Global
Challenges
Foundation

Global Catastrophic Risks

2016

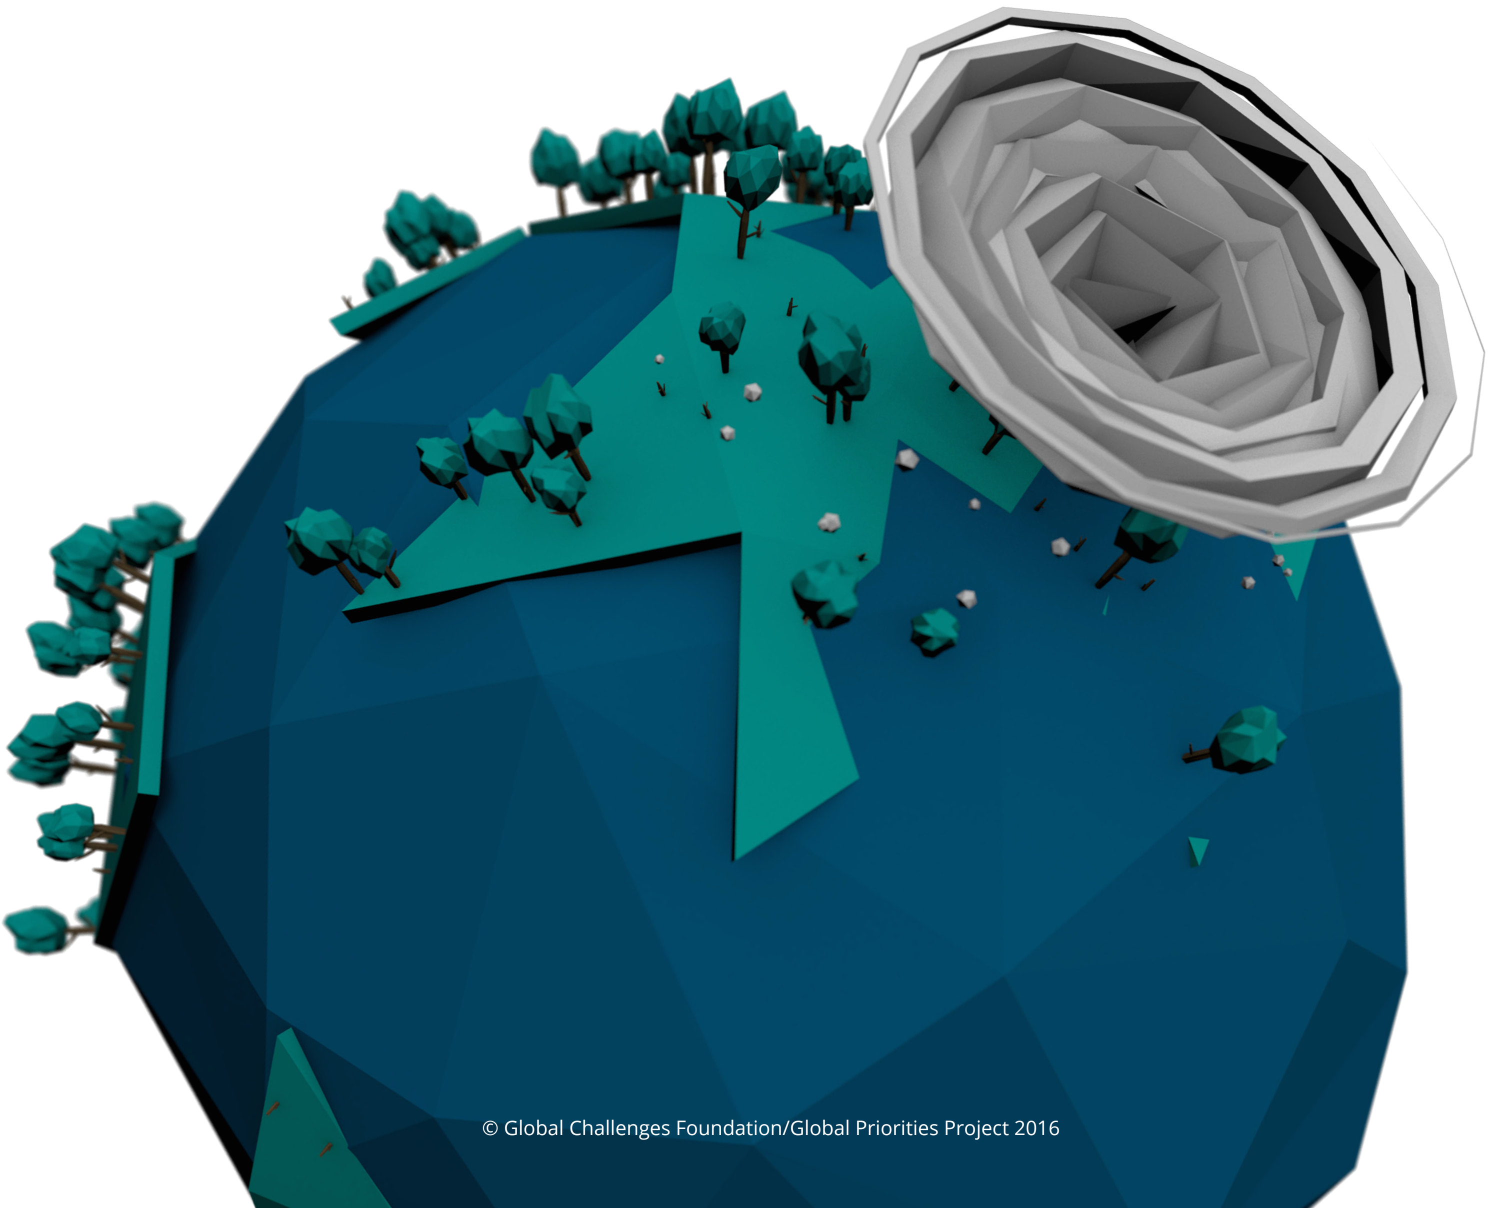

GLOBAL CATASTROPHIC RISKS 2016
The views expressed in this report are those of the authors. Their statements are not necessarily endorsed by the affiliated organisations.

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.

Contents

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 attention 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.

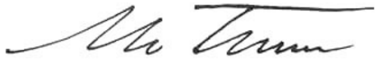

Max Tegmark
Co-founder of the Future of Life Institute
Professor of Physics at MIT

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

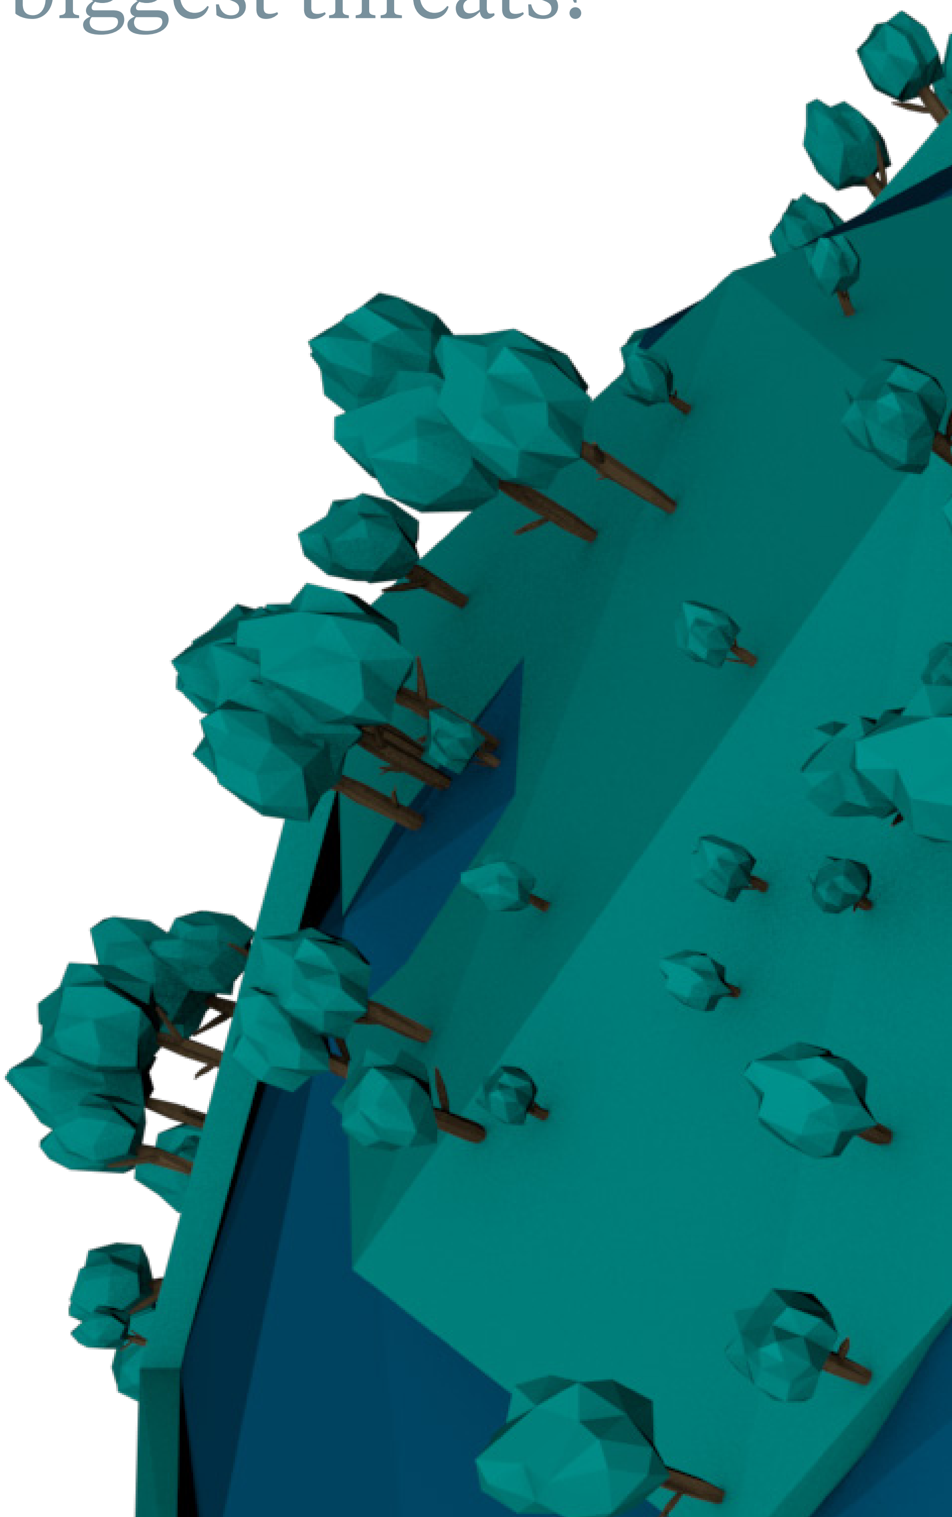

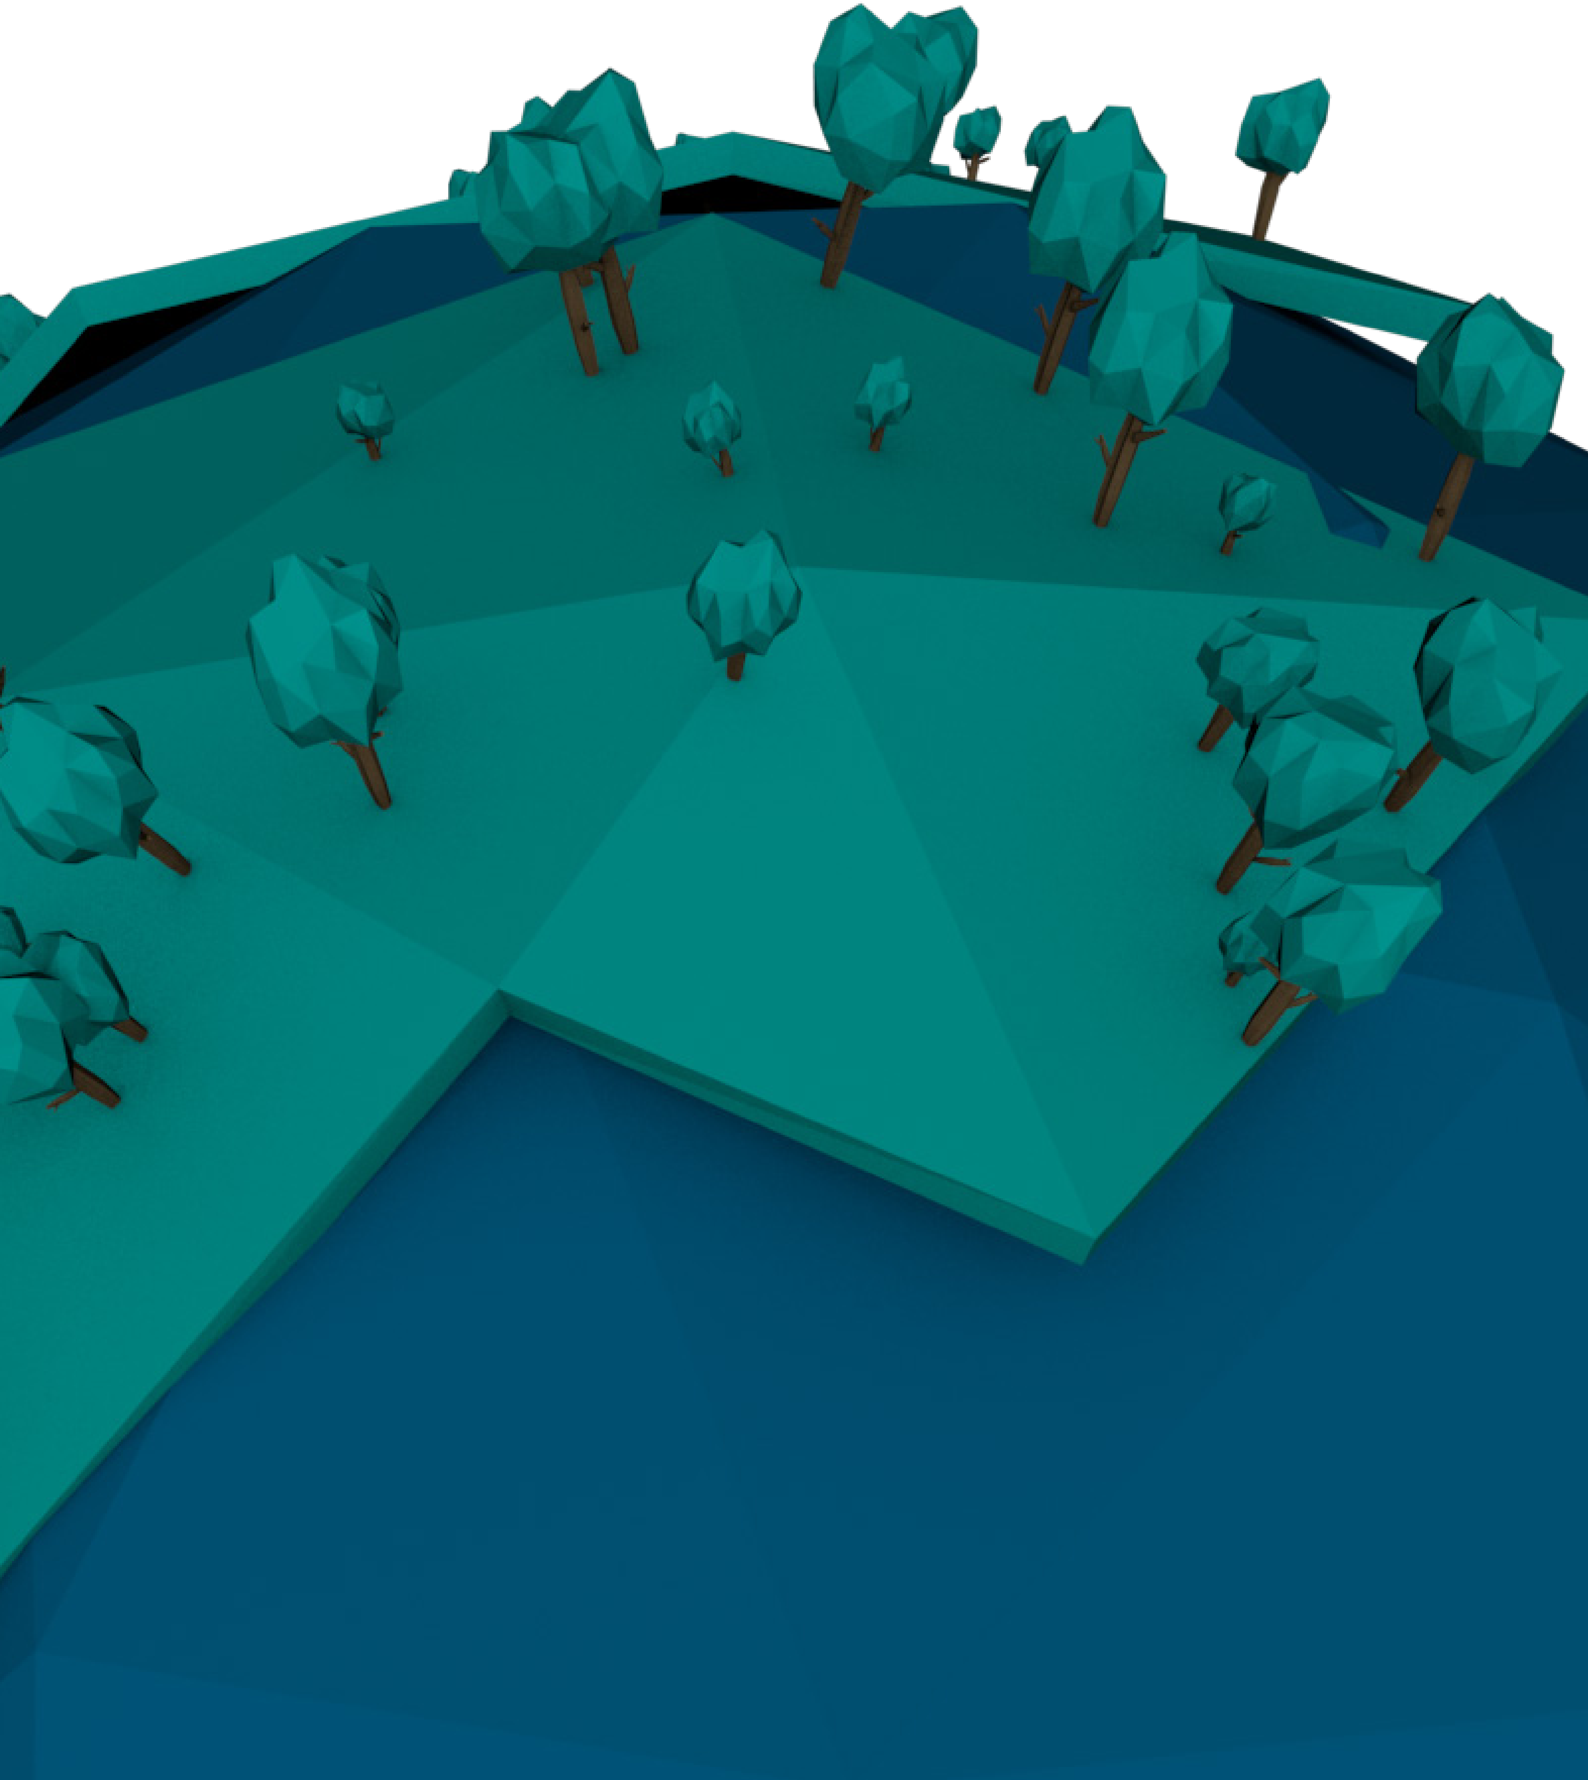

Missile Crisis, President Kennedy estimated the chance of nuclear conflict as “between one in three and even”. Tensions have eased somewhat since the Cold War, but could recur. Moreover, accidents or miscalculation with nuclear weapons continue to pose a risk.

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 Chapter 2.

FIGURE 1. OUR ASSESSMENT OF GLOBAL CATASTROPHIC RISKS

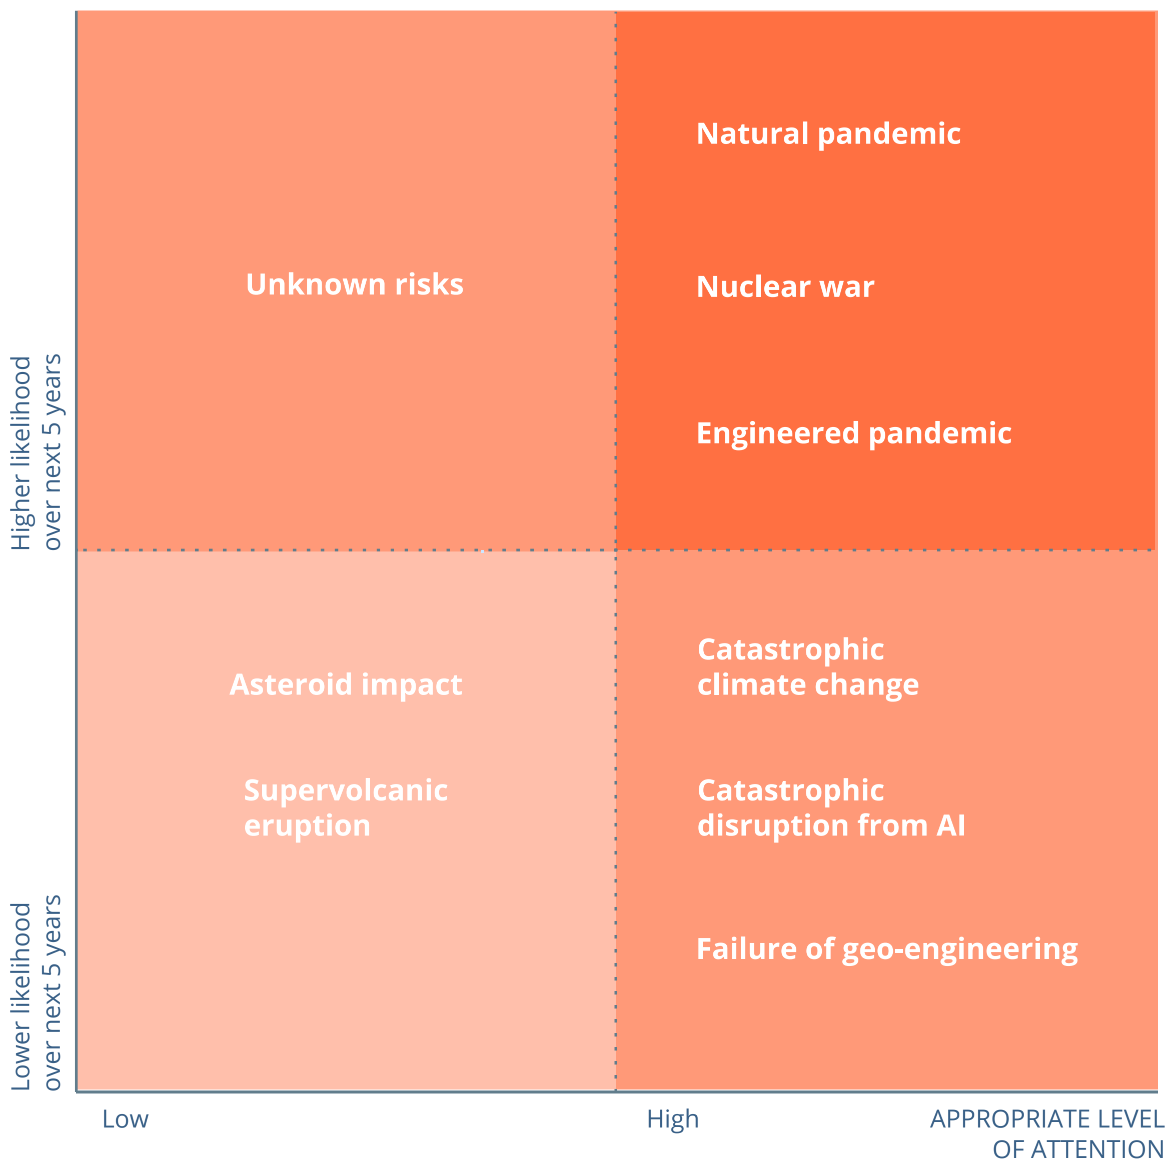

Humanity can respond to these risks

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.

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]

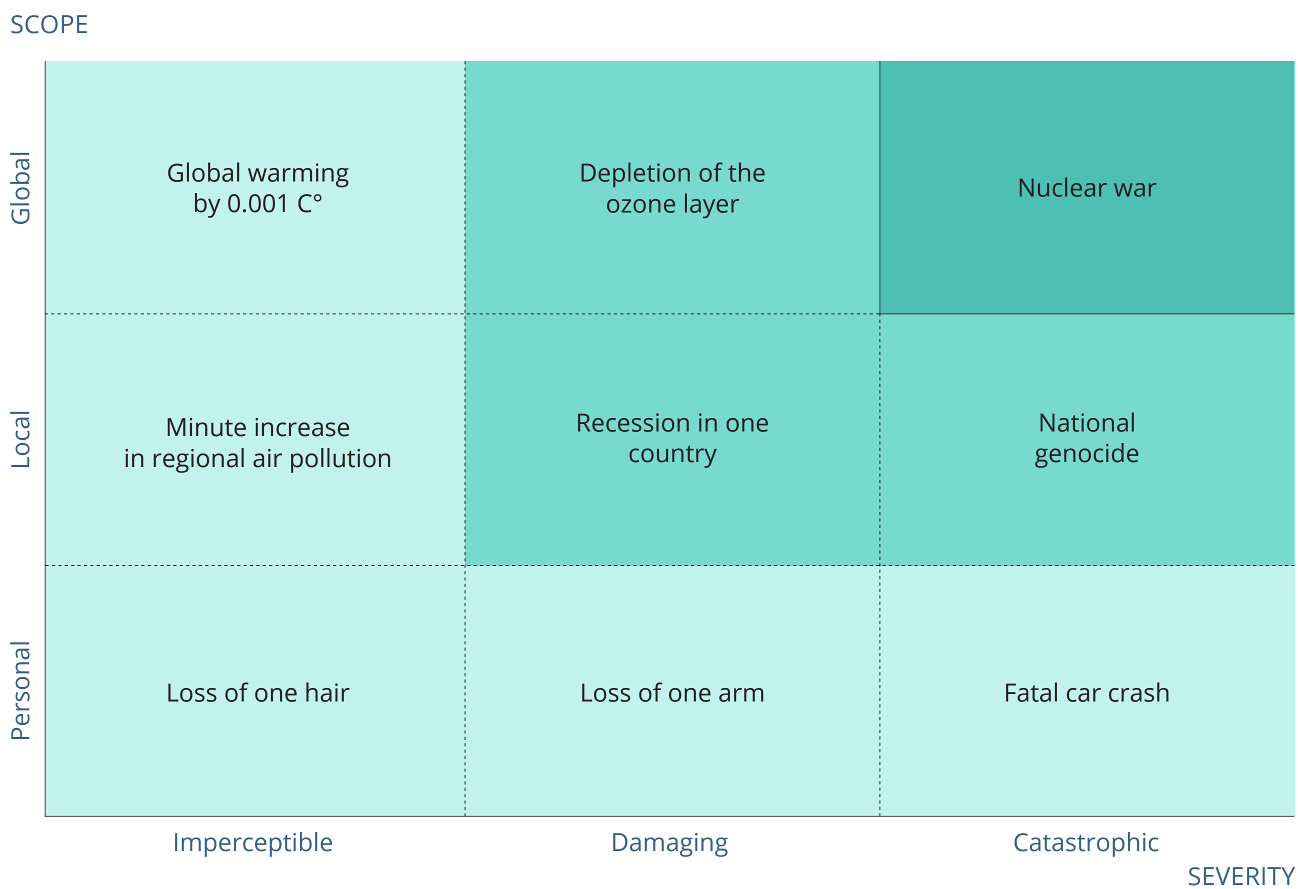

1.2. Why Global Catastrophic Risks matter

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.[14] Global catastrophic risk reduction is

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

a global public good, as even a large country would only capture a small portion of the total benefit of risk mitigation. Moreover, it is an intergenerational public good, as many of the beneficiaries are future people who have no voice in the political process. For these reasons and others, national and international actors are likely to underinvest in risk reduction.

