

Article

Are Children The Future?: Longtermism, Epistemic Discounting, and Pronatalism

James Hughes¹

¹ University of Massachusetts Boston; jamesj.hughes@umb.edu

* Correspondence: jamesj.hughes@umb.edu

Abstract: Longtermists stress the moral importance of future generations, debating how to ethically discount their interests amid uncertainty, particularly regarding pronatalism. Some effective altruists and tech elites support pronatalism, sparking controversy. This article argues for significantly discounting the value of future lives due to radical epistemic uncertainty, rendering projections beyond 2100 highly speculative. It advocates a person-affecting view as a starting point, which prioritizes the well-being of living individuals and their near-term future selves. Within this framework, progressive pronatalism—supporting childbearing through subsidies while preserving reproductive freedom—is justified for its immediate societal benefits. However, uncertainties about longevity and automation challenge large-scale pronatalist efforts. Instead of speculative, long-term strategies and an overwhelming focus on existential risk mitigation, this article calls for pragmatic policies with more proximate goals.

Keywords: Longtermism; Epistemic Discounting; Pronatalism; Existential Risk; Person-Affecting View; Reproductive Freedom; Effective Altruism

Citation: Hughes, James. 2025. Are Children The Future?: Longtermism, Epistemic Discounting, and Pronatalism. *Journal of Ethics and Emerging Technologies* 35: 2. <https://doi.org/10.55613/j eet.v35i2.173>

Received: 09/03/2025

Accepted: 09/03/2025

Published: 10/03/2025

Publisher's Note: IEET stays neutral with regard to jurisdictional claims in published maps and institutional affiliations.



Copyright: © 2025 by the authors. Submitted for possible open access publication under the terms and conditions of the Creative Commons Attribution (CC BY) license (<http://creativecommons.org/licenses/by/4.0/>).

1. Introduction

I have been talking about old age dependency ratios as an argument for universal basic income for twenty years. I have pointed to the rapid decline in fertility as a fiscal problem for the graying welfare state and as an argument for investing in anti-aging therapies to keep elders healthy longer. Healthy seniors are less expensive and will work a little longer. However, forcing seniors into the labor market is a bad idea if we soon face technological unemployment. I concluded that we must invest in healthy aging and adopt a universal basic income to even out the intergenerational inequity.

Since I'm also a father of two adult children and am hoping for grandkids, the declining fertility part of the demographic shift has been more on my mind lately. It is apparently on the minds of a growing number of people on the Right as well, ranging from those worried that feminists are pushing humanity to suicide or that there won't be enough of their kind of people in the future to those worried about the health of innovation and the economy. The reluctance by the Left to entertain any kind of pronatalism is understandable, given the reactionary ways it has been promoted. Still, I have always believed that progressive re-appropriation of a pro-family agenda is possible, including child tax credits and subsidies, free daycare, preschool and university, universal healthcare, generous family leave policies, and so on. While these policies probably won't bring us back to the replacement rate, they are good in their own right.

When the influence of futurist philosophy on AI regulation became a big story last year, a spotlight fell on the conservative, anti-feminist pronatalism promoted by people like Elon Musk, which was in turn blamed on the effective altruists and longtermists. Beyond simply trying to mitigate AI risks and x-risks more generally, futurist consequentialists had allegedly decided that the shrinking population was a catastrophe.

My intent here is first to outline the central premises of longtermist thinking and their relationship to pronatalism and the mitigation of existential risks. My initial intuitions are for a sharp temporal discounting of the future for epistemic reasons and that only the interests of the currently living should be included. Even the interests of the currently living are subject to discounting due to declining similarity with their future selves. Then, I argue that pronatalist policies that incentivize parenting while preserving reproductive freedom are in the interest of the currently living for the rest of their lives.

So, think of this as me enlisting the best minds to figure out how to convince young people to have children.

