



# AN ANTHOLOGY OF GLOBAL RISK

EDITED BY  
SJ BEARD AND TOM HOBSON



<https://www.openbookpublishers.com>

©2024 SJ Beard and Tom Hobson

Copyright of individual chapters is maintained by the chapter's authors



This work is licensed under an Attribution-NonCommercial 4.0 International (CC BY-NC 4.0). This license allows you to share, copy, distribute and transmit the text; to adapt the text for non-commercial purposes of the text providing attribution is made to the authors (but not in any way that suggests that they endorse you or your use of the work). Attribution should include the following information:

SJ Beard and Tom Hobson (eds), *An Anthology of Global Risk*. Cambridge, UK: Open Book Publishers, 2024, <https://doi.org/10.11647/OBP.0360>

Copyright and permissions for the reuse of many of the images included in this publication differ from the above. This information is provided in the captions and in the list of illustrations. Every effort has been made to identify and contact copyright holders and any omission or error will be corrected if notification is made to the publisher.

Further details about CC BY-NC licenses are available at <http://creativecommons.org/licenses/by-nc/4.0/>

All external links were active at the time of publication unless otherwise stated and have been archived via the Internet Archive Wayback Machine at <https://archive.org/web>

Digital material and resources associated with this volume are available at <https://doi.org/10.11647/OBP.0360#resources>

ISBN Paperback: 978-1-80511-114-6

ISBN Hardback: 978-1-80511-115-3

ISBN Digital (PDF): 978-1-80511-116-0

ISBN Digital eBook (EPUB): 978-1-80511-117-7

ISBN XML: 978-1-80511-119-1

ISBN HTML: 978-1-80511-120-7

DOI: 10.11647/OBP.0360

Cover image: Javier Miranda, Alien planet, June 18, 2022, <https://unsplash.com/photos/nc1zsYGkLFA>

Cover design: Jeevanjot Kaur Nagpal

# 2. Democratising Risk: In Search of a Methodology to Study Existential Risk

*Carla Zoe Cremer and Luke Kemp*

---

## Highlights:

- Existential Risk Studies is currently dominated by a Techno-Utopian Approach (TUA) that defines existential risk, not in terms of human extinction, but in terms of the loss of very large quantities of value predicated on a utilitarian, transhumanist, and longtermist set of ethical assumptions.
- The TUA is not representative of the values of most human beings. This is philosophically tenuous, given that the field often speaks of humanity as a whole, and also languorously undemocratic, given its alignment with the interests of elites who are also contributing to existential risk.
- Defining existential risk in relation to the loss of value invariably ties the field to a particular group's value system. To avoid this, we should separate the study of existential and global catastrophic risk from both extinction ethics (what is good or bad about human extinction) and existential ethics (what is good or bad about different societal forms).
- The TUA also presents methodological limitations including an apparent commitment to technological determinism, simplistic, threat-based models of risk assessment, and expected value-based approaches to decision-making under

uncertainty. These are hard to justify and push the field away from best practise in related fields such as disaster risk reduction and climate science.

- Given the above, Existential Risk Studies currently poses a high level of “response risk”, i.e. it could recommend responses that are net harmful to humanity, with many historical precedents for how this can happen. To prevent this existential risk, scholars should transparently acknowledge the moral and empirical assumptions used in risk analyses; critically embrace the latest advances in risk assessment from other fields; diversify the field’s approaches and assumptions; and democratise its judgement.

The importance of democracy in the governance of Existential Risk research is also discussed in Chapter 16, while mechanisms for engaging with the democratic process in policy making are discussed in Chapter 22.

---

## 1. Introduction

Over the past two decades, scholars have begun to methodically study human extinction and global catastrophes. This field of “Existential Risk Studies” (ERS) aims to (i) identify existential and catastrophic risks; (ii) map out the potential causes of existential catastrophes; (iii) understand the ethical implications of such calamities and (iv) devise effective strategies for mitigation and prevention.

Although the field is relatively small, it has expanded considerably, especially over the past 10 years.<sup>1</sup> It is also of increasing public interest: several popular trade-books have been published.<sup>2</sup> The ideas have been integrated into vision-setting reports from the UN Secretary General.<sup>3</sup> Institutions focusing on existential risk have received hundreds of millions of dollars in philanthropic funding.<sup>4</sup> It is a field on the rise.

It is commendable and overdue that the study of human extinction is receiving greater academic engagement and public attention. However, this field needs to be held to high standards: it is ambitious, could affect

the lives of many and attracts scholars who seek to change the trajectory of global society.

The field faces daunting challenges. How can it be inclusive of the diversity of human preferences and visions of the future? How can researchers avoid baking their subjective assumptions into risk analyses that might affect those who do not share their values? How do they conduct complex risk assessments? How do they deal with uncertainty? How do they compare risks with different quantities of evidence and degrees of plausibility? How do they ensure that the catastrophes the field studies are not misused to justify dangerous actions?

The field has not yet established the answers to such questions and we are not the first to be aware of this.<sup>5</sup> Throughout this chapter we will point to the scholars who we know have raised similar questions. The historically dominant Techno-Utopian Approach (henceforth the “TUA”) played an important role in establishing the field and drawing attention to the significance of studying human extinction. It is time to examine this approach with a critical eye. We do this to identify weaknesses, areas for further investigation and the need to also explore alternative approaches. The TUA, which relies heavily on total utilitarianism, transhumanism, and (strong) longtermism, is too morally unrepresentative, methodologically flawed, and risky an approach to rely on. It is time to diversify the definitions, tools, and frameworks of ERS. We need to develop an ERS methodology that addresses the core questions of ERS and avoids the problems of the TUA.

The question we raise is: what should the study of human futures and catastrophe look like under moral uncertainty? We suggest some solutions: a diverse range of approaches, deliberative democratic processes, and the separation of the study of catastrophe and extinction from the ethics of human existence and extinction.

We proceed in Section 2 by outlining the moral assumptions of the TUA. In Section 3 we give reasons why the TUA is not representative of wider human preferences and why representation matters to the goals of ERS. Section 4 explains how moral and empirical assumptions are often masked by abstract and ambiguous definitions and tools. Moral and empirical assumptions are important because they can distort the results of our work, whether it be in how we conceive of risks, which risks we prioritise, or what policies we recommend. In Section 5 we

focus on how studying existential risk could backfire by growing other sources of catastrophic risk. In Section 6 we suggest a democratic, risk-averse approach.

## 2. The Techno-Utopian Approach to ERS

We focus on the Techno-Utopian Approach to existential risk for three reasons. First, it serves as an example of how moral values are embedded in the analysis of risks. Second, a critical perspective towards the Techno-Utopian Approach allows us to trace how this meshing of moral values and scientific analysis in ERS can lead to conclusions, which, from a different perspective, look like they in fact increase catastrophic risk. Third, it is the original and by far most influential approach within the field.

### 2.1 Definitions and history

#### *2.1.1 The influence of the TUA*

The TUA is a cluster of ideas which make up the original paradigm within which the field of ERS was founded. We understand it to be primarily based on three main pillars of belief: transhumanism, total utilitarianism and strong longtermism. More precisely: (1) the belief that a maximally technologically developed future could contain (and is defined in terms of) enormous quantities of utilitarian intrinsic value,<sup>6</sup> particularly due to more fulfilling posthuman modes of living;<sup>7</sup> (2) the failure to fully realise or have capacity to realise this potential value would constitute an existential catastrophe;<sup>8</sup> and, (3) we have an overwhelming moral obligation to ensure that such value is realised by avoiding an existential catastrophe,<sup>9</sup> including through exceptional actions.<sup>10</sup>

Not all publications that make use of the TUA explicitly support every element of the approach, but the most widely read publications incorporate a significant number of TUA elements and share its visions of the long-term future.<sup>11</sup> The most popular definitions of existential risk are still the initial techno-utopian definitions by Bostrom and more abstract, but very similar versions based on expected value (see Section 4.1). The few attempts to put forward alternative frameworks<sup>12</sup>

are not nearly as widely cited, known, or used.<sup>13</sup> Importantly, they do not offer alternative definitions of what an existential risk is. Despite the increase in size and diversity in the field of ERS there appears to still be no coherent alternative to the TUA.

The impact of the ideas of the TUA can be seen across the field, characterising its most cited and best-known publications. Beard and Torres trace the beginning of the existing field of ERS to the early publications by Bostrom.<sup>14</sup> It was Bostrom's work in the early 2000s that first aimed to formalise the concept of an existential risk, notably via his 2002 paper "*Existential Risks: Analyzing Human Extinction Scenarios and Related Hazards*"<sup>15</sup> and the 2013 article "*Existential Risk Prevention as Global Priority*",<sup>16</sup> which have been cited a combined 1,046 times.<sup>17</sup> These two papers articulate the canonical definitions of "existential risk" and, along with Bostrom's 2003 paper "*Astronomical Waste*"<sup>18</sup> present the clearest distillation of the TUA.

The TUA also characterises almost every existential risk text with significant public profile. This includes trade-books such as *Superintelligence*,<sup>19</sup> *The Precipice*,<sup>20</sup> *Life 3.0*,<sup>21</sup> and *What We Owe the Future*.<sup>22</sup> Several culturally influential ideas such as a technological singularity<sup>23</sup> and longtermism<sup>24</sup> are influenced by the TUA. The techno-utopian worldview also appears to resonate with key funders. For example, Holden Karnofsky, the co-founder and co-chief executive of Open Philanthropy, strongly echoes the TUA in his "Most Important Century" series.<sup>25</sup>

The TUA we describe should be understood as an ideal type:<sup>26</sup> both the texts and thinkers under it may vary in specifics but converge in their broad vision. This is not an issue: despite many national variations we can speak of capitalism in general and recognise particular countries as being capitalist.<sup>27</sup>

### 2.1.2 Defining existential risk under the TUA

Bostrom provides two general formulations of existential risk. He initially defined it as "where an adverse outcome would either annihilate Earth-originating intelligent life or permanently and drastically curtail its potential".<sup>28</sup> Later, he provided a more refined definition: "one that threatens the premature extinction of Earth-originating intelligent life or the permanent and drastic destruction of its potential for desirable

future development”.<sup>29</sup> The latter paper provides a further typology of existential risks as:

- i. Premature extinction of humankind before reaching technological maturity;
- ii. Failure to reach technological maturity due to an unrecovered collapse, recurrent collapse, or plateauing;
- iii. Technological maturity being realised in an irredeemably flawed manner; or,
- iv. Ruination of Earth-originating intelligent life after technological maturity is reached but before its full potential can be realised.

According to this typology, the core feature of an existential catastrophe is a failure to attain a stable state of technological maturity, maintained for as long as is physically possible. Bostrom specifies this as a level of technological development, resource acquisition, and resource efficiency that allows for the highest feasible level of “economic productivity and control over nature”.<sup>30</sup> Failure is anything but the full exploration of possible options and the full exploitation of available matter.

A recent definitional reworking of this approach is the formulation of an existential catastrophe as “an event which causes the loss of a large fraction of expected value”<sup>31</sup> or less technically as a risk that “threatens the destruction of humanity’s longterm potential”.<sup>32</sup>

## 2.2 Transhumanism: Humans as a stepping stone

Transhumanism is the moral position that there is value in exploring posthuman and transhuman modes of being.<sup>33</sup> The results are to be beings — modified biological humans, cyborgs, androids, or digital simulants — whose lives are considered more valuable than current ones.<sup>34</sup> Transhumanists argue that these beings could achieve far longer, richer lives marked by net positive experiences.<sup>35</sup>

Achieving such lives would depend on further technological progress.<sup>36</sup> We would need three fundamental transformations: (i) protecting life; (ii) expanding cognition, and (iii) elevating well-being.<sup>37</sup> This can, for



example, take the form of achieving immortality, superintelligence, and a greater capacity for pleasure.

Although transitioning to a posthuman stage could of course entail the extinction of *Homo sapiens*, Bostrom contends that “the permanent foreclosure of any possibility of this kind of transformative change of humanbiologicalnature may itself constitute an existential catastrophe”.<sup>38</sup> Preventing existential risk is not primarily about preventing the suffering and termination of existing humans; it is focused on preserving humans so that they may give rise to a posthuman species that contains more value.

### 2.3 Total utilitarianism: Humans as containers

Total utilitarianism identifies moral rightness with the maximisation of well-being. Well-being could be interpreted in hedonistic, desire-satisfactionist, or objective-list theory terms.<sup>39</sup> People thus carry some unit of value, and the greater this value, the better. Total utilitarianism therefore demands that we maximise the total amount of value in the universe, with as many people coming to exist as possible, each person living an overall happy (i.e., net-positive) life, regardless of where or when these people come into existence. This equivalence in moral patient-hood between different “containers” of value here relies on the “impersonalist” or “non-identity” perspective, in which it is not relevant who is affected, only that someone is affected.