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.[15] 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.[16] Likewise, experts warn of the longer-term risks associated with powerful machine intelligence, which may prove hard to control safely.[17]

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,[18] 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.[19] 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

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.[20]

Chapter 2

What are the Most Important Global Catastrophic Risks?

For most of human history, humanity has had to contend with a relatively narrow range of global catastrophic risks. Super-volcanic eruptions and large asteroid impacts were possible but very unlikely, so only natural pandemics and extreme conventional warfare seriously threatened the complete destruction or permanent stagnation of human civilisation.

Today, thanks to economic and technological progress, global living standards have never been higher, but unfortunately, for the same reason, we face a number of new anthropogenic global catastrophic risks. Some of these appear to be at least as threatening to human civilisation as natural pandemics and conventional warfare. Splitting the atom brought the promise of clean power, but also led to the nuclear bomb, which has brought humanity to the brink of catastrophe on more than one occasion. The burning of fossil fuels has brought huge improvements in human welfare, but unless strong action is taken soon, there is an unacceptable chance that our children and grandchildren will face catastrophic global warming. Rapid developments in biotechnology could enable scientists to develop new therapies to reduce the global burden of disease and feed a growing population, but might also in the future give malicious groups the capacity to synthesise devastating pathogens.

This chapter surveys currently the most important global catastrophic risks by examining expert scientific opinion on the two determinants of risk: potential impact and likelihood. There is also a brief discussion of actions available to limit each of these risks, which is summarised in Chapter 5. In the final section, we give a comparative assessment of the different catastrophic risks.

Splitting the atom brought the promise of clean power, but also led to the nuclear bomb, which has brought humanity to the brink of catastrophe on more than one occasion.

2.1. Catastrophic climate change

As a result of human activity since the Industrial Revolution, atmospheric concentrations of Greenhouse Gases (GHGs) are now at their highest level for hundreds of thousands of years,[21] which has caused global surface and ocean warming. Continued increases in GHG emissions are very likely to cause future warming. The eventual level of warming depends on total GHG emissions and on the sensitivity of the climate to GHG emissions. The reports of the Intergovernmental Panel on Climate Change (IPCC) focus on the most likely levels of warming given a particular emissions pathway. It is widely agreed that the negative consequences of the most likely levels of warming will be substantial. However, it is important to take account the 'tail risk' of lower probability, but potentially catastrophic, levels of warming. Of course, what we refer to in this section as 'non-catastrophic climate change' would still have severe consequences, but it would not constitute a global catastrophe, on our definition.

POTENTIAL IMPACT OF THE CATASTROPHE

Increasing GHG emissions could potentially trigger catastrophic climate change due to high climate sensitivity or strong positive feedback loops in the carbon cycle. For example, global warming might cause the melting of arctic permafrost, which would release substantial amounts of methane – a potent GHG – into the atmosphere.[22] This process could itself trigger other positive feedback loops. Catastrophic warming could also occur even without these feedback effects, if climate sensitivity turns out to be higher than median estimates. Alternatively, it might occur simply because we are less able to coordinate internationally to reduce emissions than we expect.

It is impossible to say with confidence exactly what level of warming would bring about global catastrophe, in the sense we are interested in here. The IPCC states:

"Global climate change risks are high to very high with global mean temperature increase of 4°C or more above preindustrial levels in all reasons for concern, and include severe and widespread impacts on unique and threatened systems, substantial species extinction, large risks to global and regional food security, and the combination of high temperature and humidity compromising normal human activities, including growing food or working outdoors in some areas for parts of the year. The precise

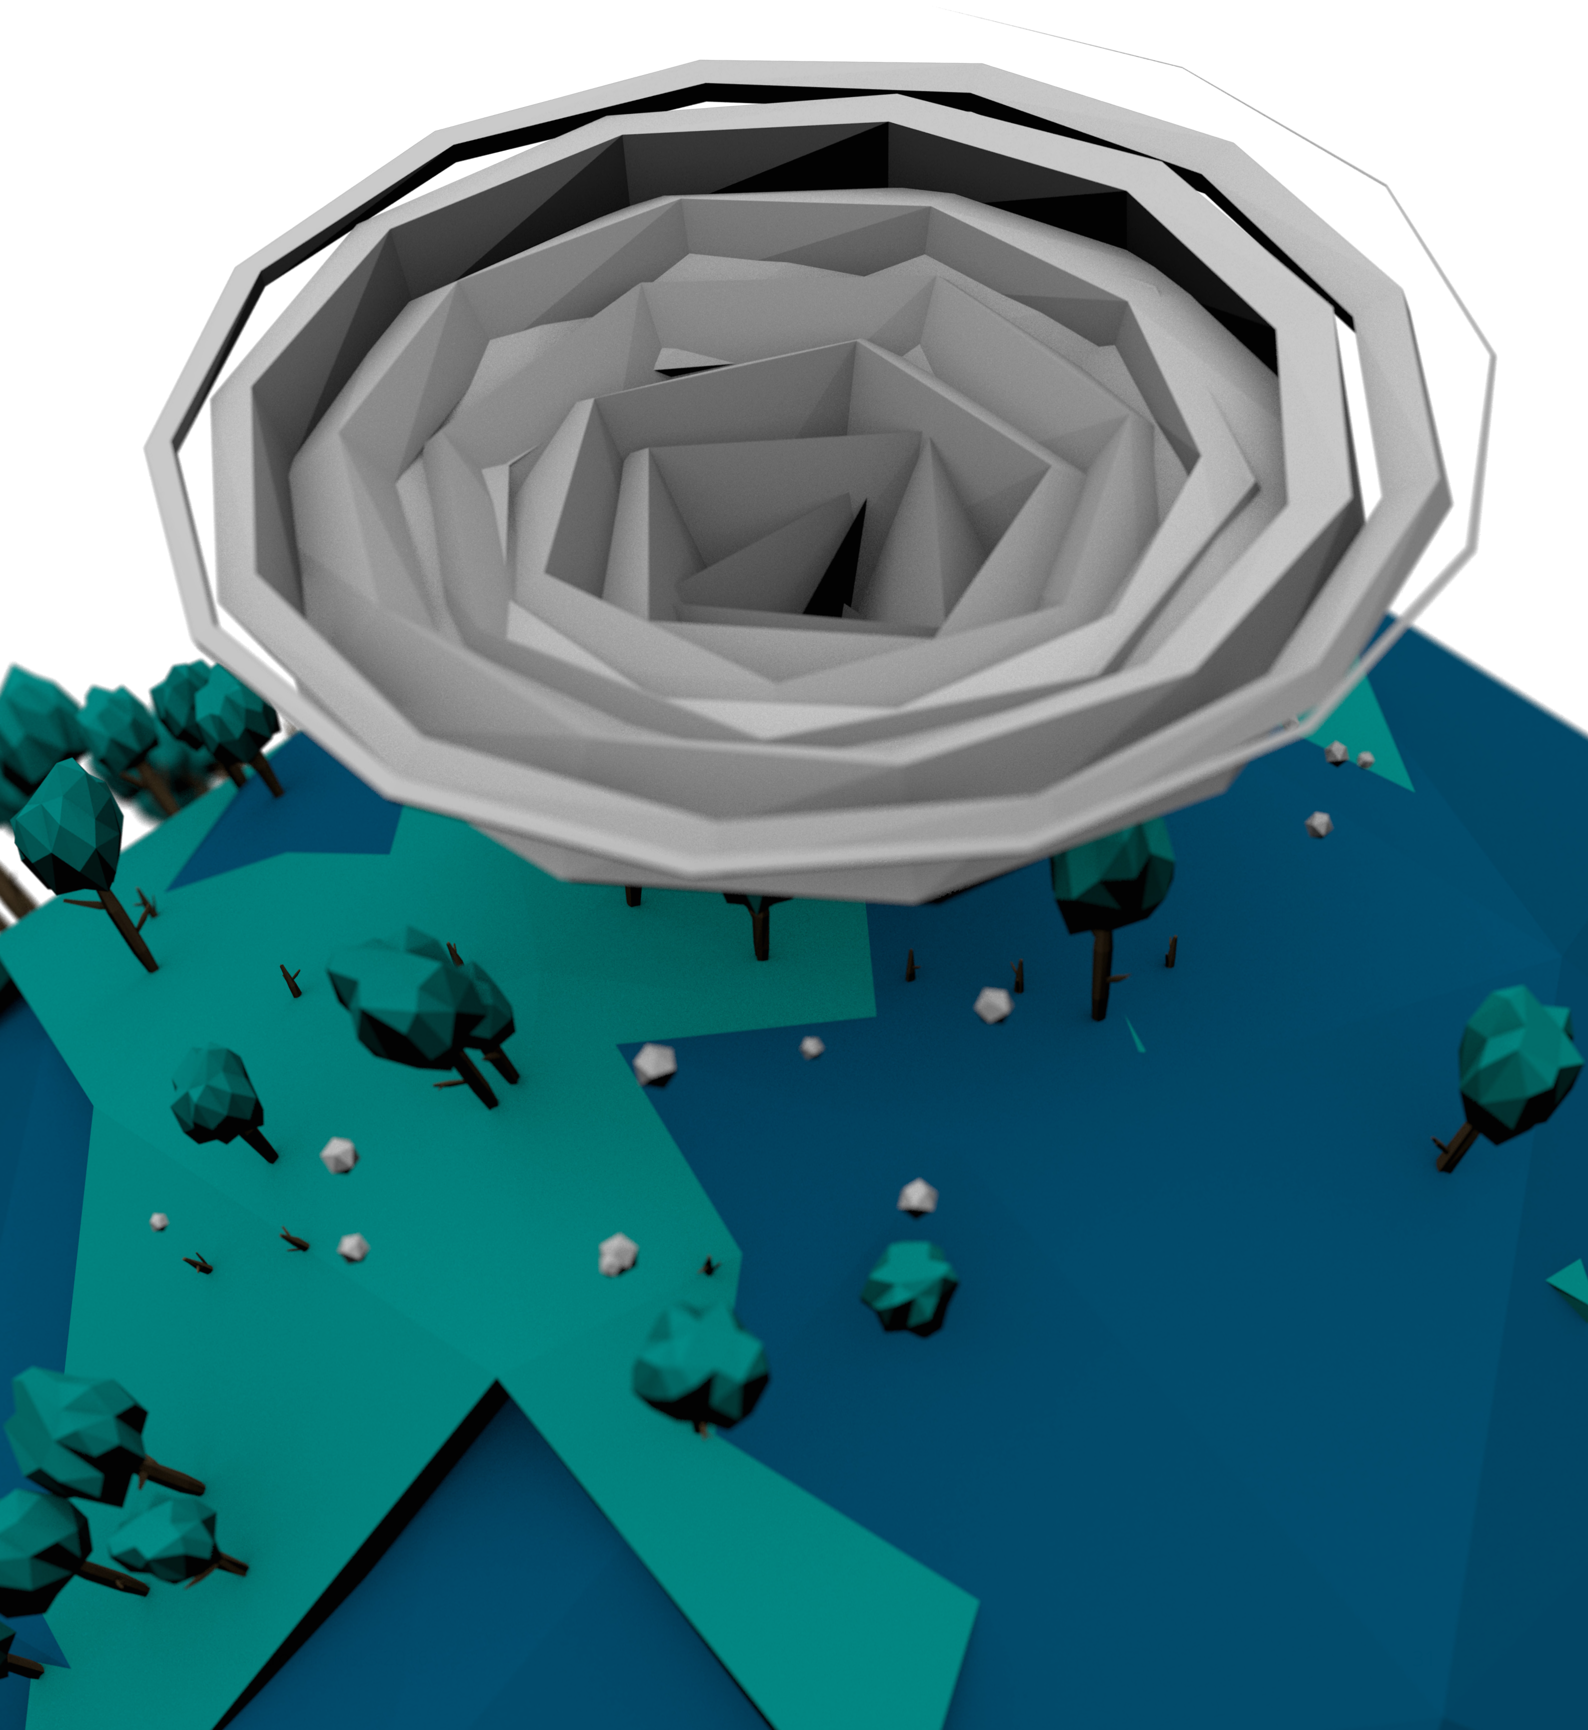

levels of climate change sufficient to trigger tipping points (thresholds for abrupt and irreversible change) remain uncertain, but the risk associated with crossing multiple tipping points in the earth system or in interlinked human and natural systems increases with rising temperature."[23]

The latest IPCC report focuses on the impacts of warming of 1°C to 4°C above pre-industrial levels, and suggests that the impacts corresponding to more extreme levels of warming are relatively unstudied.[24] However, it is likely that damages increase significantly at higher temperatures,[25] and possible that warming of 6°C or more above pre-industrial levels may be catastrophic. Warming of this magnitude is, for example, likely to render most of the tropics substantially less habitable than at present.[26]

LIKELIHOOD OF THE CATASTROPHE

The probability of catastrophic climate change depends on the level of GHGs in the atmosphere and on the sensitivity of the climate to cumulative GHGs. According to a report by King et al for the Centre for Science Policy at the University of Cambridge, if major countries and regions continue with current plans and projects, the world is most likely to follow a medium-high emissions scenario, and there is some chance we will follow a high emissions scenario.[27] However, the situation may have changed as a result of the December 2015 Paris Agreement.[28]

In estimating the impact of different emissions scenarios, the IPCC does not set out the probability of all possible levels of warming, but instead only sets out the likely range of warming, where 'likely' is defined as having a greater than 66% chance of occurring.[29] However, in order to prioritise resources effectively governments need to take into account the whole probability distribution, including lower probability but extreme levels of warming. Indeed, these may constitute the majority of the expected (probability-weighted) costs of climate change.

Some scholars have provided estimates of the probability of extreme warming. The economists Gernot Wagner and Martin Weitzman have inferred estimates of the probability of warming of more than 6°C relative to pre-industrial levels from the IPCC figures. They argue that even on a low-medium emissions scenario, there is at least a 3% chance of eventual 6°C warming (with significant uncertain-

$40
per tonne
Is the recommended global price of carbon by some leading economists. Current price is significantly less.

Governments need to take into account lower probability but extreme levels of warming. Indeed, these may constitute the majority of the probability-weighted costs of climate change.

ty). On the medium-high emissions scenario, the chance could be around 10%.[30] These figures are of course speculative, but they do provide some reason to believe that the probability of catastrophic climate change is non-negligible, unless strong action is taken on GHG emissions.

MAIN ACTIONS AVAILABLE TO LIMIT THE RISK

There are three main ways to reduce the risks from climate change: adaptation to climate change, abatement of GHG emissions, and geo-engineering. It is likely to be very costly if not impossible to to avoid many of the impacts of non-catastrophic climate change by adaptation alone,[31] and adapting to warming of 6°C or more is likely to be even more costly and difficult.

Turning to abatement, most economists agree that the best way to reduce GHG emissions is to impose a carbon tax or a cap and trade system.[32] At present, the global price of carbon is approximately $4 per tonne, whereas according to Wagner and Weitzman, to fully price in the externalities from catastrophic climate change, a price of at least $40 could be required.[33]

Scientists have demonstrated an approximately linear relationship between the total amount of carbon emitted and the resulting temperature increase.[34] The majority of the carbon still underground is in the form of coal.[35] If we were to avoid burning the remaining global coal reserves we would likely avoid catastrophic levels of climate change.[36] States could commit to building no new coal-fired power stations without carbon capture and sequestration to limit the fraction of the global coal reserves which are burned.

Geo-engineering – the deliberate use of technology to alter the world's climate – in the form of Carbon Dioxide Removal (CDR) or Solar Radiation Management (SRM), could also help to reduce the risk of catastrophic climate change, as a complement to GHG abatement. CDR techniques, such as carbon sequestration or iron fertilisation of the oceans, would remove CO2 from the atmosphere and thereby help us move towards net neutral or net negative emissions.[37] SRM techniques, such as the injection of sulphates into the stratosphere, cause global cooling by reflecting sunlight. The benefits and risks of geo-engineering are discussed in more detail in section 2.5.[38]

FIGURE 2.1. THE CHANCE OF EXTREME CLIMATE CHANGE

The probability of warming of 6°C for different atmospheric concentrations of greenhouse gases.[39]

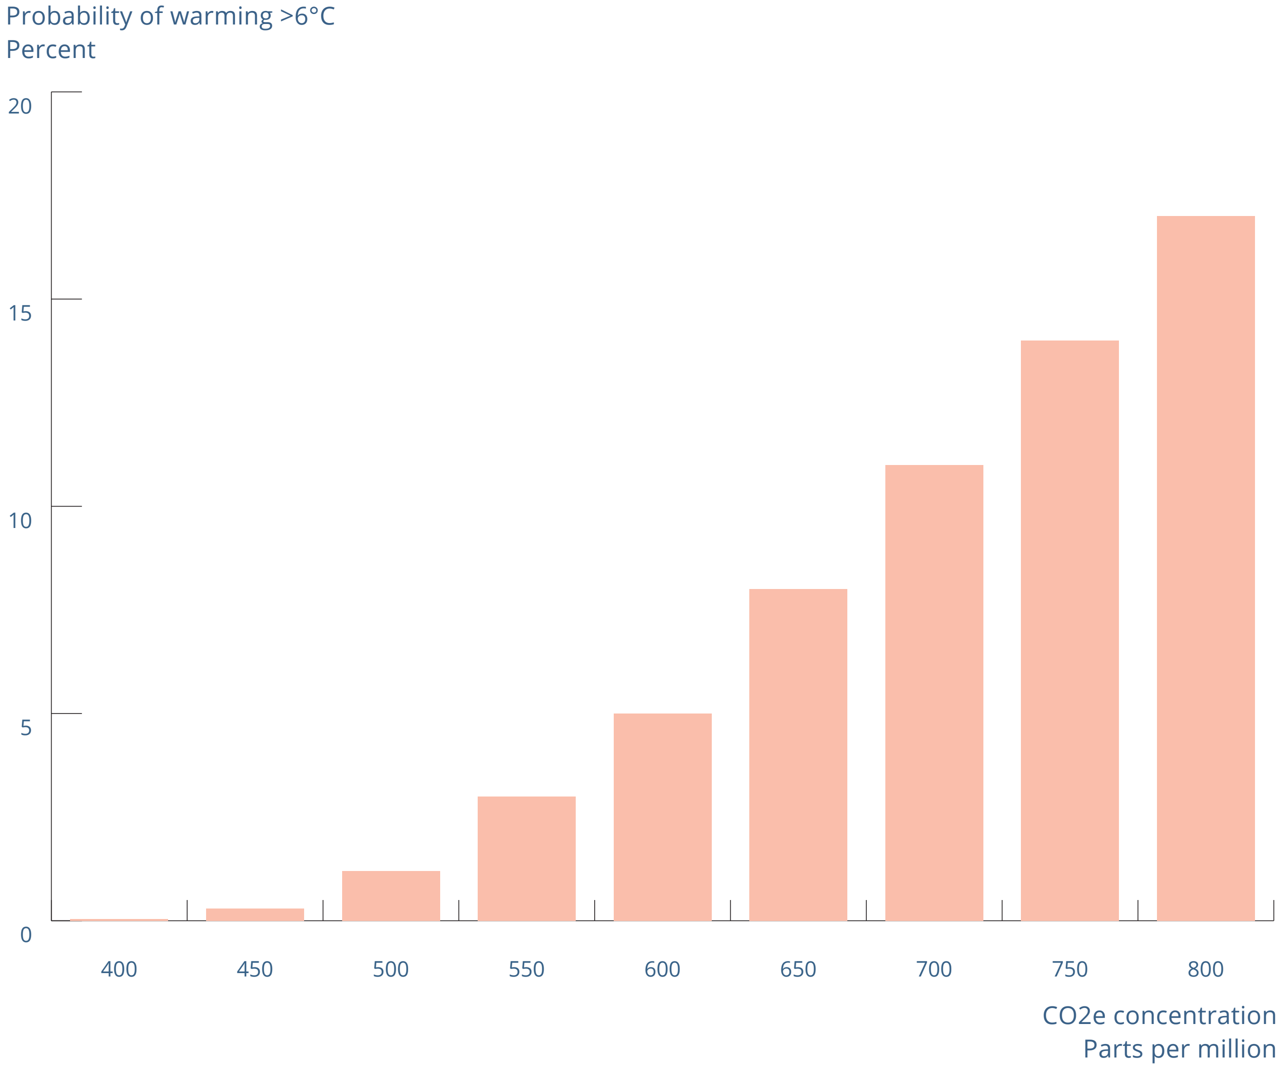

2.2. Nuclear war

The invention of nuclear weapons gave humanity the technical capacity to cause devastation on a hitherto unseen scale. Although there have been no nuclear attacks since the Second World War, we have come close to inadvertent and intentional nuclear war on a number of occasions. A worrying possible consequence of nuclear war is a nuclear winter with global climatic implications. While the chance of nuclear war may appear to have declined since the end of the Cold War, tensions between nuclear states persist. Reducing the likelihood of nuclear war is a serious ongoing global challenge.

Despite these risks, some have argued that the deterrence effects of the nuclear bomb have been, and will be, very valuable in terms of ensuring global peace.[40] How to make the trade-off between lower likelihood of nuclear or conventional conflict is unclear.

POTENTIAL IMPACT OF THE CATASTROPHE

The scale of the damage done by nuclear war obviously depends on the scale of the nuclear war itself. In 2014, there were 9,920 stockpiled nuclear warheads in the world (down from a peak of 65,000 in 1986).[41] The United States has 4,760, with yields ranging from 5 kilotons to 455 kilotons;[42] and Russia has 4,300, with yields ranging from 50 to 800 kilotons.[42] For comparison, the Little Boy bomb dropped on Hiroshima had a 15 kiloton yield.[44]

The damage from nuclear war can be divided into two main categories. Firstly, there is the damage from the blast, fire and radiation. A 1979 report by the U.S. Office of Technology Assessment estimated that, in an all-out war between the US and Russia involving thousands of nuclear weapons, 35-77% of the U.S. population and 20-40% of the Soviet population would die within the first 30 days of the attack, and millions would die globally in the following years due to the radioactive dust cloud.[45] The proportionate death toll today is likely to be lower because nuclear arsenals at that time were five times larger.[46]

The second category of damage is a possible nuclear winter, which is caused by the burning of cities, industrial facilities and other flammable materials, sending smoke into the atmosphere. Scientists have applied modern climate models to predict the scale of nuclear winter, though these predictions are uncertain. According to one model, an all-out exchange of 4,000 nuclear weapons would release 150 teragrams of smoke, leading to a 8°C fall in global temperature.[47] Due to the fall in temperature and the loss of sunlight and growing food would be virtually impossible for 4-5 years,

