Global Challenges Foundation

Executive summary

Global Catastrophic Risks

2016

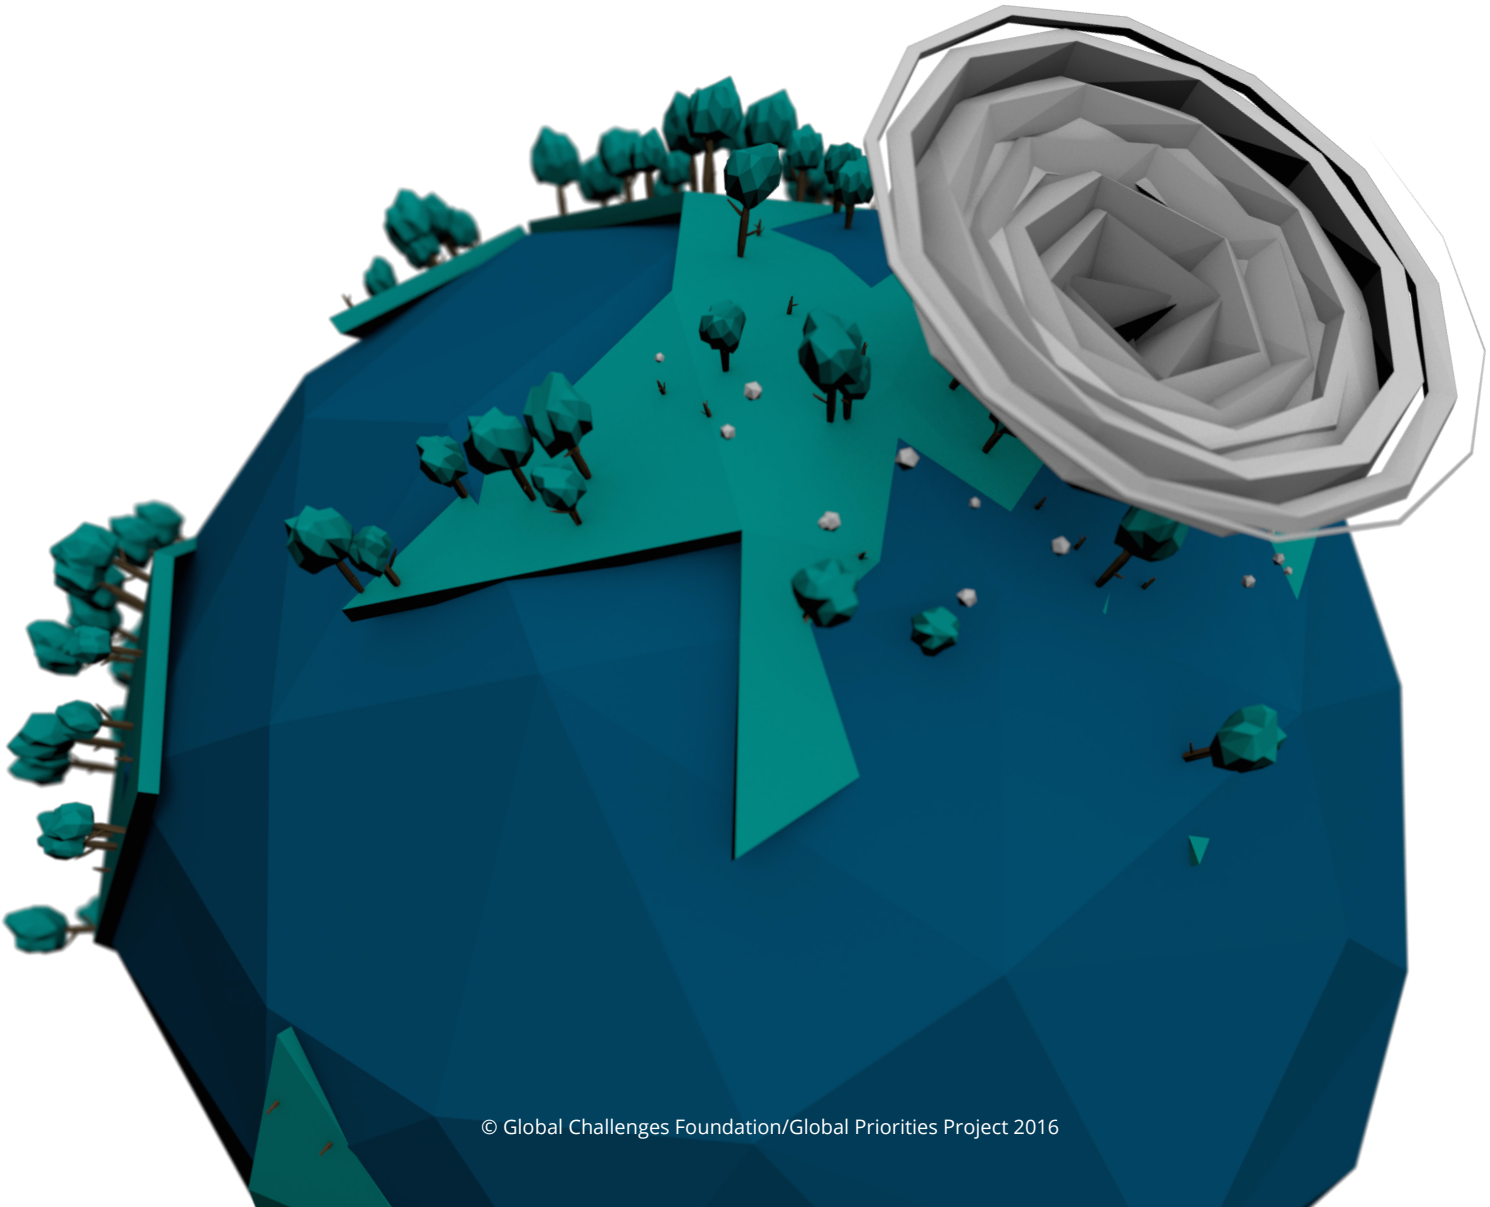

GLOBAL CATASTROPHIC RISKS 2016
The views expressed in this report are those of the authors. Their statements are not necessarily endorsed by the affiliated organisations.

Authors:
Owen Cotton-Barratt*†
Sebastian Farquhar*
John Halstead*
Stefan Schubert*
Andrew Snyder-Beattie†

* = The Global Priorities Project
† = The Future of Humanity Institute, University of Oxford

Graphic design:
Accomplice/Elinor Hägg

in association with

THE GLOBAL CHALLENGES FOUNDATION works to raise awareness of the Global Catastrophic Risks. Primarily focused on climate change, other environmental degradation and politically motivated violence as well as how these threats are linked to poverty and rapid population growth. Against this background, the Foundation also works to both identify and stimulate the development of good proposals for a management model – a global governance – able to decrease – and at best eliminate – these risks.

