsustainable, too. We could face easy-to-manufacture pathogens and other potent means of destruction without sufficient technology to defend against them. There would be a constant risk of civilization-ending catastrophe. If we stayed stuck at this unsustainable level for long enough, such a catastrophe would be essentially inevitable. To safeguard civilization, we therefore need to make sure we get beyond that unsustainable level and reach a point where we have the technology to effectively defend against such catastrophic risks. The idea of sustainability is often associated with trying to slow down economic growth. But if a given *level* of technological advancement is unsustainable, then that is not an option. We may be like a climber scaling a sheer cliff face with no ropes or harness, with a significant risk of falling. In such a situation, staying still is no solution; that would just wear us out, and we would fall

[26]MacAskill (2022: 161).

https://doi.org/10.1017/S0953820826100351 Published online by Cambridge University Press

> eventually. Instead, we need to keep on climbing: only once we have reached the summit will we be safe.[27]

This metaphor is striking and revealing. I have two main reactions to it.

First, why think that in the case of technological development, there is some "summit" that we can climb to, at which we will be safe? There is every reason to doubt this, particularly given that almost all technology can be used for both good and bad purposes. The very thing that might help us solve questions about a protein's 3D shape from its amino-acid sequence, enabling us to make progress on a host of biological and medical problems, is also the thing that will make it much easier to engineer devastating super pathogens. If anything, the history of the past 500 years is one of building more complex and advanced technology that makes it much, much easier to see how we could bring about total catastrophic events and human extinction. Genghis Khan was terrible, but it would be very hard to see how he could kill 90% or 100% of humanity on horseback. But it is relatively easy to imagine ways in which people or groups could do that now. MacAskill details many of them in the book. Why think there is safety somewhere if we just keep climbing up the technological ladder?

Second, if we've climbed up, isn't another option to climb back down? We know there is ground down there somewhere. Why should we have more confidence that safety lies in the free climb upward, rather than on the path back down?

V

Some technology has improved some aspects of well-being, but we should be wary about thinking that this will continue to be true. We should be particularly wary of a picture like MacAskill's, on which we should be continually climbing up the technological ladder.

Here is a rival picture, recommending a very different path:

1) *Limit*: given what human beings are, and what their well-being consists in, there is a limit to how much technology can contribute to the good of any one human life.
2) *Limit Reached*: for many of us in the modern post-industrial world, we have already come close to that limit; indeed, we may have gone past it.
3) *Detech*: If we are concerned about existential risks and concerned about the longterm good of possible future people, we should push to detech (but not untech).

And, to add in a somewhat related, but somewhat different dimension:

4) *Deconnect*: If we are concerned about existential risks and concerned about the longterm good of possible future people, we should push to deconnect from thin and transactional relationships with millions of people, and to connect with smaller communities.

I can't offer a full defense of this picture. But some of what I've said will have already suggested how the details might be filled in.

[27]MacAskill (2022: 145).

https://doi.org/10.1017/S0953820826100351 Published online by Cambridge University Press

In support of Limit and Limit Reached, we need a plausible theory of well-being for creatures like us. Key elements of our good, I maintain, involve deep friendships and family relationships, engagement with non-human animals and the natural world, open and curious inquiry into the world around us, joy and wonder in response to what we find, and a sense of meaning and purpose. Technology can support us in some of these, but we should be wary of ways in which it can imperil others. Borgmann (1984) and other philosophers of technology warn us about the ways in which convenience in accomplishing some task can leave us purposeless, detached, passive, isolated. Rather than working or engaging with others, we push a button. We are left with "leisure," but also with nothing to do, and no sense of purpose, achievement, or agency to inspire us to do something with our lives. Consumerist entertainment enabled by technology fills our hedonic buckets, but with only the thinnest, least satisfying, least lasting enjoyment. It needn't be this way inevitably; technology can be a genuine aid to well-being and a good life – particularly medical technology that can help us survive given the physical world we inhabit.[28] But we should be wary of the ways in which technological devices can come to replace and displace what is of genuine value, that which genuinely contributes to our well-being. And we should be wary of claims that it can somehow keep adding more and more to our well-being.