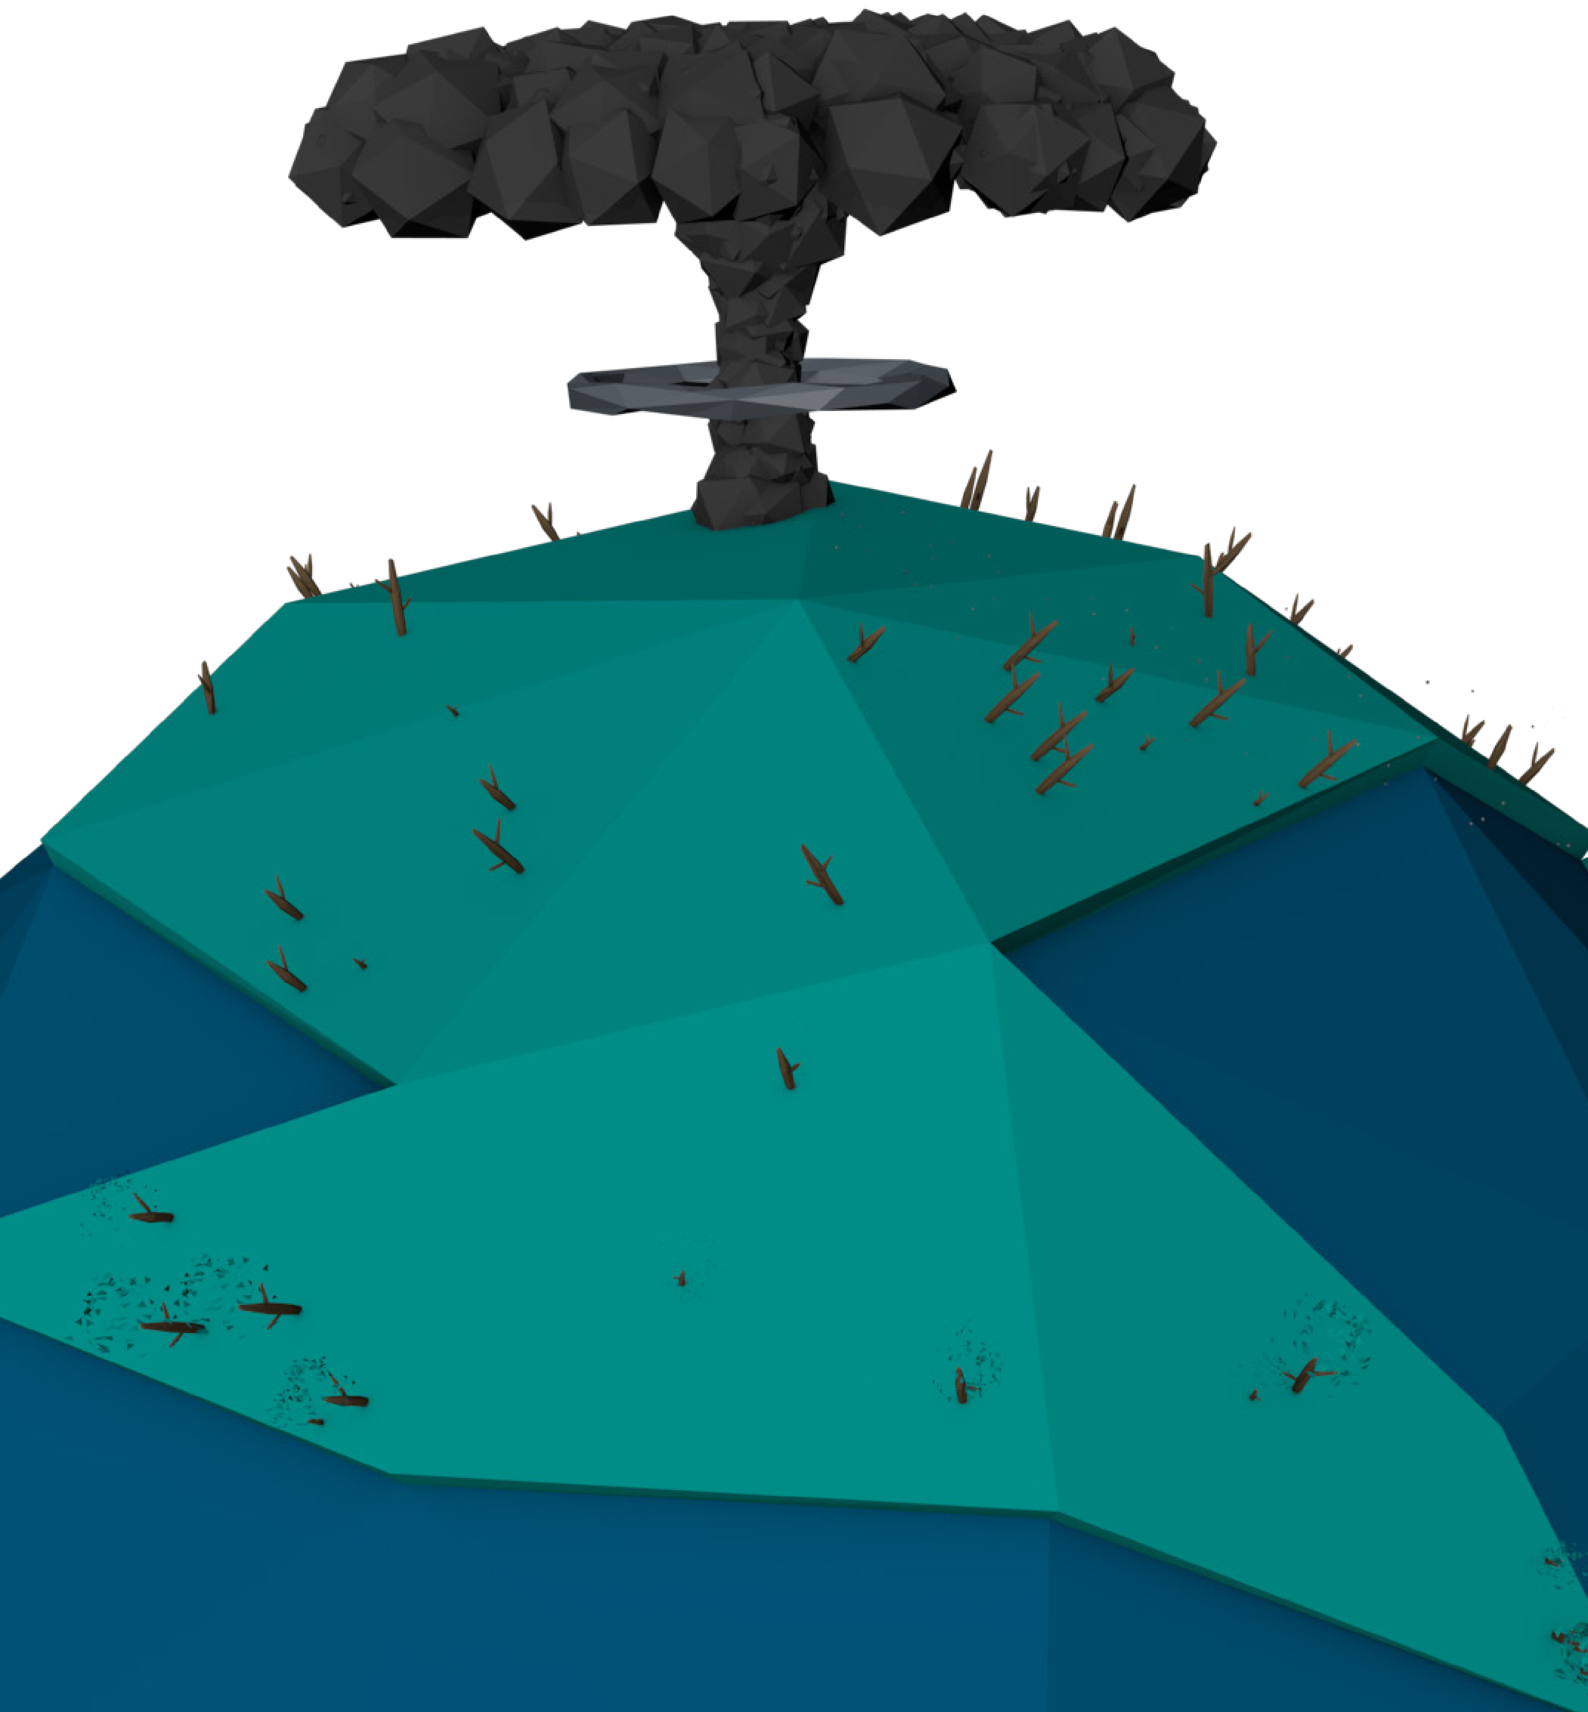

creating an unprecedented famine. According to some models, even a smaller scale regional war between India and Pakistan involving fifty 15 kiloton weapons would cause global temperatures to fall by around 1.25°C in the first year. Some have suggested that this could disrupt agriculture so significantly that one billion people would be at risk of starvation, though this has been criticised as overly pessimistic.[48]

LIKELIHOOD OF THE CATASTROPHE

It is very difficult to estimate the probability of a nuclear war in the next century. However, the chance may be too high to ignore. Over the course of the nuclear age, we have come close to inadvertent nuclear war on numerous occasions. For example, in 1995 Russian systems mistook a Norwegian weather rocket for a potential nuclear attack. Russian President Boris Yeltsin retrieved launch codes and had the nuclear suitcase open in front of him. Thankfully, Russian leaders decided the incident was a false alarm.[49] In a 2013 paper Barrett et al estimated that the 90% confidence interval of the annual probability of accidental nuclear war between the US and Russia is from 0.001% to 7%.[50] This covers a significant range, but a substantial portion of the possibilities are uncomfortably high.

We have also come close to intentional nuclear war on many occasions. President Kennedy said that the chance of nuclear war during the Cuban Missile Crisis was "somewhere between one out of three and even".[51] Although tensions have decreased since the end of the Cold War, in the wake of the crisis in Ukraine the chance of confrontation has risen,[52] and the geopolitical situation could become more unstable over the next few decades.

According to many experts, the most likely intentional nuclear war is between India and Pakistan. India and Pakistan have had numerous wars in the past and there have been various terrorist attacks against India by Pakistani groups.[53] Pakistan has pledged to meet any Indian attack on its territory with a retaliatory nuclear strike.[54]

MAIN ACTIONS AVAILABLE TO LIMIT THE RISK

There is some disagreement about how to reduce the risk of nuclear war. One policy option favoured by many is continued reduction of global nuclear arsenals.[55] There has been a great deal of progress on this front: in the 1960s 23 countries had weapons or were pursuing programmes, whereas today only nine countries have weapons.[56] However, reduction of nuclear arsenals does not guarantee safety.[57]

Weapons systems could also be altered to reduce the risk of accidental and intentional war. For instance, decision times could be increased by each side locating their weapons further from each other's borders.[58] More broadly, improved international relations would help to decrease

FIGURE 2.2.1. NUCLEAR ARSENALS OF THE US, RUSSIA, AND THE WORLD FROM 1945 UNTIL TODAY [59]

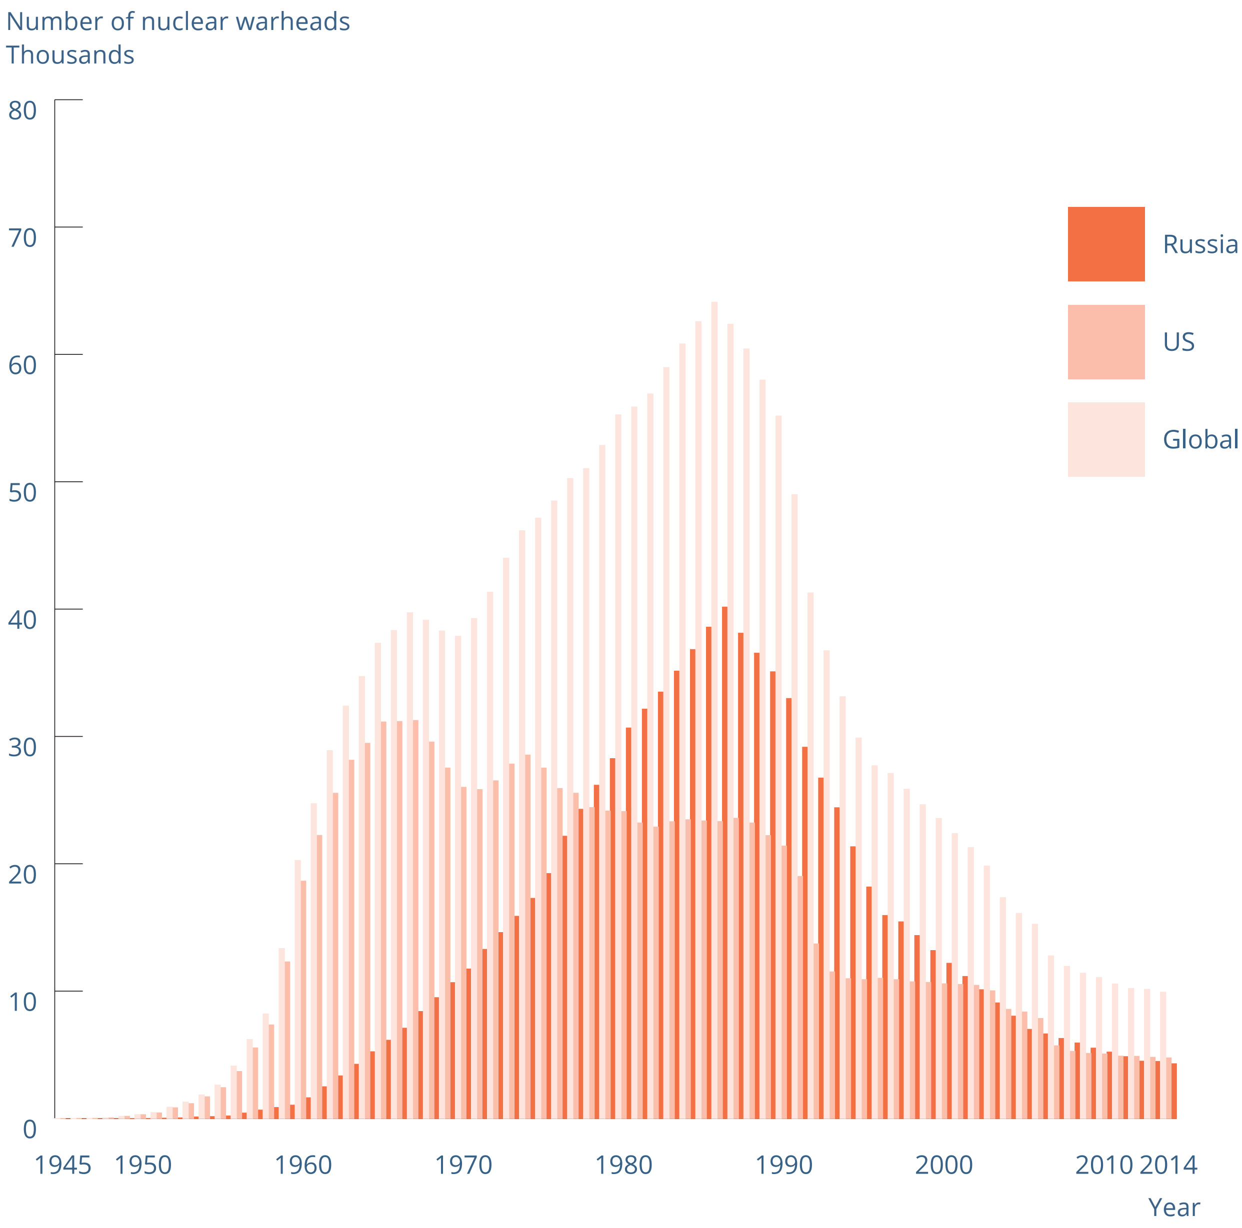

NUCLEAR NEAR MISSES

- In 1973, when Israel encircled the Egyptian Third Army, the Soviets threatened to intervene, leading to implied nuclear threats.
- In September 1983, a Soviet early warning satellite showed that the United States had launched five land-based missiles at the Soviet Union. The Soviet officer on duty, Stanislav Petrov, had only minutes to decide whether or not the satellite data were a false alarm. Since the satellite was found to be operating properly, following procedures would have led him to report an incoming attack. Going partly on gut instinct and believing the United States was unlikely to fire only five missiles, he told his commanders that it was a false alarm before he knew that to be true. Later investigations revealed that reflection of the sun on the tops of clouds had fooled the satellite into thinking it was detecting missile launches.
- The Able Archer incident of November 1983 was, in the words of former US Secretary of Defense Robert Gates, "one of the potentially most dangerous episodes of the Cold War". With talk of fighting and winning a nuclear war emanating from Washington, the Soviets reasoned that the West would mask preparations for a nuclear attack as a military exercise. The Able Archer exercise simulated the coordinated release of all NATO nuclear weapons. In response, the Soviets readied their nuclear forces and placed air units in East Germany and Poland on alert.

There have been numerous other examples of nuclear near misses.

the likelihood of international nuclear conflict. Research into nuclear winter would improve our knowledge and allow nuclear states to change their strategy accordingly. Finally, state and non-state actors could work to ensure that terrorists cannot acquire nuclear weapons by theft or on illegal markets. Global enforceable standards for nuclear weapon security would help achieve this goal.[60]

FIGURE 2.2.2. GLOBAL TEMPERATURE ANOMALY FROM NUCLEAR WINTER

Global average surface air temperature change from a release of 5 Teragram, 50 Teragram, and 150 Teragram of particular matter in the context of the climate change of the past 125 years.[61]

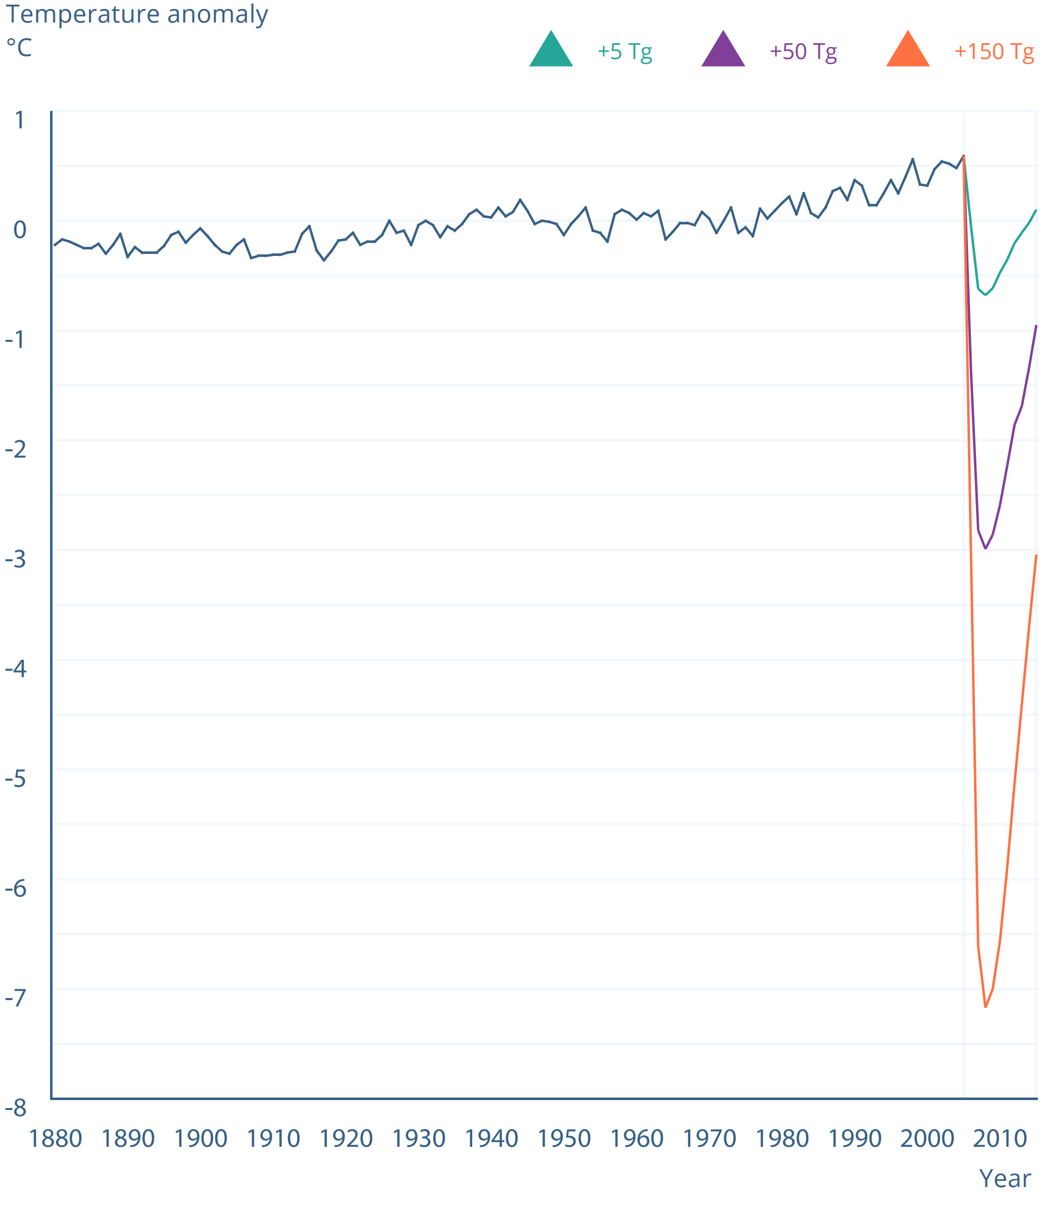

2.3. Natural pandemics

Until the discovery of nuclear weapons, pandemics were the most important global catastrophic risks, and it is plausible they still are. Plague, HIV, smallpox and other diseases have killed millions upon millions of people. There remains a serious chance that a pandemic could kill a huge portion of the world population, with pandemic influenza the most serious threat. A variety of measures, including improving international coordination and data sharing, and increasing the production of drugs and vaccines, would help to reduce this risk.

POTENTIAL IMPACT OF THE CATASTROPHE

The Black Death plague in the 14th century killed between 11% and 17% of the global population over the course of a decade.[62] In the most devastating pandemic of the last 100 years, the 1918 Spanish influenza outbreak, 50-100 million people lost their lives – between 2.5 and 5% of the global population.[63] Both of these pandemics occurred before the age of modern medicine, so it is unlikely these diseases would have similar impact today, although our more connected society may increase the spread of pandemics.

Currently, H5N1 avian influenza is thought to be the greatest pandemic threat.[64] To develop into a pandemic H5N1 would have to be easily transmissible between humans, which is not currently the case, though occasionally humans become infected through animal vectors. Influenza pandemics have previously infected about 24-38% of the world population.[65] The case fatality of a novel strain of H5N1 is unpredictable, but estimates for the H5N1 case fatality rate until today vary particularly widely – from 1% to 60%.[66] A very rough estimate of the death toll of a H5N1 pandemic can be found by multiplying the usual pandemic attack rate (24% to 38%) by the global population (roughly 7 billion) times the case-fatality ratio (1% to 60%). This produces an estimated death toll of between 16.8 million and 1.7 billion fatalities. The developing world would probably bear an overwhelming part of this burden.[67]

LIKELIHOOD OF THE CATASTROPHE

Influenza pandemics occur relatively frequently: there have been ten in the last 300 years, though none of these have killed more than 5% of the world population.[68] It is difficult to know where and when the next natural pandemic will occur. Globalisation and increased interaction between humans and animals increase the risk of pandemics, but improved health and sanitation would lessen their impact. According to the UK National Risk Register there is between a 1 in 20 and a 1 in 2 chance of a pandemic killing up

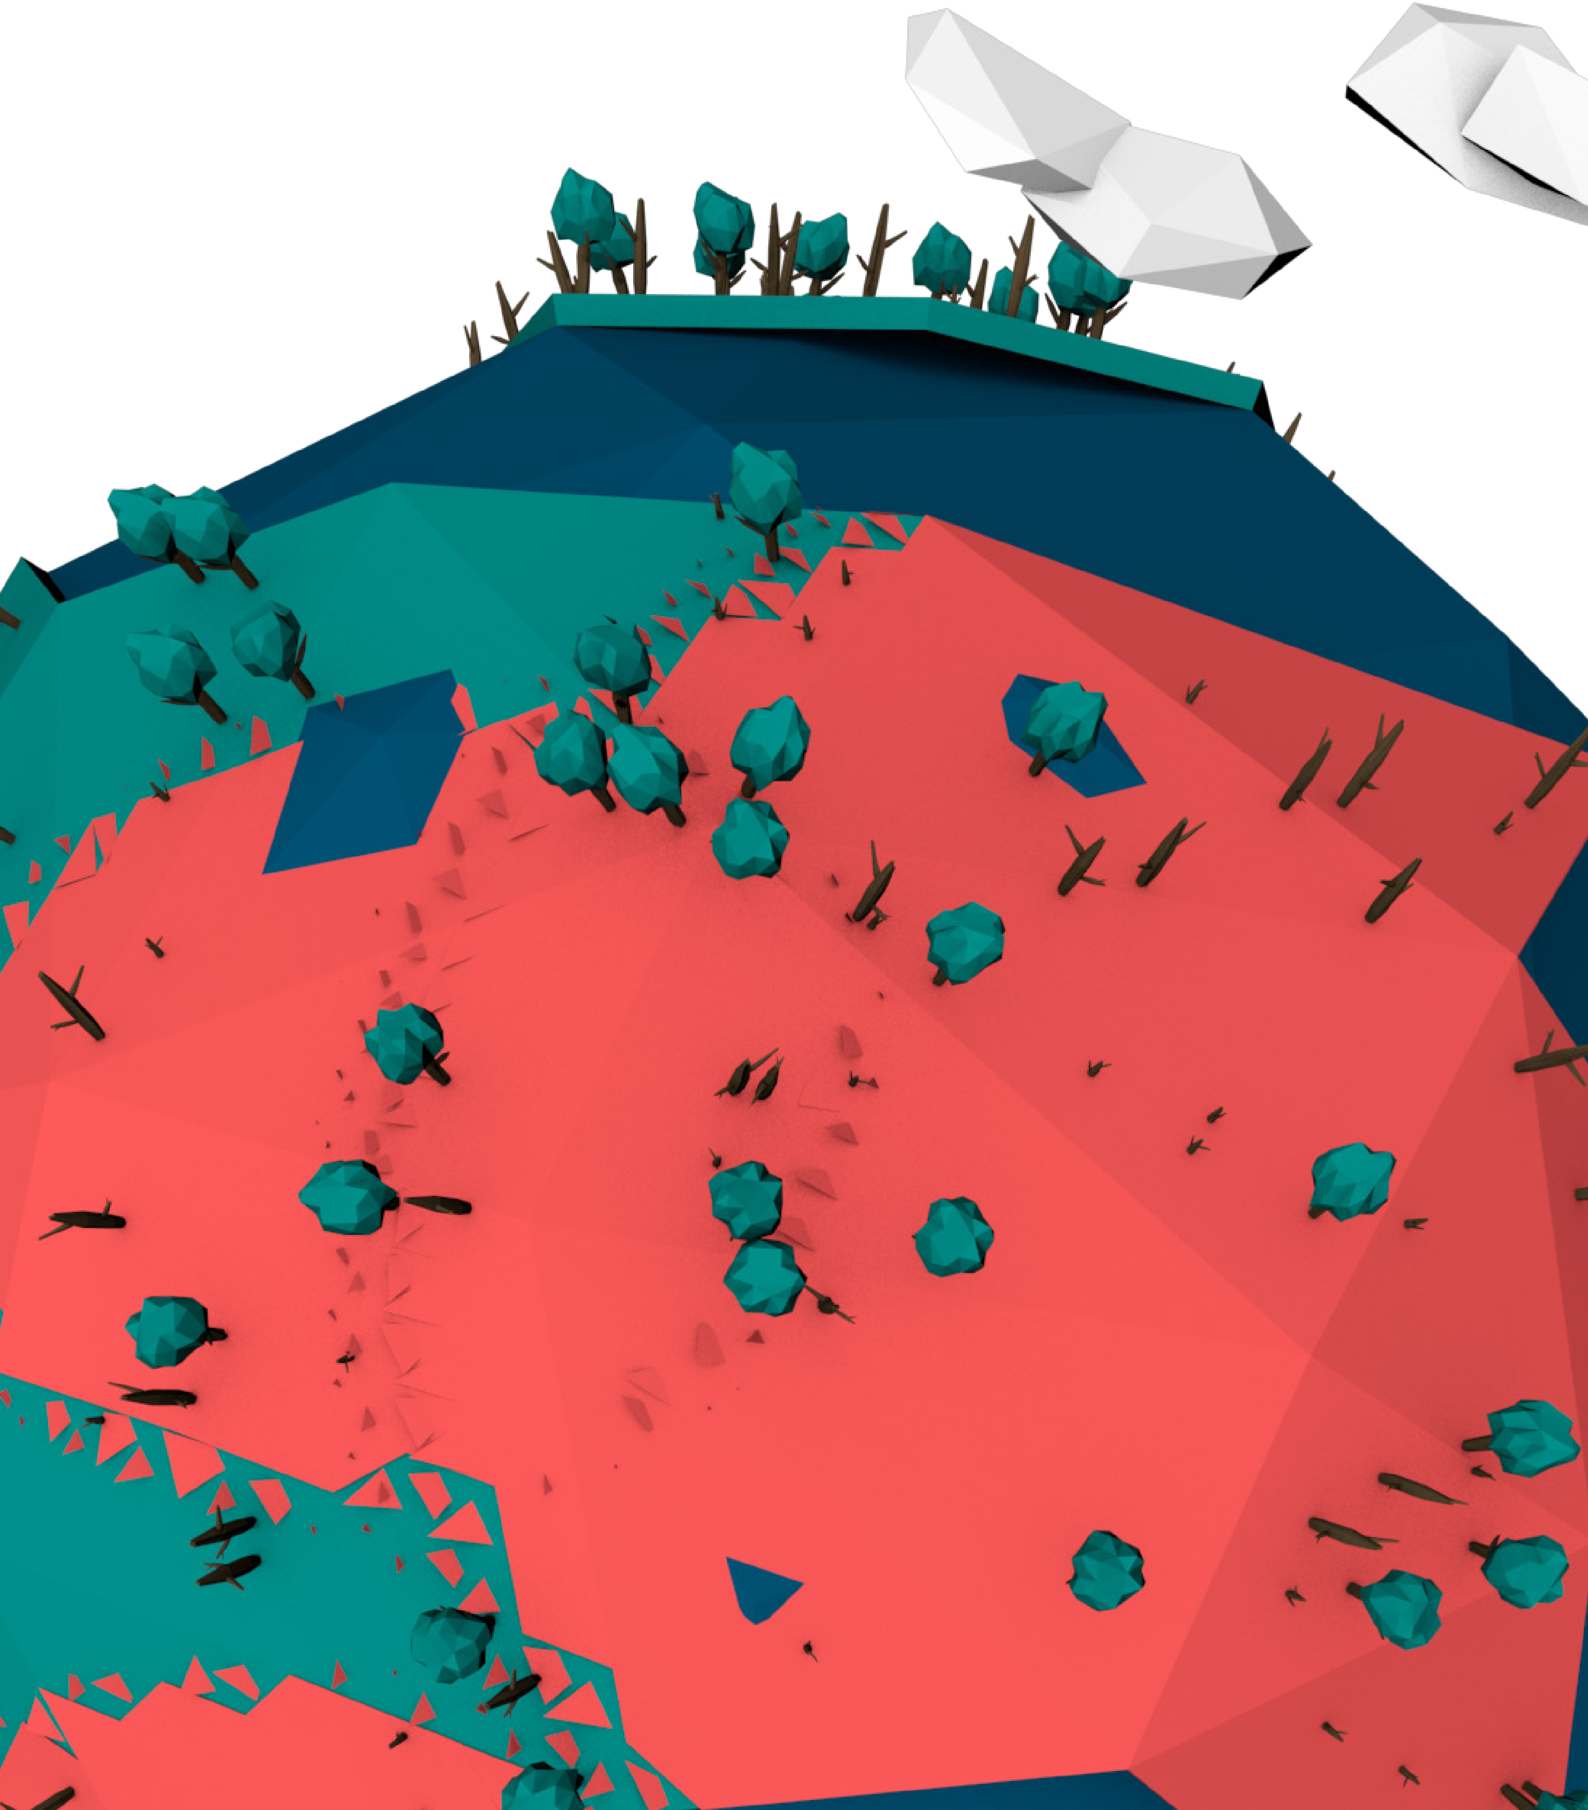

to 1% of the UK population in the next five years. Such a pandemic would go on to have global effects. Global catastrophic risk-level pandemics (killing more than 10% of the global population) are, however, further down the tail of the probability distribution. Results from expert surveys have put a 15% chance of an H5N1 pandemic over a three year period.[69]