THE GLOBAL PRIORITIES PROJECT helps decision-makers effectively prioritise ways to do good. We achieve his both by advising decision-makers on programme evaluation methodology and by encouraging specific policies. We are a collaboration between the Centre for Effective Altruism and the Future of Humanity Institute, part of the University of Oxford.

Definition: Global Catastrophic Risk – risk of events or processes that would lead to the deaths of approximately a tenth of the world's population, or have a comparable impact.

Dear Reader!

Nearly four years ago when the Global Challenges Foundation was established, we decided on a direction with two parallel strategies. The first is increasing the knowledge about Global Catastrophic Risks (GCRs), which with our terminology means threats that can eliminate at least 10% of the global population. This knowledge is an important prerequisite for the Foundation's second strategy: to encourage debates and proposals as to how we can effectively and fairly reduce – and preferably eliminate – these catastrophic risks.

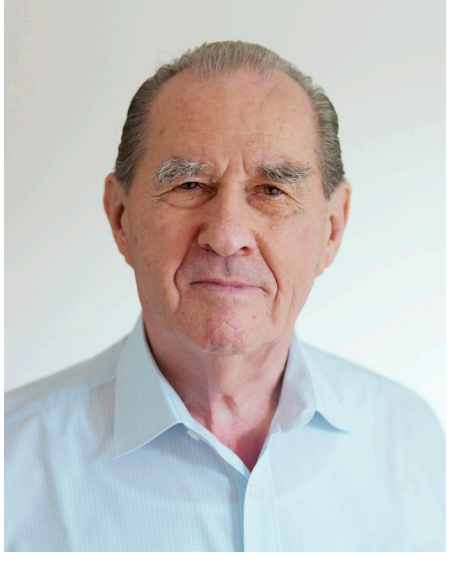

This publication, the Foundation's Annual Report for 2016, is the result of a collaboration between the Foundation and the Future of Humanity Institute (FHI) and the Global Priorities Project at Oxford University in the U.K., which has now lasted for over two years. A big group of researchers at the FHI, commissioned by the Foundation, summarized where research, focused on charting some of the greatest global risks, currently stands.

In addition to describing the risks, their effects and their likelihood of occurring, this year's Annual Report takes one step further and try's to show how different risks relate to one another, what can be done to combat the risks and who can and should do this. In addition to the risks involved in the Annual Report for 2016, the Foundation actively works with environmental degradation, weapons of mass destruction, population growth (that exacerbates several risks), and political violence which is behind many of the world's current problems.

Political violence comes in many forms. Various kinds of weapons of mass destruction represent potentially devastating weaponry. Further, political violence creates uncontrolled migration and we receive repeated reminders that there is also "digital violence" in the form of cyber-attacks. Together, this takes up a significant amount of space on the political agenda, thus stealing atten-

tion from other important risks. And above all, the defense against various forms of political violence requires a grotesquely large share of public resources. Each day, the world spends over SEK 40 billion on defence expenditure – money that would be needed to fight poverty and prevent catastrophic risks.

My personal opinion is that in order to drastically minimize GCRs we must develop a model where a majority of the world's nations, with strong support from leading nations, can make binding decisions which can be enforced in an effective and fair way. This would imply that individual nations waive their sovereignty in favor of one or more organizations that have a mandate to decide on how to mitigate GCRs.

Would this be possible? My counter question is whether there are any alternatives? To continue relying on multilateral negotiations increases the probability that decisions and actions are insufficient and executed too late. This means that the likelihood of GCRs continues to escalate.

I hope that this publication can deepen the understanding of GCRs and that these insights provide a fertile ground for both debates and proposals on how we can develop a better way of managing and addressing these risks.

Stockholm, April 2016

Laszlo Szombatfalvy
Founder of Global Challenges Foundation

Global catastrophic risks pose a pressing challenge

This report addresses one of the most important issues of our age – global catastrophic risk. Over the last decades, behavioural psychology has taught us that, as a species, we are bad at assessing scope. Issues that affect ten people do not intuitively seem ten times more important than those that affect one person. Global catastrophic risks are one area where our scope insensitivity might prove the most dangerous.

These risks can't just be treated as problem for the future, even though we might well expect them not to materialise this year or the next. At the Future of Life Institute, my team and I have been calling for global leaders to address critical global risk issues including nuclear weapons, biotechnology and artificial intelligence. This builds on existing risk reduction work led by institutions such as the United Nations.

Over the last centuries, humanity has achieved incredible things. New medical technologies save millions of lives every year. Agricultural science allows billions to be fed who might otherwise not exist. And we have begun to explore the very foundations of our universe itself – the beauty of which has inspired my own deep curiosity in cosmology.

This technological power is an enormous force for good, but carries its own risks. Although consuming fossil fuels was critical in creating the thriving and wonderful civilization we live in today, we've come to learn that there are potentially catastrophic long-term consequences from climate change. Other technologies, more powerful than combustion engines, might also offer huge benefits and carry unforeseen risks. If we fail to manage this risk well, we might be

caught out by consequences that follow from the technology more rapidly than climate change has.

As a global community, we need to win the race between the growing power of our technology and the wisdom with which we manage it. This requires a nuanced approach towards technological developments, acknowledging both that technology carries huge potential to make lives better and also that it carries some risks. Smart risk management means being realistic in weighing these factors against each other. This report offers an excellent background to the underlying issues of global catastrophic risks, and is an outstanding starting point for policy-makers developing an interest in the area or researchers considering how their own work might be brought into the study of global catastrophic risks.

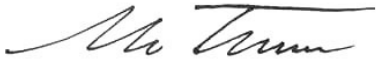

Max Tegmark
Co-founder of the Future of Life Institute
Professor of Physics at MIT

Executive Summary

Most generations never experience a global catastrophe. However, the idea of such catastrophes is not fanciful: plagues have killed over 10% of world's population and we came close to nuclear war several times in the 20th century.

Despite their scale, the risks of global catastrophes receive limited attention. One reason is that many of these risks are unlikely in any given decade. But even when the probability is low, the sheer magnitude of an adverse outcome warrants taking these risks seriously. A global catastrophic risk not only threatens everyone alive today, but also future generations. Reducing these risks is therefore both a global and an intergenerational public good.

The ever-evolving landscape of technology and society compounds these challenges. Technological and economic forces can create new global catastrophic risks, such as anthropogenic climate change and the 20th century's nuclear arms race. But technology can also reduce risk, for example through better vaccines or clean energy.

We believe the global community should work together to harness new tools to address global catastrophic risks. It is possible that, collectively, we significantly under-invest in global catastrophic risk reduction.

We believe the global community should work together to harness new tools to address global catastrophic risks.

What are the biggest threats?