(One question we might ask about this concerns the possibility of medical technology that could effectively solve our "mortality problem." It is an interesting question of whether it would be good for us to be immortal; it plausibly would be. But this is a very different way in which Limit might be rejected. And most technology doesn't promise to make a difference to our mortality in this way.)

Some evidence that we have indeed reached or gone past the limit concerns the way in which many very wealthy, technologically advanced countries also have high levels of depression, substance abuse, isolation, ideological capture, social and political division, and fear and anxiety – and this despite the fact that they are able to more widely distribute the technology that arguably does contribute to making our lives better. MacAskill himself is intrigued (if somewhat puzzled) by the high levels of measured happiness of the Hadza, one of the few remaining hunter-gatherer societies in the world.

It is worth stressing that Limit Reached is explicitly formulated to be both contextual and non-universal. There might be many people in many parts of the world who would continue to see substantial well-being improvements with expanded access to technology and perhaps even to development of new technology (such as a widely available malaria vaccine).

With respect to Detech, it is important to note that we can reject some technology and technological advance, and embrace some technological stagnation or even regression, without moving into the forest or resuming a hunter-gatherer lifestyle. It can be hard to imagine "progress" in undoing technology that is already here. But although the extant technology seems ubiquitous and inevitable for each generation born into a particular technological era, it isn't. We can see that just by looking back a little. It can seem like a one-way ratchet, or that to go back is to give up all that we have learned. But that is the result of embracing (without realizing that we are doing so) a certain false techno-capitalist ideology. We can, by values and choice, or by democratically enacted legislation, start rolling some things back, a little at a time. There is no need for it all to happen tomorrow. There is no need to jettison all technology or even most of it. We can

[28]Some are even growing skeptical about the extent of likely future benefits from the development of medicine and medical technology. See Stegenga (2018).

https://doi.org/10.1017/S0953820826100351 Published online by Cambridge University Press

gradually make choices to reject "smart" devices and "smart" phones, along with the detailed surveillance and constant capitalism-enhancing information they provide. We can reject social media and their technologies of distraction and division. We can reject devices with built-in, planned obsolescence. We could decide that AI doesn't aid us with education, expression, work, or living meaningful lives – it displaces or replaces those things. We can ask: is this improving my life, or just making it more "convenient"? We can ask: do I need this, or have I just been trained to consume, to want what is new, to use consumption as a way of temporarily distracting me from what my life is missing?

With respect to Deconnect, we should consider the way in which much of our good consists in sustained, meaningful interactions with other human beings and other creatures of the world, but that this requires attention, focus, and cultivation of specific relationships and communities, not anonymous, thin, or merely transactional interactions with potentially millions of people at once (or no people at all). As Vonnegut (1974) plausibly suggests, much of our happiness consists in feeling a part of a small community of people in which one has meaningful relationships and a sense of purpose.

Additionally, many of the extinction or collapse level threats emerge precisely because of the way in which we are so globally interconnected, via global networks of communication, trade, and transportation. These are very good for global capitalism and those who profit from worldwide markets. But they also have considerable downsides, particularly if we are thinking about possible risks to the long long term.

There are familiar ideas and movements associated with suggestions like Detech and Deconnect, but in my view we should also see them as contributing in core ways to the longtermist project, and a plausible implication of taking longtermist concerns seriously. Extinction and collapse risks from engineered pathogens, AGI, climate change, nuclear war – all plausibly go down as we work to thoughtfully detech and deconnect. There might not be carrying capacity on Earth for quite as many people, but there is reason to think that many many people will be able to exist on Earth for a much longer period of time if we go down this detech and deconnect path, rather than continuing to unthinkingly climb the treacherous technological, interconnection path. There are hard issues about whether we are already close to or past some point of no return with respect to these and other threats, so that we will need to be just as interconnected to solve these problems, given the interconnected way in which we have created them. Coordination and collective action problems might require thoughtful, organized global response. Even if that were the case, it still might be compatible with some significant steps down the Detech and Deconnect path, perhaps seeing it as an extended, gradual process or end goal.