MAIN ACTIONS AVAILABLE TO LIMIT THE RISK

There are numerous ways to limit the risk from pandemics.[70] Firstly, pharmaceutical companies could focus attention on developing vaccines and drugs for very damaging pandemics, and the international community can remove barriers to development with measures such as the establishment of a global vaccine development fund.[71] Secondly, governments and international bodies such as the WHO can stockpile drugs and vaccines to counter disease outbreaks. However, stockpiles may only be effective for some but not all pandemic pathogens.[72] Thirdly, since developing countries will probably face the highest burden from future pandemics, continued improvements in developing world health systems in accordance with the WHO's International Health Regulations, and improvements in the global distribution of drugs and vaccines, would limit the risk from pandemics.[73] Fourthly, disease surveillance and response systems could be improved so that novel threats in both animal and human populations are detected and responded to quickly.[74] Rapid dissemination of relevant data between countries is also essential for effective outbreak response, because of the exponential spread of infection in an outbreak.[75]

FIGURE 2.3.1. HISTORIC PLAGUES AND PANDEMICS [76]

PANDEMIC	DATE(S)	LETHAL IMPACT
Plague of Justinian	AD 541-542	25-33m (13-17% of the world population)
Black Death	14th Century	50-75m (11-17% of world population)
Smallpox	1520-1527	200,000 deaths within the Aztec population (75% of population in some areas)
Spanish influenza	1918 - 1919	50 - 100m (2.5-5% of population)
Smallpox	20th century	300m over the course of the 20th Century
HIV/AIDS	1981 - present	34m

2.4. Exogenous risks

Exogenous risks are those which arise independent of human activity (whereas even natural epidemics are spread by humans). They include such possibilities as super-volcanic eruptions and large asteroid and comet impacts, which are believed by some to have caused mass extinctions.[77] The likelihood of exogenous risks is better understood than that of many other global catastrophic risks because the underlying dynamics have been unchanged for a very long time. The historical evidence suggests that exogenous global catastrophic risks cannot be too frequent, and may therefore be much less likely than some of the anthropogenic risks.[78]

A. SUPER-VOLCANOES

Super-volcanoes are volcanoes capable of producing at least 1012 m3 bulk volume of fragmental material.[79] Some experts believe that the eruption of the Toba super-volcano in Indonesia around 70,000 years ago brought humanity to the brink of extinction, though there is significant disagreement about this.

POTENTIAL IMPACT OF THE RISK

The Toba eruption ejected large amounts of dust and sulphates into the Earth's atmosphere, which caused global cooling of 3-5°C for several years and led to enormous loss of plant and animal life.[80] Some have argued that Toba caused the greatest mass extinction in history and reduced the human ancestor population from around 100,000 to around 4,000 people for approximately 20,000 years, though this is controversial.[81] According to a report by the Geological Society of London:

"A layer of ash estimated at 15 centimetres thick fell over the entire Indian sub-continent, with similar amounts over much of SE Asia. Just one centimetre of ash is enough to devastate agricultural activity, at least when it falls in the growing season. An eruption of this size would have catastrophic consequences. Many millions of lives throughout most of Asia would be threatened if Toba erupted today."[82]

LIKELIHOOD OF THE CATASTROPHE

In order to assess the likelihood of super-volcano eruptions, we have to rely on a relatively limited set of observations of past super-volcanic eruptions, which makes any estimates very uncertain. Existing data suggest that there will be a super-volcanic eruption roughly every 30,000-50,000 years on average.[83]

MAIN ACTIONS AVAILABLE TO LIMIT THE RISK

At present, humanity lacks the technical capacity to prevent volcanic eruptions. Consequently, improving resilience to catastrophe is the main way to limit the risk from super-volcanoes. We discuss this in more detail

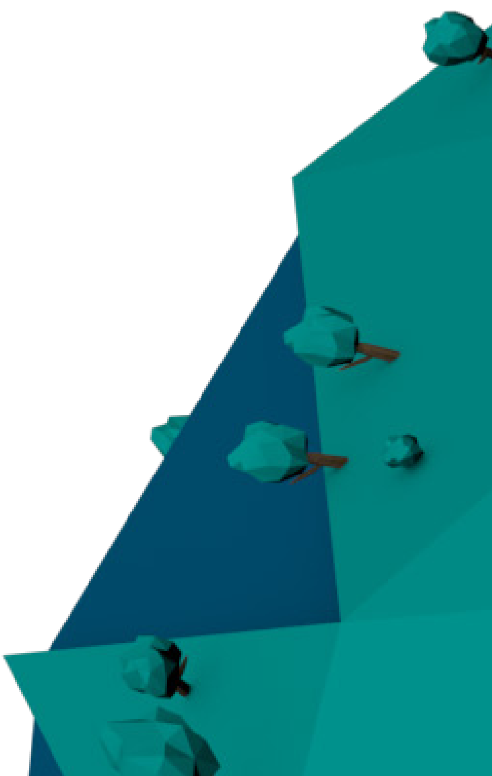

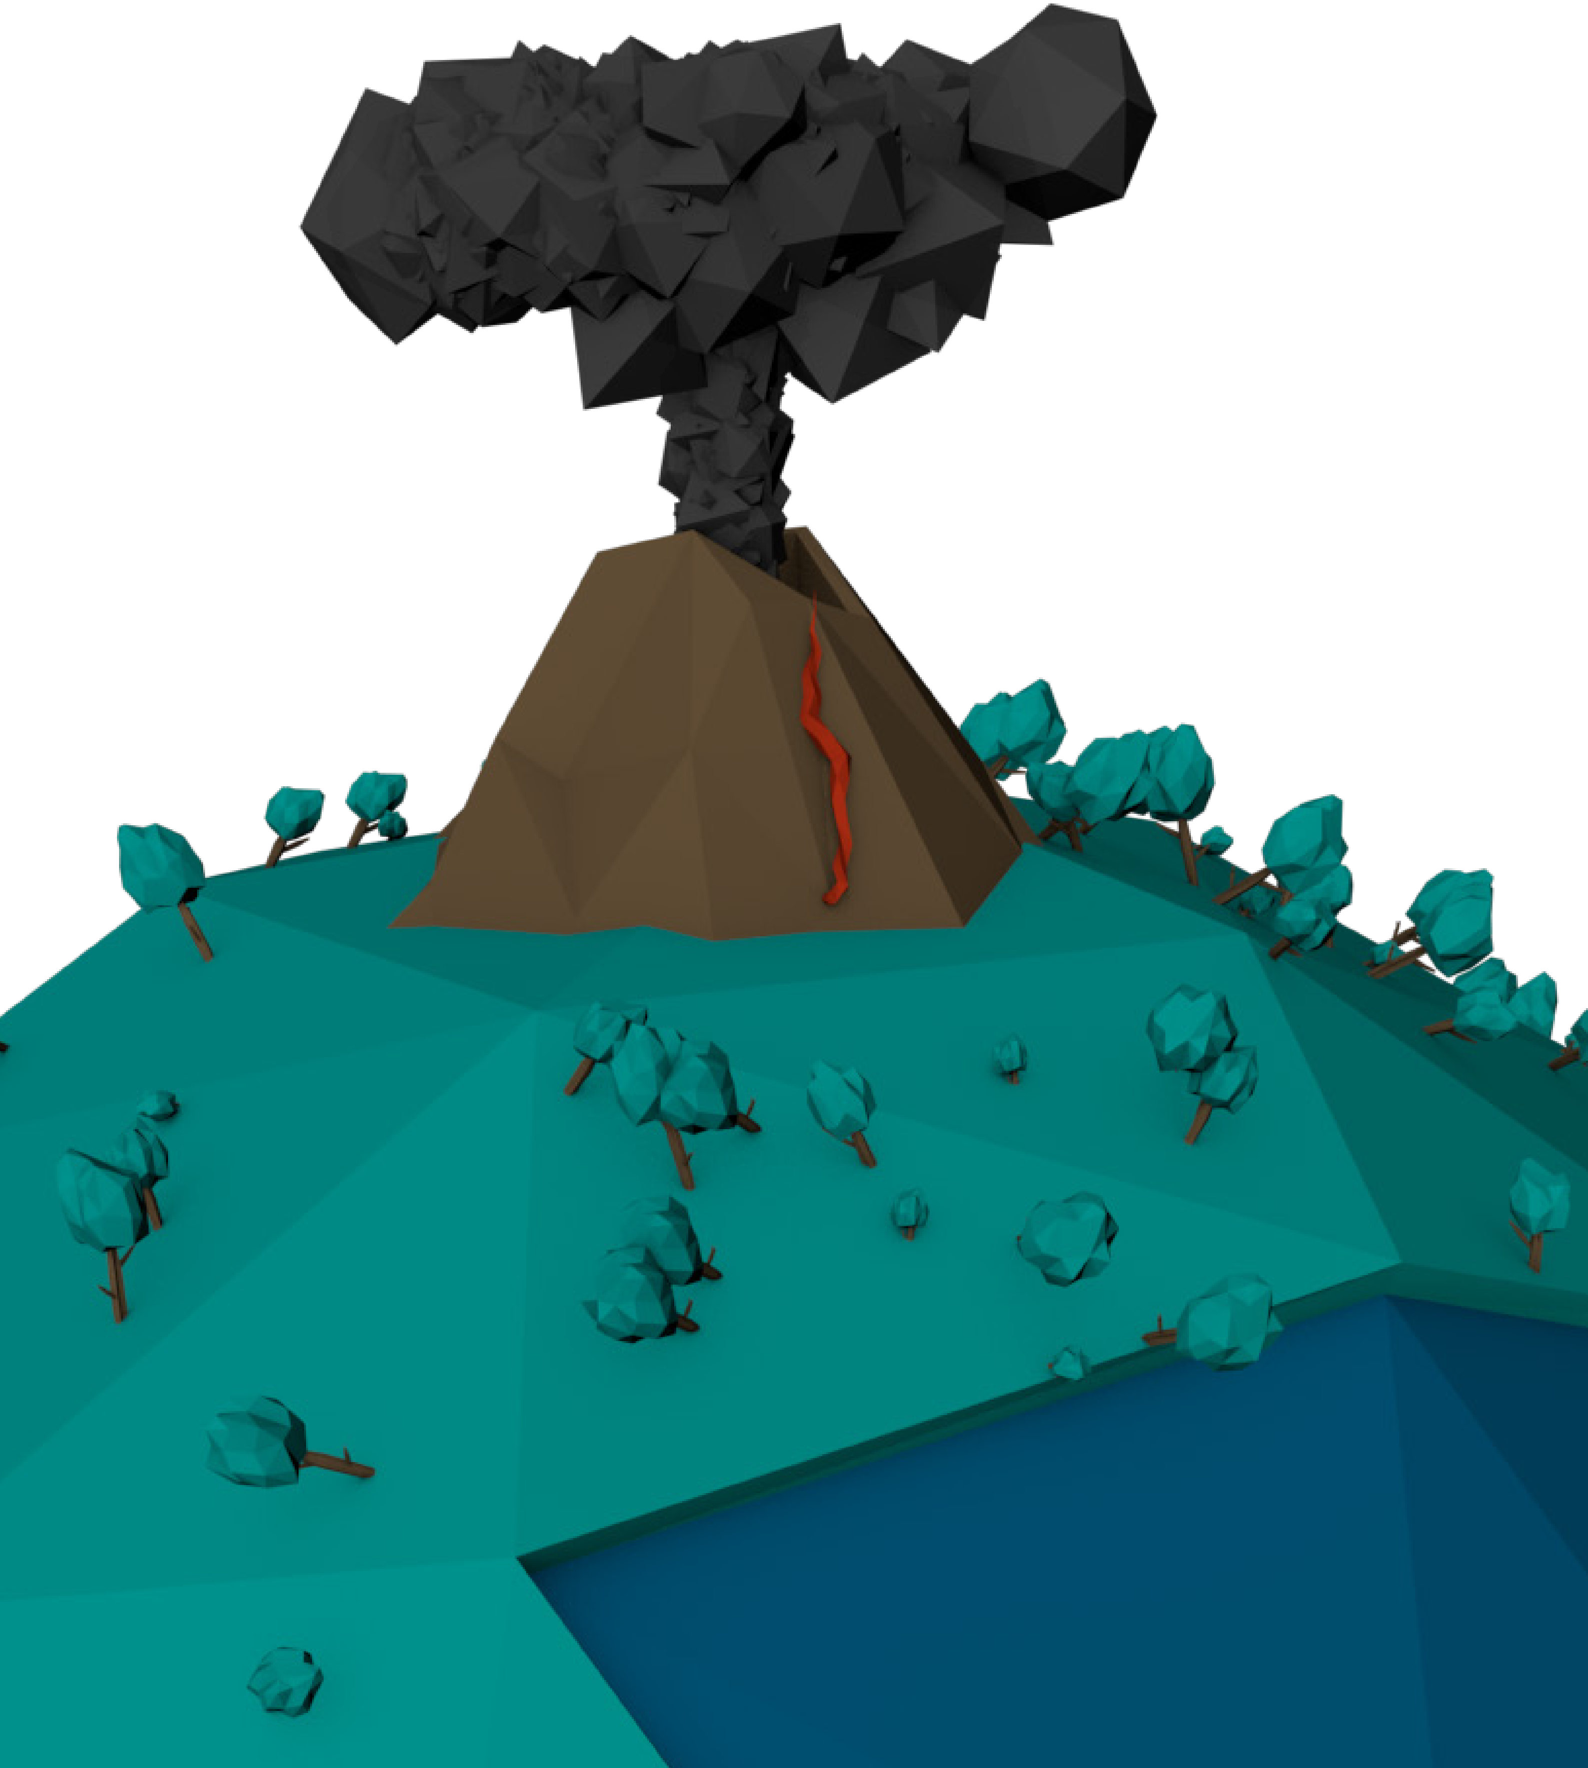

in Chapter 3. Increased investment and research might also improve our ability to predict volcanic eruptions, which would in turn improve preparedness.[84]

B. ASTEROIDS AND COMETS

Around 66 million years ago, an asteroid of around 10km in diameter struck Chicxulub in Mexico. This impact probably caused one of the three largest mass extinctions in history and may have abruptly ended the age of the dinosaurs.[85] Today, an impact by a Near Earth Object (NEO) – an asteroid or comet – larger than 1.5km in diameter would kill millions largely by causing global cooling and agricultural disruption. The likelihood of an NEO impact is reasonably well-understood and important steps have been taken to monitor the risk they pose.

POTENTIAL IMPACT OF THE RISK

Contemporary studies of asteroids and comets conservatively assume that all objects greater than 1.5km in diameter would be capable of causing damage on a global scale, via firestorms generated by impact debris and a so-called asteroid winter caused by dust and sulphates being released into the atmosphere.[86] The damage that would be caused by a multi-kilometre asteroid or comet impact is only modestly understood at present.[87] An asteroid winter would undermine agriculture at least for an entire growing season and so could cause the deaths of billions of people.[88] According to a comprehensive report by the US National Research Council, "above the conservatively assumed global catastrophe threshold from a 1.5km-diameter impactor, the number of fatalities ramps up from 10% of the world's population to the entire population for impactors above 10km in diameter".[89]

LIKELIHOOD OF THE CATASTROPHE

NEOs are a comparatively well-understood global catastrophic risk. NASA's Spaceguard system has mapped more than 90% of asteroids with a diameter of more than 1km.[90] Reinhardt et al argue that the total probability in a 100-year period of an asteroid encounter with the Earth that might cause a globally significant effect is approximately 1 in 1,250.[91]

One can calculate the expected cost of NEOs by multiplying the likely deaths per event by the frequency of events of a certain size. As of 2010, the annual expected (probability-weight-

1 in 1,250

approximate risk of an asteroid encounter with the Earth that might cause a globally significant effect in a 100-year period

FIGURE 2.4. RISK BEFORE AND AFTER THE SPACEGUARD SURVEY[92]

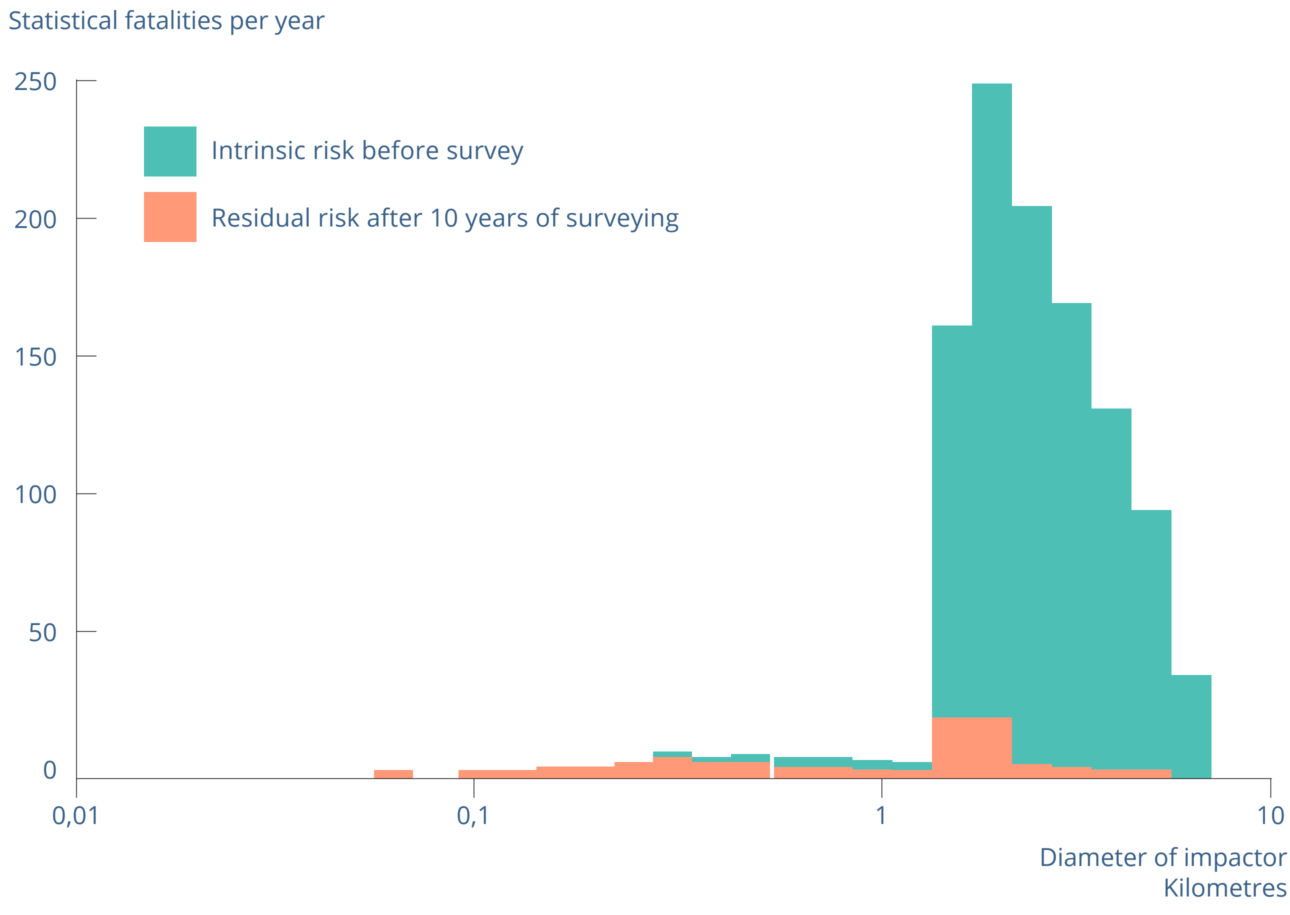

> "The likelihood of an NEO impact is reasonably well-understood and important steps have been taken to monitor the risk they pose."

ed) cost of all NEOs is currently 91 fatalities.[93]

Once ongoing asteroid surveys are completed, long-period comets – comets which take more than 200 years to orbit the Sun – may dominate the remaining unknown impact threat from NEOs.[94]

MAIN ACTIONS AVAILABLE TO LIMIT THE RISK

There are three main ways to limit the risk from asteroids. Firstly, improved detection of asteroids would enable humanity to understand the risks it faces and to take appropriate counter-measures. NASA's ongoing Spaceguard Survey has so far reduced the expected cost a large asteroid impact by more than an order of magnitude.[95] Secondly, technological research could help us to find ways to deflect threatening NEOs.[96] Thirdly, improvements in resilience could be made to ensure human survival in the event of a large asteroid or comet impact.