The utilitarian argument that the future should be an overwhelming moral priority relies on an assumption that the number of intelligent beings who could come to exist could be unimaginably large. Matheney, for example, estimates a low-range figure of  $10^{16}$ , assuming we remain Earth-bound.<sup>40</sup> Bostrom argues that if our descendants colonised as much of the universe as quickly as possible and converted celestial bodies into “computronium” for running simulations, this could result in some  $10^{38}$  simulated conscious beings per century in Earth’s supercluster alone.<sup>41</sup> Assuming wider interstellar exploration then we could even produce  $10^{58}$  happy simulations.<sup>42</sup> If we are alone in the universe and computations are run at colder temperatures towards the end of the heat death of the universe then even more could be achieved.<sup>43</sup>

The posthuman calculus, or the “astronomical value” thesis, is that the future could hold far greater value than the present due to the teeming masses of beings with higher levels of happiness than organic mortals can achieve. *Homo sapiens* is eclipsed by the towering shadow of the techno-utopian future.

## 2.4 Strong longtermism: The present in the shadow of the future

The third philosophical foundation is strong longtermism. Considering future generations as moral patients is by no means new. Notions of intergenerational justice, equity and fairness across the deep future have been extensively discussed for decades.<sup>44</sup> Strong longtermism goes a step beyond this and suggests that for some situations we may have an ethical imperative to select the choices expected to have the best effect on the long-run future,<sup>45</sup> and usually relies on a utilitarian calculus to justify this.

The best choice is often equated with the choice that has the highest expected value.<sup>46</sup> Expected value is calculated by multiplying the value of an outcome by the probability of it occurring. A calculus that numerically favours strong longtermist actions, such as reducing existential risk, rather than saving millions of today’s people, often relies on the assumption of continued technological development, happy future people, and interstellar settlements.<sup>47</sup> Thus, ensuring technological progress and maximising the quantity and expected quality of hypothetical future lives may be deemed more important than protecting current lives. We have no principled guidance about when and why a strong longtermist should prioritise living humans of today.

Transhumanism, total utilitarianism, and strong longtermism are a coherent, re-enforcing and complementary set of beliefs. Some versions of longtermism might be compatible with multiple theories of value,<sup>48</sup> but this is an ongoing area of study.<sup>49</sup> It is unclear how different those versions would be and what this would practically imply.<sup>50</sup> The reasons put forward for why other moral theories should care about existential risk include to cultivate civilisational virtues and/or to meet intergenerational obligations.<sup>51</sup> These reasons explain why different

moral theories should be concerned about human extinction, but not why they would support strong longtermism.

The non-utilitarian case for strong longtermism is, for now, weak. There have been too few attempts to understand whether different moral positions can support the belief that the vast majority of value lies in the future. There is countering evidence that this is not the case.<sup>52</sup> While it may be the case that a wide range of moral positions support caring about the future in a broad sense, it is not necessarily the case that each of these views yield the same practical implications. Hence, it is not appropriate to argue that the implications of strong longtermism follow from a diverse range of moral positions.

### 3. Representation and Existential Risk

#### 3.1 The TUA as a non-representative view

The TUA is not representative of what most humans alive now believe. Relying on the TUA, which is unrepresentative of many people's moral views today, can distort the analysis of existential risk. Representativeness itself says nothing about whether its philosophical pillars are wrong or right, but it is risky to rely exclusively on one unrepresentative approach given moral, political and empirical uncertainty. Theoretical work in ERS should be paired with and constrained by empirical studies that capture the range of existing intuitions about human extinction and longtermism. We must know the moral intuitions of the public, and when experts dismiss their moral intuitions as incorrect, they must have strong arguments to do so.

There is no consensus among philosophers on moral theory. Utilitarianism is not the most commonly held view. In one of the few surveys in the area less than a quarter (23.6%) of philosophers identified with consequentialism.<sup>53</sup> An even smaller number will be utilitarian, and a small number still will be total utilitarian. Techno-utopia offers futures of pleasurable, often virtual experiences, in which commonly valued attributes like purpose, virtue, love, and justice do not play a central role.

Transhumanism, too, is a niche perspective, and surveys reveal that those who identify as transhumanists come from a narrow demographic. The most recent high-quality survey, one that collected 760 responses from members of the World Transhumanist Association (now called Humanity+) in 2007, found that 90% of the respondents were male with a median age of 30–33 years old.<sup>54</sup> It is unknown how many people of the wider population would accept all or some of the premises of transhumanism if they were surveyed.

The implications of the TUA definition of existential risk also appear to be unrepresentative. According to the original definition of existential risk (a failure to attain technological maturity) technological plateauing — or a failure to spread beyond Earth — is an existential risk. Both near-term extinction due to nuclear war and a future in which humans persist sustainability and equitably for the next billion years without major technological progress are seen as existential risks: worst-case outcomes for humanity. Equating these outcomes as morally equivalent is likely unintuitive to many.<sup>55</sup>

The perspective that potential future lives are morally equivalent to existing lives may also be unintuitive to many. It is an active area of theoretical debate between philosophers, and we need more surveys that empirically query the moral intuitions of a wider population. Caviola et al.<sup>56</sup> find context-dependent support for adding future people. Schubert et al.<sup>57</sup> find context-specific overlap between surveyed intuitions of lay-persons and theoretical arguments by experts in ERS, although this depends heavily on survey framing. These are a commendable start. However, far more work is needed to understand what the wider population of the world wants from the far future. Ideally, this should be built not just on surveys, but also more deliberative practices (see Section 6).

The techno-utopian vision of the future, which combines three rather uncommon positions (transhumanism, total utilitarianism, and strong longtermism) and considers technological stagnation to be an existential risk, is likely a rare view among the global population. It may rise in popularity in the future, but presently it appears to be a fringe position.

### 3.2 The risk of a non-representative view

Tying the study of a topic that fundamentally affects the whole of humanity to a niche belief system championed mainly by an unrepresentative, powerful minority of the world is undemocratic and philosophically tenuous. Landemore defines the term “elite” as a group of people that would not likely be selected at random from its wider population and that is granted decision-making powers.<sup>58</sup> Under this definition, the field of existential risk is decidedly elitist at present. There are ways to mitigate against elitist research projects: diversifying the field and thus its policy recommendations, and democratising the evaluation of policies that are proposed by the field (see Section 6).

An obvious retort here would be that these are scholars, not decision-makers, that any claim of elitism is less relevant if it refers to simple intellectual exploration. This is not the case. Scholars of existential risk, especially those related to the TUA, are rapidly and intentionally growing in influence. To name only one example noted earlier, scholars in the field have already had “existential risks” referenced in a vision-setting report of the UN Secretary General. Toby Ord has been referenced, alongside existential risks, by UK Prime Minister Boris Johnson. Dedicated think-tanks such as the Centre for Long-Term Resilience have been channelling policy advice from prominent existential risk scholars into the UK government.<sup>59</sup>

The field also appears to be elitist in the more common-sense notion of being representative of a small stratum of people with disproportionate economic and political power. The main research centres are clustered in a few of the most elite universities in the world, with most of the field located in Oxford, Cambridge, or the San Francisco Bay Area. As noted earlier, the field is disproportionately supported by billionaires and millionaires. They not only hold financial power, but often advisory positions as well. The ideas of the TUA also closely echo popular Silicon Valley ideology.<sup>60</sup> Indeed, in 1995 Barbrook and Cameron referred to the techno-utopian ideas of the “Californian Ideology”, which distinctly echoes the TUA. The Californian Ideology began in Silicon Valley during the tech boom of the 90s and was

underpinned by a commitment to technological determinism and neoliberal economics.<sup>61</sup>

The point is a broader one: it is highly risky to grant privileged influence over the fate of *Homo sapiens* to a tiny minority. This is true for the study of existential risk, but more so for the implementation of policies that are meant to reduce existential risk, which will need to balance trade-offs between different interests. An attempt to reduce elitist interference in the study and implementation for existential risk mitigation is important due to moral uncertainty.<sup>62</sup> For every moral theory, there exist recommendations which fail to match our intuitions.<sup>63</sup> It is thus all the more important to know what choices would empirically be preferred by widening the range of people that are allowed to decide what risks are worth and not worth taking.

Some scholars associated with the TUA have written about moral uncertainty. This is excellent and should continue. They advance theories such as expected moral value (ranking alternative axiologies by their expected value,<sup>64</sup> or expected moral choice-worthiness (weighting moral theories by credence and combining them)<sup>65</sup> to navigate moral uncertainty. Issues remain for these approaches, including whether moral theories are comparable in this manner, whether empirical and moral uncertainty are equivalent, how such approaches are scaled up to collectives, as well as technical problems.<sup>66</sup> Using the suggested approaches tends to lead to total utilitarian perspectives outweighing others once large populations in the future are assumed.<sup>67</sup> There is room for considerably more research, and approaches to dealing with moral uncertainty have yet to be consistently and practically applied to existential risk.

### 3.3 Existential risk: Who is threatened?

The original definition of the techno-utopian paradigm is not concerned with humans *per se*. Instead, it is focused on Earth-originating intelligent life, and enhanced posthumans. This is a different inquiry to studying human extinction, and it is not obvious that it should be conducted under the banner of existential risk. The existing species *Homo sapiens* can be approached empirically, offering the opportunity to develop a science of existential risk; by taking an interest in the future of *Homo*

*sapiens*, scholars can approach existential risk reduction as a communal project, able to engage with the subject of inquiry — existing humans — and consider their individual preferences and visions of what a good future would look like.

## 4. Flawed Definitions, Frameworks and Tools: How Ambiguity and Unverified Assumptions Undermine Research into the Apocalypse

### 4.1 Ambiguous definitions

This section will look at the problem of defining existential risk. We look at three definitions that can be considered part of the TUA, where an existential risk is one that:

- a) “threatens the premature extinction of Earth-originating intelligent life or the permanent and drastic destruction of its potential for desirable future development”;<sup>68</sup>
- b) “threatens the destruction of humanity’s long-term potential” (through extinction, unrecovered collapse, or permanent dystopias);<sup>69</sup>
- c) “causes the loss of a large fraction of expected value”.<sup>70</sup>

The three definitions share a core feature: they are all fixated on future value. The tragedy to be averted is not the suffering or loss of existing humans, but rather the loss of future value or potential. All the definitions are motivated by future, long-term value. There are also some differences. In his extended typology (see Section 2.1) Bostrom’s definition enshrines a particular moral view by specifying desirable futures as technological maturity. The other two definitions are, at least in theory, more abstract and value-agnostic.

By leaving “value” and “potential” undefined, these latter definitions theoretically avoid the charge of existential risk as being a project of a niche philosophical view. According to Ord, our potential should be the entire set of possible futures. He suggests that we should first reduce existential risks to a minimum level to achieve “existential security” before undertaking a “Long Reflection”: a patient, collective discussion

of what exactly humanity's potential is. There are significant problems with this approach.

First, in practice, value is still expressed in techno-utopian terms. For example, the last chapter of *The Precipice* expands on a vision of humanity's potential: transhumanist space expansion receives ample attention and adoration. Unrecoverable civilisational collapse (a state in which technological progress is not ensured) is described as an existential risk. Here, "civilisational collapse" refers to a permanent reversion back to non-agricultural ways of living. It is not explained why the presence of agriculture, or many of the commonly assumed trappings of "civilisation", such as urbanism, writing, and states (although these rarely came as a coherent package; see Graeber and Wengrow)<sup>71</sup> would increase the likelihood of reaching our potential. For a techno-utopian, it does. For others who value virtue, freedom, or equality, it is unclear why a long-term future without industrialisation is abhorrent: it all depends on one's notion of potential. The definition is seemingly agnostic in the abstract, but in practice there are numerous signals that it expresses the same commitment to total utilitarianism and transhumanism.

Secondly, we need to define what our potential is before we can identify threats to it. How else would we know which risks to address? This is an inherent tension within *The Precipice* since we are supposed to achieve existential security before undertaking the Long Reflection. It is difficult to know if we have achieved existential security if we haven't defined what an existential risk is, since we haven't undertaken the Long Reflection to define our potential. A reasonable counter could be that, in theory, there are certain futures that almost no one would like to live in (such as nuclear winter), and that there may be certain risks (for instance, an asteroid strike) that would take lots of plausibly good options off the table. Extinction may indeed be an outcome which we could assume most people would agree we should avoid. Beyond this point of convergence, there may be far more disagreement on what futures are worth protecting.

Indeed, agreement on the badness of extinction and disagreement over our potential is evident in wider philosophical debates. Philosopher Elizabeth Finneron-Burns argues that extinction is wrong due to the suffering and psychological traumas it could cause, not



due to the prevention of millions or billions of people yet to be born.<sup>72</sup> Similarly, an intrinsic end value of humanity can be grounds to want to ensure the long-term survival of humanity, but not the potential for many additional future lives.<sup>73</sup> Others have argued that the badness of extinction is a generation-centred issue and neither the future masses or appeals to the natural lifespan and shape of humanity (there isn't one) provide sufficient grounds.<sup>74</sup> There is a strong case that different value theories should be concerned about extinction. There is not a compelling case that many are compatible with longtermist concerns about the deep future.