2. Longtermism

Longtermists start with the assumption of "temporal impartiality," that future lives have the same value as today's. Longtermists then determine what actions to take today to have the maximum optimistic influence on the long-term future, given the vast number of potential future lives. Strong longtermists accept that this logic may oblige significant sacrifices today for hypothetical future lives, which Parfit called the "excessive sacrifice" argument, "any small increase in benefits that extends far into the future might demand any amount of sacrifices in the present because in time the benefits would outweigh the cost." Weak longtermists argue for some discounting of future lives so that future interests don't overwhelm current interests. Longtermism is associated with the work of Nick Bostrom (2002, 2003), Nick Beckstead, William MacAskill (2016), Hillary Greaves (2017, 2021) and Toby Ord (2020), and the Global Priorities Institute and the Future of Humanity Institute at Oxford University. However, long-termists draw on a long tradition of work on utilitarianism, from Mill and Sidgwick to Derek Parfit and Peter Singer.

Sociologically, longtermists were also a product of the effective altruism movement that had grown around tech culture for the last twenty years. Effective altruists attempt to apply consequentialism systematically to charitable giving and public policy, drawing on ideas like the Quality-Adjusted Life Year. Effective altruists argue we should commit our resources to reduce suffering and apply a consequentialist logic to the causes and organizations likely to reduce the most suffering. For effective altruists who included future lives in their equations, mitigating existential risks, risks that could extinguish all future value, become a very high priority. Some effective altruism groups like Open Philanthropy have programs devoted to reducing the risks of bioweapons, for instance.

It should be noted that there are non-consequentialist arguments for the interests of future generations. Scheffler (2018), for instance, argues that our concern for future generations is not merely derived from utility maximization but from deeper normative commitments about the continuation of humanity. Similarly, contractual and deontological arguments can be made for protecting future people, independent of aggregative calculations. Nonetheless, while intuitions about the goodness of future life may influence how we debate the coming century, most of the debate can be reduced to consequentialism.

Although effective altruists have addressed the risks to humanity from nuclear war, engineered pandemics, and climate-induced collapse (Ord, 2020), runaway AI is the most popular existential risk in these tech and philosophy circles. Many see the rapid advance of artificial intelligence as a hinge point of history, in which our actions have enormous long-term consequences. For those convinced that artificial intelligence is a threat that can extinguish all future value or provide unimaginable benefits for a very long future, effective altruism, longtermism, and Singularitarianism have fused into a quasi-religious mission.

2.1. Existential Risks and Pascal's Mugging

Even if we see the future as uncertain, we can safely assume that people will always prefer to exist and that avoiding extinction risks today is good. For longtermists, it is

evident that there is an enormous difference between no one living in the future and every other scenario. Even if existential risks are low probability, longtermist calculations give them more weight than most people think is prudent. For many in this milieu, runaway AI is the most significant existential risk, and ensuring a safe transition to the AI future is the most important thing we can do now for the far future.

Some call this result "Pascal's Mugging" (Bostrom, 2009), prioritizing low-probability, high-impact risks leading to misplaced resource allocation. Bostrom highlights the problem of Pascal's Mugging as a weakness in expected value calculations when dealing with astronomically large future populations. If longtermists justify extreme resource allocation based on speculative, low-probability scenarios, they risk over-prioritizing existential risks without robust probability constraints.

One way to navigate this is through risk-weighted epistemic discounting, which incorporates probability assessments into decision frameworks rather than dismissing existential risk mitigation outright. For people outside of this milieu it seems evident that we can argue that mitigating existential risk is valuable without reference to 1025 people in the year 3 billion.

Ironically, some longtermists have pointed out that the maths can make other interventions appear more critical than mitigating x-risks if their benefits, on population or well-being, are astronomical. So, there can also be statistical traps or attractors in the positive intervention ledger (Alexandrie and Eden, 2023).

2.2. Political Influence and Backlash

Very few advocates of these ideas are billionaires, but when advocated by people like Elon Musk and Sam Bankman-Fried, they got much more attention. Hundreds of millions of dollars and crypto coins have flowed into philanthropic support of effective altruist causes and charities, especially those associated with "AI safety." OpenAI and Anthropic were founded by people with this "superalignment" agenda. About a hundred "AI safety" experts have been embedded in federal agencies and Congress, and they actively lobby for their brand of AI regulation in the UK and European Union.