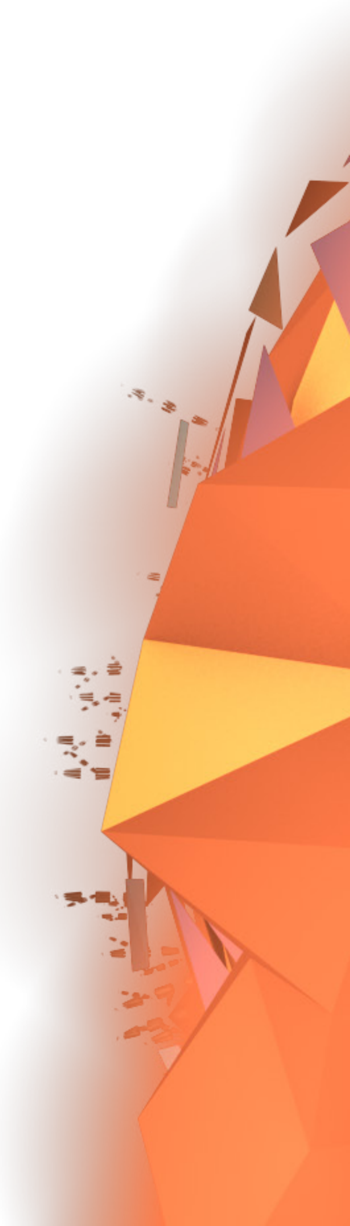

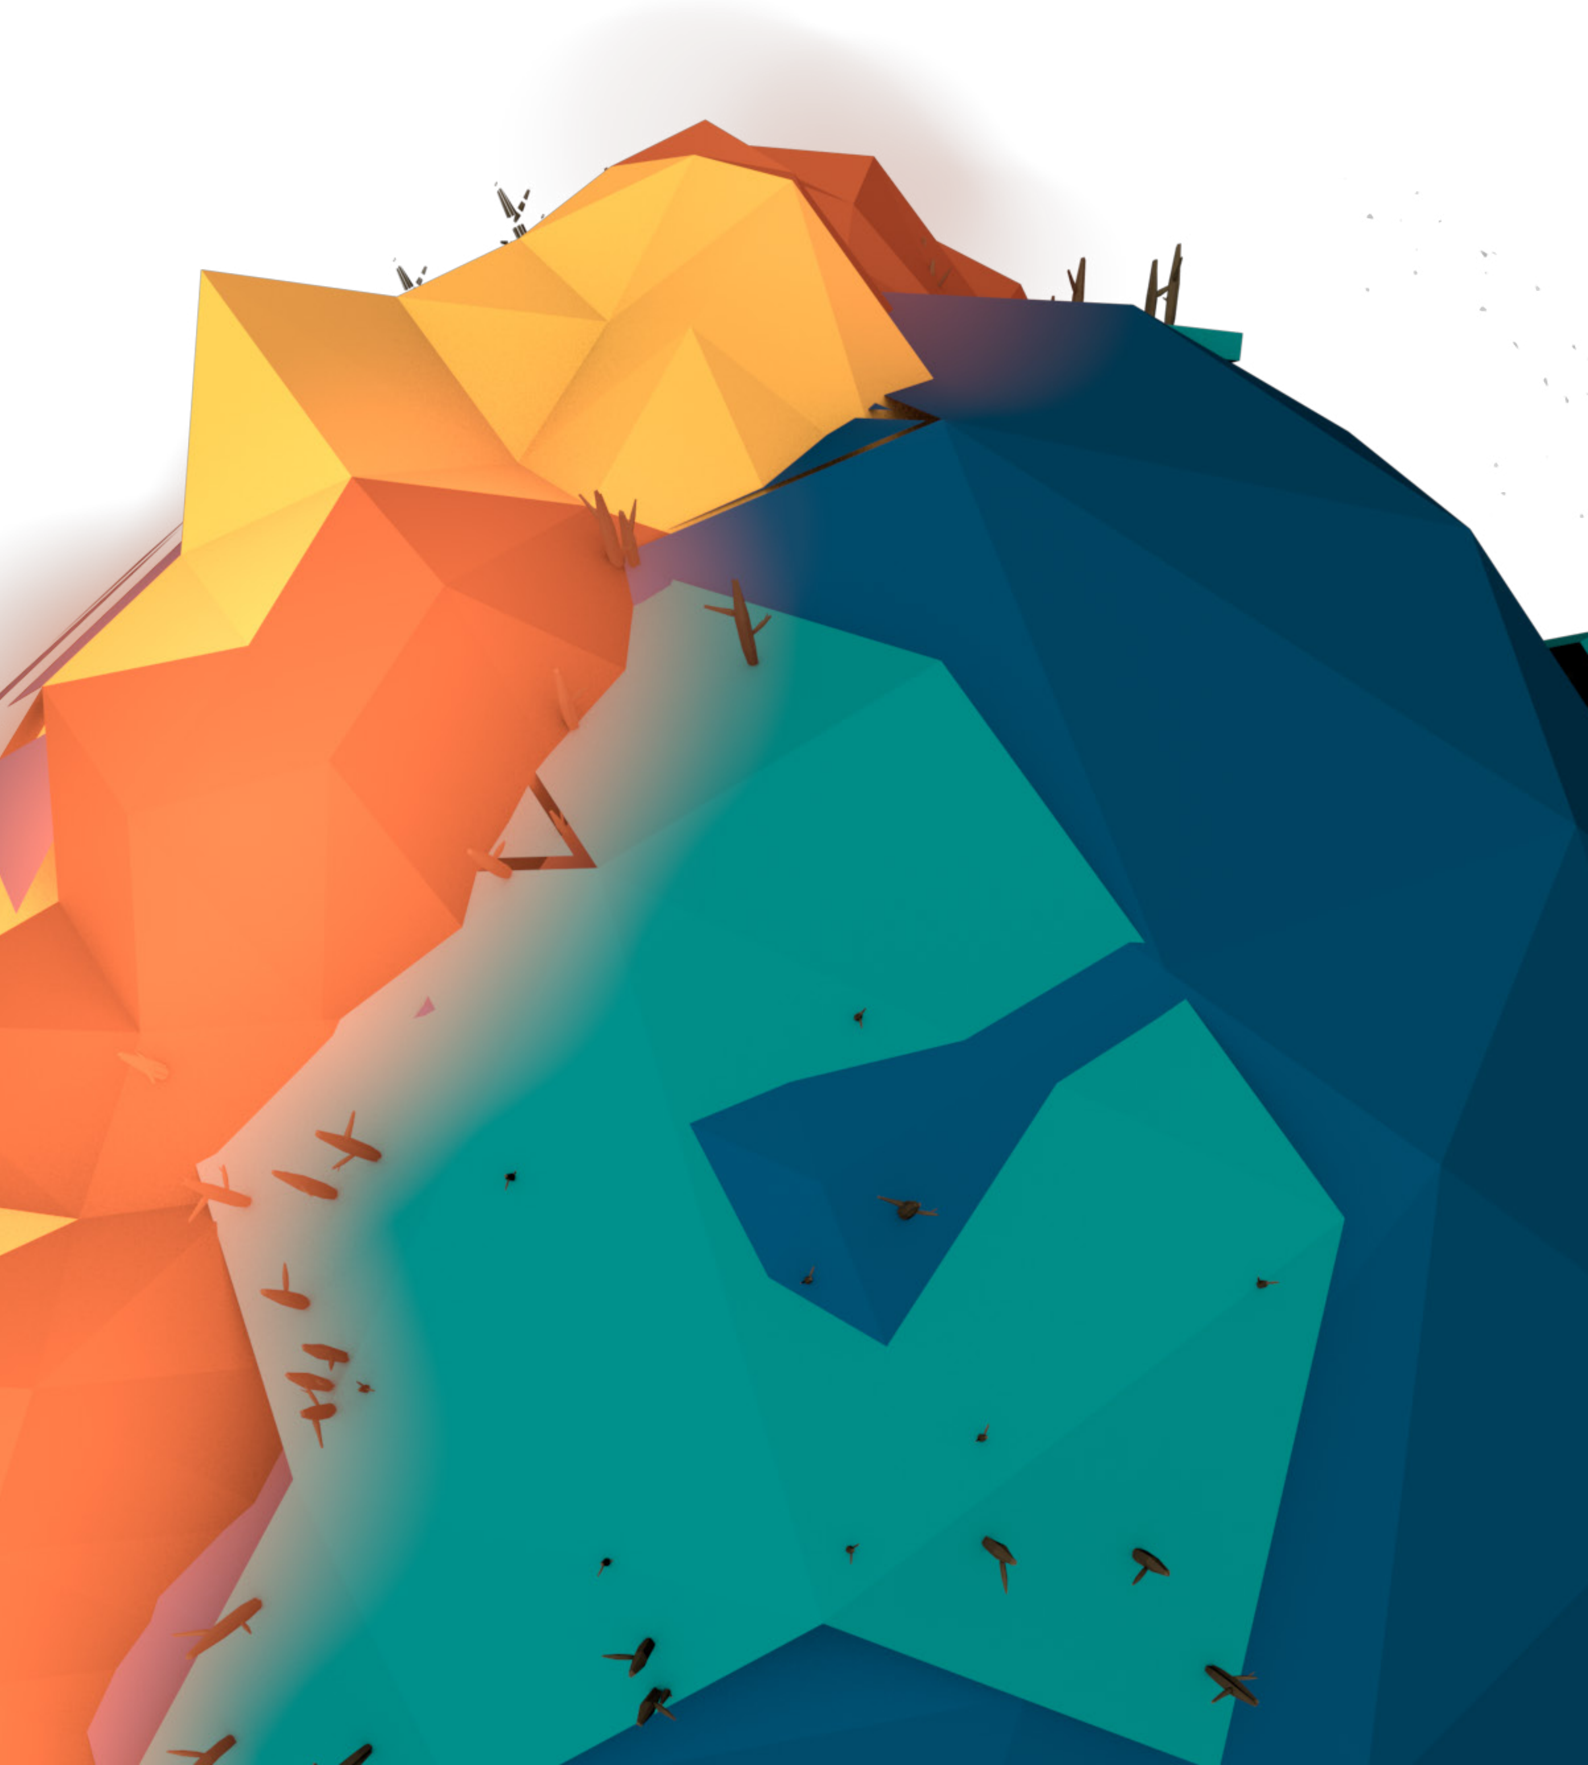

2.5. Emerging risks

In the coming decades, emerging technologies will provide major benefits to society, but they may also create significant and unprecedented risks. Certain types of biotechnology, if more widely accessible, could give terrorist groups the access to pathogens as dangerous as smallpox. Geoengineering technologies could give single countries the power to unilaterally alter the earth's climate. Finally, artificial general intelligence could, if developed, leave human control.

Technological risks could emerge very quickly and give certain groups large and perhaps unprecedented destructive power. Moreover, some emerging technologies may be particularly difficult to control because barriers to access the technology may be quite low.

A. ENGINEERED PANDEMIC

The past decades have seen rapid advances in biotechnology, in part due to the falling costs of gene sequencing and synthesis.[97] Improvements in ease-of-use of certain specific kinds of biotechnology bring increased concerns about bioterrorism. Gene synthesisers have the capacity to turn digital sequence data into physical genetic sequences, enabling individuals to create viruses from digital files (as was done with the 1918 Spanish Flu virus).[98] Should gene synthesis become increasingly streamlined, the tools enabling such risks may become widely accessible.[99] Previously benign digital information, such as the widely available online genetic data for smallpox, will become more hazardous.

POTENTIAL IMPACT OF THE RISK

Evolutionary pressure generally constrains the lethality of pathogens, as most highly lethal pathogens fail to spread far before killing their host. There is thus some evidence for an inverse relationship between a pathogen's lethality and transmissibility, thereby limiting the damage from a naturally occurring pandemic. Biotechnology has the potential to break this correlation, allowing organisms with extraordinarily high lethality and transmissibility. In 2001 Australian researchers accidentally created a highly lethal and vaccine resistant form of mousepox.[100] Similar techniques could potentially be applied to smallpox.[101] Two recent controversial papers have shown how to create a version of H5N1 which is potentially transmissible between humans.[102] Basic calculations based on H5N1 parameters suggest that a single release of these modified viruses could cause hundreds of millions of casualties.[103] Engineered pathogens with dangerous features could be released accidentally from a lab or intentionally by states or terrorist groups. The legal

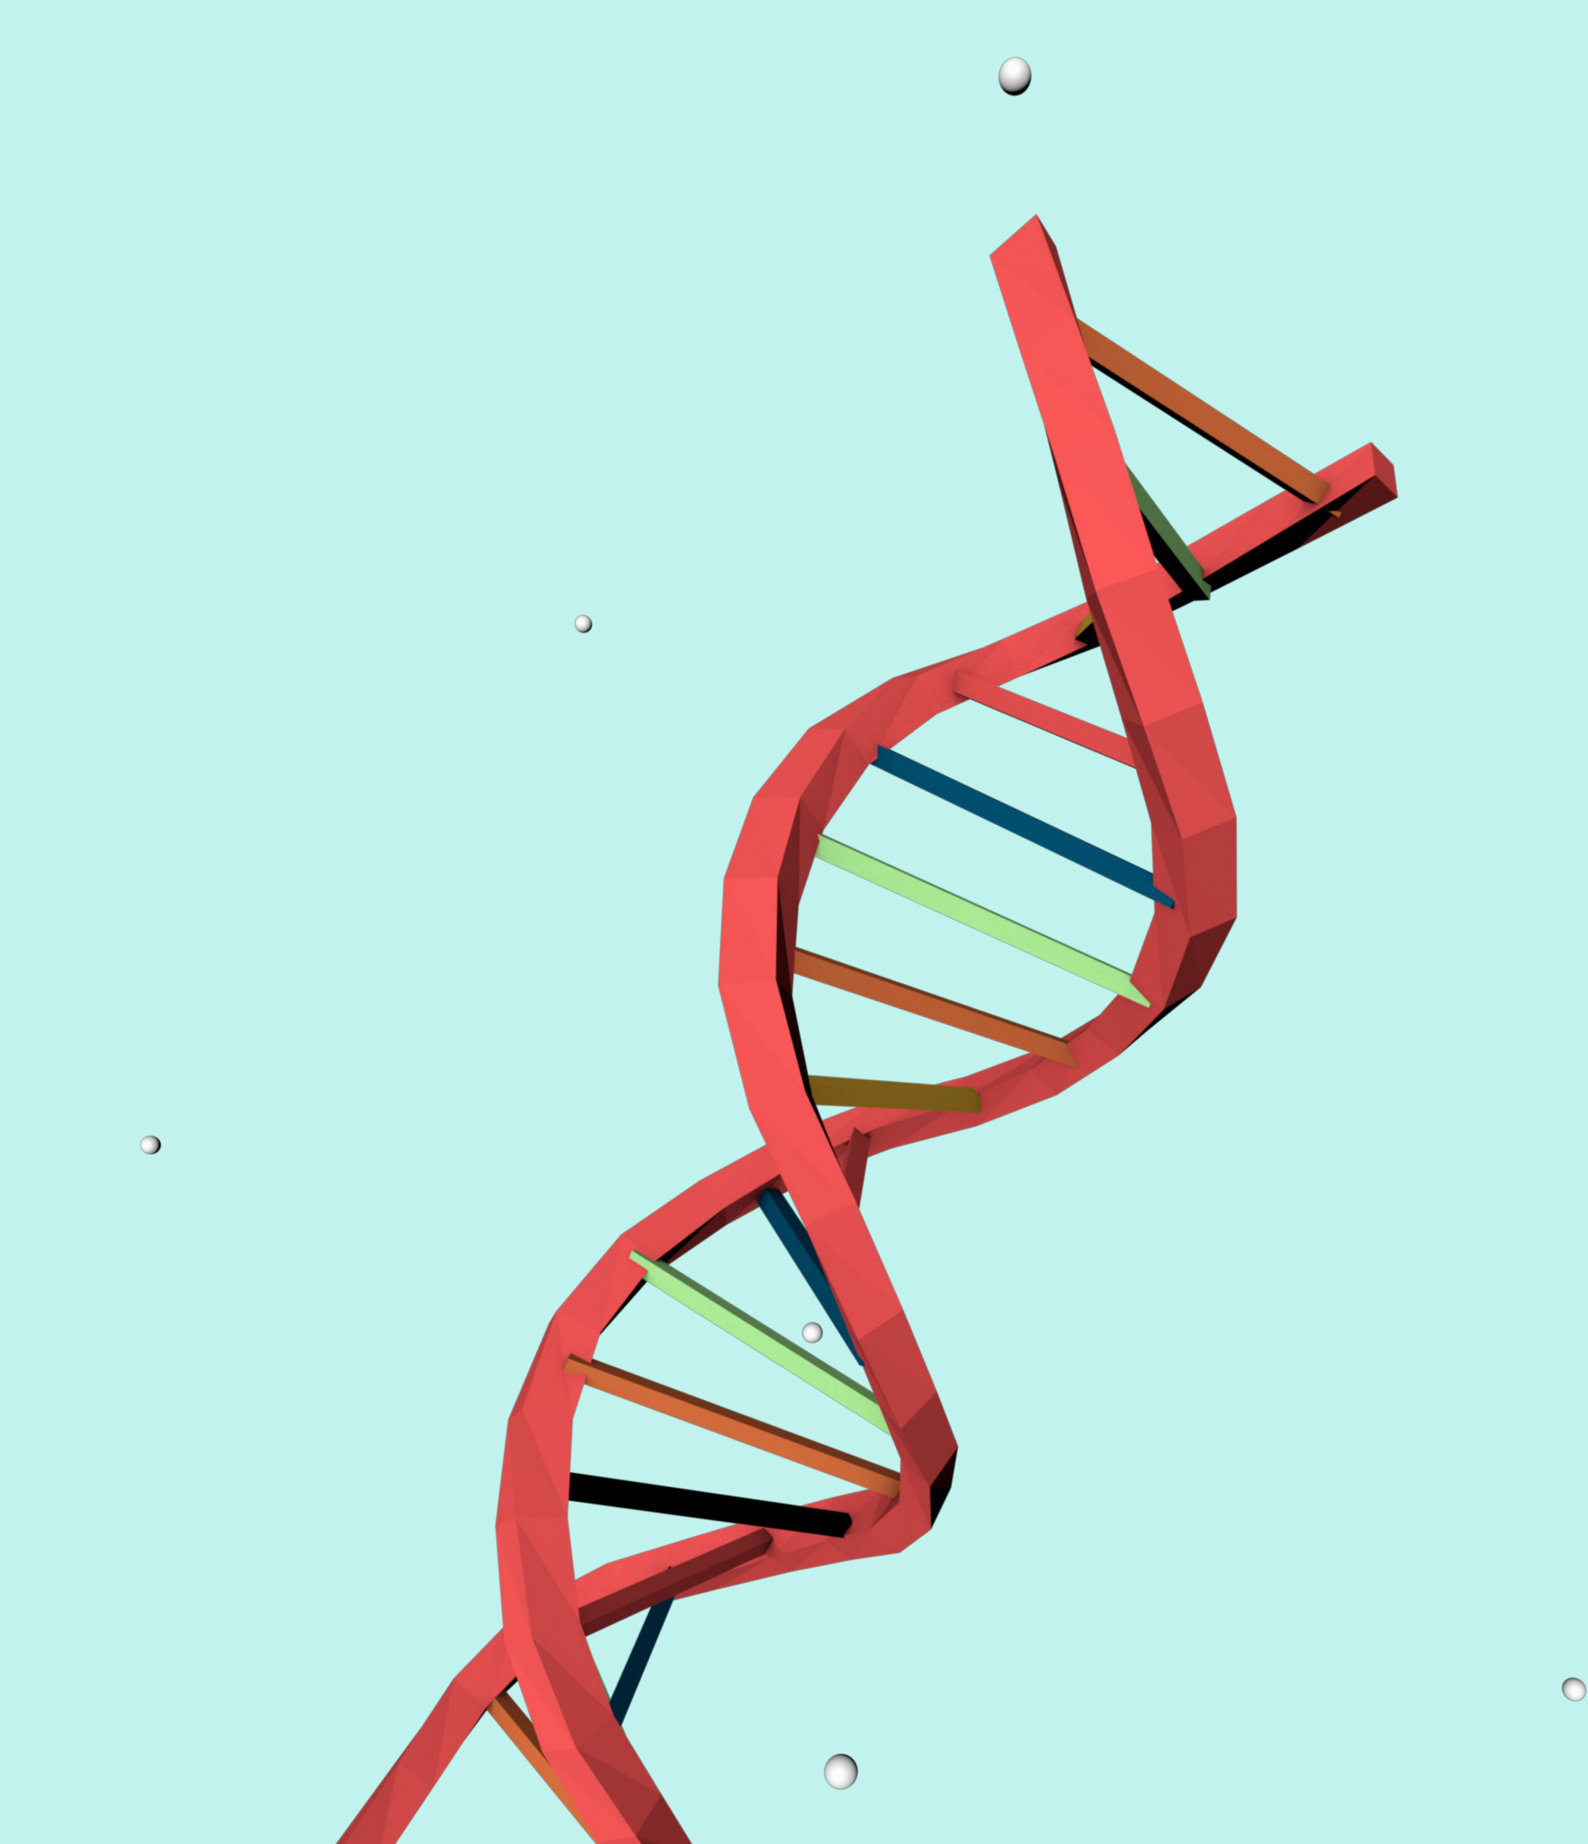

theorist Richard Posner has described a scenario in which terrorists synthesise a version of smallpox that is incurable, immune to vaccine and kills all its victims, which is then released in a large city via aerosol.[104] Although such scenarios are not currently feasible, eminent figures including George Church,[105] Nathan Myhrvold,[106] and Martin Rees[107] have all argued that biotechnology poses serious risks.

LIKELIHOOD OF THE CATASTROPHE

There is a real possibility that a dangerous engineered pathogen could be released by accident. The H1N1 influenza strain, responsible for significant morbidity and mortality around the world from 1977 to 2009, is thought to have originated from a laboratory accident.[108] As of 2012, there were at the very least 42 laboratories engaged in live research on potential pandemic pathogens, though this will have been affected by the US moratorium on 'gain of function' research in 2014.[109] Lipsitch and Inglesby estimate that "work with a novel, transmissible form of influenza virus carries a risk of between 0.01% and 0.1% per laboratory-year of creating a pandemic… or between 0.05% and 0.6% per full-time worker-year".[110] These estimates are illustrative of the probability that existing 'gain of function' influenza research would produce a pandemic, though there is significant uncertainty about them.[111]

As biotechnology develops, the level of expertise required to create dangerous pathogens will fall. In the longer term, if biotechnology matures sufficiently and gene synthesis is not well-regulated, states and small groups will find it increasingly easy to synthesise and alter dangerous pathogens.[112] This poses a serious risk of a global catastrophe.

MAIN ACTIONS AVAILABLE TO LIMIT THE RISK

Many of the measures which limit the risk from natural pandemics would also limit the risk from synthesised or modified pathogens. However, other measures would help specifically with biotechnology risks. Fostering a culture of safety in the relevant research and technical communities is probably very valuable. Researchers could be encouraged to take significant precautions with this research and to avoid disseminating research when doing so brings major risks. With regard to regulation, a licensing regime for DNA synthesis could be one first step,[113] while mandatory liability insurance for dual-use research - which helps medical progress but could be used maliciously - would ensure that researchers have incentives to maintain high levels of laboratory biosafety.[114] Governments could also require journal editors to consider whether publication of research could lead to adverse outcomes.[115]

FIGURE 2.5.1. COST PER GENOME[116]

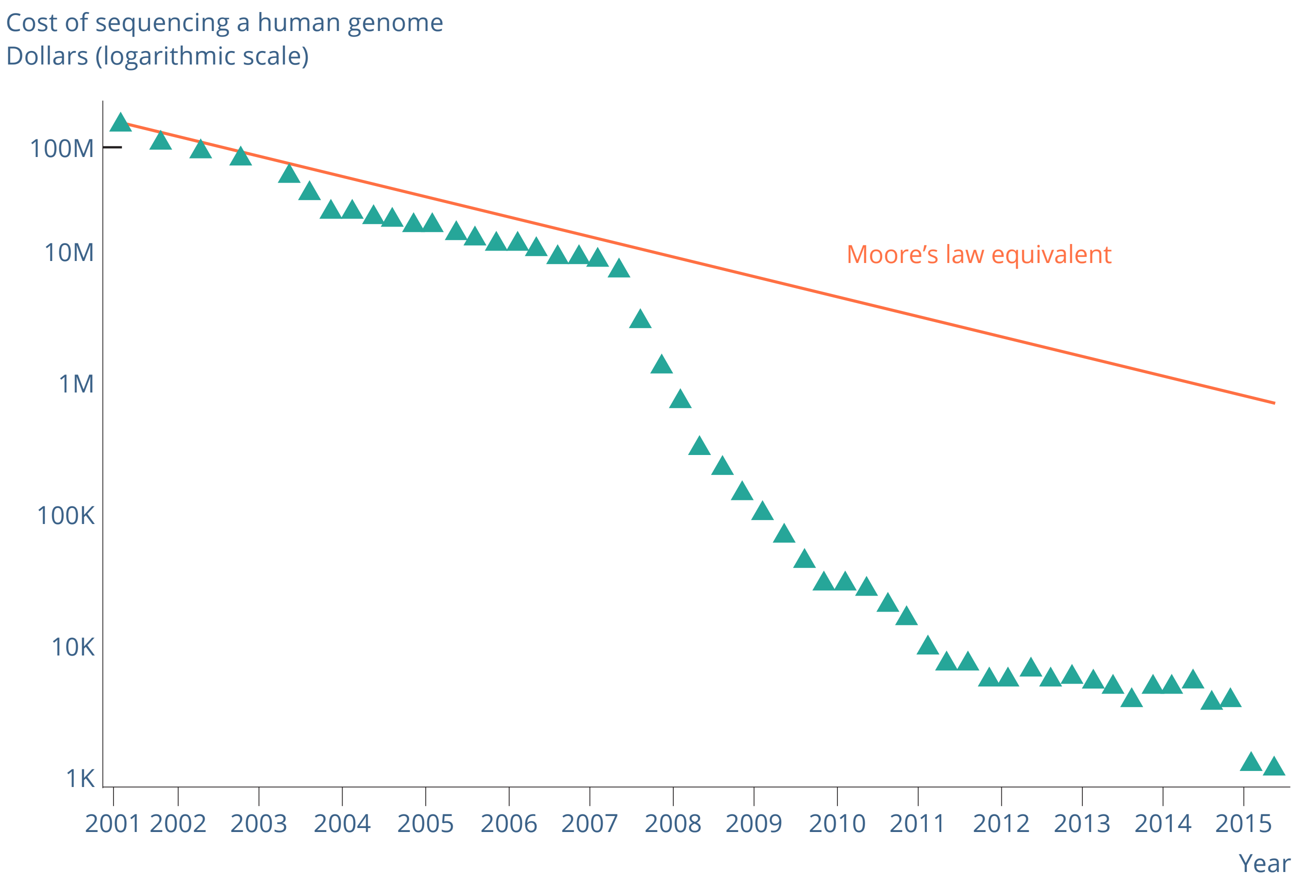

B. ARTIFICIAL INTELLIGENCE

Prominent figures uch as Stuart Russell, Professor of Computer Science at Berkeley, Peter Norvig, Director of Research at Google, and Nick Bostrom, author of Superintelligence, recently signed a letter warning of the risks posed by Artificially Intelligent (AI) systems.[117] Past experience demonstrates that AI systems can go from significantly subhuman to superhuman relatively quickly in narrow domains, such as recently in the game Go. In the coming decades, we may create AI systems which surpass humans in all relevant domains. If this were to happen, the effects would be uniquely transformative. However, AI also has huge potential benefits, as well as risks.[118] It could, for example, greatly reduce the cost of many goods, and allow us to solve other global problems.

POTENTIAL IMPACT OF THE RISK

Some risks of AI likely fall short of posing a global catastrophic risk. Widespread automation could cause significant economic and social disruption, which has only a relatively small chance of leading deaths at the scale of other global catastrophic risks.[119]

In the longer term, AI may enable important new capabilities, perhaps extremely quickly if it turns out that we can automate AI development.[120] If we have powerful and generalisable automated systems, the goals they are programmed with may exert significant influence over the future. It has been argued that if these goals were not aligned with human values, the consequences could be truly catastrophic.[121] It is difficult not only to specify human values in a robust, machine-interpretable way, but also to agree on human values in the first place. Moreover, if even if the values of very powerful artificially intelligent systems can be aligned with their creators', such systems might destabilise the geopolitical balance in a destructive way.

LIKELIHOOD OF THE CATASTROPHE

It is widely accepted that we will be able to create AI systems that are capable of performing most tasks as well as a human at some point ('human-level AI'). Experts disagree about when this will occur. According to the median surveyed expert, there is a roughly 50% chance of such AI by 2050.[122] The median surveyed expert believes that there is at least a 5% chance of superintelligent AI within two years after human-level AI, and a 50% chance within thirty years.[123] Assuming that human-level AI is developed, its long-term social impact is unclear. According to the median surveyed expert, there is around a 7% chance that it would be 'extremely bad'.[124] The extreme uncertainty surrounding these estimates should, however, be strongly emphasised.