Third, the "Long Reflection" and Ord's definition assumes we can reduce risks to humanity's potential without choosing between conflicting values. This is almost certainly not the case. Maybe Ord means that we should reach existential security without closing off any possible futures, but retaining option value is presumably restricted to future options that are morally valuable. Indeed, the very notion of existential risk presumes that certain futures, such as dystopias, are to be avoided. We must define what is morally valuable to identify what is dystopian, and we cannot wait until a "Long Reflection" to do so.

Beyond the simpler domain of extinction, things get far murkier. Different plausibly good futures will often involve trade-offs against each other. Steering the world through an age of perils will involve difficult choices. Choices that will frequently have divergent answers depending on one's values. Look no further than the COVID-19 pandemic to see how comparatively smaller crises lead to clashes in values and understandings of what defines a good society. The field should provide clear delineations between risks that were identified as threatening across a broad swathe of value assumptions, and risks that are only threatening given a particular notion of potential. Some scholars may choose to stick with studying extinction risk, rather than trying to specify all possible good futures and the existential threats to them.

Deciding what risks are worth taking and which risks should be taken seriously will need to be a matter of reflection and collective decision-making if it is to respect moral uncertainty and diverse preferences.

Ord writes: “If we steer humanity to a place of safety, we will have time to think”,<sup>75</sup> but who is “we”? What appears a risk worth accepting to some will not be considered a risk worth taking by others. For instance, slowing technological progress is a risk to a transhumanist who sees our potential in overcoming human biological limitations, but accelerating technological progress may be perceived as a risk by others. In the following sections, we will list several examples of proposed mitigation efforts which show that judgements of strategies to mitigate risk depend on subjective notions of value. Generally, it is easier to identify universal extinction risks than universal existential risks.

Fourth, in practice the attempt at neutrality can end up masking rather than eliminating values from the analysis. Explicit commitments to transhumanist values in Bostrom’s definition of existential risk have the advantage of transparency. It is easier to reject or counterbalance a researcher’s perspective when their underlying values are clear. Abstract definitions can end up implicitly incorporating moral assumptions. This can happen unbeknownst to the researcher.

Fifth, the definitions are sufficiently ambiguous to render them unfit for use in a rigorous or replicable risk assessment. If we conceptualise an existential risk as the “permanent and drastic destruction” of desirable future development,<sup>76</sup> human potential,<sup>77</sup> or expected value,<sup>78</sup> then how severe does the destruction need to be? How significant must be the loss of human potential or expected value? How high does the likelihood of its permanency need to be? Even if a risk is judged sufficiently impactful, how large does the probability of the risk occurring need to be, before it can be considered as a legitimate existential risk? How should we compare such risk being incurred by some action against potential benefits of said action? Should risks based on speculation or thought experiments or naive technological extrapolation be treated seriously (see Section 5.2)? These decision-relevant ambiguities have not yet been clarified.

These are not minor, theoretical quibbles. Empirical and moral assumptions determine what is considered an existential risk and what mitigation efforts are recommended. Whether or not these choices are considered reasonable doesn’t depend on a replicable

framework. Instead, it relies on whether the judge shares the same assumptions.

An example: Ord<sup>79</sup> considers Artificial Intelligence (AI) to be the biggest contributor to existential risk within the next 100 years. The author's probability estimate relies on a survey of machine learning researchers,<sup>80</sup> a study with questionable methodological value for determining whether AI is an existential risk.<sup>81</sup> The policy recommendations for mitigating such risks in *The Precipice* support R&D into aligned Artificial General Intelligence (AGI), instead of delaying, stopping, or democratically controlling AI research and deployment.<sup>82</sup> These policy recommendations were echoed in the "Future Proof" report by the Centre for Long-Term Resilience,<sup>83</sup> which was aimed at UK policymakers. This recommendation seems to assume a kind of technological determinism (see below), or an implicit advocacy for building advanced technology for instrumental purposes. Either explanation echoes the TUA and leads to recommendations that support a particularly existentially risky course of action: developing advanced AI. Despite an explicit acknowledgment that AI could be a major contributor to risk and that slowing, delaying or halting development can help avoid the risk, the book recommends R&D, rather than e.g., citizen surveys, moratoria, or transparency measures.

What appears to be a risk worth accepting to some will appear to be a risk not worth taking to others. Take, for example, Bostrom's "Vulnerable World Hypothesis",<sup>84</sup> which argues for the need for extreme, ubiquitous surveillance and policing systems to mitigate existential threats, and which would run the risk of being co-opted by an authoritarian state.<sup>85</sup> It is a solution that some may find appealing and others appalling. Without a deliberation between different moral views, it should not be assumed that this risk is acceptable. Given disagreements about risk-taking, the field needs to ask who gets to put forward recommendations and who gets to choose between them (see Section 6).

Many definitions face the challenge that they are too abstract to allow for robust, replicable analysis, but the TUA's definitions of existential risk are particularly faulty. These definitions conflate the study of global catastrophe or human extinction with that of the

longtermist ethics of existential risk. The question of what futures are worth taking which risks for will always rely heavily on a notion of value. For this reason, we suggest scholars consider separating the following areas:

- *Extinction Ethics*: the study of the ethics of human extinction; the badness or goodness of human extinction given different ethical considerations.
- *Existential<sup>86</sup> Ethics*: the study of the ethical implications of different societal forms. This includes not just our potential in the deep future, but also how societies should be structured in the present. This in turn provides ground for defining what would constitute what the field currently calls an *existential* risk. An easier initial step here may be to specify dystopias that most value theories wouldn't want, rather than develop a widely shared notion of potential.
- *Catastrophic Risks*: the study of contributors to the occurrence and probability of global catastrophic events.
- *Extinction Risks*: the study of contributors (which include global catastrophic events) to the probability of human extinction.

While each inquiry can have both empirical and theoretical elements, the latter two lend themselves to scientific risk analyses, while the former two are more philosophical inquiries by nature. All the existing definitions currently conflate all of these. For many analyses, this is unnecessary and counterproductive. The media repeatedly makes the mistake of equating Ord's estimated one in six chance of existential catastrophe with extinction risk,<sup>87</sup> but these are meant to be drastically different concepts.<sup>88</sup> The study of extinction and existential (or longtermist) ethics does not need to be resolved to scientifically study risks, and it requires a different set of skills and procedures (see Section 6). There will be some necessary and fruitful areas of overlap. For instance, ensuring that measures to prevent and mitigate risks are proportional and not dangerous, we will need to have some debate on existential ethics.

Risk analysis will always be at least partly subjective and value-laden. Yet, it can be made more objective and scientific through precise definitions, the transparent statement of assumptions, and where possible, separating risk assessment from the study of extinction and existential ethics.

## 4.2 Arbitrary categorisations

ERS currently lacks a framework or methodology to categorise risks consistently, comprehensively, and rigorously. The TUA does not provide such a framework. It purports to distinguish between existential and catastrophic risks, but the distinction is hazy and arbitrary.

ERS currently distinguishes between existential risks and “global catastrophic risks”. All existential risks are global catastrophic risks, but not all global catastrophic risks are existential risks. Definitions of global catastrophic risks proposed in the literature vary but tend to focus on a significant global loss of human life (such as a loss of 10%).<sup>89</sup> Most include the great catastrophes of the past, such as the Black Death.

Under the earliest TUA definition, a catastrophe that does not jeopardise the attainment of technological maturity is assigned comparatively little moral consideration. This is due to the belief that such disasters have not influenced our long-term fate, and thus do not constitute existential risks,<sup>90</sup> and that limited resources would be wasted if they were directed towards global catastrophes that did not threaten technological maturity.<sup>91</sup> This extreme prioritisation of existential risks downgrades the importance of addressing other global catastrophes. What under Bostrom’s view would plausibly be considered “feel-good projects of suboptimal efficacy”, and what falls short of “efficient philanthropy”,<sup>92</sup> could in fact very much be worth the attention of those who study human extinction.

Ord<sup>93</sup> presents a more moderate version of this original TUA framework. He considers it justified to spend some resources on large catastrophes, because those catastrophes could indirectly amplify (but not directly cause) existential risk. His framework still draws clear distinctions between existential and catastrophic risks framing some

as direct “risks” and others as indirect “risk factors”, prioritising the reduction of direct risks.

We lack the necessary understanding of how human society and the Earth system operate to make such neat, surgical distinctions. Whether a single global catastrophe can or would fundamentally alter the trajectory of humanity is one of the great unanswered questions of history. For now, we struggle with forecasting GDP even a year ahead<sup>94</sup> and in fact reliably make incorrect forecasts about population growth, despite knowing all essential variables.<sup>95</sup> In a complex adaptive system it may be impossible to forecast how small changes will affect the *longue durée*. Neat distinctions between GCRs and existential risks presume a level of systemic knowledge and certainty we currently do not possess. The TUA thus far does not offer the tools to make any fine-grained, credible separation between existential and non-existential global catastrophes. It has not offered explanations of how such events affect extinction and existential risk.

This is not to say that some persistent, long-term, societal trends cannot be identified and understood. For instance, the historian Walter Scheidel has put forward a compelling empirical case that (intra-country) wealth inequality increases inexorably until a great leveller (a state collapse, pandemic, revolution, or mass mobilisation warfare) resets the playing field.<sup>96</sup> Investigating these trends will be critical to foreseeing how catastrophic risks are produced and could unfold deep into the future. Such an analysis does not require or justify the crude split between global catastrophes and extinction risks.

Under the TUA, an existential risk is understood as one with the potential to cause human extinction directly or lead us to fail to reach our future potential, expected value, or technological maturity. This means that what is classified as a prioritised “risk” depends on a threat model that involves considerable speculation about the mechanisms which can result in the death of all humans, their respective likelihoods, and a speculative and morally loaded assessment of what might constitute our inability to reach our potential.

Imagining pathways to human extinction (kill mechanisms) invariably requires some creativity and speculation. This has meant that some areas of risk (e.g. AI) which are not as empirically constrained are often prioritised above others for which we have far more empirical

data (e.g. climate change). This has led some (non-peer-reviewed) publications (including a Google Doc) to conclude that climate change is not an existential risk.<sup>97</sup>

This is the wrong question to ask. Adding up the predicted impacts of a selection of hazards and asking “Will the total of these kill everyone?” is a simplistic and ineffective way of conducting risk analysis. This is not how risk unfolds in reality: hazards interact with networks of societal vulnerabilities and responses as well as each other, and can trigger cascading failures.<sup>98</sup> We need to consider different pathways and ways in which climate change (or any other source of risk) can contribute to the overall level of extinction or catastrophic risk we face.<sup>99</sup> The question of “Is this an existential risk?” is naive. We should instead ask: in a given world-state (with structure, vulnerabilities, and the capacity for change) how much will a given process or event increase the overall likelihood of human extinction, and what are the plausible<sup>100</sup> pathways for it to contribute to extinction risk?

A field looking for the *one hazard to kill them all* will end up writing science fiction. More speculative risks are prioritised because a seemingly more complete story can be told and speculative mechanisms by which AI could kill every human can seemingly not yet be ruled out.

In practice this could be addressed by lowering the threshold of what risks should be treated as relevant to extinction and include more of what Ord calls “risk factors”, such as those commonly thought of as Global Catastrophic Risks (GCRs). Global catastrophes and responses to GCRs can give vital insight into vulnerabilities that should be mended and resilience factors that should be enhanced.

The distinction between “direct” and “indirect” risk factors also depends on speculation. Direct risks appear to be those for which we can tell a story about how they might “directly” (presumably limited to third or fourth order effects) lead to the extinction of *Homo sapiens*. Unaligned AI, for example, is often considered a direct risk, but the story about extinction from AI is far from complete.<sup>101</sup> Strong expert disagreement regarding risks from AI<sup>102</sup> is testament to how debatable the empirical foundations of “direct” risk pathways of AI still are. Additionally, risks originating from AI could come in many forms, each relying on speculation about different kill mechanisms and assumptions about the nature and use of a system that has never been

built. Surely, not all these pathways are equally direct, and yet AI is prioritised as a seemingly homogenous hazard-cluster across key texts within ERS.<sup>103</sup>

A risk perception that depends so strongly on speculation and yet-to-be-verified assumptions will inevitably (to varying degrees) be an expression of researchers' personal preferences, biases, and imagination. If collective resources (such as research funding and public attention) are to be allocated to the highest priority risk, then ERS should attempt to find a more evidence-based, replicable prioritisation procedure.