Last year, the failed coup at OpenAI was partly attributed to the influence of effective altruist board members. Elon Musk and Sam Bankman-Fried's reputational collapse, and the enormous wealth trying to shape AI regulation, spotlighted the growing influence of these kooky ideas. These ideas and organizations first came under attack from the Left, who saw the promotion of science fiction scenarios as an attempt to distract from proximate concerns about algorithmic bias and technological unemployment. But longtermism and "AI safety" are also now challenged by libertarian titans of industry counseling an "effective accelerationism" that removes all regulatory constraints.

2.3. Effective Altruism, Pronatalism, and the Right

Anxiety about shrinking populations is now widespread, and consequentialists have their rationale for pronatalism. If lives, on average, are worth living, and we consider future lives, then, all else being equal, more lives will produce more aggregate utility (Mogensen, 2022). Parfit added that even if population growth reduces average utility, a society producing maximum utility would have more lives and lower average utility than we would find attractive, which he called the "repugnant conclusion." While there is no wisdom in repugnance, we must be explicit about how we think about population. We are not near the repugnant conclusion threshold, where fewer people would improve aggregate or individual utility.

Assumptions about future population growth or decline also dramatically impact how important future people's interests are compared to the billions of current people. Thorstad (2023) points out that a steadily declining population eventually zeroes out so that there will only be a couple tens of billion more people ever.

On standard population models, even a constant population of ten billion humans lies out of reach, and there is a good chance that most of the humans who have ever lived have already been born. Standard population models exert very strong downward pressure on the value of existential risk mitigation. (Thorstad, 2023)

Responding to scenarios of declining population, conservatives in tech culture have adopted compelling altruist arguments for pronatalism, adding critiques of feminism, reproductive rights, and liberalism as drivers of population decline, adding racist "great replacement" fears based on ethnic groups' different fertility patterns.

3. Priors: Eudaemonia, Capabilities Theory, Liberalism

Turning to my intuitions, I have four priors. First, I believe consequentialism is the appropriate logic for public policy. For instance, I advocated for a QALY approach to human enhancement therapies in Citizen Cyborg. Some critics have broadly attacked any use of consequentialism, and I think it should be defended and made the default policy calculus in democratic institutions.

Second, my approach has evolved from Millian utilitarianism in the last two decades towards a more eudaemonic view. In particular, the Sen/Nussbaum "capabilities approach" attempts to operationalize public policy aimed at flourishing lives, focusing on people's freedom and capability to pursue life projects. (See Yoon, 2021, for an excellent presentation of capabilities theory consistent with my "technoprogressive" perspective.)

Third, liberal and social democracies in general, and reproductive freedom in particular, are best at maximizing human flourishing and capabilities, at least so far. Reproductive freedom guarantees everyone the ability to pursue life projects, especially women.

Fourth, I accept that we should adopt impartiality and reject kinship bias. At least in these calculations, we should not value our descendants over others. (See Mogensen 2022 for why we should reject kinship partiality for a pure rate of time discounting).

4. Discounting Future Lives

4.1. Radical Uncertainty and Epistemic Discounting

My intuitions on longtermism start with my profound uncertainty about the future. Even if we believe that we should consider future interests, we are not obliged to try if we are completely uncertain about what we could do to improve things, as Sidgwick said 140 years ago.