The global catastrophic risks in this report can be divided into two categories. Some are ongoing and could potentially occur in any given year. Others are emerging and may be very unlikely today but will become significantly more likely in the coming decades. The most significant ongoing risks are natural pandemics and nuclear war, whereas the most significant emerging risks are catastrophic climate change and risks stemming from emerging technologies. Even where risks remain in the future, there are things we can do today to address them.

The Spanish influenza pandemic of 1918 may have killed as much as 5% of the world population. Some outbreaks since then infected over a third of the world's population (e.g., pandemic influenza), whereas others killed over half of people infected (e.g., Ebola or SARS). If a disease were to emerge that was as transmissible as the flu and as lethal as Ebola, the results could be catastrophic. Fortunately, this rarely transpires, but it is possible that it could, for example with the H5N1 influenza virus.

The invention of nuclear weapons ushered in a new era of risks created by human action. A large nuclear war between major powers would likely kill tens or hundreds of millions in the initial conflict, and perhaps many more if a nuclear winter were to follow. During the Cuban

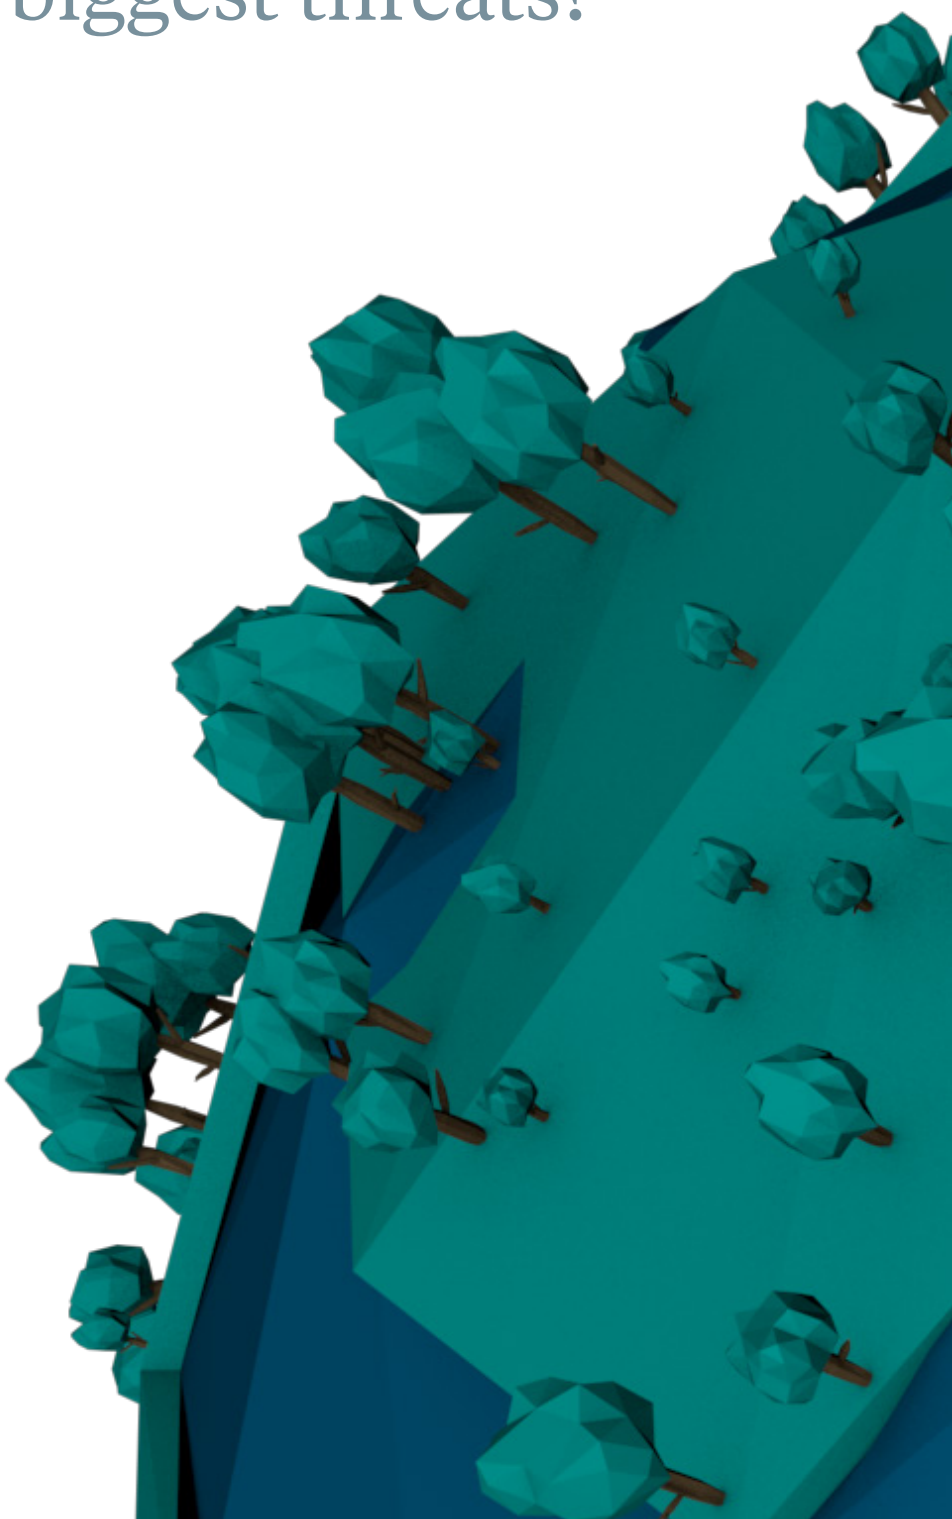

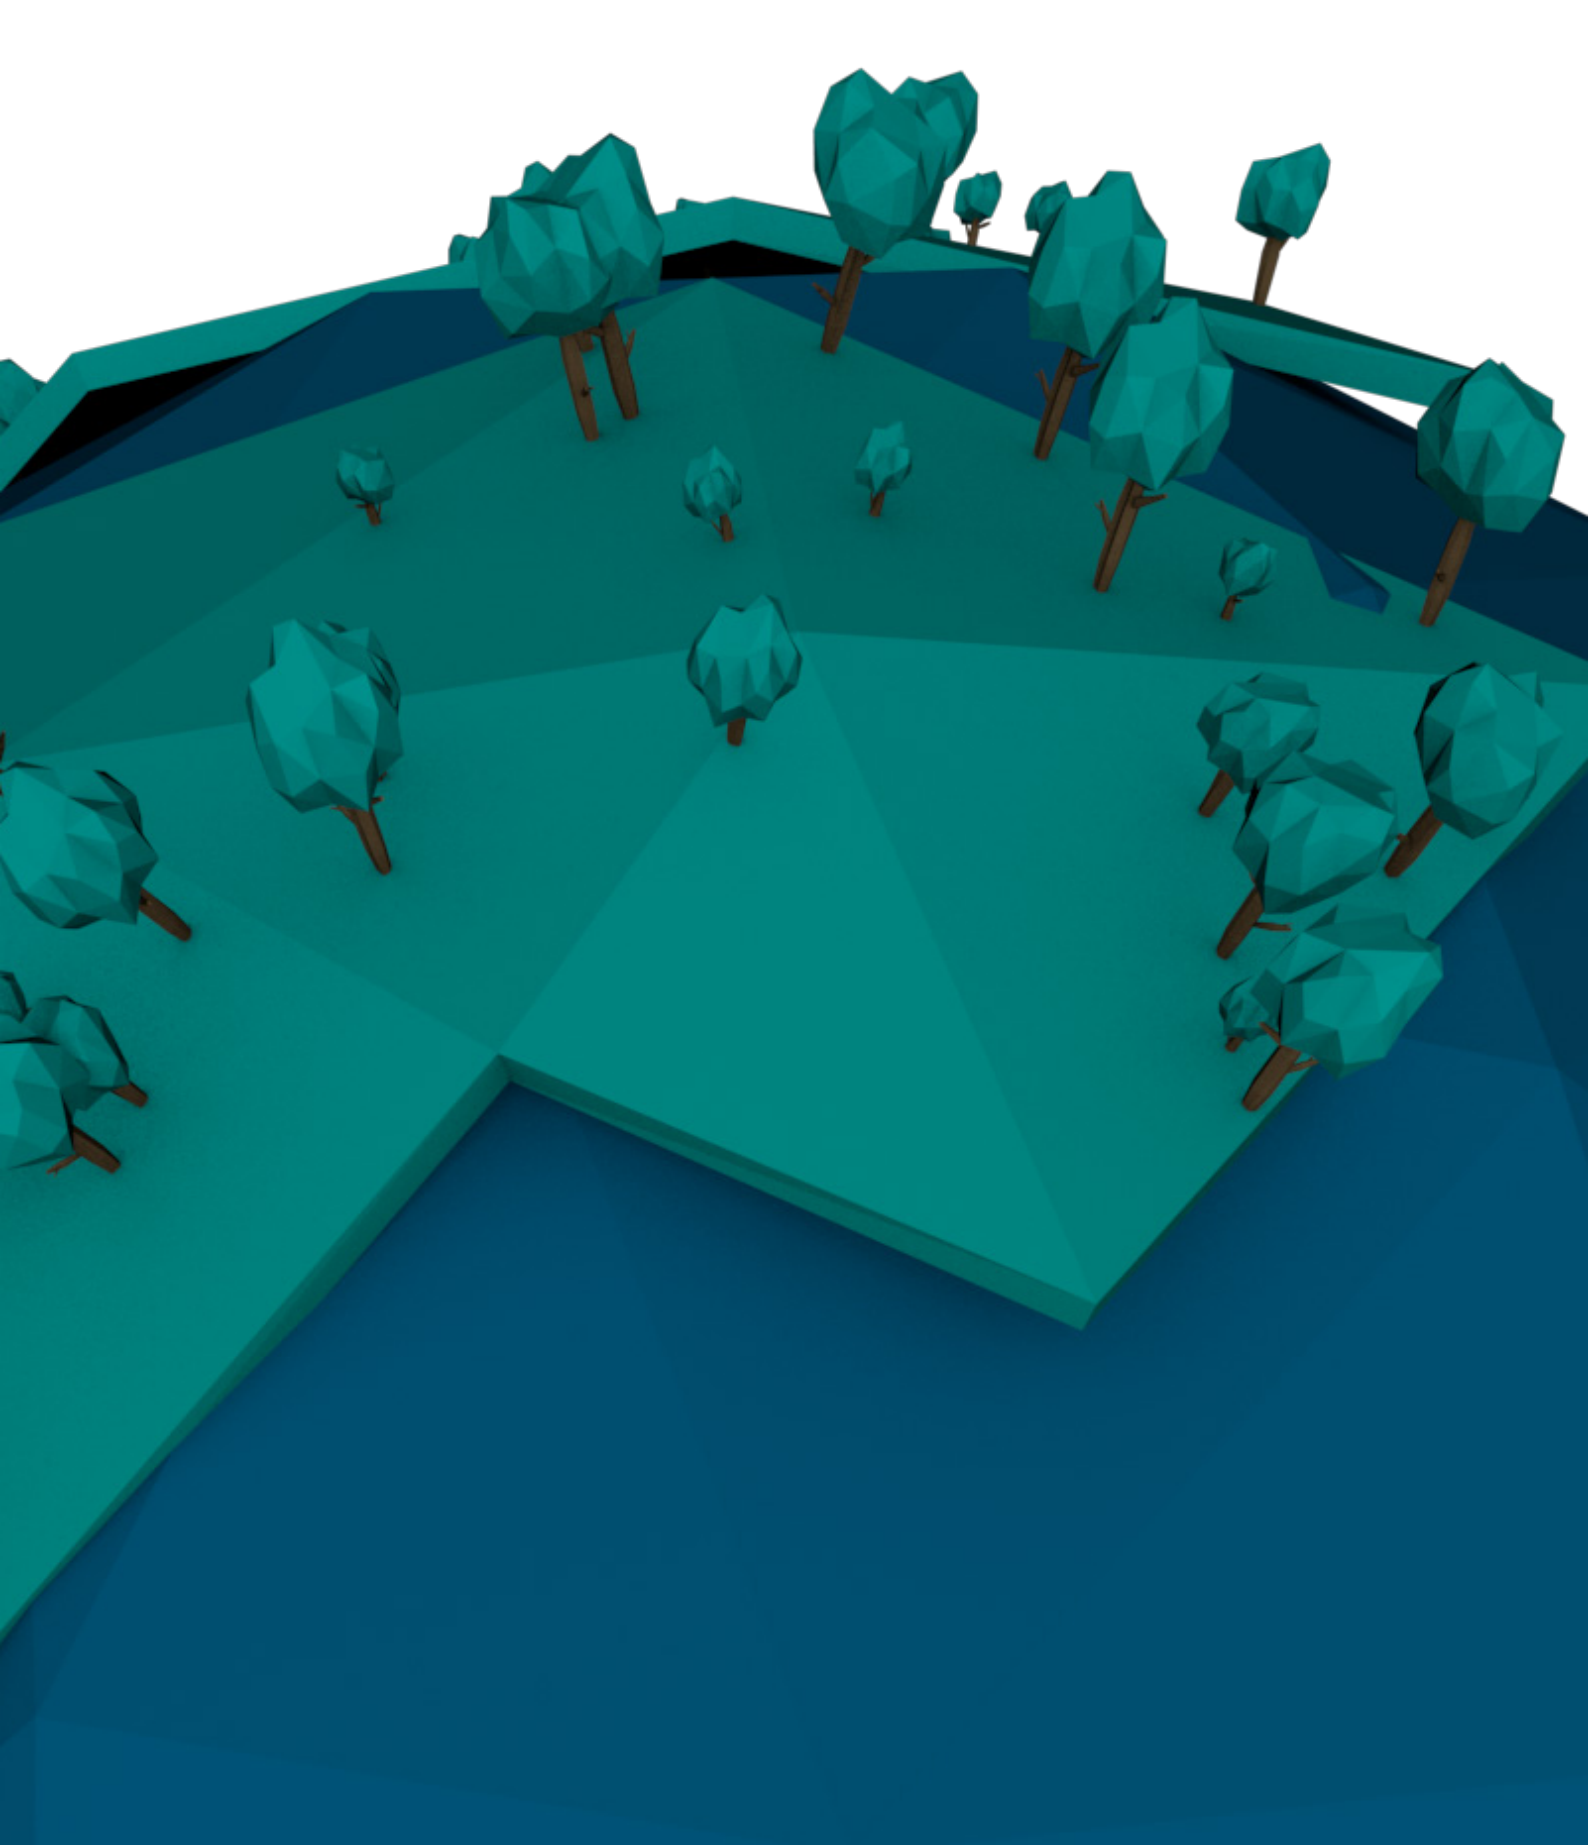

Missile Crisis, President Kennedy estimated the chance of nuclear conflict as "between one in three and even". Tensions have eased somewhat since the Cold War, but could recur. Moreover, accidents or miscalculation with nuclear weapons continue to pose a risk.

Climate change is a well-known anthropogenic risk. Even if we succeed in limiting emissions, scientists expect significant climate change to occur. This could bring a host of global challenges including environmental degradation, migration, and the possibility of resource conflict. But this is not the worst-case scenario. Although it receives far less attention, scientists also acknowledge the possibility of catastrophic climate change. There is a small likelihood that warming could even exceed 6 °C, leaving large swathes of the planet dramatically less habitable. This could occur if emissions are not cut sufficiently, if the sensitivity of the climate system is different from what is expected or if positive environmental feedback loops occur.