### 4.3. Simplistic risk models

#### *4.3.1 Complex vs. crude risk assessments*

Risk assessment has evolved dramatically in past decades. Scholars now commonly analyse systemic risk (the ability for a single disruption to cascade into systems failures),<sup>104</sup> how risks can cascade across borders and sectors,<sup>105</sup> and how failures in critical systems can synchronise and reinforce each other.<sup>106</sup> This has led to new forms of complex risk assessment, particularly in climate science and disaster risk reduction. The Intergovernmental Panel on Climate Change (IPCC) sees risk as composed of vulnerabilities, hazards, and exposures, as well as response risks.<sup>107</sup> Similarly, others have suggested that a complex risk assessment needs to consider four determinants of risk (hazard, vulnerability, exposure, response) as well as how risks link and cascade. Understanding the common drivers across each of these determinants is critical to mitigation efforts.<sup>108</sup>

How we assess risk is fundamental to what we consider as an existential or catastrophic risk. For instance, Ord focuses on super volcanoes as a potential existential catastrophe. However, lower magnitude volcanic eruptions could have catastrophic impacts due to their cascading effects and the vulnerable nature of critical infrastructure systems.<sup>109</sup> Similarly, stratospheric aerosol injection<sup>110</sup> does not appear to pose direct risks that would classify it as a global catastrophic threat. Yet, this depends on how it is deployed, the world in which it operates, and the level of warming it is masking. If another



calamity, such as a volcanic eruption, solar flare, or nuclear war, destroys the mitigation system for a prolonged period, the ensuing “termination shock” (rapid global warming over a short timeframe) would likely result in catastrophic effects.<sup>111</sup> A more complex risk assessment will be more difficult to do, but it will be more accurate, realistic, and informative.

Most existential risk texts take a simpler, hazard-centric approach. They tend to focus on a few selected hazards: biologically engineered pandemics, Artificial General Intelligence (AGI), nuclear war, climate change, and asteroid strikes.<sup>112</sup> As currently framed, TUA equates risk with hazard and ignores the wider literature on risk assessment in fields such as disaster risk reduction. It is also unclear how the TUA suggests systematically clustering, prioritising, or analysing these hazards: current attempts rely on simplistic categories of “Natural”, “Anthropogenic”, and “Future” hazards<sup>113</sup> or presenting the selected hazards as the ones worth discussing without explanation. There have been some recent attempts to provide alternative frameworks<sup>114</sup> but these have found little application thus far, and still do not consider response risks and many other relevant areas. Research efforts are often split across the lines of these different hazards. Working across them as part of a more complex risk assessment could offer novel insights.

#### 4.3.2 *Technological determinism*

The choice to structure risk assessment this way has not been explained or defended. It may have been chosen due to an implicit techno-determinist threat-model: the TUA often appears to assume an exogenous threat model in which existential hazards naturally and apolitically arise from inevitable and near-autonomous technological progress. The TUA rarely examines the drivers of risk generation. Instead, key texts contend that regulating or stopping technological progress is either deeply difficult, undesirable, or outright impossible.<sup>115</sup> Bostrom proposed a “Technological Completion Conjecture”: if technological developments do not cease, then all important, basic technological capabilities will be obtained in the long run.<sup>116</sup> Others offer a more sophisticated view, in which military-economic competition exerts a powerful selection pressure on technological development.

This “military-economic adaptionism” constrains sociotechnical change to deterministic paths. Technologies that gift a strong strategic advantage will almost certainly be built.<sup>117</sup> Many in the related Effective Altruism community disregard controlling technology on the grounds of a perceived lack of tractability.<sup>118</sup>

Whether it is technological determinism, the more nuanced military-economic adaptationism model, or concerns around tractability, the result is the same: regressing, relinquishing, or stopping the development of many technologies is often disregarded as a feasible option.

The proposed alternative is “differential technological development”: speeding up and slowing down different technologies to ensure they occur in the safest order possible. Why this is more tractable or effective than bans, moratoriums, and other measures has not been fully explained and defended (see Section 5.1 for further discussion). This could be interpreted as hard-nosed pragmatism, an argument that we can stop technologies from being built, but it will not be an efficient or prudent use of resources. Again, a compelling analysis has not been made for why this is the case, and in practice this ends up looking identical to technological determinism. The irony is that if the world is locked into the development of dangerous technologies, then we are already in a “lock-in scenario” so dreaded by many within the TUA.<sup>119</sup> In the eyes of the TUA, the range of future options available to humanity is already greatly restricted.

It is unclear whether technological determinism in the TUA is descriptive or prescriptive. It could be a genuine belief that controlling technology is infeasible. It could also be that under the TUA unabated technological progress is vital to achieve technological maturity and avoid existential risk.

In any case, the assumption of technological determinism leads scholars to focus on hazards, rather than, say, exposure, maybe because there appears to be no point in trying to reduce humanity’s exposure to a technology since the development of the technology is assumed inevitable. It is then merely a question of whether benevolent technologies are built first. Attempts to change political or economic drivers of different risks get less attention. This unstated threat model leads to sharp divergence from modern developments in risk analysis towards a crude hazard-centrism.

Importantly, assumptions around technological determinism are highly contested. Indeed, technological determinism is largely (for better or worse) derided and dismissed by scholars of science and technology studies.<sup>120</sup> We have historical evidence for collective action and coordination on technological progress and regress. One example is weather modification. Early attempts were made by the US during the Vietnam War to use weather modification technologies to extend the monsoon season and disrupt enemy supply chain.<sup>121</sup> The introduction of the 1976 Convention on the Prohibition of Military or Any Other Hostile Use of Environmental Modification Techniques (ENMOD Convention) seems to have successfully curtailed further research into the area.<sup>122</sup> The assumptions on technological progress must be thoroughly examined, empirically and theoretically, before they should be used to determine policy and mitigation actions.

There are enough exceptions to doubt that any strategically powerful technology will be due to competition between (largely) rational actors (usually states). Many important technologies such as glass and steam engines were used for ceremonial purposes for centuries before being redirected towards practical purposes.<sup>123</sup> During the early industrial revolution, water mills were more reliable and efficient than coal-fired steam engines. The latter were adopted not because of their inherent superiority, but because they could be located in urban areas with a large and desperate population, which appealed to early capitalists.<sup>124</sup>

Any approach to existential risk will struggle to find frameworks that are comprehensive yet elegant and practical. It will need to be transparent about its empirical assumptions, including on how risks are created. For now, the hazard-centrism and opaque technological determinism of the TUA provides a framework that is overly simplistic, unduly curtails the available mitigation options, and provides no compelling method to understand or address the common drivers behind risk determinants. It is inadequate for the grand challenge of understanding and mitigating extinction risks.

There is room for different empirical worldviews, different frameworks, and different moral positions to be considered in the study of existential risks. We do not at all recommend that hazards should no longer be studied. Similarly, speculation will always be a part of assessing

unseen risks, but a science of extinction should adopt frameworks which minimise the need for this aspect. Different frameworks will make different predictions about what policy and research efforts appear to plausibly reduce risk.

#### 4.4 Inappropriate translations of theory into practice

Mitigating existential risk requires decision-making under uncertainty. Decision theory in the context of ERS and longtermism is an active area of research. For now, we want to caution against applying idealised decision-theoretic results to the evaluation of risky choices in practice. This is because empirical uncertainty can affect the applicability of results that hold in theory and because expressing subjective notions of risks and benefits numerically can provide a false sense of certainty.

Take Expected Value (EV), defined as the value of the outcome multiplied by the probability of it occurring. The TUA extensively uses Expected Value calculations to justify its own approach and prioritisations. For example, Bostrom argues that existential risk mitigation should be prioritised over other altruistic acts: if there is just a 1% chance of  $10^{54}$  people coming to exist in the future, then “the expected value of reducing existential risk by a mere one billionth of one billionth of one percentage point is worth a hundred billion times as much as a billion human lives”.<sup>125</sup> According to Bostrom, “even the tiniest reduction of existential risk has an expected value greater than that of the definitive provision of any “ordinary” good, such as the direct benefit of saving 1 billion lives”.<sup>126</sup> Elsewhere, Millett and Snyder-Beattie use EV to argue for reducing risks from biological pathogens.<sup>127</sup>

While EV is a useful theoretical tool in a range of contexts, in practice, it is hard to apply rigorously when working with the generally low and highly uncertain probabilities characteristic of existential risks. The TUA applies expected value theory to the very areas where it faces the most pitfalls, that is, situations of deep uncertainty and low information about probabilities.<sup>128</sup> Ord attempts to estimate probabilities of existential catastrophes caused by various hazards for communicative purposes, most notably the aforementioned one in six figure.<sup>129</sup> While this has been successful in terms of public communications (the estimate has

been widely reported in the press), it is unclear how to evaluate the accuracy of these estimates or whether his methodology for arriving at them is sufficient to warrant the sense of scientific credibility that numbers inevitably imply to the lay person. In addition, to evaluate the EV of mitigating an event, one must decide upon a particular conception of value. Thus, personal intuitions and non-representative moral preferences are at risk of being captured and made to appear objective by numbers. How can policy recommendations be considered robustly good and replicable if they are evaluated on subjective assessments of probabilities and hidden ethical assumptions? EV is still unsuitable to be relied upon in practice for estimating human impact on the value of the long-term future.<sup>130</sup>

Furthermore, EV and decision theories more widely are affected by Pascal's Mugging,<sup>131</sup> as well as what has been called *fanaticism*.<sup>132</sup> We know of no pragmatic and consistent response to those challenges yet. Fanaticism describes how we may be required to put considerable effort into mitigating terrible events with an arbitrarily minuscule probability of occurring. Similarly, Pascal's Mugging describes how vast or near-infinite quantities of value can overwhelm even the most minuscule probabilities: the term comes from a thought-experiment in which Pascal is conned by a self-proclaimed wizard who promises to magically grant 1,000 quadrillion happy days in exchange for his wallet.<sup>133</sup> The probability that the grifter is a powerful wizard is infinitesimally low, but not zero, and outweighed by the expected utility of so many happy days. The Pascal's Mugging problem arises when applying EV to existential risks as defined within the TUA. Any risk that could prevent technological maturity, no matter how small, should be taken seriously and acted on due to the sheer amount of expected value at stake. Scholars have suggested that Pascal's mugger returns to swindle another unsuspecting victim, but this time using existential risk studies and longtermism.<sup>134</sup>

Both these challenges are usually evaded by claiming that hazards like AGI have an unambiguously high enough probability of occurring this century to merit considerable action.<sup>135</sup> This only side-steps the question in the case of AGI in particular (assuming that these doubtful estimates can be trusted) and does not tell us whether highly unlikely or speculative risks should be acted on. There are good reasons to defend

fanaticism<sup>136</sup> and further theoretical work might resolve the challenges presented here, but while the integration of theory and empirical work is still ongoing, we should consider drawing pragmatic lines between mere speculation and risks humanity should significantly focus on. For instance, a simple threshold or plausibility assessment<sup>137</sup> could protect the field's resources and attention from being directed towards highly improbable or fictional events.

## 5. The Risks of Studying Existential Risk

How could the study of global risks and longtermism contribute to catastrophe? The worst-case outcome is not that existential risk remains unaffected, or that resources are wasted on incorrect speculations (although these are problems). Instead, it is that these risks are aggravated by research into them. This issue must be addressed for any approach to ERS. Unfortunately, the TUA appears to be particularly prone to both ignoring response risks and aggravating them.

### 5.1 A risky road to safety

Mitigating risks incurs a risk of its own. As noted earlier, this is a fundamental part of sophisticated risk assessments, including those used by the IPCC.<sup>138</sup> Not all approaches incur the same risks. While we will not compare the techno-utopian approach against alternatives in this chapter, we think the TUA is prone to an especially high level of response risk.

The zealous pursuit of technological development, according to proponents of the TUA, accounts for the vast majority of risk over the coming centuries. The risk of human extinction from natural hazards is likely low, with an upper-bound of less than one in 14,000<sup>139</sup> and a best guess of around 0–0.05% per century.<sup>140</sup> In contrast, several scholars of existential risk place the likelihood of an existential catastrophe far higher at 1/6,<sup>141</sup> or >1/4<sup>142</sup> over the coming century. Rees puts the chance of collapse or extinction by 2100 at 1/2.<sup>143</sup> This discrepancy is mainly due to anthropogenic risks arising from climate change, nuclear weapons, synthetic biology, and artificial intelligence (AI).

If the lion's share of extinction risk stems from emerging technologies, why do we rarely ask how to stop dangerous developments? This option is usually considered infeasible,<sup>144</sup> or outright impossible.<sup>145</sup> This may be due to the exogenous threat model and technological determinism of the TUA. Since halting the technological juggernaut is considered impossible, an approach of differential technological development is advocated.<sup>146</sup> This involves trying to develop beneficial and protective technologies (aligned Artificial Intelligence<sup>147</sup> is often used as an example) first, before proceeding to riskier options.

It is not clear how scholars plan to reliably determine which non-existing technologies will be more or less risky years or decades in advance. Even if we did have such a refined vision of the future, it is unclear why a precise slowing and speeding up of different technologies (which are interlinked and presumably require a set of fine-tuned regulatory tools) across the world is more feasible or effective than the simpler approach of outright bans and moratoriums.

More importantly, in the TUA the stark choice between one of only two destinies — technological maturity or existential catastrophe — is a *fait accompli*. The path to techno-utopia appears to be the only one available, despite its risks. From a techno-utopian perspective, a failure to build these dangerous, powerful technologies is an existential risk. Bostrom, aware of the tension arising from recommending the (albeit careful) development of technologies, warns: "We should not blame civilization or technology for imposing big existential risks. Because of the way we have defined existential risks, a failure to develop technological civilization would imply that we had fallen victims of an existential disaster. [...] Without technology, our chances of avoiding existential risk would therefore be nil".<sup>148</sup> The TUA dramatically restricts the options available for avoiding existential catastrophe.