There is no abstract reason why the interest of future generations should be less considered than that of the now-existing human beings; allowance being made for the greater uncertainty that the benefits intended for the former will actually reach them and actually be benefits. (Sidgwick, 1883)

Thorstad calls this the "cluelessness problem: it is tough to know how our acts will affect the future" (Thorstad, 2023). Thorstad points out that the longtermists think they have dodged the cluelessness problem by focusing on mitigating existential risks: preventing nuclear war is good no matter what. As Parfit (1984) noted, existential risks differ from ordinary risks because they eliminate all future potential value. While we may be uncertain about how many people may exist in the future and what their interests are, we can be very confident they will want to exist. Suppose the discounting is based on our inability to know precisely what actions to take to mitigate existential risks. In that case, our certainty about moral action still declines over time, even when we agree that an existential threat would drive all future uncertainty-adjusted values to zero from that point on.

Existential risks will still loom disproportionately in a consequentialist calculus, but the farther out the risk is in time, the more our cluelessness will eventually reduce our horizon for meaningful action to zero. The best-known rejoinder to the uncertainty about mitigating risks is probably the "butterfly effect" time travel and alternate history stories in science fiction. An example is the Star Trek episode "City on the Edge of Forever," in which Spock and Kirk realize they have to allow a pacifist to die in the 1930s so she doesn't keep the USA out of World War II. History can be perverse.

It is amusing that longtermists are the familial descendants of the concept of the Singularity, which equates the advent of artificial general intelligence with entering the event horizon of a black hole. The outcomes imagined after a "takeoff" scenario range from apocalypse to utopia, but the levers to pull for one outcome or the other were beyond our capacity to comprehend. Our best hope was to build a friendly, super-intelligent AI in the hope that they would understand the best course of action and take over. Today, the Singularitarians' grandchildren are the ones trying to statistically quantify whether the population counts at the end of time would be higher from a one year AI pause versus fighting fascism.

By contrast, I want to argue for radical uncertainty on several grounds. For me, the sources of radical uncertainty include

- Chaos and non-linear complexity in the natural and social world
- Uncertainties about the nature and moral standing of future persons (animals, posthumans, machines, hybrids, collectivities)
- Axiological uncertainty – values will change rapidly
- Eschatological bias - attempting to imagine far futures is always somewhat biased by our eschatological beliefs and fears

These sources of radical uncertainty are very different. Chaotic complexity in socio-political and technological developments creates unpredictability in specific outcomes, leading to a shorter epistemic horizon. Axiological uncertainty, the evolution of moral values over time, suggests a fundamental unpredictability in long-term ethical considerations. Eschatological bias is a heuristic warning against overly deterministic future predictions rather than a source of uncertainty itself. Epistemic discounting for the likelihood of eschatological bias is like reminding oneself that optimistic or pessimistic biases skewed previous predictions.

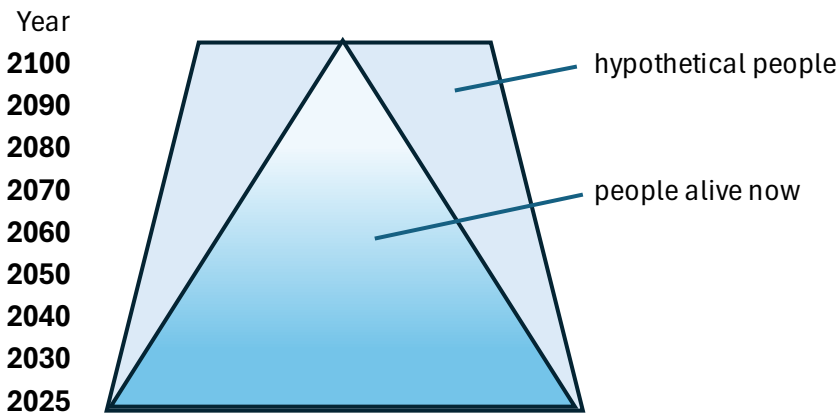
For me, given these various sources of cluelessness, the epistemic horizon for moral action is restricted to this century. Even if we have solid reasons to assume some things about the far future, axiological change means our descendants will bring different values and better information to their decisions.