MAIN ACTIONS AVAILABLE TO LIMIT THE RISK

Different challenges from AI systems will require different responses, and

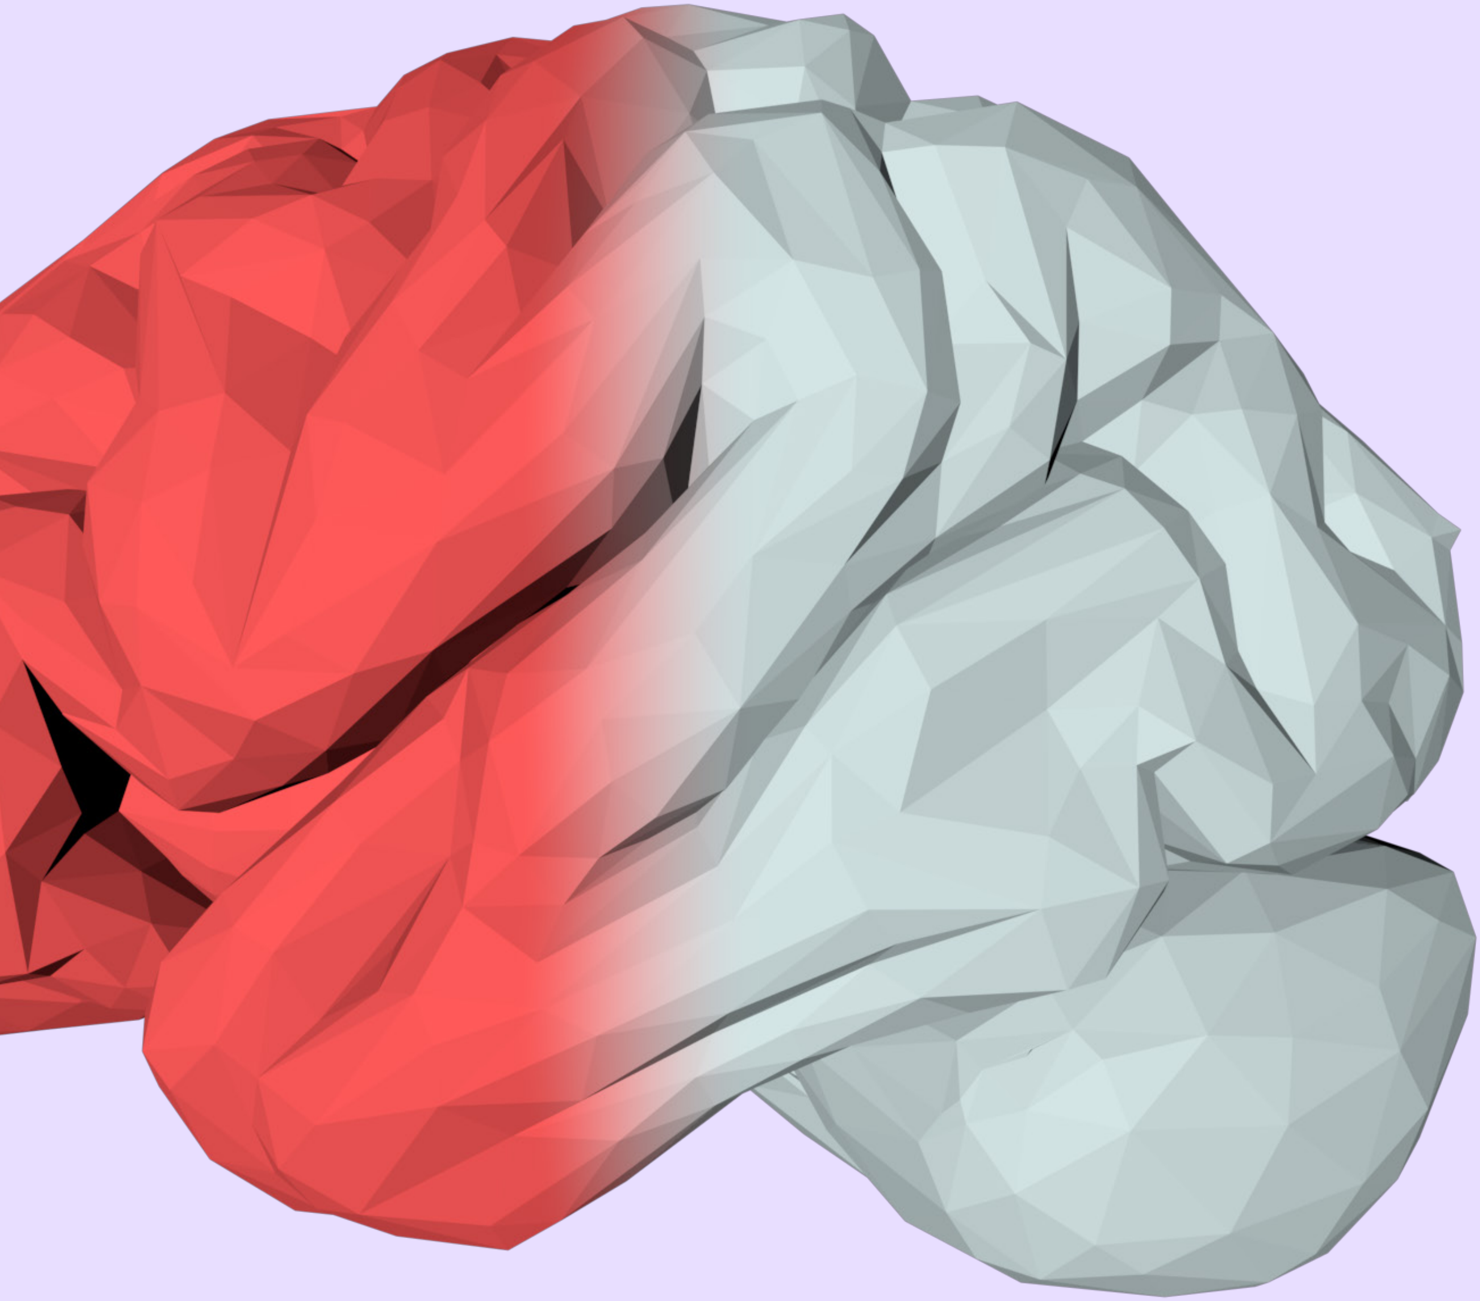

in many cases it will be more appropriate to respond closer to the time the technologies mature. However, there is a strong case for early work preparing for possibly rapid changes brought about by AI systems because a reactive response won't succeed when the change is very rapid. In particular, important research could be done on how to give AI systems desirable goals. Foresight work could be done to better understand which other issues require advance preparation, or where there may be an eventual role for other responses. Work could also be done to encourage understanding of the risks among AI developers, especially around automated AI system development, which might enable very fast transitions. However, the benefits of AI could be very great so it is important not to unnecessarily impede AI development.

C. GEOENGINEERING

As mentioned previously in this chapter, geo-engineering – in the form of Carbon Dioxide Removal (CDR) or Solar Radiation Management (SRM) – could help to limit the risks of catastrophic climate change. CDR techniques, such as ocean fertilisation or carbon sequestration, remove Carbon Dioxide from the atmosphere, whereas SRM techniques, such as cloud brightening or the injection of sulphates or other particles into the stratosphere, reflect the sun's light and heat back into space. Certain forms of CDR could carry major risks. For example, ocean fertilisation using iron or urea could pose major risks to marine ecosystems.[125] However, most forms of SRM are thought to carry much greater risks than most forms of CDR,[126] and worries about civilisation-threatening consequences have generally focused on SRM (and in particular on currently the leading form of SRM: the injection of sulphate particles into the stratosphere).[127] The remainder of this section will therefore focus on SRM only.

SRM is the only known technique for quickly stopping (or even reversing) the rise in global temperatures. This means that it could be used as a complement to GHG reduction, to manage temperatures while the world phases out fossil fuels. Some have proposed that SRM could provide insurance against a 'climate emergency', such as unexpected abrupt and extreme warming.[128] Thus, on some scenarios, failing to use SRM could constitute a global catastrophic risk. Moreover, SRM has the potential to reduce the costs of warming at very low cost. Some estimate that the annual cost of stratospheric aerosols could be less than $10 billion per year, which is orders of magnitude less than the costs of climate change mitigation strategies.[129]

POTENTIAL IMPACT OF THE RISK

Four main arguments have been given for the view that SRM brings global catastrophic risks. Firstly, while existing models suggest that SRM could reduce the catastrophic effects of climate change and will not bring their

2.5.2. GLOBAL TEMPERATURE ANOMALY FROM GEO-ENGINEERING FOLLOWED BY TERMINATION

Global average surface air temperature change from business as usual emissions, injection of 3 megatons/annum (Mt/a) of SO2 in the Arctic, 5 Mt/a of SO2 in the tropics, and 10 Mt/a SO2 in the tropics.[130]

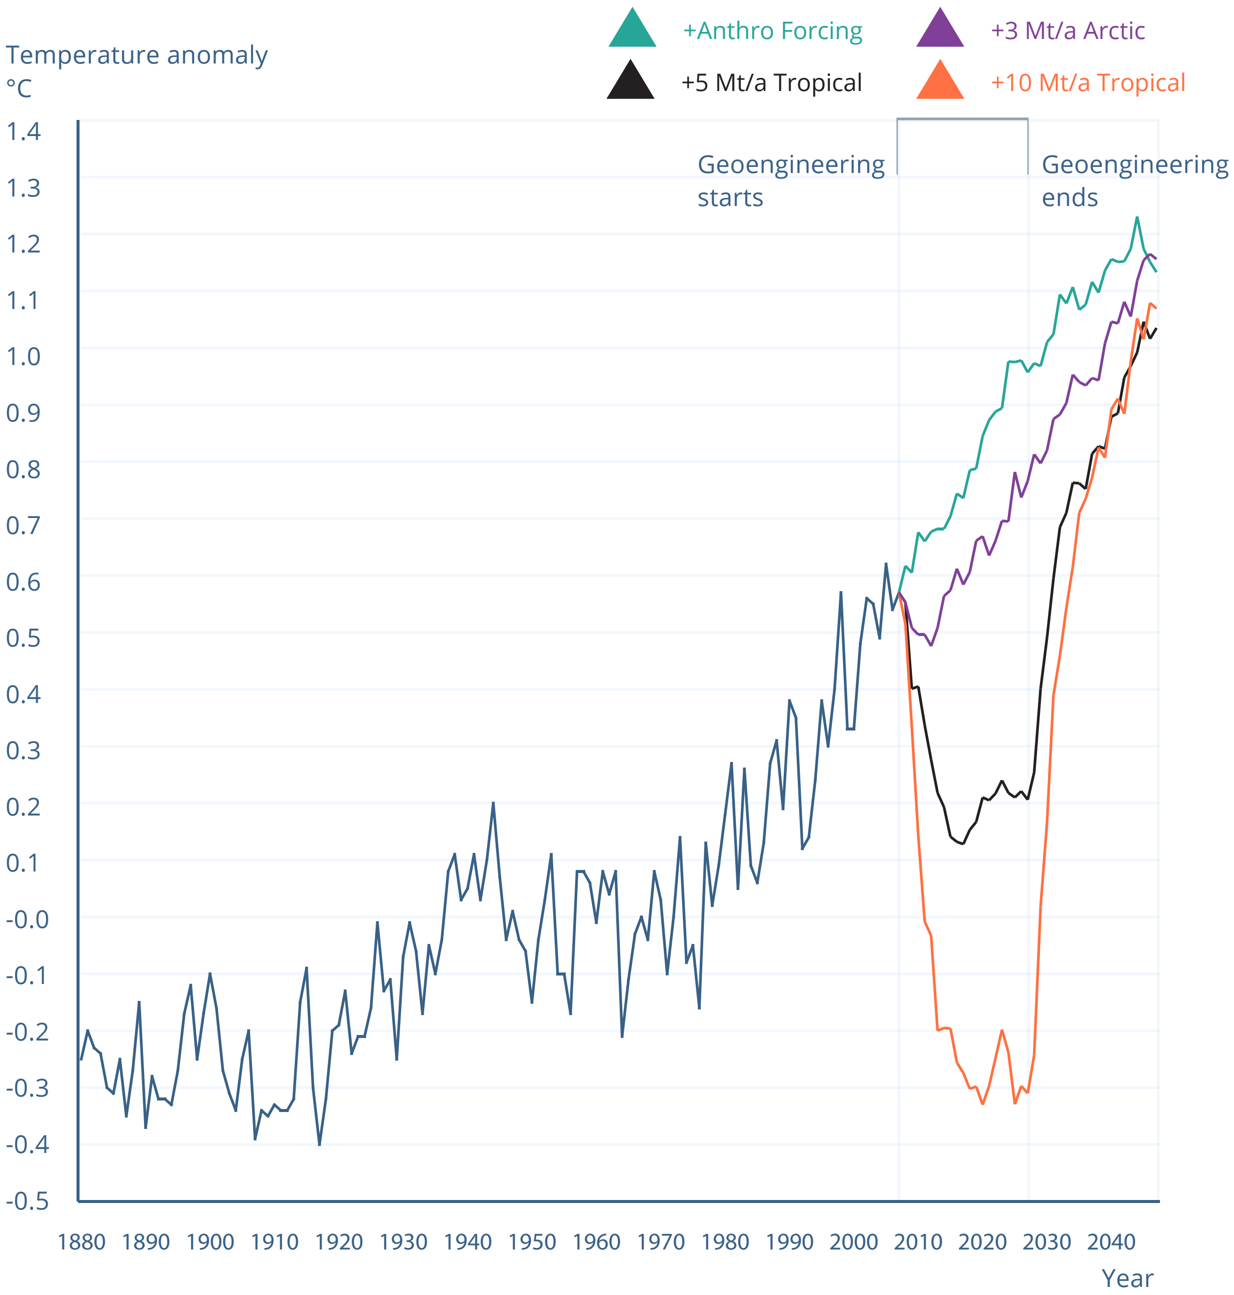

own catastrophic impacts, it may nevertheless bring currently unknown risks, particularly through impacts on global precipitation.[131] The climate system is imperfectly understood and the deployment of a novel technology with global effects is inherently risky. Whether the use of SRM is more risky than allowing the planet to warm is one of the key questions for anyone considering using SRM. Secondly, because SRM is so cheap, individual states could feasibly deploy it and unilaterally bring about global climatic impacts. Moreover, individual states might also have the incentive to do this because they could be particularly badly affected by climate change.[132] Individual states acting alone may be less likely to properly take into account the interests of other states and may be concerned about catastrophic consequences in other regions.

Thirdly, sudden termination of SRM would lead to rapid and severe global warming.[133] There are some reasons to think that an SRM system could be very resilient against external shocks and termination. If the current deployer were to suddenly stop SRM for some reason, every other country would have strong incentives to resume SRM.[134] Thus, sudden termination might only be likely in the event of a severe global catastrophe which undermines the capacity of all countries to use SRM.[135]

Finally, research into geo-engineering or advocacy for geo-engineering could present a 'moral hazard' in that it could discourage GHG reduction efforts.[136] All major reports studying SRM have concluded that it does not present an alternative to emissions reductions, as it only masks the effects of GHGs and does nothing to counter ocean acidification.[137] Therefore, it could be problematic if geo-engineering drew a disproportionate amount of policy attention. However, it is important to bear in mind that there are countervailing reasons in favour of SRM research and advocacy.

LIKELIHOOD OF THE CATASTROPHE

The probability that there will be a global catastrophe brought about by SRM depends on many factors, including the timing, speed and severity of global warming; the state of SRM technology and our knowledge of the climate system; the response of the climate system to SRM; the form of SRM deployed and how it is used; and how well the world does at organising governance of SRM. There is obviously very large uncertainty about all of these factors.

We argued above that unless strong action is taken soon, there is a sizeable chance of catastrophic warming. Catastrophic warming would create very strong incentives to use SRM. Therefore, unless strong GHG reduction action is taken soon, the chance that SRM is used will increase. The greatest chance of catastrophe probably comes from poorly planned and governed the use of SRM, perhaps by an individual state or small group of states. Since most of the risks of a

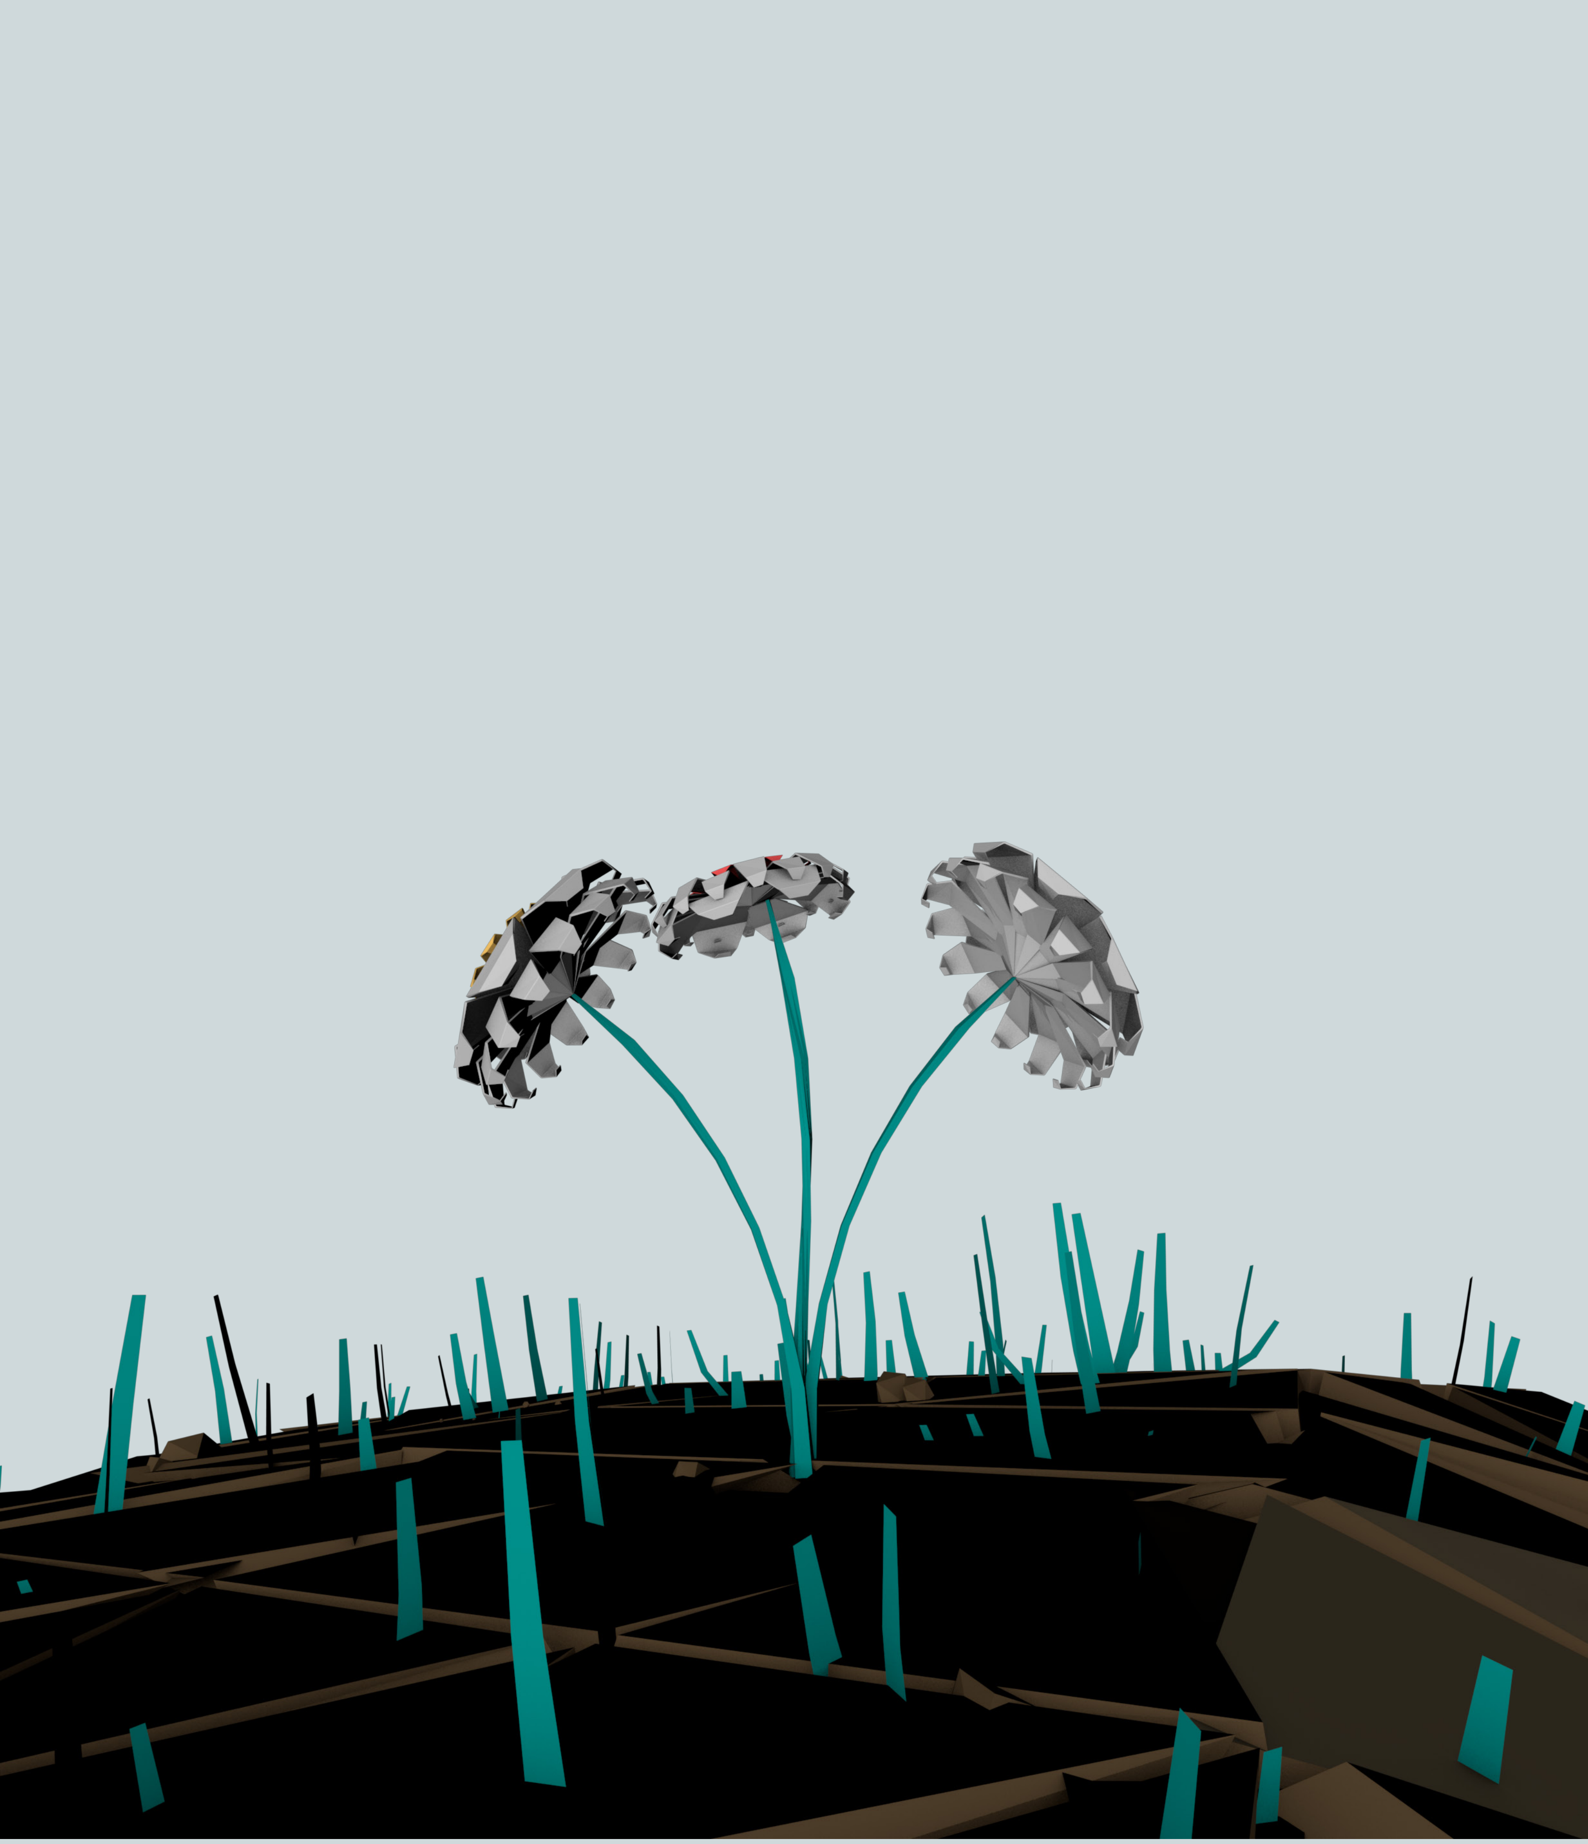

well-governed and well-planned form of SRM are unknown, it is very difficult to say how likely that form of SRM is to produce a global catastrophe.

Finally, it is unclear the impact research into, advocacy for, or preparation for, geo-engineering would have on global willingness to cut greenhouse gas emissions. Therefore, it is not clear to what extent geo-engineering is a moral hazard.

MAIN ACTIONS AVAILABLE TO LIMIT THE RISK

The main way to reduce the risks from SRM would be to take strong action on GHGs to reduce the expected costs of climate change and thereby reduce the incentives to geo-engineer. Further research into the different kinds of SRM and into the response of the climate to them would reduce the unknown risks of SRM. Finally, working to develop geo-engineering governance through climate treaties and through global institutions might limit the risks of unilateral SRM.[138] However, both of these actions might also have moral hazards. We cannot settle the moral hazard debate here, but it is important for policymakers to be aware of these issues.

THE CHALLENGE OF EMERGING RISKS

The catastrophic risks from emerging technology are particularly challenging. Firstly, we have no track record of dealing with these emerging technological risks and, as we discuss in chapter 4, existing national and international institutions are not designed to deal with them. It is therefore less likely that our eventual responses will be effective.

Secondly, these technologies might, like nuclear weapons, reach maturity more quickly than expected. If so, the nature of the problem would only become fully apparent over a short time-frame and it would be difficult to make an appropriate response. Thirdly, some of these technologies could be harder to control than nuclear weapons. Nuclear weapons require the rare and controllable resources of uranium-235 or plutonium-239. In contrast, if some of these technologies reach full maturity, they could be accessible to small countries or even terrorist groups.

In the coming decades, emerging technologies will provide major benefits to society, but they may also create significant and unprecedented risks.

2.6. Other risks and unknown risks

There is a broad spectrum of possible global catastrophic risks, from very salient and well-understood threats to those which are extremely low probability and highly speculative. In compiling this report, we had to make decisions about which risks to investigate carefully, and which to ignore. Our guiding principle was to include risks which could clearly cause a global catastrophe in our definition if they were to happen, which could be speculative but not too speculative, and which were not too low in probability.

There are several risks we didn't include for these reasons. For example, gamma ray bursts are fairly well understood and would be catastrophic, but have a very low annual likelihood.[139] A global totalitarian state is a possible scenario, but it is unclear what proportion of the scenarios would meet the threshold for a global catastrophe and how likely it would be to arise given the current geopolitical environment.[140] Conventional and chemical warfare are real threats, but quite unlikely to reach the scale of a global catastrophe.

It is important to remember, however, that nearly all of the most threatening global catastrophic risks were unforeseeable a few decades before they became apparent. Forty years before the discovery of the nuclear bomb, few could have predicted that nuclear weapons would come to be one of the leading global catastrophic risks. Immediately after the Second World War, few could have known that catastrophic climate change, biotechnology, and artificial intelligence would come to pose such a significant threat.

These risks emerged due to rapid economic and technological development, which looks set to continue apace in the coming century. That might create a number of new risks. Therefore, it seems likely that some future global catastrophic risks are at present unknown.

Detailed planning for unknown risks is of course impossible, but steps can be taken to improve our preparedness. Bodies specifically tasked with horizon scanning and the discovery of new global catastrophic risks would give the international community more time to craft an appropriate response. Measures can also be taken to improve general societal resilience to catastrophe. These are discussed in chapter 5.

