



# What is Techno-Optimism?

Peter Königs<sup>1</sup>

Received: 3 June 2022 / Accepted: 11 June 2022 / Published online: 13 July 2022  
© The Author(s) 2022

**Keywords** Techno-optimism · Risk

Should we be optimists or pessimists about technology? This question goes to the heart of how we should think about and relate to technology. Curiously enough, there are few systematic discussions of this question in the philosophical literature. In his nuanced and wide-ranging essay ‘Techno-optimism: an analysis, an evaluation and a modest defence’ (1), John Danaher remedies this situation by providing a long-overdue systematic exploration of techno-optimism. His essay does four things: it provides an analysis of how techno-optimism should be understood. It offers a general framework for assessing when techno-optimism is warranted. It explains why strong forms of techno-optimism are difficult to justify. And it defends a modest, agency-based version of techno-optimism.

There is a lot in the paper that I find compelling, and the general framework he puts forth will prove useful for everyone thinking about the topic. In this discussion note, I want to engage with one part of Danaher’s analysis that I find less compelling, namely his account of how techno-optimism should be understood. I will first explain why we should reject Danaher’s account and then outline an alternative account.

Danaher defends the preponderance account of techno-optimism. Essentially, this account defines techno-optimism as “the stance that holds that technology plays a key role in ensuring that the good prevails over the bad.” (2022, p. 8) He rejects the improvement account, according to which techno-optimism is, roughly, the stance that holds that technology makes things better, perhaps substantially so. Put differently: on the improvement view, techno-optimists believe that technology makes the world a *better* place. On Danaher’s preponderance view, they believe that it makes the world not only a better place but also a *good* place.

---

Discussion note on John Danaher’s “Techno-optimism: an analysis, an evaluation and a modest defence”.

---

✉ Peter Königs  
p.koenigs@fs.de  
<https://peterkonigs.mystrikingly.com/>

<sup>1</sup> Philosophy & Law Department, Frankfurt School of Finance & Management, Adickesallee 32-34, 60322 Frankfurt am Main, Germany

I find Danaher's case for the preponderance account unpersuasive. The preponderance account has a couple of implications that, in my view, speak strongly against it. For instance, someone who asserts 'I firmly believe that technology will make the world a much better place' should definitely count as a techno-optimist. In order to classify her as a techno-optimist, we need not know whether she thinks that technology will tip the overall balance of good and bad in the world. In fact, I would say that she need not have any beliefs about this question at all to qualify as a techno-optimist. According to the preponderance account, however, we cannot classify this person as a techno-optimist until we have verified that (1) this person has any beliefs about the overall balance of good and bad in the first place, and (2) that she believes that technology contributes to tipping the overall balance in favour of the good. I find this implausible. The preponderance account has even more counterintuitive implications when applied to individual technologies rather than to technology as a whole. Surely, someone who is optimistic about, say, artificial intelligence or CRISPR need not believe that these technologies will contribute to ensuring that the good prevails over the bad. She need only have positive expectations about their net impact on human welfare. To be sure, one could reject the preponderance view for individual technologies and apply it only to optimism about technology as a whole. But this move would be ad hoc.

Danaher's preponderance account thus seems to be at odds with ordinary language. Deviating from ordinary language is, of course, legitimate if doing so serves the interest of one's inquiry. Danaher's goal is to provide a useful *ameliorative* analysis of techno-optimism. As such, it need not accord with ordinary language usage. But the reason why we are interested in techno-optimism in the first place does not seem to support the preponderance account either. Arguably, the reason why it matters whether we may be optimists about technology is because we want to know whether we ought to unleash the forces of technological progress, or whether we should rather inhibit, regulate or even halt or reverse technological development. But to answer this question, we need not determine whether technology will make the good prevail over the bad. We need only know whether technology can be expected to make things better or worse. There might be people who are interested in whether technology makes the good prevail over the bad. But this seems like a philosophical question of relatively little real-world relevance. The more important question, from a social and political point of view, is whether technology can be expected to improve the human condition or not. I strongly suspect that this is also the question that many of the authors cited by Danaher, tech people, technology critics, and politicians are most interested in. The preponderance account seems to be the wrong account of techno-optimism *given what we are interested in*.

The improvement account yields more plausible results. It correctly classifies someone as a techno-optimist who firmly believes that technology (or *a* technology) will make the world a much better place irrespective of her beliefs regarding the overall balance of good and bad in the world. It also fits better with the motivation underlying our interest in techno-optimism. However, while capturing some important truths, the improvement account is too simple. The principal problem with it is that it fails to account for risk and uncertainty. Typically, we do not know with certainty what the effect of a technology, or of technology *tout court*, is going to be. Many of our beliefs about technology are not of the form 'Technology *T* will make

things better'. Rather, we are aware that different outcomes are possible and try to assign rough probabilities to these outcomes. The simple improvement account, according to which to be a techno-optimist is to hold that technology makes things better, is too coarse-grained to deal with more nuanced assessments of the impact of technology.