There might be a concern that "ought implies can" and that we cannot halt or reverse technological development. But surely, we can. For the past 300 years, we haven't gone backward to a previous technological level, but that doesn't mean that it is impossible to do so. 300 years is a blip even in human history. We might need to get better at learning how to do this. It might require complex social and political coordination, as we witnessed with global regulation of nuclear and chemical weapons. But the difficulty involved in that is not impossibility, and it is likely considerably easier than what would be required to prevent catastrophic uses (by individuals, rogue nations, terrorists, and the like) of all the technology that might be developed on our continual climb upwards.

https://doi.org/10.1017/S0953820826100351 Published online by Cambridge University Press

VI

I am not *sure* that the Detech and Deconnect path is going to be better for the long long term than MacAskill's continuing to push for technological advancement path. As noted at the outset, it is very hard to be sure about the way in which any of our actions and choices now might affect the very very long term. There are heroic stories on which things all work out just pushing forward with tech, we learn how to colonize space, trillions and trillions of lives become possible beyond the death of the Sun and Earth, and perhaps this doesn't happen (or is less likely to happen) on the Detech and Deconnect pathways. MacAskill seems to have his sights on stories of this kind. I'm skeptical of the optimism on this path; it is much easier to see disaster.

Even if one isn't fully convinced by what I've said here, however, there seems to at least be a strong case for hedging in the Detech and Deconnect direction. We might also consider ways in which we should be enabling and promoting substantial social and political space for experimentation in these directions, allowing and enabling small group autonomy, and considering how we could keep some of the medical technology of the modern world while also enabling some of the more happily human social structures of the past. We should take seriously that the risk to human life on Earth has gone up dramatically over the past 500 years and take seriously that reducing that risk might mean changing what we are doing and building as well as changing our fundamental orientation to technology and human well-being.

References

Albert Borgmann, *Technology and the Character of Contemporary Life* (Chicago: University of Chicago Press, 1984).

Anne Case and Angus Deaton, *Deaths of Despair and the Future of Capitalism* (Princeton: Princeton University Press, 2020).

R.J. Dwyer, K. Kushlev, and E.W. Dunn, 'Smartphone Use Undermines Enjoyment of Face-to-Face Social Interactions', *Journal of Experimental Social Psychology*, **78** (2018), 233–239.

Hilary Greaves and William MacAskill, 'The Case for Strong Longtermism', in *Essays on Longtermism*, eds. J. Barrett, H. Greaves, and D. Thorstad (Oxford: Oxford University Press, 2025).

Jonathan Haidt, *The Anxious Generation: How the Great Rewiring of Childhood is Causing an Epidemic of Mental Illness* (New York: Penguin Press, 2024).

Serene J. Khader, *Adaptive Preferences and Women's Empowerment* (Oxford: Oxford University Press, 2011).

William MacAskill, *What We Owe the Future* (New York: Basic Books, 2022).

S. Misra, L. Cheng, J. Genevie, and M. Yuan, 'The iPhone Effect: The Quality of In-Person Social Interactions in the Presence of Mobile Devices', *Environment and Behavior*, **48** (2014), 275–298.

Max Roser, 'Mortality in the Past: Every Second Child Died.' 2023. <https://ourworldindata.org/child-mortality-in-the-past>

Amartya Sen, *Commodity and Capabilities* (Amsterdam: North Holland, 1985).

Jacob Stegenga, *Medical Nihilism* (Oxford: Oxford University Press, 2018).

Nikhil Venkatesh, 'Capitalism and the Very Long Term', *Moral Philosophy and Politics*, **12** (2025), 33–58.

Kurt Vonnegut, 'Commencement Address, Hobart and William Smith College'. 1974. <https://www.hws.edu/news/transcripts/vonnegut.aspx>

Florence Williams, *The Nature Fix* (New York: Norton, 2017).

Cite this article: Guerrero, Alexander. 2026. Longtermism, Technology, and Well-Being. *Utilitas* 1–15. https://doi.org/10.1017/S0953820826100351

https://doi.org/10.1017/S0953820826100351 Published online by Cambridge University Press