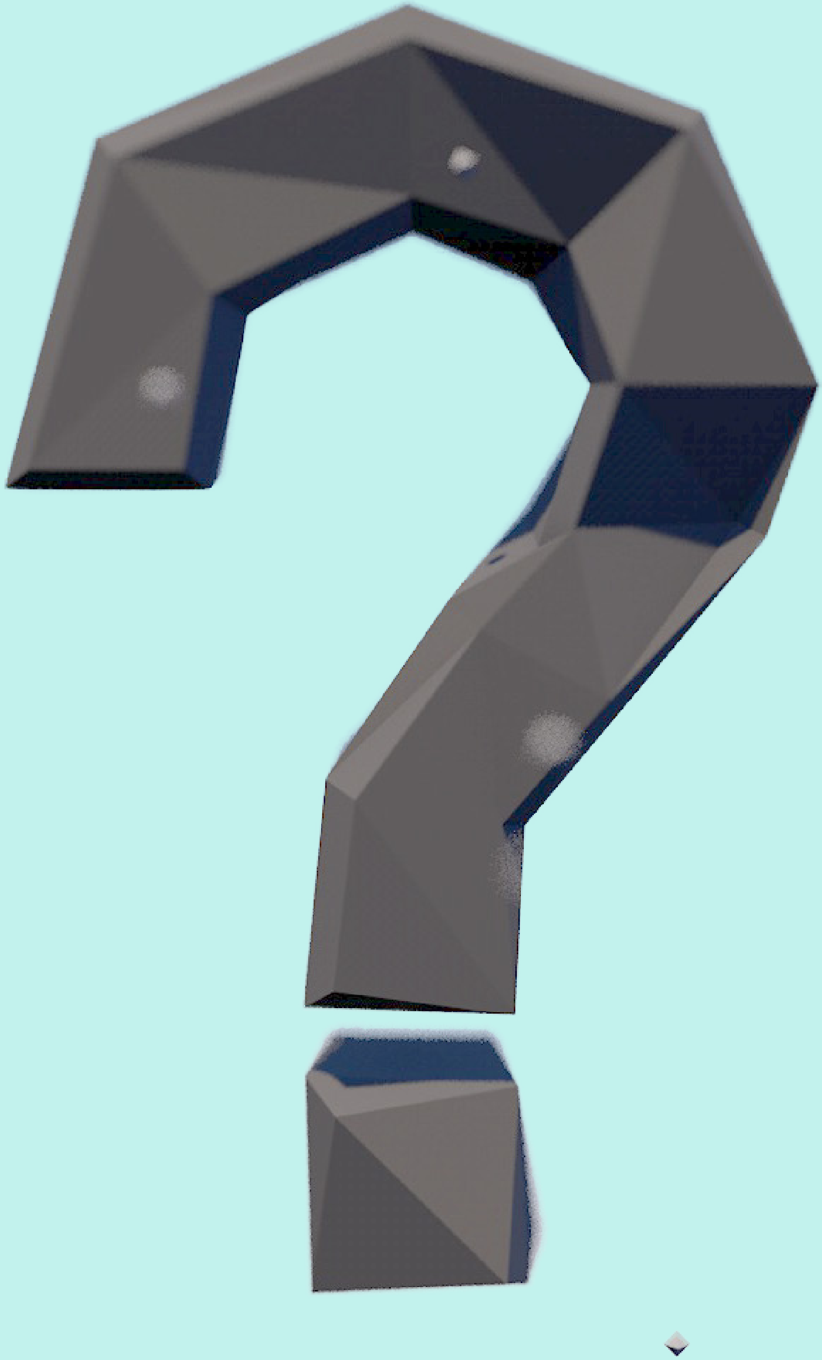

2.7. Our assessment of the risks

We've given qualitative descriptions of each of these risks. It is useful to consider comparisons between them as well. Unfortunately there is a great deal of uncertainty around the level of many of the risks, and around how much can usefully be done. Nonetheless, some attempt must be made in order to prioritise the world's scarce resources, and there are some areas where we can be more confident than others.

In this section, therefore, we make some comparative assessments of the risks. Because of the high level of uncertainty we categorise the risks into broad bands rather than attempt precise assessment. The categorisation remains our subjective judgement, and it is possible that some risks should be in different bands. But we believe that it is better to offer even such an imperfect assessment than nothing.

A. OUR METHOD

We have assessed the risks on two different dimensions. First, the current likelihood: how likely is it to materialise in the next few years? Second, how much work should be given to reducing the risk in the next few years (according to our judgement)? These two dimensions are linked, but they can come apart, as in the case of climate change whose main harms are not likely to be felt soon, but which demands action today. We think that this split should be explicit.

To assess current risk, we considered the likelihood over the next five years. Our upper band consists of events which we consider to be distinct possibilities. The lower band consists of risks which appear low in absolute likelihood, either because the base rate is low enough or because five years is too short a timescale for the risk to develop. Because of the large scale of the potential catastrophes, even risks in the lower category may still be significant over this time period.

It is even harder to assess how large the response should be. We considered how much risk was posed, both at present and in years to come. We also considered how good the opportunities to reduce the risk appear, and for emerging risks how much they benefit from early responses. Our upper band indicates risks where we think a significant global response is likely appropriate, perhaps involving thousands of people, or billions of dollars. There is a wide range within this band, and we do not think all

Bodies specifically tasked with horizon scanning and the discovery of new global catastrophic risks would give the international community more time to craft an appropriate response.

of the risks there should receive the same amount of attention. For risks in the lower band, we still think responding is appropriate, but perhaps at a smaller scale for the time being.

Our assessment is summarised in this diagram:

B. HOW WE MADE THESE ASSESSMENTS ABOUT PARTICULAR RISKS

Catastrophic climate change would likely have effects decades out, and we consider the likelihood in the next five years to be small unless the scientific community has significantly mis-modelled climate dynamics. Nonetheless, because of cumulative nature of greenhouse gas emissions, and the relatively well understood dynamics, we believe a large response is justified.

Natural pandemics have been responsible for past global catastrophes, and are seen high on lists of national risk assessments. We are not confident assigning a particular probability to a pandemic, but we do think it is among the most likely risks in the next five years. International work in reducing the risk from pandemics seems both important and effective, and we think a large response is appropriate.

Although the likelihood may have decreased since the Cold War, nuclear war appears to remain a real possibility. We have correspondingly placed it in the upper category for likelihood. It is the only catastrophe that is definitely within the reach and control of humanity today, and we think it is therefore appropriate that the response be high.

Asteroid impacts and supervolcanic eruptions are caused exogenously to human actions. This means we have better estimates of their rates, through knowledge of historical incidents. A supervolcanic eruption (which may or may not cause a global catastrophe) is estimated to occur very approximately every 30,000 years. An asteroid impact large enough to cause a global catastrophe is estimated every 120,000 years. We are therefore reasonably confident that they belong to the lower category of current risk. Because of this relatively low risk, we have also put them in the low category for attention.

We are confident that this is justified for supervolcanic eruptions, where there are few clear actions to reduce the risk. Asteroids are exceptional in that we have relatively well-understood ways to reduce the risk, and this could mean that the attention category should be higher.

Engineered pandemics are an

1 every 30,000 years

is the estimated occurrence of a supervolcanic eruption (which may or may not cause a global catastrophe).

FIGURE 2.7. OUR ASSESSMENT OF GLOBAL CATASTROPHIC RISKS

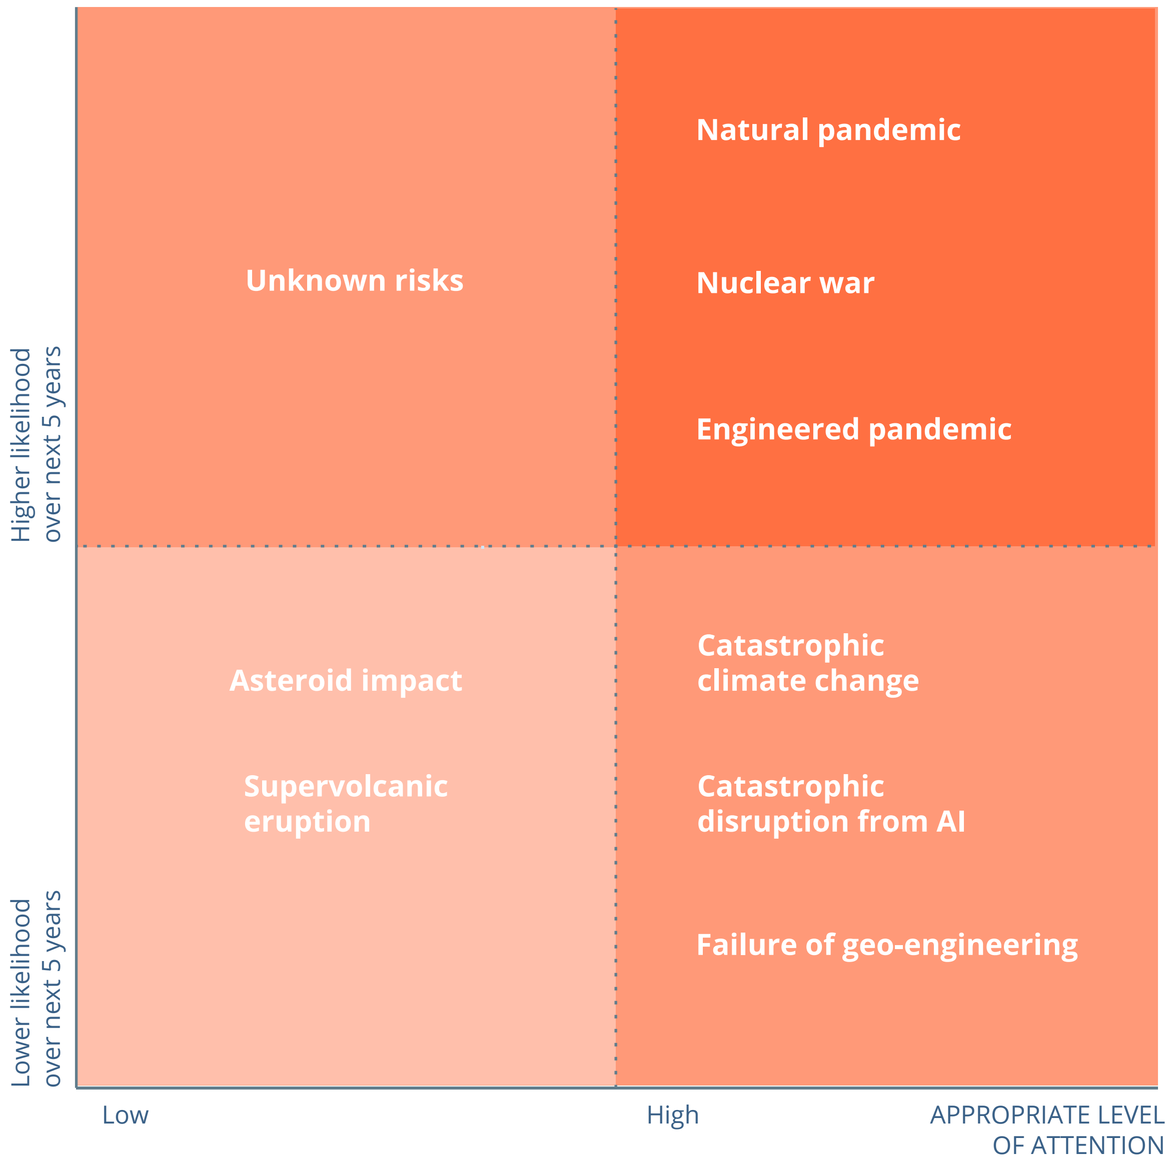

emerging risk in that future technology will make them more plausible. But even existing experiments have been claimed to pose real risk of sparking a pandemic, so we put this in the upper band for likelihood. We also think it is appropriate that it gets significant amounts of attention, particularly as the unmanaged risk may be increasing.

Risks from artificial intelligence appear to be low in the short-term, since we most likely are some way off the kind of capabilities that might cause a global catastrophe. We have therefore put them in the lower category of current risk, although it is hard to predict the speed of technological progress. In the longer term the risks could be extremely significant, and experts have identified useful work that can be done today, so we have put it in the upper category of attention.

Finally, risks from geoengineering are another emerging technological risk where the current risks seem low. But because geoengineering may be employed to tackle climate change, it seems important to build a more comprehensive understanding of the risks it poses in itself. For this reason we have put it in the higher category for attention.

There is a great deal of uncertainty around many of the risks, and around how much can usefully be done. Nonetheless, some attempt to estimate must be made in order to prioritise the world's scarce resources.

Chapter 3

Risk factors and interactions between risks

In this chapter we discuss some of the factors that affect the likelihood of global catastrophic risks. We first discuss some of the key factors influencing individual risks, before turning to factors that the risks have in common, and ways in which catastrophic events could trigger further catastrophes.

Global catastrophic risks are parts of extremely complicated causal networks, and the drivers of each one could quite sensibly be the subject of multiple books. Here, we provide a brief overview where we highlight a few key drivers. This serves as a background to the more action-oriented discussions of Chapters 4 and 5.

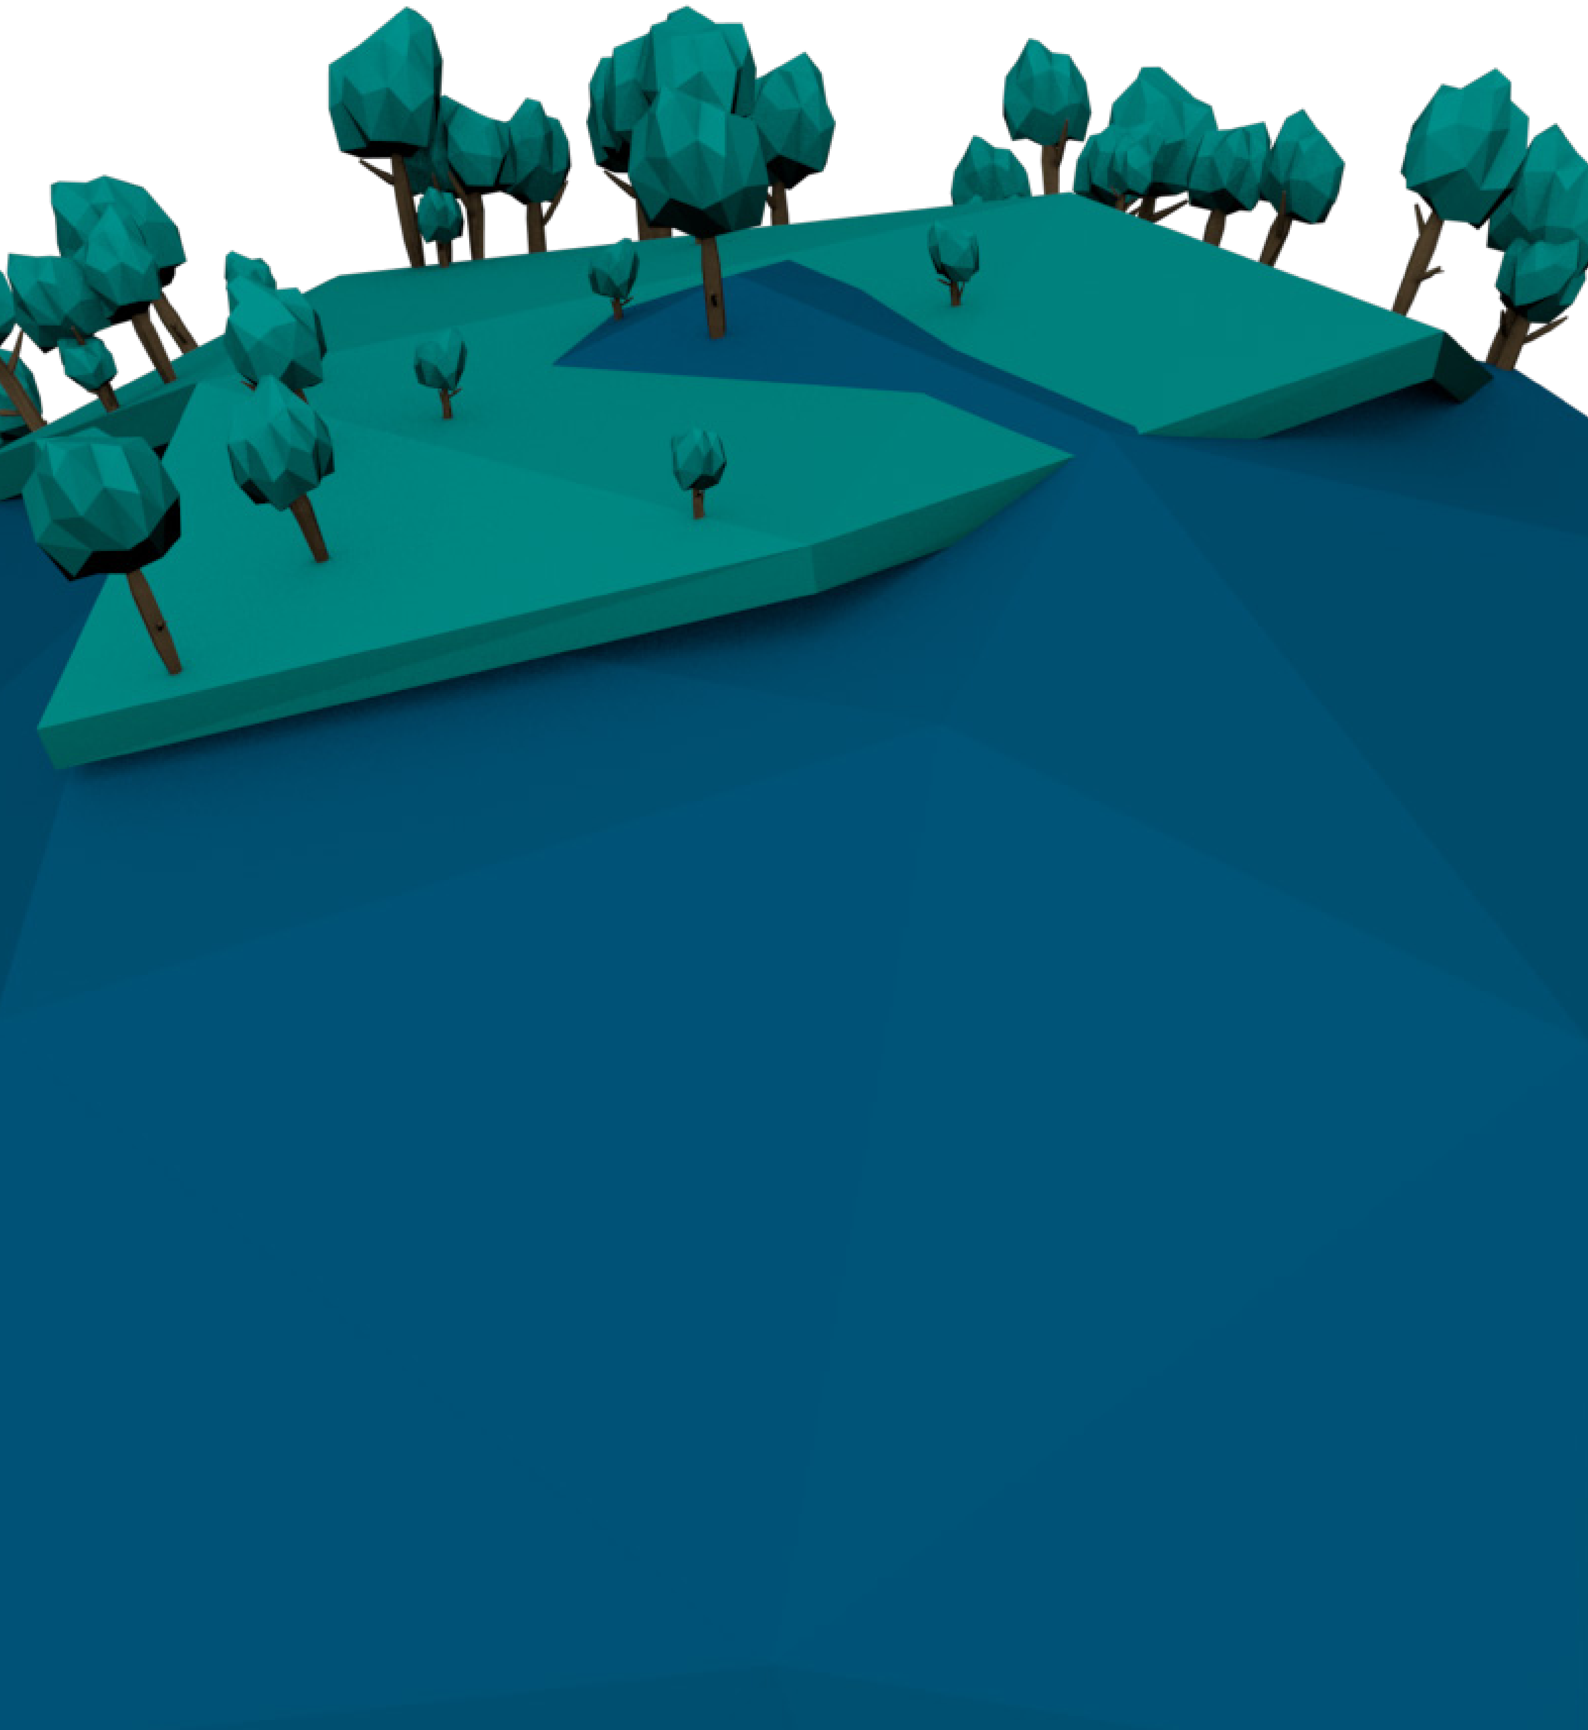

3.1. Drivers of individual risks

NUCLEAR WAR

Increasing the number of nuclear-armed actors -- horizontal proliferation -- probably increases the likelihood of nuclear conflict. Increasing the size and sophistication of their arsenals -- vertical proliferation -- increases the amount of damage that could occur in a given conflict.[141] These factors are, in turn, influenced by political factors such as leading nations' ability to control rogue actors, and the level of conflict between nuclear states.[142]

The possibility of false alarms or miscalculation between nuclear powers is another driver of risk, exacerbated further when geopolitical tensions are high.[143] There is also the unresolved question of exactly how sensitive the climate system is to dust and sulphates. The higher the sensitivity, the greater the risk of a nuclear winter scenario.[144]

Our response to climate change may also have significant effects on the risk of nuclear war. Many authorities on climate change argue that it will be hard to reduce GHG emissions sufficiently without nuclear power being part of the energy mix.[145] Unfortunately, rising use of nuclear power may increase the chance of nuclear weapon proliferation.[146] Moreover, the mass migration and resource scarcity which could result from significant climate change may increase the geopolitical tensions that drive nuclear risk.[147]

PANDEMICS

The factors driving pandemic risk divide into those affecting the likelihood of potentially pandemic pathogens arising, and those affecting how much society will be affected by such pathogens. Whether a given pathogen could cause a pandemic depends on its natural parameters, such as transmission vector, resistances, and lethality, over which we generally have little control.[148] However, concentrated populations of poultry or other animals can increase the probability of zoonotic spillover, raising the risk.[149] The transmission of pathogens through society is increased by global travel and dense populations, but decreased by factors such as improved hygiene.[150] The effectiveness of our countermeasures constitutes another important factor, ranging from disease surveillance to healthcare access in the developing world.[151]

SUPER-VOLCANOES, ASTEROIDS, AND COMETS

The probability of super-volcanic eruptions and asteroid and comet impacts are driven by natural process-

> “If ‘positive feedback loops’ prove to be worse than anticipated, the risk of catastrophic climate change will be higher.”

es outside of human control. However, our ability to predict, and, in the case of asteroids and comets, to prevent these catastrophes could have a significant bearing on their eventual impact. Because we have less control over the likelihood of an adverse event with these than with most global catastrophic risks, resilience is a more central part of the risk management strategy. A particularly important factor is the level of food security, including stockpiles and the ability to switch to less sunlight-dependent food sources.[152] Better food security means a lower chance that a particulate winter would have catastrophic consequences.[153]

CATASTROPHIC CLIMATE CHANGE & GEO-ENGINEERING

The risk of catastrophic climate change and the risk of a geoengineering catastrophe are strongly influenced by the level of GHG emissions. These depend on factors including the amount of energy the world consumes and the proportion of this which comes from fossil fuels. The total energy consumption is shaped by, among other things, population growth, economic development, and energy efficiency. The proportion of energy which comes from fossil fuels depends on other factors such as the rate of progress in clean energy technology development and the prevalence of economic incentives to adopt clean energy technologies.[154]

Another important set of factors concern how the climate is likely to react to increased levels of GHG emissions. If the climate turns out to be generally more sensitive to GHG emissions than expected, or if "positive feedback loops" prove to be worse than anticipated, the risk of catastrophic climate change will increase.[155]

Finally, the development of new geoengineering techniques might either decrease or increase global catastrophic risk. They could decrease global catastrophic risk if they prove to be an effective tool to mitigate catastrophic climate change.[156] But they could increase risk if they have a high chance of causing a catastrophe of their own, or if they prove to be ineffective while at the same time leading countries to avoid emission abatement (as discussed in Chapter 2).[157]

EMERGING TECHNOLOGY

The factors that impact the risks of emerging technology are surrounded by a high degree of uncertainty. However, one plausible risk factor is misaligned incentives, which might cause nation-states to engage in arms races or under-regulate risky technology for economic reasons.[158] The speed and suddenness of technological breakthroughs could also be a risk factor, as sudden breakthroughs might leave inadequate time for social and political institutions to adjust their risk management mechanisms.[159]

Formal regulation of technology has an unclear impact—in some cases it could prevent risky activities, but in other cases it could stifle critical innovation or, if implemented before the issues are properly understood, even increase risk.[160] On the positive side, technology developed in a culture of responsible innovation will likely be safer.

For some of the emerging technological risks we may also hope to develop solutions. It may reduce risk to make progress on developing safety aspects of a technology faster relative to progress on the technology itself. For example having excellent diagnosis and treatment tools to contain outbreaks before there is significant risk from artificial pandemics could reduce the total risk significantly. Similarly, a solution to the problem of aligning advanced AI with human values before we are able to create advanced AI would reduce risk.[161] The proposal of aiming for safety-enhancing technologies ahead of risk-increasing technologies in order to reduce global catastrophic risk is called 'differential technological development'.[162]

UNKNOWN RISKS

It is of course impossible to provide a detailed account of the drivers of currently unknown risks. Unknown risks could in principle include both exogenous and anthropogenic risks. We have access to a long historical

> The speed and suddenness of technological breakthroughs could also be a risk factor, as sudden breakthroughs might leave inadequate time for social and political institutions to adjust their risk management mechanisms.

record, which might have brought many exogenous risks to our attention already. The historical record offers us little or no guidance on anthropogenic risks. They will often be driven by social or technological change, so a higher rate of change could mean more currently unknown risks in the future (although technological progress could also reduce global catastrophic risk; see section 3.2.). Another important factor influencing these unknown risks is the quality of our foresight work, since the ability to plan for a risk could help both avoidance and mitigation.

3.2. Shared risk factors and interactions between risks

So far we have mainly discussed each risk in isolation. But the global catastrophic risks share numerous risk factors, and also interact with each other. In this section, we discuss some salient examples of shared risk factors - variables associated with an altered level of risk - and interactions between risks.

One important shared risk factor is governance. Good governance may help timely and appropriate preparation and response to risks. By the same token, bad governance, nationally or internationally, could increase the likelihood and potential impact of every risk. Lax oversight might lead to the accidental release of a dangerous engineered pathogen from a laboratory. Dysfunctional governments would be less able to intervene early in an outbreak to prevent it becoming a pandemic. Poor international coordination might radically worsen our prospects to avoid catastrophic climate change, and international tensions could increase the risk of nuclear war. Bad governance could in principle also lead to over-prioritising catastrophic risk reduction, but because of the political distortions discussed in Chapter 4 we think this is less likely.

While bad governance can cause global catastrophes, the causality might also be reversed. A global catastrophe could lead to a breakdown of social and political institutions, which in turn could cause an outbreak of violence. In a 2008 paper, Nel and Righarts argued that even smaller natural disasters such as earthquakes, tsunamis, and heat waves significantly increase the risk of violent civil conflict in the short and medium term, especially in low- and middle-income countries.[163] Similarly, there has been significant attention on the relationship between climate change and political conflict. Many scholars and security experts have argued that even moderate climate change could increase the risk of political violence because of conflicts over dwindling natural resources (such as food and water), massive international migration and a range of other factors.[164] Greater disasters are likely to lead to even more upheaval, and might significantly weaken the defenses against further catastrophic events. In this way, one global catastrophe might trigger another.