Furthermore, pursuing a techno-utopian future is dangerous and may come with considerable cost. It has already been noted that the attempt at colonising space and an expansion of technological capabilities could end in catastrophe if it foments a new arms race and large-scale warfare.<sup>149</sup> Upgrading the human body could construct a biological caste system, where, an enhanced, genetic elite could oversee a subjugated, unenhanced, "inferior" class.<sup>150</sup> These are not far-flung

speculations. Researchers already actively monitor, evaluate and debate near-term bio-engineering enhancements and their ethical implications.<sup>151</sup> Enhanced inequality would not only be unjust, but also amplify many social ills.<sup>152</sup> Similarly, a horizon scan under the WHO Science and Research Division has noted that the pursuit of advances in the life-sciences could produce many technologies that could be easily misused to cause harm. These range from using bioregulators for the delivery of bioweapons to the use of deep learning algorithms to identify novel biological pathogens.<sup>153</sup>

Existential risk does not need to be defined in reference to technological maturity, nor does it need to be accompanied by these response risks of accepting or even speeding up disruptive technologies. A different vision for a good future could lead to dramatically different policy recommendations. A less determinist view of technological change and pessimistic view of political change would open up a plethora of other interventions. A democratic approach to ERS will provide ample room for different moral and empirical assumptions to affect the assessment, discussion, and negotiation of collective risk-taking. Other paths and approaches may be more risk-averse and must be explored if humanity wants to safely reduce existential risk.

## 5.2 High stakes: Existential exceptions and a risk-averse approach to existential risk

### 5.2.1 *The Stomp Reflex*

There is a long history of security threats being used to enable draconian emergency powers. Emergency powers are intended to be conservative: to protect existing legal and political structures in a period of tumult. The logic is that drastic times call for drastic measures. To protect institutions, emergency powers allow governments to disregard existing laws and exempt themselves from judicial or democratic restrictions and oversight. Rather than protect, such measures are often abused to erode and transform fundamental political structures; when trying to centralise and extend state powers, fear is a powerful



justification. The larger the fear, the easier it is to justify more potent emergency powers. If the perceived threat is human extinction, then the measures could be extreme.

Recent examples abound. The Patriot Act, adopted by the US just 45 days after the 9/11 attack, allowed for a range of draconian actions, including indefinite detainment of migrants without criminal prosecution. The provisions broadly underpin the current US surveillance network, most notably through the NSA's PRISM program, which Edward Snowden exposed in 2013. The *War on Terror* became a useful cover for the creeping power of the US security apparatus. This despotic drift is not a purely historical threat. Clauses for states of emergency have spread over past decades<sup>154</sup> and 2020 marked the highest ever use of such measures.

The prolific use of emergency powers can lead to the creation of a *state of exception* in which the sovereign transcends the regular rule of law. Temporary measures become permanent, and spill into the operation of the legal system.<sup>155</sup> The transition from the Roman Republic to the Roman Empire, the fall of the Weimar Republic in the 1920s and 1930s into the Nazi regime, and many other political declines were underpinned by the normalisation of emergency powers.<sup>156</sup>

The irony is that emergency powers rely on an inaccurate understanding of human nature and disaster risk. Emergency powers inevitably empower those atop hierarchies, despite abundant evidence from disaster risk reduction and other fields that while mass panic is a myth, the risk of elite panic and elite co-option of catastrophes is real.<sup>157</sup> There is even evidence that such a response worsens crisis. One study of natural disasters found that the larger the number of emergency provisions used by an executive, the higher the fatalities (controlling for disaster severity and size).<sup>158</sup> This is a "*Stomp Reflex*": governments using emergency powers to reassert and veil systems of authority. Such a response is counter-productive and ultimately shifts power into the shadows, away from transparency and public accountability.<sup>159</sup>

Existential risk is the perfect excuse for enacting the Stomp Reflex. Indeed, catastrophic hazards such as nuclear weapons have already been used to justify anti-democratic shifts. In the US, the accumulated

nuclear stockpile and threat of sudden war justified a profoundly autocratic move: a single individual — the President — was given the ability to launch nuclear attacks. Richard Nixon once boasted “I can go into my office and pick up the telephone and in twenty-five minutes seventy million people will be dead”. As sociologist Elaine Scarry has argued, the nuclear decision-making apparatus violates constitutional rights, the deliberative nature of democracy, and any social contract. The world lives in the shadow of a “thermonuclear monarchy” rather than a democracy.<sup>160</sup>

Nuclear weapons also saw a revolution in secrecy in the US. The threat of thermonuclear war was used as a justification to construct unprecedented levels of secrecy in the military and intelligence communities. These were of dubious efficacy in preventing the spread of nuclear weapons, but they did have the effect of eroding transparency and democratic control over the military industrial complex.<sup>161</sup>

The best empirical example of policy responses to an existential threat do not inspire confidence. In the US at least, the threat of thermonuclear war spurred dramatic reforms that made for a less democratic and open state, but not necessarily a safer one.

### *5.2.2 Survival through security, surveillance and suppression*

Any approach to understanding and mitigating existential risks runs the risk of becoming *securitised*. Securitisation refers to a discursive manoeuvre that moves an issue from the arena of normal politics to that of national security, making it more likely to permit emergency powers and be placed under the control of unelected military and intelligence officials. Moves towards thermonuclear monarchy and elevated secrecy were largely underpinned by neorealist foreign policy and game theory developed in such a context.<sup>162</sup> This is not to say that all securitisation approaches are equally dangerous; some are far more likely to enable authoritarian responses.

There are reasons to expect that the TUA is particularly vulnerable to misuse. As philosophers such as Peter Singer and Phil Torres have noted, if the world is viewed from the TUA’s lens of existential

risk, then we run the risk that almost any action is justified if it is believed to improve our chance of surviving to expand beyond Earth.<sup>163</sup> Problems which are not considered to be an existential risk dwindle into irrelevance, as other values are sacrificed on the altar of expected astronomical value.

This is not to say that most believers of the TUA are intent on using it to justify morally abhorrent actions, nor that they are unaware of these weaknesses. Rather, we argue that the basic logic of protecting a high-tech future of astronomical value could be easily co-opted. Marx never intended for communism to justify brutal dictatorships. Nonetheless, it was easily twisted by Stalin and others to do so.

Scholars of existential risk have already shown some proclivity for invoking security, whether it be for “existential security”<sup>164</sup> or “epistemic security”.<sup>165</sup> Extreme emergency responses have also been raised. Bostrom’s “Vulnerable World Hypothesis”<sup>166</sup> identified the combination of ubiquitous surveillance, preventative policing, and global governance (understood to be “a world government; a sufficiently powerful hegemon or a highly robust system of inter-state cooperation”)<sup>167</sup> as a comprehensive agenda to protect against possible technological hazards that could devastate civilisation. He proposes a typology of four potential threats: “easy nukes” (readily accessible and easy to use weapons of mass destruction), “safe first strike” (the ability to safely destroy others with impunity), “surprising strangelets” (experiments that could harbour an unforeseen, or foreseen but low-probability catastrophe); and those in which the accumulation of minor damages by individuals eventually accumulate into global catastrophe.

Bostrom’s preferred solution — extreme preventative policing and widespread surveillance — could involve the mandatory use of ironically-named “freedom tags” fitted with multiple cameras and microphones to continuously track individual behaviour. These would be distributed to all citizens and monitored by state employees — “freedom officers” who themselves are watched by artificial intelligence to prevent misuse — who can order preventative interventions using drones or police.<sup>168</sup>

Both the journal paper (published in the journal *Global Policy*), as well as public-facing spin-off articles<sup>169</sup> about the Vulnerable World Hypothesis feature clear policy recommendations. A box in the 2019 paper titled “Policy Implications” includes the recommendations that dealing with “black balls” (technological innovations that by default destroy the world) would require “a system of ubiquitous real-time worldwide surveillance. In some scenarios, such a system would need to be in place before the technology is invented”.<sup>170</sup> The public-facing article from 2021 asks: “If you find yourself in a position to influence the macroparameters of preventive policing or global governance, you should consider that fundamental changes in those domains might be the only way to stabilise our civilisation against emerging technological vulnerabilities”. It is not difficult to foresee how such ideas could provide grounds for aspiring autocrats to subvert democratic institutions in the face of global threats.

Bostrom<sup>171</sup> also includes a discussion of pre-emptive strikes. He describes the responsibility and need for nations to (on some occasions) enact pre-emptive, unilateral infringements of sovereignty. If extinction threatening technologies (he imagines biosphere-destroying nanobots) are not controlled under international treaties, “the mere decision to go forward with development of the hazardous technology [...] must be interpreted as an act of aggression” and would justify pre-emptive infringement of national sovereignty.<sup>172</sup>

There is a clear danger in authoritative recommendations based on speculative thought experiments. Scholars using the TUA providing recommendations for surveillance and pre-emptive measures in the name of avoiding catastrophe could contribute to birthing the very dystopias they fear.

There is little evidence that the push for more intrusive and draconian policies to stop existential risk is either necessary or effective. It is empirically dubious to think that we cannot halt or delay the development or spread of dangerous technologies. Nor is it convincing that surveillance measures would prove effective. We have little to no evidence that the use of mass surveillance has been effective at preventing terrorist attacks in the US.<sup>173</sup> Moreover, the main creators of such hazards — the *Agents of Doom* — are often the very people who control the surveillance apparatus: military industrial complexes,

enormous technology firms, and powerful states.<sup>174</sup> At worst, the knee-jerk reaction of surveillance and preventative policing to prevent a speculative calamity could simply create one of its own: entrenched authoritarianism.

The obvious option to discontinue certain technological developments — if we assume that the Vulnerable World Hypothesis is true — is considered “hardly realistic” and “extremely costly, to the point of constituting a catastrophe in its own right”.<sup>175</sup> Ord, too, warns of so-called “desired dystopias”, in which an ideology (or manipulation and surveillance) has corrupted our choices to the extent that, for example, we “completely renounce further technological progress”.<sup>176</sup>

Under the TUA we appear to be trapped. We either develop technologies which carry immense risk, such that ubiquitous surveillance becomes necessary, or we cease technological development only to manifest the existential risk of failing to reach a technologically mature utopia. This trap is an idiosyncratic feature of the TUA.

There are more options for humanity than merely picking between two highly risky paths. Ord confidently asserts that ceasing any further technological development would “ensure our destruction at the hands of natural risks”,<sup>177</sup> but we have seen no convincing analysis that shows we could not safeguard our survival with current technologies and re-directed resources. Indeed, it is unclear why we cannot develop technology to address threats from asteroids and supervolcanoes without indulging in the entire range of dangerous inventions. The opinion that all technological progress must continue at all costs is a dangerous one.

Scholars of existential risk need to be vigilant to response risks, and risk-averse in their own suggested interventions. This means we must be truthful about uncertainty, consider the worst-case outcomes of our actions, and verify the acceptability of proposed interventions by subjecting them to democratic oversight.

## 6. Democratising Risk

There is an intimate and neglected relationship between existential risk and democracy. Democracy must be central to efforts to prevent and mitigate catastrophic risks. It is also an antidote to many of the problems manifest in the TUA. Do those who study the future of humanity have good grounds to ignore the visions, desires, and values of the very people whose future they are trying to protect? Choosing which risks to take must be a democratic endeavour.

We understand democracy here in accordance with Landmore as the rule of the cognitively diverse many who are entitled to equal decision-making power and partake in a democratic procedure that includes both a deliberative element and one of preference aggregation (such as majority voting).<sup>178</sup> Decision-making procedures are not either democratic or non-democratic, but instead lie on a spectrum. They can be more or less democratic, inclusive, and diverse.

We posit three reasons for why we should democratise research and decision-making in existential risk: the nature of collective decision-making about human futures, the superiority of democratic reason, and democratic fail-safe mechanisms.

Avoiding human extinction, or crafting a desirable long-term future, is a communal project. Scholars of existential risk who take an interest in the future of *Homo sapiens* are choosing to consider the species in its entirety. If certain views are excluded, the arguments for doing so must be compelling.

Democracy will improve our judgments in both the governance and the study of existential risks. Asking how our actions today influence the long-term future is one of the most difficult intellectual tasks to unravel, and if there is a right path, democratic procedures will have the best shot at finding it. Hong and Page<sup>179</sup> demonstrate both theoretically and computationally that a diverse group of problem-solving agents will show greater accuracy than a less diverse group, even if the individual members of the diverse group were each less accurate. Accuracy gains from diversity trump the gains from improving individual accuracy. Landmore<sup>180</sup> builds on this work to advance a probabilistic argument that inclusive democracies will, in expectation, make epistemically superior choices to oligarchies or even the wise few. This is supported

by promising results in inclusive, deliberative democratic experiments from around the world.<sup>181</sup> In the long run, democracies should commit fewer mistakes than alternative decision-making procedures. If this is true, it should improve the accuracy of research efforts and decision-making. We are more likely to make accurate predictions about the mechanisms of extinction, probable futures, and risk prevention if the field invites cognitive diversity, builds flat institutional structures, and avoids conflicts of interest.