4.2 Person-Affecting View and Non-Identity

Given radical uncertainty, it doesn't make sense to me to include future populations' welfare beyond this century, at least in the first estimation. However, I am also inclined to two additional restrictions on calculating future utility: the person-affecting view and the declining identity of current people with their future selves. The person-affecting view argues that, for moral and pragmatic reasons, we should focus on the interests of existing people and their interests in the future and not try to calculate the partial probabilistic utilities of hypothetical future people. (Narveson, 1973). Narveson's famous formulation is that morality is about "making people happy, not making happy people" (Narveson, 1973: 80).

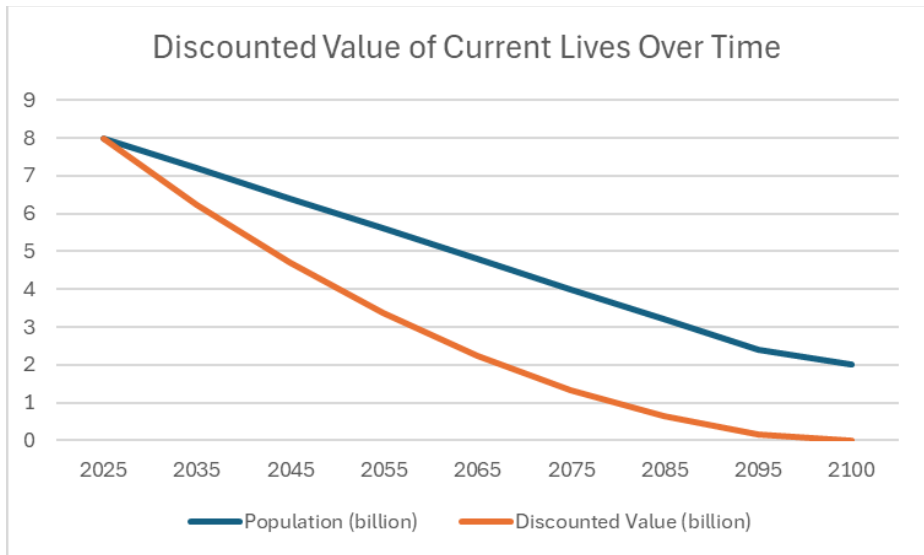
The second limitation I would place on modeling future utility is Parfit's probabilistic decline in identity over time. Parfit also robustly rejected the person-affecting view because it is incompatible with having moral intuitions about the fate of unborn people. I will return below to an account that includes the unborn. However, Parfit's account of our declining identity with ourselves over time does much of the same work I want to do here, epistemically discounting future lives. In his 1984 *Reasons and Persons*, Derek Parfit argued that we should discount our lives over time since the longer we live, the less

identical we are to our current selves. Even just by focusing on the future lives of people alive now, we can be certain that the longer they live, the less they will have in common with their current selves. By 2100, we are not only ignorant of what decisions will need to be made and what values will be in play but also whether the young people today are predictive of the elders they will become.



4.3. Implications

Suppose we roughly model the life expectancies of current persons (person-affecting view) and apply a steady discount rate that reduces to 0 by 2100, due to epistemic uncertainty and non-identity. In that case, it yields something like this weighting.