Catastrophic risks from emerging technology are less well understood. Emerging technologies promise significant benefits, but a handful could also create unprecedented risks to civilisation and the biosphere alike. Biotechnology could enable the creation of pathogens far more damaging than those found in nature, while in the longer run, artificial intelligence could cause massive disruption.

The relative likelihood and urgency of the different risks matters when deciding how to respond. Even though the level of uncertainty is extreme, rational action requires explicit assessments of how much attention the different risks deserve, and how likely they are. The views of the authors on these vexed questions, based on our reading of the scientific evidence, are summarised in the following table. More information can be found in the full version of this report.

FIGURE 1. OUR ASSESSMENT OF GLOBAL CATASTROPHIC RISKS

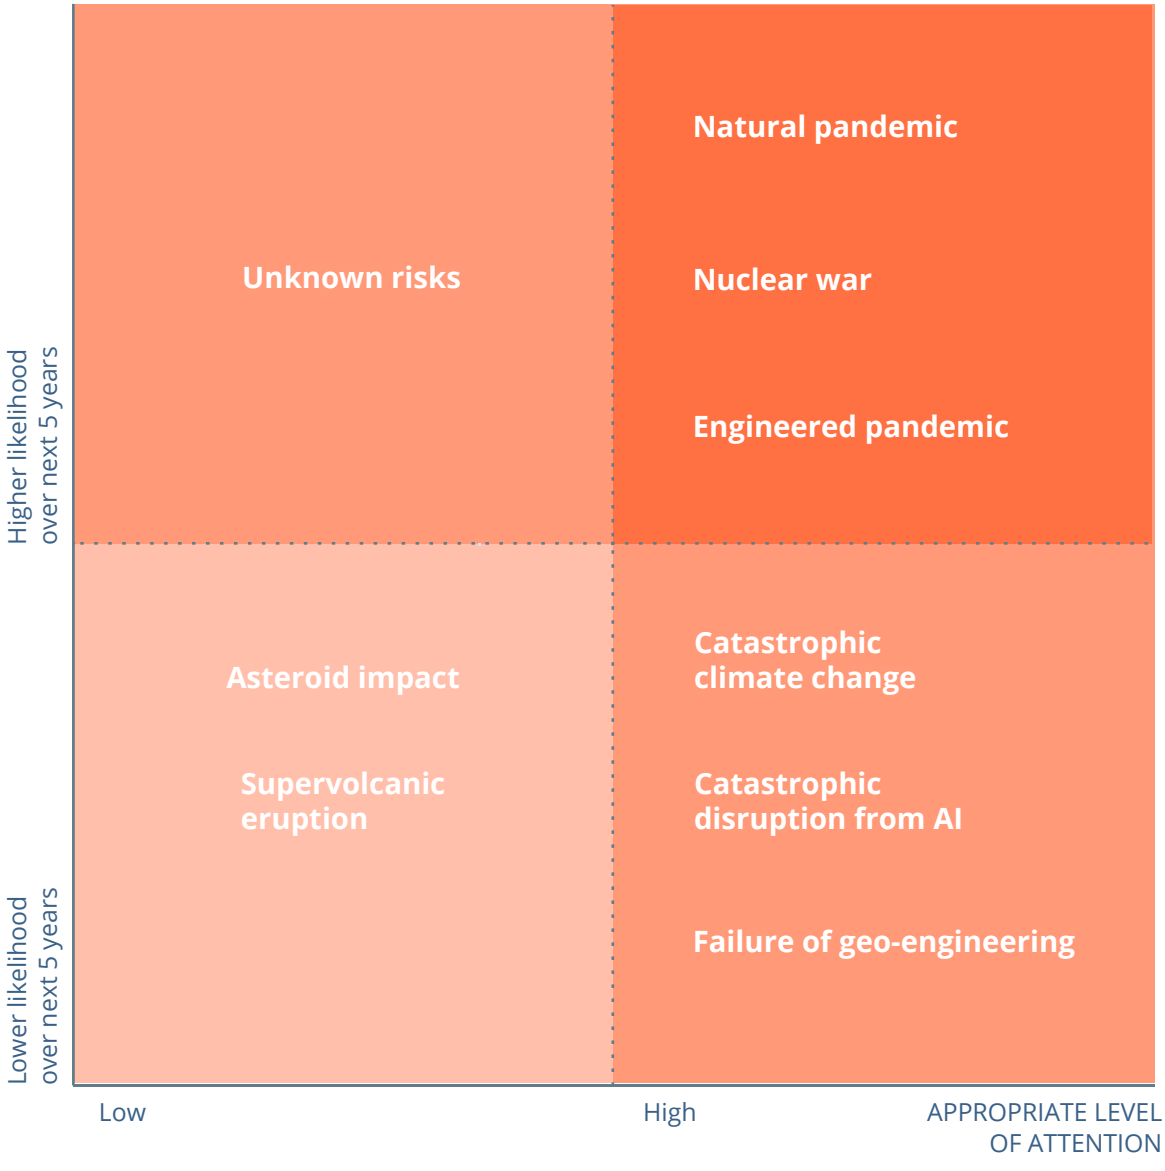

Humanity can respond to these risks

For each of the risks in this report, we consider actions available to avoid or mitigate the risk, and which actors are best-placed or responsible for taking that action. For the most significant risks, some of the most promising opportunities are listed here. The full version of this report lists further recommendations which might apply to these and other domains.