There are many ways to consider the interests of the many. Democratic assemblies could allow global citizens to deliberate about the futures they prefer, citizens could be surveyed, and the field of ERS itself could be diversified. At the moment, the field is, as many academic disciplines are, unrepresentative of humanity at large and variably homogenous in respect to income, class, ideology, age, ethnicity, gender, nationality, religion, and professional background. The latter issue is particularly true of existential risk, which, despite being an inherently interdisciplinary endeavour, is at the highest levels dominated by analytic moral philosophers. We need to be vigilant to what perspectives are not represented in the study of existential risk. An awareness of bias will go some way towards mitigating its negative effects. To get close to replicating the cognitive diversity found among humans, we must begin by inviting different thinkers with different values and beliefs into the field.

Democracies can limit harms. Any approach to mitigating existential threats could create response risks, and the TUA seems particularly vulnerable to this. Despite good intentions and curiosity-driven research, it could justify violence, dangerous technological developments, or drastically constrain freedom in favour of (perceived) security. If we hope to explore ideas but minimise harms, democracies can be used to moderate the measures taken in response to harmful ideas. It seems, for example, vanishingly unlikely that a diverse group of thinkers or even ordinary citizens would entertain the idea of sacrificing one billion living, breathing beings for an infinitesimal improvement in reaching an intergalactic techno-utopia. In contrast, the TUA could recommend this trade-off.

The democratic constraint of extreme measures may simply be a form of collective self-interest. Voters are unlikely to tolerate global

catastrophic risks (GCRs), which incur the death of a sizeable portion of the electorate, if they know they themselves could be affected. We expect that scholars who do not support sacrificing current lives in the name of abstract calculations, but would still like to explore the use of expected value theory in existential risk, will be in support of democratic fail-safe mechanisms.

Empirically, this fail-safe mechanism seems to work. Even deeply imperfect democracies, like the ones we inhabit now, often avert detrimental outcomes. Democracies prevent famines<sup>182</sup> (although not malnutrition).<sup>183</sup> They make war — a significant driver of GCRs — less likely.<sup>184</sup> The inclusion of diverse preferences in democracies, such as those achieved through women's suffrage, further decreases the likelihood of violent conflict.<sup>185</sup> Citizens often show a significant risk aversion in comparison to their government. While surveys are notoriously difficult to collect and interpret, existing data suggest that the public has little support for nuclear weapons use,<sup>186</sup> but strong support for action against climate catastrophe.<sup>187</sup> We can further show that when citizens deliberately engage with the subject at hand, their concern and readiness for action often increases.<sup>188</sup> For example, citizen assemblies on climate change have recommended widespread policy-changes across sectors, amendments to incentive structures and laws against ecocide to reach emissions targets.<sup>189</sup> Indeed, many lament that when it comes to genetically modified organisms and nuclear power, citizens are far too risk-averse.<sup>190</sup> The problem is not that the public is riddled with cognitive biases that make them unconcerned about global catastrophes.

Democratic debate cannot be an afterthought. Navigating humanity through crises will involve many value-laden decisions under deep uncertainty. Democratic procedures can deal with such hard choices. Greater cognitive diversity should be represented amongst scholars of ERS. Recommendations on policies that would reduce risk should be passed through deliberative assemblies and await the approval of a wider pool of ordinary citizens, as they will be the ones who will bear this risk. A homogenous group of experts attempting to directly influence powerful decision-makers is not a fair or safe way of traversing the precipice.



## 7. Conclusion

The case for the importance of studying existential risk has been made. ERS must now converge on a trustworthy methodology.

The general problems in ERS we identify are that (i) an inclusive definition of existential or extinction risk will require some ambiguity, (ii) any categorisation of risks will be at least partly arbitrary, (iii) any risk assessment will not be entirely comprehensive, and (iv) any study of existential risks and proposed interventions could also increase risk. This field of study is inseparable from a moral inquiry. Definitions of what are catastrophic or desirable futures are inextricably intertwined with questions of value. Dealing with risk is not restricted to risk analysis but includes the question of what risks are worth taking. We believe these challenges are not insurmountable and are worthy of our attention.

The original and influential Techno-Utopian Approach (the TUA) faces specific, daunting problems. These include idiosyncratic and non-representative moral visions of the future, and the use of definitions that are excessively ambiguous and founded on opaque, questionable assumptions. The categorisation problem is exacerbated by combining the study of catastrophe with longtermist ethics. The frameworks are crude and do not include recent advances in risk assessment, or even basic knowledge from directly relevant fields. The TUA does not yet consider response risks and at times advocates for risky policies. It is susceptible to being misused to justify exceptional emergency actions.

We suggest some initial, modest steps for improving the field. First, extinction risks should be analysed separately from extinction ethics and existential ethics. Some research questions will still require combining these inquiries, but the attempt to separate the science of risk from the moral evaluation of risk will benefit each endeavour. Second, existential risk scholars should transparently acknowledge the moral and empirical assumptions used in risk analyses. Third, we must critically embrace the latest advances in risk assessment from other fields, such as climate science and disaster risk reduction. Fourth, and most importantly, existential risk must be cognitively diversified, and the judgement of its recommendations democratised. We can't afford to wait for a "Long

Reflection". Open democracy and collective deliberation need to be central to reducing existential risk and navigating the future.

We encourage existing scholars to enrich the diversity of available frameworks by revising or abandoning the TUA. We encourage researchers to find entirely new approaches and take a more inclusive, participatory approach to thinking about and shaping our responses to potential catastrophes. To have good judgement, represent the interests of the vulnerable, and avoid dangers, the study of existential risk needs to be democratised.

## Acknowledgements

We would like to thank the almost two dozen reviewers who took the time to peruse this draft and provide invaluable suggestions. The chapter is much stronger due to your help.

## Notes and References

- 1 Shackelford, G. E. et al. 'Accumulating evidence using crowdsourcing and machine learning: A living bibliography about existential risk and global catastrophic risk', *Futures*, 116 (February 2020): 102508. <https://doi.org/10.1016/j.futures.2019.102508>
- 2 Rees, M. J. *Our Final Century: Will Civilisation Survive the Twenty-First Century?* Arrow (2004); Bostrom, Nick. *Superintelligence: Paths, Dangers, Strategies*. Oxford University Press (2016); Tegmark, M. *Life 3.0: Being Human in the Age of Artificial Intelligence*. Penguin Books (2018); Russell, S. J. *Human Compatible: Artificial Intelligence and the Problem of Control*. Penguin Books (2020); Ord, T. *The Precipice: Existential Risk and the Future of Humanity*. Hachette Books (2020); Christian, B. *The Alignment Problem: How Can Machines Learn Human Values?* Atlantic Books (2021).
- 3 UN. *Our Common Agenda — Report of the Secretary-General*. United Nations (2021).
- 4 The following is not a comprehensive summary, but an indicator of the origin of resources in the field: The Future of Life Institute began with a \$10 million donation from billionaire Elon Musk. It launched a \$25 million program on AI safety with financial support from billionaire Vitalik Buterin. The technology millionaire Jaan Tallinn was a co-founder of CSER, while Elon Musk sits on the Advisory Board. The Future of Humanity Institute has received funding from Elon Musk, as well as over \$14 million in grants from Open Philanthropy. Open Philanthropy, founded by billionaires Cari Tuna and Dustin Moskovitz (a cofounder of Facebook), funded the establishment of the Centre for Security and Emerging Technology (CSET) with a \$55 million dollar donation. Open Philanthropy has, to date, given over \$37 million to projects related to global catastrophic risks, \$99 million to biosecurity and pandemic preparedness, and \$196 million worth of grants for potential risks from AI (not all of which are related to existential risk) (calculated on the 01/12/2021 using Open Philanthropies online database).

- 5 Greaves, H. et al., *A Research Agenda for The Global Priorities Institute*. Global Priorities Institute (2019).
- 6 Ord (2020); Matheny, J. G., 'Reducing the risk of human extinction', *Risk Analysis* 27(5) (7 December 2007): 1335–44. <https://doi.org/10.1111/j.1539-6924.2007.00960;> Bostrom, Nick. 'Existential risk prevention as global priority', *Global Policy*, 4(1) (February 2013): 15–31. <https://doi.org/10.1111/1758-5899.12002>; Bostrom, Nick. 'Astronomical waste: The opportunity cost of Delayed technological development', *Utilitas*, 15(3) (November 2003): 308–14. <https://doi.org/10.1017/S0953820800004076>; Bostrom, Nick (ed.). 'Existential risks: Analyzing human extinction scenarios and related hazards', *Journal of Evolution and Technology*, 9 (2002).
- 7 Bostrom, Nick. 'Letter from Utopia', *Studies in Ethics, Law, and Technology*, 2(1) (9 January 2008). <https://doi.org/10.2202/1941-6008.1025>; Bostrom, Nick and Philosophy Documentation Center. 'Transhumanist values', *Journal of Philosophical Research*, 30(9999) (2005): 3–14. [https://doi.org/10.5840/jpr\\_2005\\_26](https://doi.org/10.5840/jpr_2005_26)
- 8 Ord (2020); Bostrom (2013); Bostrom (2002).
- 9 Greaves, Hilary and Will MacAskill. *The Case for Strong Longtermism*. Global Priorities Institute (2019).
- 10 Bostrom (2013); Bostrom, N. 'The vulnerable world hypothesis', *Global Policy*, 10(4) (November 2019): 455–76. <https://doi.org/10.1111/1758-5899.12718>
- 11 The failure modes we describe in our chapter will thus apply to the elements of the TUA which different publications lean on.
- 12 Avin, Shahar et al. 'Classifying global catastrophic risks', *Futures*, 102 (September 2018): 20–26. <https://doi.org/10.1016/j.futures.2018.02.001>; Liu, Hin-Yan, Kristian Cedervall Lauta and Matthijs Michiel Maas. 'Governing boring apocalypses: A new typology of existential vulnerabilities and exposures for existential risk research', *Futures*, 102 (September 2018): 6–19. <https://doi.org/10.1016/j.futures.2018.04.009>
- 13 Liu, Lauta and Maas (2018) suggest examining GCRs in terms of their hazards (what harms humanity), vulnerabilities (how and why a given hazard could cause humanity to come to harm) and exposure (how and why humanity is in harm's way). Avin et al. (2018) classify GCRs by the critical system under threat, the mechanism for spreading risk globally, and failures to mitigate or prevent the risk. See Beard, SJ et al. 'Assessing climate change's contribution to global catastrophic risk', *Futures*, 127 (March 2021): 102673. <https://doi.org/10.1016/j.futures.2020.102673> for a brief summary, as well as a survey of existing definitions of GCRs.
- 14 Beard, Simon and Phil Torres. 'Ripples on the great sea of life: A brief history of existential risk studies', *SSRN Electronic Journal* (2020). <https://doi.org/10.2139/ssrn.3730000>
- 15 Bostrom (2002).
- 16 Bostrom (2013).
- 17 Calculated with Google Scholar on the 10th May, 2021. As an imperfect comparison, one of the few papers proposing an alternative way of thinking about existential risks (Liu, Lauta and Maas, 2018) has a total of 24 citations (calculated on the 10th May, 2021).
- 18 Bostrom (2003).
- 19 Bostrom (2016).
- 20 Ord (2020).

- 21 Tegmark (2017).
- 22 MacAskill, Will. *What We Owe the Future*. Basic Books (2022).
- 23 Kurzweil, Ray. *The Singularity Is Near: When Humans Transcend Biology*. Duckworth (2018); Chalmers, David J. 'The singularity: A philosophical analysis', *Journal of Consciousness Studies*, 17(9–10) (2010): 7–65.
- 24 Beckstead, Nick. *On the Overwhelming Importance of Shaping the Far Future*. Rutgers University (2013).
- 25 Karnofsky, Holden. 'Most Important Century' Series: Roadmap (2021). <https://www.cold-takes.com/most-important-century-series-roadmap/>
- 26 A Platonic ideal of a phenomenon that abstracts its essential features, but in reality is observed in many variations.
- 27 Milanović, Branko. *Capitalism, Alone: The Future of the System That Rules the World*. The Belknap Press of Harvard University Press (2021).
- 28 Bostrom (2002).
- 29 Bostrom (2013).
- 30 Bostrom.
- 31 Cotton-Barrat, Owen and Toby Ord. 'Existential risk and existential hope: Definitions', *Future of Humanity Institute* (n.d.).
- 32 Ord (2020).
- 33 Bostrom and Philosophy Documentation Center (2003).
- 34 Bostrom (2013).
- 35 Bostrom (2008).
- 36 Bostrom and Philosophy Documentation Center (2003); Bostrom, Nick. *What Is Transhumanism?* (2001). <https://www.nickbostrom.com/old/transhumanism.html>
- 37 Bostrom (2008).
- 38 Bostrom (2013).
- 39 Parfit, Derek. *Reasons and Persons*. Oxford University Press (1986). <https://doi.org/10.1093/019824908X.001.0001>
- 40 Matheny, J. G. 'Reducing the risk of human extinction', *Risk Analysis*, 27(5) (October 2007): 1335–44. <https://doi.org/10.1111/j.1539-6924.2007.00960.x>
- 41 Bostrom (2013).
- 42 Bostrom (2014).
- 43 Sandberg, Anders, Stuart Armstrong and Milan M. Cirković. 'That is not dead which can eternal lie: The aestivation hypothesis for resolving Fermi's Paradox', *ArXiv:1705.03394 [Physics]* (27 April 2017). <http://arxiv.org/abs/1705.03394>
- 44 Brown-Weiss, Edith. *In Fairness to Future Generations: International Law, Common Patrimony, and Intergenerational Equity*. Transnational Publishers (1989).
- 45 Greaves, Hilary and Will MacAskill. 'The case for strong longtermism', *GPI Working Paper*, 5. University of Oxford, Global Priorities Institute (2019).
- 46 Tarsney, Christian. 'The epistemic challenge to longtermism', *GPI Working Paper*, 10. University of Oxford, Global Priorities Institute (2019).