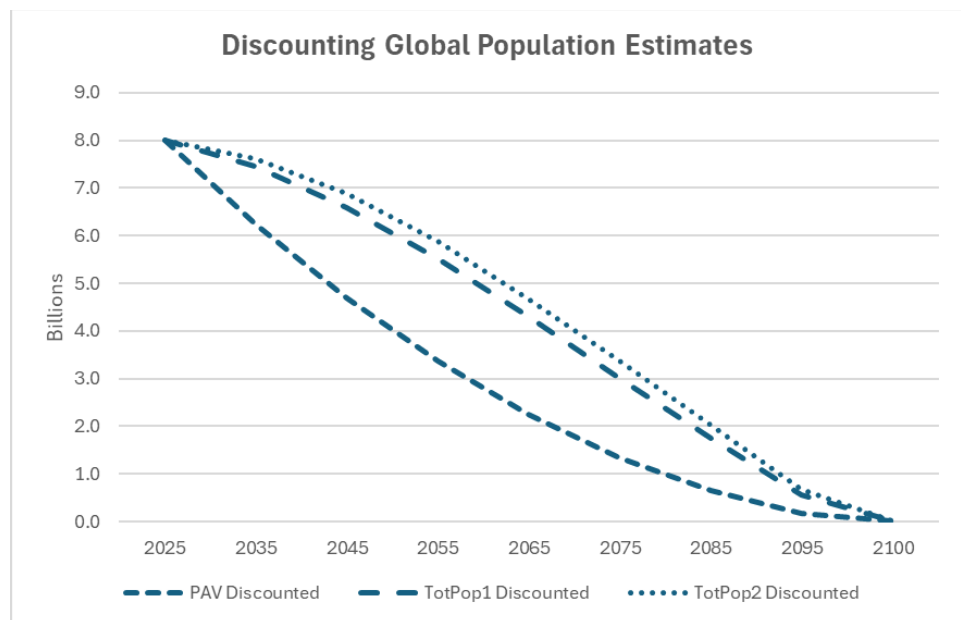
In this accounting, the same intervention is worth 6 times more in 2035 than in 2080. In other words, we are six times more certain of what actions to take for people's interests in 2035 than in 2080. For instance, in the case of pronatalism, we can be six times more certain about the social benefit of a baby born in 2035 than one born in 2080. Again, all babies are equally precious, and all people alive in 2080 are of equal moral worth, but right now we are much less confident about what to do for them today.

It turns out that the hard epistemic barrier of 2100 is doing almost all the work in this model. If we include those born between now and 2100, the shape of the curve will only change somewhat. In the chart below, I have included the median and pessimistic estimates of the global population developed by the International Institute for Applied

Systems Analysis (IIASA). The IIASA is an Austrian think tank that includes trends in family planning and education in its modeling and is used by the Intergovernmental Panel on Climate Change.

If we use IIASA's low growth model to predict population decline without pronatalist intervention, it predicts a peak in the 2060s, declining to 8 billion by 2100. More optimistically, if pronatalist interventions succeed or anti-natalist conditions abate, the IIASA's median model (K.C. et al., 2024) predicts global population peaking in the 2080s and declining to less than 10 billion by 2100.

However, once the continuous epistemic discount rate is applied to both models, they generate roughly equal discounted cluelessness-weighted guesses about welfare. The difference between these top two lines and the bottom line, which is based on discounted current lives (e.g. person-affecting view), is the additional moral weight of the people likely to be; population aging and decline, while not necessarily catastrophic, carries predictable economic and social costs born between now and 2100. From this perspective, the argument for pronatalism is still positive but very weak compared to all the other things we could do to improve future welfare.



5. Person-Affecting Population Decline

Suppose we accept radical epistemic uncertainty and adopt the person-affecting view with declining identity. In that case, we have obligations to try to do the best for the lives of existing adults and children, our own and all others, with a sharp discount rate on our ability now to act in their interests over the coming century. One aspect of their well-being will be determined by whether the population is shrinking or growing and how quickly it is aging. We have good economic reasons to assume that the future well-being of existing people will be better; the slower the population ages, the less it shrinks (Yglesias, 2021). The interests of the currently living then warrant at least some attention to near-term pronatalist interventions. I presume that in the year 2050 when the current cohort is aged 25 to 100, they will be better off in societies with more young people, while the impacts of pronatalist interventions on the well-being of everyone in 2100, including the dwindling number of currently living people, is unknowable.

5.1. Life Expectancy and Automation

Two factors that will change how we think about well-being and population are how long people live and how much labor we can replace with automation. If children alive today live to 2100, they may have technically unlimited life expectancies. Radically extended life expectancy would be so different that we have an even stronger rationale for discounting. But radical healthy life extension by 2050 would mean a quicker aging society and probably a faster fertility decline. I guess another prior for me is that radical life extension would be a net good and force us to adapt to a grayer, smaller future more quickly.