To reduce the risk of global catastrophe caused by pandemic:

- The World Health Organisation, nation states, and other bodies should increase their planning for extremely bad pandemics.
- The global health community should improve developing world capacity for response, for example by ensuring that vaccine production facilities are well-distributed around the world.

To reduce the risk of global catastrophe caused by climate change:

- Research communities should increase their focus on understanding the pathways to and likelihood of catastrophic climate change, and possible ways to respond.
- Nations should continue to implement and improve mechanisms for emissions abatement such as carbon taxes or tradable emissions quotas, as for non-catastrophic climate change.

To reduce the risk of global catastrophe caused by nuclear war:

- The international community should continue the policy of nuclear non-proliferation, and nuclear states can continue to reduce stockpiles.
- Nuclear-weapon states should continue to work to reduce the chance of accidental launch or escalation.

To reduce the risk of global catastrophe caused by emerging technologies:

- Research communities should further investigate the possible risks from emerging capabilities in biotechnology and artificial intelligence, and possible solutions.
- Policymakers could work with researchers to understand the issues that may arise with these new technologies, and start to lay groundwork for planned adaptive risk regulation.

To reduce global catastrophic risk in a cross-cutting way:

- Research communities should focus greater attention on strategies and technologies for resilience to and recovery from global catastrophe, for example by developing alternate food sources.
- Nations should work to incorporate the interests of future generations into their decision-making frameworks.

Research communities should further investigate the possible risks from emerging capabilities in biotechnology and artificial intelligence, and possible solutions.

Chapter 1

An Introduction to Global Catastrophic Risks

Over the course of history, the world has suffered disasters of such magnitude that human civilisation itself has been threatened. Warfare and pandemics have caused especially significant damage. Originating in 541-542, the initial outbreak of the 'Great Plague of Justinian' killed 25-33 million people – between 13% and 17% of the world population at the time.[1] The plague had trans-generational consequences: many historians believe that it weakened the Byzantine Empire at a crucial time, undermining its attempts to reconquer Europe.[2]

In recent times, humanity has not endured catastrophic events on the proportionate scale of Plague of Justinian. However, the risk of global catastrophe, which is determined by the potential damage of the event and its probability of occurring, has at times been uncomfortably high.

Throughout the Cold War, the threat of nuclear warfare loomed large. The United States and the Soviet Union possessed tens of thousands of high yield nuclear warheads, and their retaliatory strike systems were programmed to respond to any attack within minutes. The world has come close to the nightmare scenario on a number of occasions. Perhaps the narrowest escape came on 27th October 1962. Two Russian B-59 submarine commanders off the coast of Cuba gave the order to launch a nuclear strike against the United States, on the mistaken assumption that war had already started. The launch of a nuclear torpedo required the consent of all three officers on board; the second in command, Vasili Arkhipov, was alone in refusing permission.[3]

While the nuclear threat has receded since the

As many as 17%

of the world's population was killed in the initial outbreak of the 'Great Plague of Justinian'. That equals 25-33 million people

> The most deadly event of the 20th century was probably the Spanish influenza pandemic of 1918-1920 which killed between 2.5% and 5% of the world population... Our focus here is on even more extreme possibilities which receive less attention.

end of the Cold War, the risk remains. In addition, ongoing economic and technological developments bring, alongside their benefits, a range of new unprecedented anthropogenic risks: for example, catastrophic climate change, pandemics of global proportions, and the potential for machine intelligence which could behave in a manner incompatible with human values.

However, our governments and institutions, whose primary focus is understandably on more day-to-day concerns, may systematically be neglecting global catastrophic risks.

1.1. Defining Global Catastrophic Risk

Global catastrophes are events or processes that would inflict serious damage to humanity on a global scale, such as all-out nuclear war or a pandemic killing hundreds of millions. The severity of a risk is a function of its scope (the size of the population at risk), intensity (how badly this population would be affected), and probability (how likely the disaster is to occur).[4]

A fatal car crash is a personal catastrophe: a small number of unfortunate victims suffer a severe harm. Most genocides are examples of local catastrophes: thousands or millions of people within a country or region lose their lives. The focus of this report is catastrophic risks with global scope. We define a global catastrophe as a possible event or process that, were it to occur, would end the lives of approximately 10% or more of the global population, or do comparable damage. Extinction risks are a subset of global catastrophic risks, which would end the human race.

It is important to put the scale of global catastrophic risks in context. None of the various humanitarian disasters of the 20th century killed more than 10% of the world population. Around 1% of the world population died in the First World War, while up to 3% died in the Second World War.[5] The most deadly event of the 20th century was probably the Spanish influenza pandemic of 1918-1920 which killed 50 - 100 million people – between 2.5% and 5% of the world population.[6] Although these were huge tragedies, our focus here is on even more extreme possibilities which receive less attention. In Chapter 2 we give an overview of the major risks of global catastrophe, and in Chapter 3 we look at the factors which may increase or decrease these risks.

Limited historical evidence makes it very difficult to provide a definitive list of past global catastrophes. There have been at least two in the past two millennia – the Plague of Justinian and the Black Death. Some scholars have argued that more than 10% of the world population lost their lives in pre-industrial wars, though this is heavily disputed.[7]

FIGURE 1.1. QUALITATIVE RISK CATEGORIES [8]