- 47 Bostrom (2013); Tarsney (2019).
- 48 Ord (2020).
- 49 Riedener, Stefan. 'Existential risks from a Thomist Christian perspective', *GPI Working Paper*, 1. University of Oxford, Global Priorities Institute (2021); Tarsney, Christian and Teruji Thomas. 'Non-additive axiologies in large worlds', *GPI Working Paper*, 9. University of Oxford, Global Priorities Institute (2020); Greaves, Hilary and Toby Ord. 'Moral uncertainty about population axiology', *Journal of Ethics and Social Philosophy*, 12(2) (2 October 2017): 135–67. <https://doi.org/10.26556/jesp.v12i2.223>
- 50 Greaves and MacAskill (2019).
- 51 Ord (2020).
- 52 Finneron-Burns, Elizabeth. 'What's wrong with human extinction?', *Canadian Journal of Philosophy*, 47, (2–3) (2017): 327–43. <https://doi.org/10.1080/00455091.2016.1278150>; Frick, Johann. 'On the survival of humanity', *Canadian Journal of Philosophy*, 47(2–3) (2017): 344–67. <https://doi.org/10.1080/00455091.2017.1301764>; Lenman, James. 'On becoming extinct', *Pacific Philosophical Quarterly*, 83(3) (September 2002): 253–69. <https://doi.org/10.1111/1468-0114.00150>
- 53 Bourget, David and David J. Chalmers. 'What do philosophers believe?', *Philosophical Studies*, 170(3) (September 2014): 465–500. <https://doi.org/10.1007/s11098-013-0259-7>
- 54 Hughes, James. *Report on the 2007 Interests and Beliefs Survey of the Members of the World Transhumanist Association*. World Transhumanist Association (2008).
- 55 Torres, Phil. *Were the Great Tragedies of History "Mere Ripples"? The Case Against Longtermism* (2020). <https://www.xriskology.com/mini-book>
- 56 Caviola, Lucius et al. 'Population ethical intuitions', *Cognition*, 218 (January 2022): 104941. <https://doi.org/10.1016/j.cognition.2021.104941>
- 57 Schubert, Stefan Lucius Caviola and Nadira S. Faber. 'The psychology of existential risk: Moral judgments about human extinction', *Scientific Reports*, 9(1) (December 2019): 15100. <https://doi.org/10.1038/s41598-019-50145-9>
- 58 Landemore, Hélène. *Open Democracy: Reinventing Popular Rule for the Twenty-First Century*. Princeton University Press (2020).
- 59 Ord, Toby, Angus Mercer, and Sophie Dannreuther. *Future Proof: The Opportunity to Transform the UK's Resilience to Extreme Risks*. Centre for Long-Term Resilience (2021). <https://www.longtermresilience.org/futureproof>
- 60 Schuster, Joshua and Derek Woods. *Calamity Theory: Three Critiques of Existential Risk*. University of Minnesota Press (2021); Stoecker, Christian. 'Ist "Longtermism" Die Rettung — Oder Eine Gefahr?', *Spiegel Wissenschaft* (2021). <https://www.spiegel.de/wissenschaft/longtermism-was-ist-das-rettung-oder-gefahr-kolumne-a-983e60ba-6265-40a8-8c65-8f2668e4e9ff>
- 61 Barbrook, Richard and Andy Cameron. 'The Californian ideology', *Mute Magazine* (1995). <https://www.metamute.org/editorial/articles/californian-ideology>
- 62 MacAskill, Will, Krister Bykvist and Toby Ord (eds.). *Moral Uncertainty*. Oxford University Press (2020).
- 63 Ord (2020).
- 64 Greaves and Ord (2017).

- 65 MacAskill, Bykvist and Ord (2020).
- 66 Kernohan, Andrew. 'Descriptive uncertainty and maximizing expected choice-worthiness', *Ethical Theory and Moral Practice*, 24(1) (March 2021): 197–211. <https://doi.org/10.1007/s10677-020-10139-3>
- 67 Greaves and Ord (2017).
- 68 Bostrom (2013).
- 69 Ord (2020).
- 70 Cotton-Barrat and Ord (2015).
- 71 Graeber, David and David Wengrow. *The Dawn of Everything: A New History of Humanity*. Allen Lane (2021).
- 72 Finneron-Burns (2017).
- 73 Frick (2018).
- 74 Lenman (2002).
- 75 Ord (2020).
- 76 Bostrom (2013); Bostrom (2002).
- 77 Ord (2020).
- 78 Cotton-Barrat and Ord (2015).
- 79 Ord (2020).
- 80 Grace, Katja et al. 'Viewpoint: When will AI exceed human performance? Evidence from AI experts', *Journal of Artificial Intelligence Research*, 62 (31 July 2018): 729–54. <https://doi.org/10.1613/jair.1.11222>
- 81 Cremer, Carla Zoe. 'Deep limitations? Examining expert disagreement over deep learning', *Progress in Artificial Intelligence*, 10(4) (December 2021): 449–64. <https://doi.org/10.1007/s13748-021-00239-1>; Cremer, Carla Zoe and Jess Whittlestone, 'Canaries in technology mines: Warning signs of transformative progress in AI', *1st International Workshop on Evaluating Progress in Artificial Intelligence* (24 September 2020). <https://doi.org/10.17863/CAM.57790>
- 82 Note this unwillingness to stop or delay progress in technology or research is interestingly often not observed in biorisk research in ERS. It is unclear why this is the case. One explanation could be that banning or stopping technologies and practices in biorisk is seen to be more feasible than doing the same for other areas. We have yet to see a convincing argument as to why this should be the case. Another explanation could be that this is about desire, not tractability. AI is central to the techno-utopian vision of the future. Gain-of-function experimentation is not. This suggests that the techno-determinism of the TUA is prescriptive, not descriptive (see Section 4.3 for further discussion).
- 83 Ord, Mercer and Dannreuther (2021).
- 84 Bostrom (2019).
- 85 While some may choose to read Bostrom's discussion as a merely speculative proposal for a hypothetical scenario, it exemplifies a willingness to consider, and even justify, measures that would normally be deemed unconscionable in the name of averting supposed existential catastrophes. Moreover, the fact that the essay was published in a journal for policy (with policy implications considered), rather than philosophy, suggests that the discussion should not be considered as entirely theoretical and not

- actionable. Additionally, intentional outreach through a TED Talk (with over 2.5 million views), the Sam Harris podcast, and a piece in Aeon all stress this aspect of the article. These actions signal that the argument is not intended to be an innocuous thought experiment.
- 86 The term *existential ethics* is already sometimes used in relation to the philosophy of existentialism. We are aware of this; however, it seems difficult to move away from the use of the existential given the field's reliance on the term existential.
  - 87 Purtil, Corinne. 'How close is humanity to the edge?', *The New Yorker* (2020). <https://www.newyorker.com/culture/annals-of-inquiry/how-close-is-humanity-to-the-edge>
  - 88 We owe this point to Joshua Teperowski Monrad.
  - 89 Beard et al. (2021).
  - 90 They are, in Bostrom's words, "mere ripples on the surface of the great sea of life" (Bostrom, 2002).
  - 91 Bostrom (2013).
  - 92 Bostrom.
  - 93 Ord (2020).
  - 94 Silver, Nate. *The Signal and the Noise: The Art and Science of Prediction*. Penguin Economics (2013).
  - 95 Smil, Vaclav. *Growth: From Microorganisms to Megacities*. The MIT Press (2020).
  - 96 Scheidel, Walter. *The Great Leveller: Violence and the History of Inequality: From the Stone Age to the Twenty-First Century*. Princeton University Press (2017).
  - 97 Halstead, John. 'Is climate change an existential risk?', *Google Doc* (blog) (2019). <https://docs.google.com/document/d/1qmHh-cshTCMT8LX0Y5wSQm8FM BhaxhQ8OI0eRLkXIF0/edit#>; Piper, Kelsey. 'Is climate change an "existential threat" — or just a catastrophic one?', *Vox* (2019). <https://www.vox.com/future-perfect/2019/6/13/18660548/climate-change-human-civilization-existential-risk>. Halstead does not provide a definition of existential risk, so it is unclear what this means.
  - 98 Pescaroli, Gianluca and David Alexander. 'Understanding compound, interconnected, interacting, and cascading risks: A holistic framework for understanding complex risks', *Risk Analysis*, 38(11) (November 2018): 2245–57. <https://doi.org/10.1111/risa.13128>
  - 99 Pescaroli and Alexander (2018); Tang, Aaron and Luke Kemp. 'A fate worse than warming? Stratospheric aerosol injection and catastrophic risk', *Frontiers in Climate*, 3(720312) (2021). <https://doi.org/10.3389/fclim.2021.720312>
  - 100 Defined as being consistent with our background knowledge (Betz, 2016).
  - 101 "The case for existential risk from AI is clearly speculative. Indeed, it is the most speculative case for a major risk in this book" (Ord, 2020, p.149).
  - 102 Cremer (2021); Grace et al. (2018).
  - 103 Bostrom (2016); Ord (2020).
  - 104 Centeno, Miguel A. et al. 'The emergence of global systemic risk', *Annual Review of Sociology*, 41(1) (14 August 2015): 65–85. <https://doi.org/10.1146/annurev-soc-073014-112317>
  - 105 Challinor, Andy J. et al. 'Transmission of climate risks across sectors and borders',

- Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 376(2121) (13 June 2018): 20170301. <https://doi.org/10.1098/rsta.2017.0301>
- 106 Homer-Dixon, Thomas et al. 'Synchronous failure: The emerging causal architecture of global crisis', *Ecology and Society*, 20(3) (2015): art6. <https://doi.org/10.5751/ES-07681-200306>
  - 107 IPCC. *The Concept of Risk in the IPCC Sixth Assessment Report: A Summary of Cross-Working Group Discussions*. Intergovernmental Panel on Climate Change (2020). [https://www.ipcc.ch/site/assets/uploads/2021/02/Risk-guidance-FINAL\\_15Feb2021.pdf](https://www.ipcc.ch/site/assets/uploads/2021/02/Risk-guidance-FINAL_15Feb2021.pdf)
  - 108 Simpson, Nicholas P. et al. 'A framework for complex climate change risk assessment', *One Earth*, 4(4) (April 2021): 489–501. <https://doi.org/10.1016/j.oneear.2021.03.005>. In disaster risk management even more risk components have been proposed, such as capacity. We highlight this formulation due to the prominence and rigour of the IPCC.
  - 109 Mani, Lara, Asaf Tzachor and Paul Cole. 'Global catastrophic risk from lower magnitude volcanic eruptions', *Nature Communications*, 12(1) (December 2021): 4756. <https://doi.org/10.1038/s41467-021-25021-8>
  - 110 The injection of aerosols into the stratosphere to reflect sunlight and mitigate some of the effects. of climate change.
  - 111 Tang and Kemp (2021).
  - 112 Rees (2004); Ord (2020).
  - 113 Ord (2020).
  - 114 Avin et al. (2018); Cotton-Barratt, Owen, Max Daniel and Anders Sandberg. 'Defence in depth against human extinction: Prevention, response, resilience, and why they all matter', *Global Policy*, 11(3) (May 2020): 271–82. <https://doi.org/10.1111/1758-5899.12786>
  - 115 Rees (2004); Ord (2020); Bostrom (2013).
  - 116 Bostrom, Nick. 'The future of humanity', in *New Waves in Philosophy of Technology*, ed. Jan Kyrre Berg Olsen, Evan Selinger and Søren Riis. Palgrave Macmillan (2009), pp.186–215. [https://doi.org/10.1057/9780230227279\\_10](https://doi.org/10.1057/9780230227279_10)
  - 117 Dafoe, Allan. 'On technological determinism: A typology, scope conditions, and a mechanism', *Science, Technology, & Human Values*, 40(6) (November 2015): 1047–76. <https://doi.org/10.1177/0162243915579283>
  - 118 Balwit, Avital. 'Response to recent criticisms of longtermism', *Effective Altruism Forum* (blog) (2021).
  - 119 Ord (2020).
  - 120 Dafoe (2015).
  - 121 Jacobsen, Annie. *The Pentagon's Brain: An Uncensored History of DARPA, America's Top Secret Military Research Agency* (1st ed.). Little, Brown and Company (2015).
  - 122 Fleming, James Rodger. 'The pathological history of weather and climate modification: Three cycles of promise and hype', *Historical Studies in the Physical and Biological Sciences*, 37(1) (1 September 2006): 3–25. <https://doi.org/10.1525/hsps.2006.37.1.3>
  - 123 Mumford, Lewis. *Technics and Civilization*. The University of Chicago Press (2010).
  - 124 Malm, Andreas. *Fossil Capital: The Rise of Steam-Power and the Roots of Global Warming*. Verso (2016).