As to automation, much of the argument for the impacts of population decline on well-being rests on its economic impacts, from the burden of pensioners on tax-payers to labor shortages and reduced economic vitality. Automation and robotics will partly address the decline in young workers. We may find that an automated economy can provide a high level of well-being for everyone without population growth. If we assume that a fully automated economy will be achieved over this century, as I do, then there is declining salience for pronatalist interventions on person-affecting grounds.

5.2. The Value of Freedom and Democracy

My rejection of the conservative framing of pronatalism is based on my commitment to reproductive freedom, feminism, and liberal individualism, which is grounded in capabilities consequentialism. Historically, liberal and social democracies have been best for promoting individual freedom and collective empowerment, thus maximizing individual and aggregate flourishing. In particular, the freedom to choose whether to have children and how many is a significant constraint on pursuing all life's options and is thus an essential liberty. While there are lifeboat situations where we may agree, "You three have to have kids," only policies that encourage but do not require childbearing would be consistent with maximal future well-being.

Liberal and social democracies are also the best systems we have found for aggregating collective interests into policy. The folks who make it to 2080 may have better mechanisms for social choice, but from the vantage of 2025, more liberal, egalitarian, and democratic societies still look like their best bet. Insofar as liberal and social democracy are on the defensive, the defense of these regimes should also be a moral priority, now and for the future.

5.3. Progressive Pronatalism

Suppose we accept that there are moral grounds for pronatalism, and we reject any coercion to reach the replacement rate. In that case, we are left with creating more generous incentives to become parents. Most of the experiments so far trying to encourage fertility with financial incentives, from Denmark to Japan, have had only marginal impacts on birth rates, leading to demographic pessimism (North, 2023).

I believe we just haven't found the price point yet. Child subsidies could be more generous, and daycare and university can be free. A society that acknowledged parenting as a part-time job worthy of compensation and each child as a precious resource to be nurtured to full potential would have a higher birth rate, albeit probably not at a replacement rate.

Progressives, such as folks at the United Nations Population Fund, believe that calling these policies "pronatalist" implies that governments should have population targets, which opens the door to restrictions on women, sexuality, marriage, and reproductive freedom. Since child subsidies, access to reproductive health services, including fertility assistance, and expanded social services for families, are desirable in their own right, progressives argue that we should instead campaign for everyone who wants to have children to be able to. While I share the anxiety about providing ammunition for reactionary attacks on reproductive freedom, as the world becomes more

focused on the problem of falling birth rates, progressives should not cede the argument to reactionaries.

6. Longtermism: Right Questions, Suspect Answers

I am in awe of the rapid translation of this bundle of futurist philosophy into an influential and controversial political movement, which may now reach into the Trump administration through several longtermist oligarchs. I am hostile to the reactionary and elitist, even racist, appropriations of these ideas by self-interested tech bros. However, I am also convinced that the work the diverse "effective altruist" and "AI safety" communities have attempted is vital and that the majority of their members are basically liberal and center-left in their politics (Wenar, 2024). Being more explicit about the expected impacts of different distributions of power and wealth and different political and economic arrangements on future well-being would help clarify that many disagree with the shallow appropriation of futurist philosophy by problematic celebrities and billionaires. In other words, what I have called the "technoproggressives" should divide from and critique the right-wing appropriators of futurist philosophies.

Longtermism asks essential questions about how we think about a universalistic morality that includes future people. However, as with many fields of putatively secular futurism, it reads more like eschatology than policy analysis to this sociologist. Even if my assumptions are too strict, the field could benefit from more epistemic humility.

As to my takeaways, I think there are pragmatic, epistemic, and political reasons to limit our futurism to about a century, focusing on the well-being of the currently living. The future well-being of living people in this century will probably be better in a world that makes it as easy as financially possible to become a parent while preserving reproductive choice, even if this doesn't get us back to replacement. We can systematically examine existential risks and demographic challenges without far-future fairy tales.