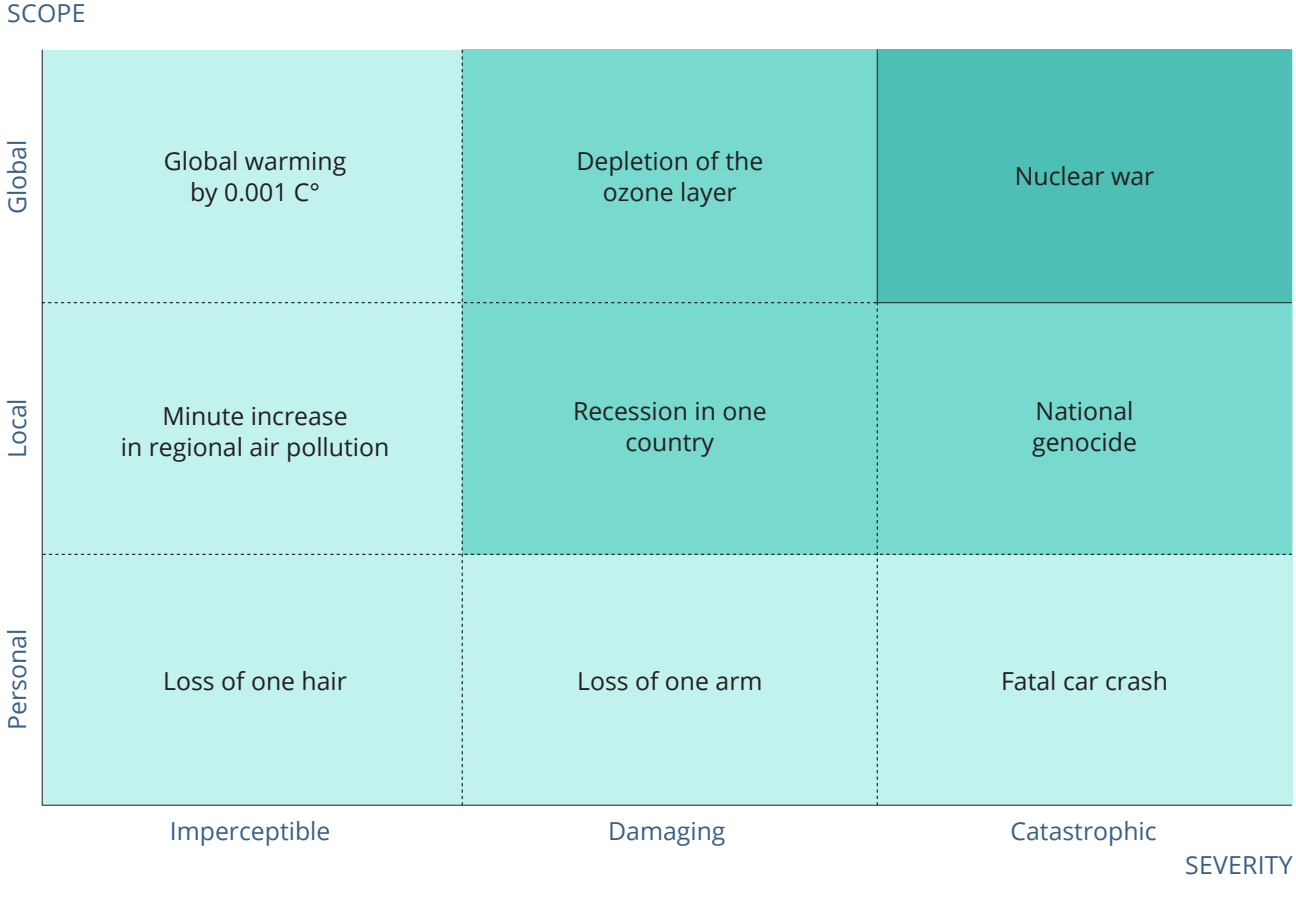

1.2. Why Global Catastrophic Risks matter

Although the chance of dying in a car crash is small, we each take steps to mitigate the risk such as wearing seat belts and driving safely. National governments take steps to mitigate the risk of rare natural disasters, such as earthquakes and hurricanes. Similarly, it is important that the global community works to reduce the risk of catastrophic events which would have a global scope.

The probabilities of these catastrophic events are low but not negligible. Moreover, small annual probabilities compound significantly over the long term.

We do not know of a robust estimate of the annual probability of global catastrophic risk. Nor do we believe that we are able to create a robust estimate because the uncertainties in key parameters are so large. However, for extinction risks some experts have suggested that a 0.1% annual chance of extinction is within the range of plausible orders of magnitude. A 2008 Oxford survey of expert judgement on the topic implied an average annual extinction risk over the next century of around 0.2%.[9] The UK's Stern Review on the Economics of Climate Change used 0.1% as an upper bound modeling assumption for annual extinction risk.[10]

Now let's suppose that the chance of extinction were 0.1% per year and consider the consequences. It may seem at first glance that this would be an acceptable level of risk. However, that would mean an individual would be more than five times as likely to die in an extinction event than a car crash.[11] Moreover, these small annual probabilities add up, so that the chance of extinction within the next century under this scenario is 9.5%.[12] A global catastrophe, which involves the death of 10% of the global population, is more likely than an event that involves human extinction. As a result, even if 0.1% were on the high side for extinction risk, it might be of the appropriate order of magnitude for global catastrophic risk.

Reducing these risks has obvious humanitarian benefits for those alive today. But we should also consider the welfare of future generations. A global catastrophe could reduce the standards of living for many generations to come, while outright human extinction denies existence to all future generations. Many leading moral philosophers have argued that the welfare of these future generations is of utmost importance.[13]

Global catastrophic risks are also likely to be politically neglected, for

> The probabilities of these catastrophic events are low but not negligible. Moreover, small annual probabilities compound significantly over the long term.

several reasons. First, global catastrophic risk mitigation is a global public good. Governments pursuing the national interest are therefore likely to attempt to free ride on the efforts of others to reduce global catastrophic risk, which means that global catastrophic risk reduction is likely to be undersupplied. Global catastrophic risk reduction is also an intergenerational public good, so even if there is effective international cooperation, mitigation policy will tend to pay insufficient attention to the interests of future generations. As a result, our descendants may be exposed to unnecessarily high levels of global catastrophic risk, or be too likely to have to live in the aftermath of one. Special interests are likely to work against global catastrophic risk mitigation because many regulations will impose concentrated costs on specific industries, whereas the beneficiaries are dispersed. Finally, because most global catastrophic risks are unprecedented, they are not salient to voters or policymakers, and our regulatory response cannot rely on learning from experience.

1.3. Why Global Catastrophic Risks are especially relevant today

Now more than ever before, global catastrophic risks deserve attention. Prior to the 20th century, the main global catastrophic risks that humankind faced were natural pandemics and conventional warfare. However, economic and technological development have brought a range of new anthropogenic risks.

The first of these new risks was nuclear weapons, which gave states unprecedented destructive power and emerged very rapidly: the bombings of Hiroshima and Nagasaki came only six years after Einstein's letter to Roosevelt warning of the dangers of nuclear fission.[14] Other anthropogenic risks might also mature quickly giving us little time to prepare. Advances in certain kinds of biotechnology, for example, might at some point in the next few decades give states, or even terrorist groups, the capacity to create devastating designer pathogens.[15] Likewise, experts warn of the longer-term risks associated with powerful machine intelligence, which may prove hard to control safely.[16]

Indeed, experience over the last century suggests that many of the most important future risks may be at present unknown. Just as in the early 20th century it would have been impossible to predict nuclear weapons, catastrophic climate change, or biotechnology risks, it may be that many of the future leading global catastrophic risks are not yet within sight.