What I want to suggest, therefore, is that whether a person qualifies as a techno-optimist depends on the impact/probability distribution she assigns to a technology, or to technology as a whole, and on whether this impact/probability distribution is thought to be favourable or not. By 'impact/probability distribution', I mean the distribution of probabilities and impacts of the various possible outcomes: what are the possible outcomes, and how likely and desirable are they? This idea yields the following tentative definition, which could be understood as a more sophisticated version of the improvement account: a techno-optimist is someone who believes that technology's impact/probability distribution is favourable. An optimist about a particular technology is someone who believes that this particular technology's impact/probability distribution is favourable. Techno-pessimism is the view that the impact/probability distribution of technology is unfavourable. A pessimist about a particular technology is someone who believes that this particular technology's impact/probability distribution is unfavourable. Let us call this the refined improvement account.

To better understand the refined improvement account, consider the following three examples and what the suggested definition implies about them. The examples feature individual technologies, but the same applies to technology *tout court*.

### Scenario 1

We can be certain that the net impact of some technology is positive. There is no significant risk of it having a net negative impact. This impact/probability distribution is favourable.

=> The assessment of the technology warrants optimism about this technology. Someone who holds these beliefs is an optimist about this technology.

### Scenario 2

There is a high probability of some technology having a significant net positive impact, and only a small probability of it having a negligible net negative impact. This impact/probability distribution is favourable.

=> The assessment of the technology warrants optimism about this technology. Someone who holds these beliefs is an optimist about this technology.

### Scenario 3

There is a high (say, ~80%) probability of some technology having a net positive impact, and a relatively small but significant (~20%) probability of it resulting in the extinction of humankind. This impact/probability distribution is unfavourable.

=> The assessment of the technology does not warrant optimism about this technology. Someone who holds these beliefs is a pessimist about this technology.

Like the simple improvement account, the refined improvement account classifies someone as a techno-optimist who is certain that the impact of technology is positive (scenario 1). But scenario 1, involving no risk or uncertainty, is rather atypical. Scenario 2 involves risk, and it is easy to see that the overall impact/probability distribution is still favourable. The most interesting case is scenario 3. It highlights that the refined improvement account can classify someone as a pessimist even if she believes that the effects of some technology are probably going to be positive. This, I think, is a plausible implication of the account. What matters is the overall impact/probability distribution and how favourable it is. Someone who believes that there is a significant chance of some technology destroying human civilization hardly qualifies as an optimist, even if she believes that the technology is more likely to have a positive impact. If somebody were to say ‘I think we can be optimistic about AI. It will probably make our lives better. There is only a 20% chance of it destroying civilization’, we would think that this person is making a cynical joke. Her assessment clearly does *not* support optimism about AI.

My talk of ‘favourable’ has admittedly remained somewhat vague. I suggest that the impact/probability distribution of some technology is favourable if it is such that we should prefer the introduction or arrival of this technology over the status quo. We should welcome the introduction of the technologies in scenarios 1 and 2, but not in 3, given the significant existential risk. Note that whether the good ends up prevailing over the bad seems irrelevant for whether we should prefer the arrival of some technology over the status quo. Ultimately, what matters is whether technology will make the world a better or a worse place, and how likely it is going to do so. This is why this account is best understood as a refined version of the improvement account.

On the refined improvement account, much like on Danaher’s preponderance account, the techno-optimist stance is not easy to justify. As Danaher plausibly points out, the way technologies will develop in the future and what their impact is going to be is extremely difficult to anticipate. The refined improvement account introduces an additional normative source of uncertainty, on top of this factual uncertainty. Even when we can determine the impact/probability distribution, which is difficult enough, we might struggle to determine whether it is favourable or not. While the above scenarios were easy to decide, many impact/probability distributions are much more ambivalent normatively, as any risk ethicist will be able to confirm. For instance, how small must the risk of human extinction be in order for super-intelligent AI to be a technology that we should welcome? What we can say, though, is that the refined improvement account is less demanding than Danaher’s preponderance account. For it need not be the case that the good prevails over the bad for some impact/probability distribution to be favourable and thus for optimism to be warranted.

I am not suggesting that the refined improvement account is the only legitimate or plausible way of defining (techno-)optimism. But it does seem to fit well with what I take to be the motivation underlying our interest in techno-optimism and whether it is warranted. Is not what we are interested in precisely whether we should welcome new technology, or *a* new technology, in light of its estimated impact/probability distribution? For this tells us how we should practically relate to it.

**Acknowledgements** The author is a member of the project group “Regulatory theories of Artificial Intelligence” within the “Centre Responsible Digitality” (ZEVEDI).

**Author Contribution** Not applicable.

**Funding** Open Access funding enabled and organized by Projekt DEAL.

## Declarations

**Ethics Approval** Not applicable.

**Consent to Participate** Not applicable.

**Consent to Participate** Not applicable.

**Competing Interests** The author declares no competing interests.

**Open Access** This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

## Reference

Danaher, J. (2022). Techno-optimism: an analysis, an evaluation and a modest defence. *Philosophy & Technology*, 35(54), 8. <https://doi.org/10.1007/s13347-022-00550-2>

**Publisher's Note** Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.