- 125 Bostrom (2002).
- 126 Bostrom (2013).
- 127 Millett, Piers and Andrew Snyder-Beattie. 'Existential risk and cost-effective biosecurity', *Health Security*, 15(4) (August 2017): 373–83. <https://doi.org/10.1089/hs.2017.0028>
- 128 Singh, Riddhi, Patrick M. Reed and Klaus Keller. 'Many-objective robust decision making for managing an ecosystem with a deeply uncertain threshold response', *Ecology and Society*, 20(3) (2015): art12. <https://doi.org/10.5751/ES-07687-200312>
- 129 Ord (2020).
- 130 Morgensen, Andreas. 'Maximal cluelessness', *GPI Working Paper*, 2. University of Cambridge, Global Priorities Institute (2019); Morgensen, Andreas and David Thorstad. 'Tough enough? Robust satisficing as a decision norm for long-term policy analysis', *GPI Working Paper*, 15. University of Cambridge, Global Priorities Institute (2020).
- 131 Kokotajlo, Daniel. 'Tiny probabilities of vast utilities: Bibliography and appendix', *Effective Altruism Forum* (blog) (2018). [https://forum.effectivealtruism.org/s/ji4eCGYmpjWY6yp/p/nMfhXPiqPRDtEdauY#Academic\\_literature\\_](https://forum.effectivealtruism.org/s/ji4eCGYmpjWY6yp/p/nMfhXPiqPRDtEdauY#Academic_literature_); Bostrom, Nick. 'Pascal's mugging', *Analysis*, 69(3) (2009): 443–45.
- 132 Wilkinson, Hayden. 'In defense of fanaticism', *Ethics*, 132(2) (1 January 2022): 445–77. <https://doi.org/10.1086/716869>
- 133 Bostrom (2009).
- 134 Balfour, Dylan. 'Pascal's mugger strikes again', *Utilitas*, 33(1) (March 2021): 118–24. <https://doi.org/10.1017/S0953820820000357>
- 135 Based on personal communications with several scholars in the field.
- 136 Wilkinson (2022).
- 137 Betz, Gregor. 'Accounting for possibilities in decision making', in *The Argumentative Turn in Policy Analysis* (vol. 10), ed. Sven Ove Hansson and Gertrude Hirsch Hadorn. Springer International Publishing (2016), pp.135–69. [https://doi.org/10.1007/978-3-319-30549-3\\_6](https://doi.org/10.1007/978-3-319-30549-3_6)
- 138 IPCC (2020); Simpson et al. (2021).
- 139 Snyder-Beattie, Andrew E., Toby Ord and Michael B. Bonsall. 'An upper bound for the background rate of human extinction', *Scientific Reports*, 9(1) (December 2019): 11054. <https://doi.org/10.1038/s41598-019-47540-7>
- 140 Ord (2020). These figures are not water tight: they assume that the vulnerability of humanity in the Palaeolithic is identical to today.
- 141 Ord (2020).
- 142 Bostrom (2002).
- 143 Rees (2004).
- 144 Ord (2020).
- 145 Bostrom (2019).
- 146 Ord (2020); Bostrom (2002).
- 147 That is, a superintelligent AI system that is aligned with human values and will aid rather than harm humanity.

- 148 Bostrom (2002).
- 149 Deudney, Daniel. *Dark Skies: Space Expansionism, Planetary Geopolitics, and The Ends of Humanity*. Oxford University Press (2020).
- 150 Harari, Yuval Noah. *Homo Deus: A Brief History of Tomorrow* (revised edition). Vintage (2017); McKibben, Bill. *Enough: Staying Human in an Engineered Age*. Owl Books (2004).
- 151 Kemp, Luke et al. 'Bioengineering horizon scan 2020', *Elife*, 9 (2020): e54489. <https://doi.org/10.7554/eLife.54489>; Wintle, Bonnie C. et al. 'A transatlantic perspective on 20 emerging issues in biological engineering', *Elife*, 6 (14 November 2017): e30247. <https://doi.org/10.7554/eLife.30247>
- 152 Wilkinson, Richard G. and Kate Pickett. *The Spirit Level: Why Equality Is Better for Everyone [with a New Chapter Responding to Their Critics]*. Penguin Books (2010); Wilkinson, Richard G. and Kate Pickett. *The Inner Level: How More Equal Societies Reduce Stress, Restore Sanity and Improve Everyone's Well-Being*. Allen Lane (2018).
- 153 Kemp, Luke et al. *Emerging Technologies and Dual-Use Concerns: A Horizon Scan for Global Public Health*. World Health Organization (2021).
- 154 Humphreys, S. 'Legalizing lawlessness: On Giorgio Agamben's state of exception', *European Journal of International Law*, 17(3) (1 June 2006): 677–87. <https://doi.org/10.1093/ejil/chl020>; Keith, Linda Camp and Steven C. Poe. 'Are constitutional state of emergency clauses effective? An empirical exploration', *Hum. Rts. Q.* 26 (2004): 1071.
- 155 Ferejohn, John and Pasquale Pasquino. 'The law of the exception: A typology of emergency powers', *International Journal of Constitutional Law*, 2(2) (2004): 210–39.
- 156 De Wilde, Marc. 'Just trust us: A short history of emergency powers and constitutional change', *Comparative Legal History*, 3(1) (2 January 2015): 110–30. <https://doi.org/10.1080/2049677X.2015.1041728>
- 157 Clarke, Lee. 'Panic: Myth or reality?', *Contexts*, 1(3) (August 2002): 21–26. <https://doi.org/10.1525/ctx.2002.1.3.21>; Clarke, L. and C. Chess, 'Elites and panic: More to fear than fear itself', *Social Forces*, 87(2) (1 December 2008): 993–1014. <https://doi.org/10.1353/sof.0.0155>; Bregman, Rutger. *Humankind: A Hopeful History*. Bloomsbury Publishing (2021); Solnit, Rebecca. *A Paradise Built in Hell: The Extraordinary Communities That Arise in Disasters*. Penguin Random House (2009).
- 158 Bjørnskov, Christian and Stefan Voigt. 'More power to government = more people killed? On some unexpected effects of constitutional emergency provisions during natural disasters', *SSRN Electronic Journal* (2018). <https://doi.org/10.2139/ssrn.3189749>
- 159 Kemp, Luke. 'The "stomp reflex": When governments abuse emergency powers', *BBC Future* (2021). <https://www.bbc.com/future/article/20210427-the-stomp-reflex-when-governments-abuse-emergency-powers>
- 160 Scarry, Elaine. *Thermonuclear Monarchy: Choosing Between Democracy and Doom* (1st ed.). W. W. Norton & Company (2014).
- 161 Wellerstein, Alex. *Restricted Data: The History of Nuclear Secrecy in the United States*. The University of Chicago Press (2021).
- 162 It is worth noting that both neorealist assumptions about international relations and game theoretic approaches appear frequently in research under the TUA.

- 163 Singer, Peter. 'The hinge of history', *Project Syndicate* (2021). <https://www.project-syndicate.org/commentary/ethical-implications-of-focusing-on-extinction-risk-by-peter-singer-2021-10>; Torres, Phil. 'Against longtermism', *Aeon* (2021). <https://aeon.co/essays/why-longtermism-is-the-worlds-most-dangerous-secular-credo>
- 164 Ord (2020).
- 165 Seger, Elizabeth et al. *Tackling Threats to Informed Decision-Making in Democratic Societies: Promoting Epistemic Security in a Technologically-Advanced World*. Apollo — University of Cambridge Repository (14 October 2020). <https://doi.org/10.17863/CAM.64183>
- 166 Bostrom (2019).
- 167 Bostrom, Nick and Matthew Van Der Mewe. 'How vulnerable is the world?', *Aeon* (2019). <https://aeon.co/essays/none-of-our-technologies-has-managed-to-destroy-humanity-yet>
- 168 Bostrom (2019).
- 169 Bostrom and Van Der Mewe (2019).
- 170 Bostrom (2019).
- 171 Bostrom (2002).
- 172 Bostrom.
- 173 Granick, Jennifer Stisa. *American Spies: Modern Surveillance, Why You Should Care, and What To Do About It*. Cambridge University Press (2017); Kirchner, Lauren. 'What's the evidence mass surveillance works? Not much', *Propublica* (2015). <https://www.propublica.org/article/whats-the-evidence-mass-surveillance-works-not-much>
- 174 Kemp, Luke. 'Agents of doom: Who is creating the apocalypse and why', *BBC Future* (2021). <https://www.bbc.com/future/article/20211014-agents-of-doom-who-is-hastening-the-apocalypse-and-why?ocid=twfut>
- 175 Bostrom and Van Der Mewe (2019).
- 176 Ord (2020).
- 177 Ord.
- 178 Landemore (2020); Landemore, Hélène. *Democratic Reason: Politics, Collective Intelligence, and the Rule of the Many*. Princeton University Press (2017).
- 179 Hong, L. and S. E. Page. 'Groups of diverse problem solvers can outperform groups of high-ability problem solvers', *Proceedings of the National Academy of Sciences*, 101(46) (16 November 2004): 16385–89. <https://doi.org/10.1073/pnas.0403723101>; Hong, Lu and Scott E. Page. 'Problem solving by heterogeneous agents', *Journal of Economic Theory*, 97(1) (March 2001): 123–63. <https://doi.org/10.1006/jeth.2000.2709>
- 180 Landemore (2017).
- 181 OECD. *Innovative Citizen Participation and New Democratic Institutions: Catching the Deliberative Wave* (2020). [https://www.oecd-ilibrary.org/governance/innovative-citizen-participation-and-new-democratic-institutions\\_339306da-en](https://www.oecd-ilibrary.org/governance/innovative-citizen-participation-and-new-democratic-institutions_339306da-en)
- 182 Burchi, Francesco. 'Democracy, institutions and famines in developing and emerging countries', *Canadian Journal of Development Studies/Revue Canadienne d'études Du Développement*, 32(1) (March 2011): 17–31. <https://doi.org/10.1080/02255189.2011.576136>

- 183 Massing, Michael. 'Does democracy avert famine?', *The New York Times* (2003). <https://www.nytimes.com/2003/03/01/arts/does-democracy-avert-famine.html>
- 184 Dafoe, Allan, John R. Oneal, and Bruce Russett. 'The democratic peace: Weighing the evidence and cautious inference', *International Studies Quarterly*, 57(1) (March 2013): 201–14. <https://doi.org/10.1111/isqu.12055>
- 185 Barnhart, Joslyn N. et al. 'The suffragist peace', *International Organization*, 74(4) (2020): 633–70. <https://doi.org/10.1017/S0020818320000508>
- 186 Kull, Steven. 'Survey says: Americans back arms control', *Arms Control Today*, 34(5) (2004): 22; Løvold, Magnus. 'Lessons from the ICRC's "Millennials on War" survey for communication and advocacy on nuclear weapons', *Journal for Peace and Nuclear Disarmament*, 3(2) (2 July 2020): 410–17. <https://doi.org/10.1080/25751654.2020.1859216>; The Simons Foundation, *Global Public Opinion on Nuclear Weapons*. The Simons Foundation (2007).
- 187 UNDP. *The G20 Peoples' Climate Vote 2021*. United Nations Development Programme and the University of Oxford Department of Sociology (2021). <https://www.undp.org/publications/g20-peoples-climate-vote-2021>; European Commission. *Citizen Support for Climate Action: 2021 Survey*. European Commission (2021). [https://ec.europa.eu/clima/citizens/citizen-support-climate-action\\_en](https://ec.europa.eu/clima/citizens/citizen-support-climate-action_en); Marlon, Jennifer et al. *Yale Climate Opinion Maps 2020*. Yale Program on Climate Change Communication (2020). <https://climatecommunication.yale.edu/visualizations-data/ycom-us/>
- 188 OECD (2020).
- 189 Convention Citoyenne pour le Climat. *Les Propositiones de La Convention Citoyenne Pour Le Climat* (2021). <https://propositions.conventioncitoyennepourleclimat.fr/pdf/ccr-rapport-final.pdf>
- 190 Lynas, Mark. *The God Species: How Humans Really Can Save the Planet*. Fourth Estate (2012).