Moreover, to reduce these new anthropogenic risks we may need levels of international coordination that existing institutions are not designed to produce, something we discuss in Chapter 4. A good illustration of this is the threat of catastrophic climate change. Atmospheric concentrations of greenhouses gases are now at their highest level for hundreds of thousands of years,[17] and if the international community fails to take strong action soon there is a worryingly high chance of warming in excess of 6°C (compared to pre-industrial levels) by the end of the century.[18] But because unilateral action is costly for any state, and the benefits are felt by everyone regardless of their contribution, action on greenhouse gas emissions has been slow in coming.

1.4. What can be done?

When dealing with global catastrophic risks we cannot generally rely on historical experience or trial and error. Given the severity of global catastrophes, learning from experience would be extremely costly or, in the event of human extinction, impossible. But policy-makers, industries, research communities, and citizens can take preemptive steps to limit global catastrophic risks. This report outlines the key features of the world's most significant global catastrophic risks and identifies, in Chapter 5, some strategies for limiting them.

1.5. How to read this report

These ideas, and the current state of the science underlying them, are explored in significantly more detail in the full version of this report. Chapter 2 gives a comprehensive overview over the main global catastrophic risks, from catastrophic climate change and nuclear war to risks associated with emerging technologies such as biotechnology, artificial intelligence and geo-engineering. We discuss the potential impact and the likelihood of each risk, as well as the main actions to limit them. Finally, we attempt to compare the risks in terms of how likely they are and in terms of how much attention we ought to pay them at present.

Chapter 3 discusses the causes of global catastrophic risks. These include both factors that increase the likelihood or impact of individual risks, as well as factors that affect multiple risks, such as poor governance.

Chapter 4 discusses why we currently collectively underinvest in global catastrophic risk. It also discusses what actors are best placed to overcome this neglect.

Finally, in Chapter 5 we briefly discuss a number of concrete steps to reduce global catastrophic risk.[19]

Endnotes

1. William Rosen, *Justinian's Flea : Plague, Empire, and the Birth of Europe* (New York: Viking, 2007), 3, 209. Percentages deduced from Demographic Internet Staff US Census Bureau, "International Programs, World Population," accessed February 12, 2016, https://www.census.gov/population/international/data/worldpop/table_history.php. Some earlier estimates put the death toll as high as 100 million. See B. Lee Ligon, "Plague: A Review of Its History and Potential as a Biological Weapon," Seminars in Pediatric Infectious Diseases, Return of the Gram-Positives, 17, no. 3 (July 2006): 161–70.
2. Rosen, Justinian's Flea, 3–5, 321–324.
3. Edward Wilson, "Thank You Vasili Arkhipov, the Man Who Stopped Nuclear War," *The Guardian*, October 27, 2012, http://www.theguardian.com/commentisfree/2012/oct/27/vasili-arkhipov-stopped-nuclear-war.
4. This is adapted from Nick Bostrom, "Existential Risk Prevention as Global Priority," *Global Policy 4*, no. 1 (February 1, 2013): 16.
5. Steven Pinker, *The Better Angels of Our Nature : A History of Violence and Humanity* (London: Penguin, 2012),195. Note that although the fatalities are too low, the non-fatality impacts mean it could be argued that the Second World War was just about large enough to count as a global catastrophe in our sense.
6. Niall Johnson and Juergen Mueller, "Updating the Accounts: Global Mortality of the 1918-1920 'Spanish' Influenza Pandemic," *Bulletin of the History of Medicine 76,* no. 1 (2002): 105–15.
7. For further discussion see the sections on natural pandemics, nuclear war and super-volcanic eruptions in Chapter 2.
8. Adapted from Nick Bostrom, "Existential Risk Prevention as Global Priority," Global Policy 4, no. 1 (February 1, 2013): 15–31, doi:10.1111/1758-5899.12002.
9. Sandberg, A. & Bostrom, N. (2008): "Global Catastrophic Risks Survey", Technical. Report #2008-1, Future of Humanity Institute, Oxford University: pp. 1-5.
10. N. H. Stern, *The Economics of Climate Change : The Stern Review* (Cambridge: Cambridge University Press, 2007), 47.
11. "Institute for Health Metrics and Evaluation (IHME). GBD Compare. Seattle, WA: IHME, University of Washington, 2015. Available from http://vizhub.healthdata.org/gbd-compare" from 'ibid.'
12. Adapted from Nick Bostrom, "Existential Risk Prevention as Global Priority," Global Policy 4, no. 1 (February 1, 2013): 15–31, doi:10.1111/1758-5899.12002.
13. Derek Parfit, *Reasons and Persons* (Oxford: Oxford University Press, 1984); Nick Bostrom, "Astronomical Waste: The Opportunity Cost of Delayed Technological Development," *Utilitas* 15, no. 03 (November 2003): 308–14; Nick Bostrom, "Existential Risk Prevention as Global Priority".
14. http://www.atomicarchive.com/Docs/Begin/Einstein.shtml
15. Ali Nouri and Christopher F. Chyba, "Biotechnology and Biosecurity," in *Global Catastrophic Risks*, ed. Nick Bostrom and Milan M. Ćirković (Oxford University Press, 2008).
16. Future of Life Institute, "AI Open Letter," accessed February 3, 2016, http://futureoflife.org/ai-open-letter/.
17. NASA, "The Relentless Rise of Carbon Dioxide," *Climate Change: Vital Signs of the Planet*, accessed February 2, 2016, http://climate.nasa.gov/climate_resources/24/.
18. Gernot Wagner and Martin L. Weitzman, *Climate Shock : The Economic Consequences of a Hotter Planet* (Princeton: Princeton University Press, 2015).
19. Clearly, enormous literatures exist in each of these areas which can inform policy. We suggest only a small selection of the policy options which are available.

Acknowledgements

This report has greatly benefited from the advice and feedback of many people. For detailed comments on a large portion of the report, we are very grateful to Niel Bowerman, Martina Kunz, Catherine Rhodes. For their suggestions and comments we are grateful to Seán Ó'heigeartaigh, Shahar Avin, and Toby Ord.

For very helpful advice on individual risks, we would like to thank Dan Bernie, Sean Brocklebank, Clark Chapman, Daniel Dewey, Simon Driscoll, Marius Gilbert, David Graff, Marc Lipsitch, Mark Lynas, Piers Millet, Oliver Morton, Andrew Parker, Alan Robock, Jonathan Skaff, Stephen Sparks, and Naomi Standen.

CONTACT INFO

Authors:
Global Priorities Project
16-17 St. Ebbes Street
Oxford, OX1 1PT
United Kingdom

contact@globalprioritiesproject.org

The Global Challenges Foundation:
Strandvägen 7A
114 56 Stockholm
Sweden

info@ globalchallenges.org
+46 (0) 709 98 97 97