References

- Alexandrie, G., & Eden, M. (2023). Is Existential Risk Mitigation Uniquely Cost-Effective? Not in Standard Population Models. In GPI Working Paper (Issue 5).
- Bostrom, N. (2002). Existential Risks: Analyzing Human Extinction Scenarios and Related Hazards. *Journal of Evolution and Technology*, 9(1). <http://www.nickbostrom.com/existential/risks.html>
- Bostrom, N. (2003). Astronomical Waste: The Opportunity Cost of Delayed Technological Development. *Utilitas*, 15(3), 308–314. <https://doi.org/10.1017/S0953820800004076>
- Bostrom, N. (2009). Pascal's Mugging. *Analysis*, 69(3), 443–445
- Bostrom, N. (2013). Existential Risk Reduction as Global Priority. *Global Policy*, 4(3), 15–31. <https://doi.org/10.1111/1758-5899.12002>
- Bostrom, N. (2014). *Superintelligence: Paths, Dangers, Strategies*. Oxford University Press.
- Burgis, B. (2022, December 10). Effective Altruism Is No Substitute for a Better Society. *Jacobin*, 1–6.
- Danaher, J., & Hopster, J. (2022). The normative significance of future moral revolutions. *Futures*, 144. <https://doi.org/10.1016/j.futures.2022.103046>
- Gheaus, A. (2016). The right to parent and duties concerning future generations, *Journal of Political Philosophy* 24,4: 487–508.
- Greaves, H., Macaskill, W., & Thornley, E. (2021). The Moral Case for Long-Term Thinking. In N. Cargill & T. M. John (Eds.), *The Long View* (pp. 19–8). First.
- Greaves, H., & Ord, T. (2017). Moral uncertainty about population axiology. *Journal of Ethics and Social Philosophy*, 12(2), 135–167.
- K.C., S., Dhakad, M., Potančoková, M., Adhikari, S., Yildiz, D., Mamolo, M., Sobotka, T., Zeman, K., Abel, G., Lutz, W., & Goujon, A. (2024). Updating the Shared Socioeconomic Pathways (SSPs)

-
- Global Population and Human Capital Projections. IIASA Working Paper. Laxenburg, Austria: WP-24-003. <https://pure.iiasa.ac.at/id/eprint/19487/>
- MacAskill, W. (2016). *Doing Good Better: Effective Altruism and a Radical New Way to Make a Difference*. Avery.
- Mogensen, A. L. (2022). The only ethical argument for positive δ ? Partiality and pure time preference. *Philosophical Studies*, 179(9), 2731–2750. <https://doi.org/10.1007/s11098-022-01792-8>
- Narveson, J. (1973). Moral Problems of Population. *The Monist*, 57(1), 62–86.
- North, A. (2023). You can't even pay people to have more kids. *Vox*, Nov 27. <https://www.vox.com/23971366/declining-birth-rate-fertility-babies-children>
- Ord, T. (2020). *The Precipice: Existential Risk and the Future of Humanity*. Bloomsbury Publishing.
- Parfit, D. (1984). *Reasons and Persons*. Oxford University Press.
- Scheffler, S. (2018). *Why Worry about Future Generations?* Oxford University Press.
- Sidgwick, H. (1883). *The Principles of Political Economy*. New York: Macmillan
- Tarsney, C. (2024). The Epistemic Challenge to Longtermism. *Synthese*, 1–39.
- Thorstad, D. (2023). Three mistakes in the moral mathematics of existential risk. In GPI Working Paper (Issue 7).
- Wenar, L. (2024, March 29). The Deaths of Effective Altruism. *Wired*, 1–27. <https://www.wired.com/story/deaths-of-effective-altruism/?s=09>
- Yglesias, M. (2021). "The economic case for pronatalism," *Slow Boring*. July 21. <https://www.slowboring.com/p/the-economic-case-for-pro-natalism>
- Yoon, I. S. (2021). Amartya Sen's Capabilities Approach: Resistance and Transformative Power in the Age of Transhumanism. *Zygon*, 56(4), 874–897.