Another set of factors which affect many global catastrophic risks is technological. Technological advances change the system we live in by giving actors new powers. These in turn could increase or decrease global catastrophic risk, even aside from cases

Food stockpiles and the ability to rapidly increase production of alternate sources of food would increase resilience to a broad range of risks.

where the technology directly poses risk. It is hard to fully predict the effects of technological trends, but we will highlight three trends which may be relevant: economic productivity gains, surveillance, and distributed manufacturing.

Global productivity has increased dramatically in the last few decades. Partly, this is because of technological progress and organisational efficiency growth. Partly, it is due to automation, a form of technological progress which we expect to continue.

This growth may help us reduce risk because it lets us spend more on prevention and resilience. Measures like clean energy and food stockpiles are costly, and more likely to receive investment when people are wealthy. In the tail case, with very powerful artificial intelligence, there could be a radical improvement in our ability to manage other global catastrophic risks. However, economic development may exacerbate some risks, at least in the short term, for example by increasing greenhouse gas emissions.

The global surveillance disclosures by Edward Snowden and others revealed that state surveillance is now extensive. In addition, the use of CCTV is expanding, and police and other law enforcement agencies are making increasing use of cameras in their day-to-day work.

Increased surveillance has the potential both to exacerbate and reduce global catastrophic risk. Extensive surveillance could make it easier for unscrupulous states to control their citizens, which makes it easier for states to engage in extreme behaviour. It might also weaken trust between states, which could cause political tension. However, surveillance may make it more difficult for malicious states or terrorist groups to act in secret. This will become increasingly important as barriers to access destructive weapons fall. Surveillance between countries could even facilitate international cooperation by making the actions of states more transparent. It is not clear what overall effect surveillance has on levels of global catastrophic risk.

Distributed manufacturing is a set of technologies that allow products to be designed and built without centralised factories. They offer many benefits, and may increase resilience to catastrophe by spreading out production capacity. However, they also bypass some government controls designed to prevent the construction of destructive weapons. 3D printing, an early form of such technology, has already generated security risks by allowing people to create functional homemade firearms. More powerful forms of distributed manufacturing could increase risk caused by malicious actors, by increasing access to powerful weaponry such as bioweapons.

Nuclear war, geo-engineering, super-volcanoes, asteroids, and comets all pose global catastrophic risk in significant part because of the ‘particulate winter’ scenarios they might produce. By ejecting large amounts

FIGURE 3.2. A WORKING GUN MADE FROM PLASTIC ON A 3D PRINTER[165]

of smoke, dust, and/or sulphates into the stratosphere they could cause global cooling, sunlight loss, ozone loss, and subsequent agricultural disruption.[166] Because so many risks share this mechanism, many strategies for resilience that address one risk address several. Food stockpiles and the ability to rapidly incerase production of alternate sources of food would increase resilience to a broad range of risks.[167]

Chapter 4

Do institutions collectively underinvest in global catastrophic risk?

In Chapters 2 and 3 we gave an overview of several global catastrophic risks, the main factors that affect their likelihood and impact, and some of the levers available to influence them. In the first half of this chapter, we argue that market and political distortions mean that these risks are likely to be systematically neglected by many actors. This increases the importance of attention to the risks, and suggests some mechanisms for reducing risk by countering the distortions.

In the second half of the chapter, we examine the kinds of actors or institutions that may be well-placed to act on global catastrophic risks, or have a responsibility to do so. We look at how they can help to correct the market and political failures we consider in the first half.

> We should expect global catastrophic risks to systematically receive less attention than they merit.

4.1. Market and political failures

It may seem surprising that relatively little effort has gone into global catastrophic risk mitigation. However, we should expect global catastrophic risks to systematically receive less attention than they merit, for structural reasons described below.

4.1.1. GLOBAL PUBLIC GOODS

Global catastrophic risk reduction is a global public good – the benefits of reduction spill over to other countries and it is hard (and arguably unethical) to only protect the countries that contribute to risk reduction. Many global public goods tend to be under-provided, primarily because countries try to free-ride.[168] Every country might hope that other countries invest in risk reduction, so they can get protected for free. If everyone shares this hope, no one will invest. To overcome this, there is a need for good coordination between actors.

In well-functioning states, national public goods, such as defence and clean air,[169] are provided despite the difficulties of collective action. Sometimes, central regulation supplies these goods, sometimes market mechanisms are employed, and sometimes cooperative institutions for supplying the goods emerge.[170] Supra-national institutions are generally weaker and international communities less cohesive, and therefore are less able to implement these solutions to collective action problems. Sometimes, such as with the World Trade Organisation, nations do collectively give up some of their autonomy in order to provide a public good, but this process requires a great deal of negotiation and trust. Thus, institutions for aligning the incentives of nation-states to provide global public goods are typically less mature and less effective than those that provide national public goods.

4.1.2. INTERGENERATIONAL PUBLIC GOODS

Many of the benefits of global catastrophic risk reduction accrue to future generations. However, the interests of future generations tend to be systematically neglected because they cannot vote and have no direct voice in the political process.[171] For example, reducing the risk from catastrophic climate change may provide only small benefits to many people alive today, but could be very valuable to people who will exist in sixty years' time If we do not have formal processes for taking the interests of future generations into account, we may under-invest in risk reduction.

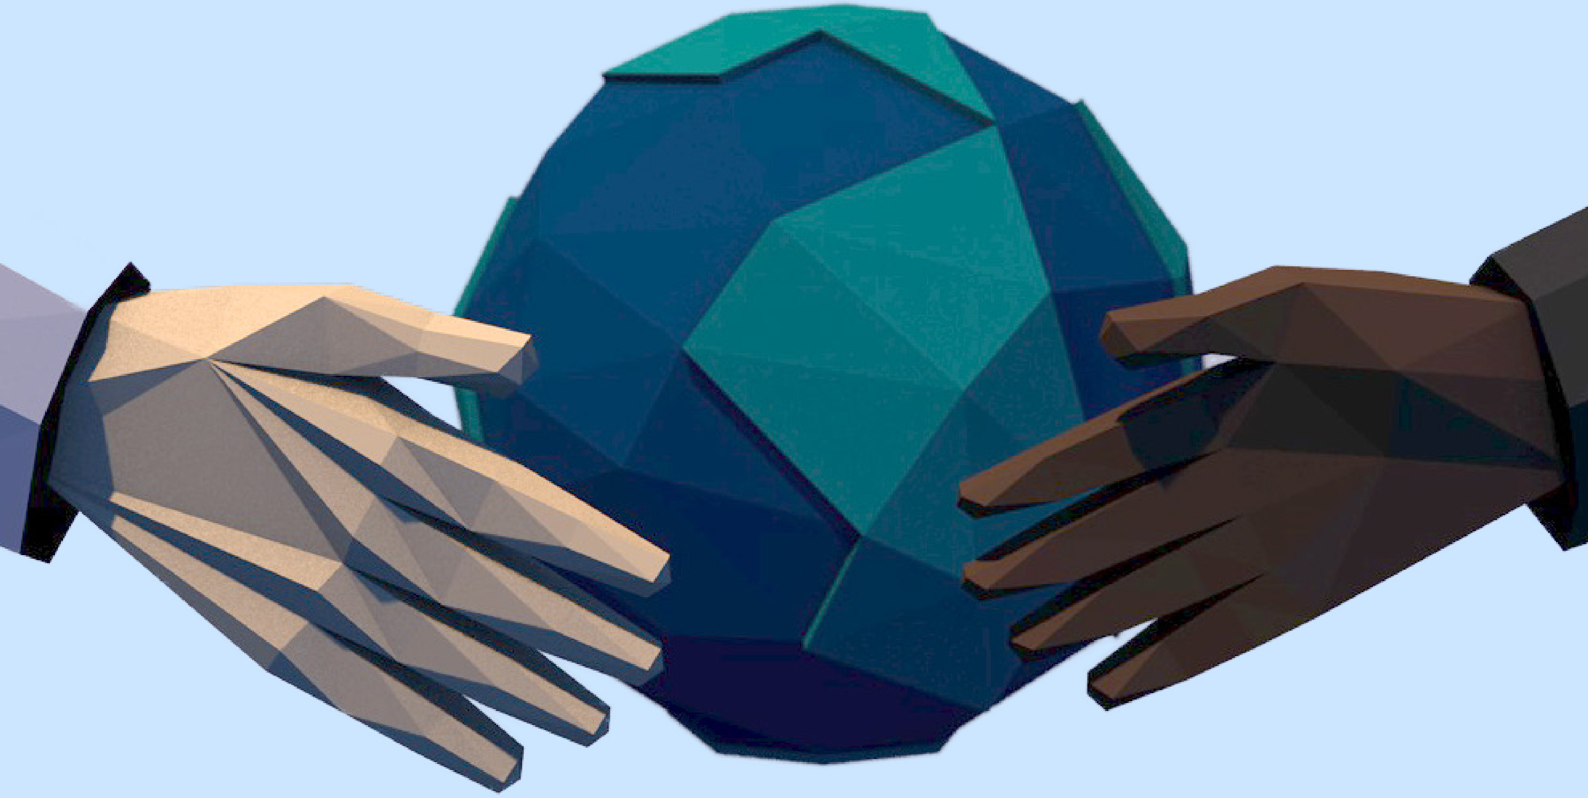

4.1.3. SPECIAL INTERESTS

The costs of regulatory actions to reduce global catastrophic risks will typically be concentrated on particular industries, whereas the benefits are dispersed. The small affected groups have strong incentives to lobby and campaign, while those gaining the benefits will not regard those benefits as a voting priority. This means that industry lobbying could wield disproportionate power over the regulatory process.[172] Consequently, the trade-offs made between economic profit and risk reduction could be skewed towards underprovision of risk reduction.

4.1.4. UNPRECEDENTED RISKS

Global catastrophic events occur very infrequently. In the last two millennia, there may have only been two such events – the Plague of Justinian and the Black Death. Anthropogenic global catastrophes are completely unprecedented. Because unprecedented events are typically less salient, it is less likely that governments and voters will pay appropriate attention to them, in spite of their very high costs in expectation.[173]

FIGURE 4.1. PREPAREDNESS IN THE AFRICA REGION

Multi-hazard national public health emergency preparedness & response plan in the African Region. [174]

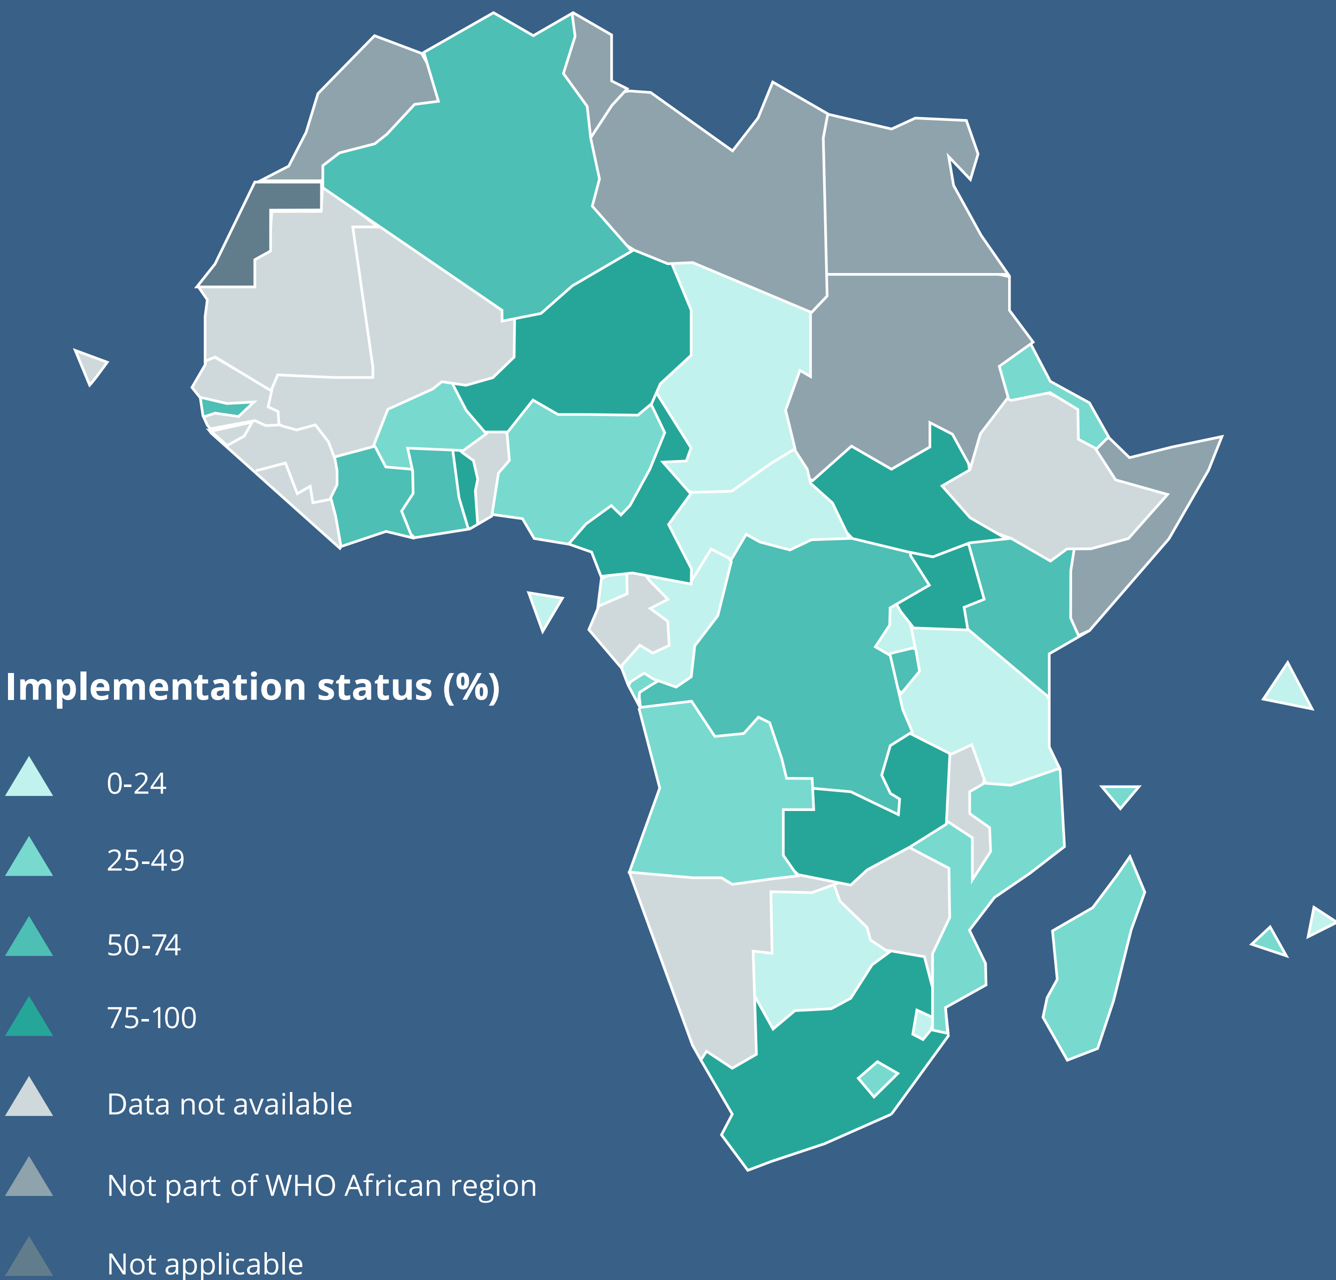

4.2. Which actors can help reduce global catastrophic risk?

There is a broad range of actors that can help us overcoming the current neglect of global catastrophic risk. In this section, we give an overview of some ways in which different groups, ranging from the international community to individuals, can contribute to the mitigation of global catastrophic risks.

4.2.1. THE INTERNATIONAL COMMUNITY

As discussed in section 4.1, reduction of global catastrophic risk is a global public good. This means that the international community will probably need to play a major role in reducing global catastrophic risk. In some cases, the global level of risk is significantly dependent on the capabilities of the countries that are the weakest links.[175] For example, to safeguard against the risk of a global pandemic, the least prepared countries should be enabled to strengthen their health systems.[176] In other cases, the risk of a catastrophe is mainly dependent on a few major actors. For example, the nuclear-weapon states have a disproportionate influence on the risk of nuclear war. International coordination can aim to reduce the risk of nuclear conflict between these major actors.[177]

There have been many agreements intended to reduce global catastrophic risk. Two of the more important examples[178] are the Treaty on the Non-Proliferation of Nuclear Weapons[179] and the 1992 United Nations Framework Convention on Climate Change, with its objective of preventing "dangerous" anthropogenic climate change.[180] Similarly, there are several permanent international bodies, such as the International Atomic Energy Agency and the World Health Organisation, which are (partly) concerned with reducing global catastrophic risks (nuclear war[181] and pandemics, respectively). The international community often collaborates productively with relevant parts of the academic community, as it did when it set up the Intergovernmental Panel on Climate Change.[182] Many of the global catastrophic risks - not least those associated with emerging technologies - require a deep level of technical expertise to be properly managed.

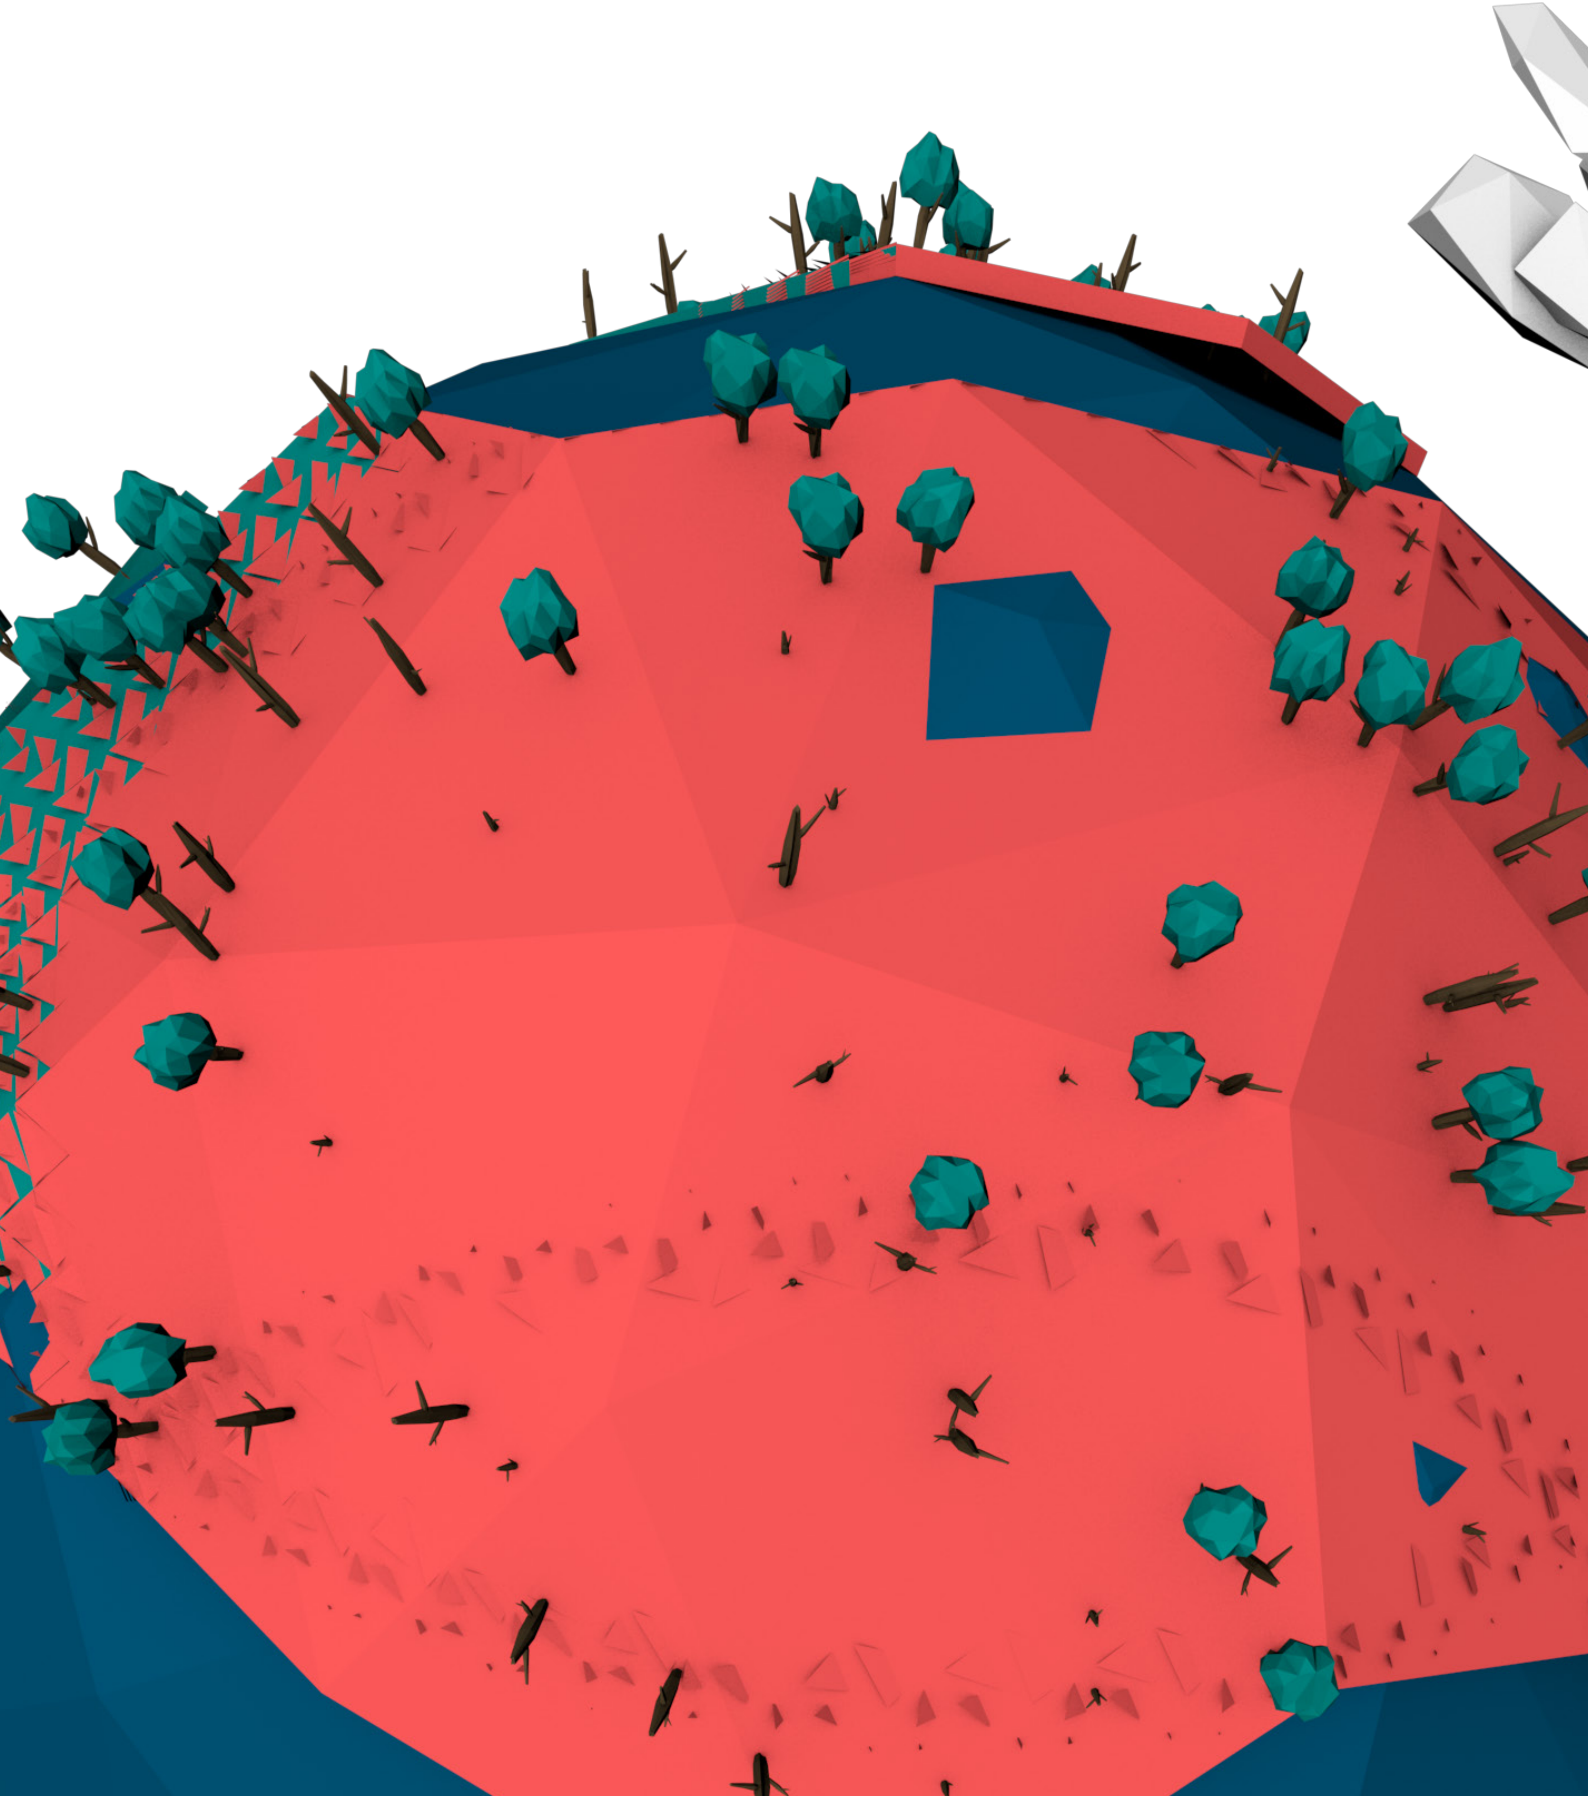

4.2.2. NATION-STATES

Although the international community has a crucial role in the reduction of global catastrophic risk — because this is a global public good — much power currently resides in the hands of nation-states. In practice, international bodies have limited power unless they are backed up by nation-states (especially the more powerful ones). This means that nation-states are essential to galvanising action. They also need to lead the process of implementation. For instance, it is up to nation-states to see that emissions actually are cut to the extent required by the Paris agreement. Similarly, it is the responsibility of nuclear-weapon states to guard against the possibility of accidental launch.

Nation-states can also reform their internal political processes in ways that are conducive to reduction of global catastrophic risks. Ensuring that decisions take account of the interests of future generations is, for instance, likely to lead to a more appropriate degree of focus on global catastrophic risk. Similarly, constraining the excessive power of special interests could decrease distortions on decision-making and so in turn reduce global catastrophic risk.

4.2.3. THE RESEARCH COMMUNITY

Many of the global catastrophic risks are not well-understood, and more research would allow more appropriate decisions about when to act and how to respond. For instance, we would benefit from more resources devoted to the study of global catastrophic risks such as tail risk climate change[183] and particulate winter scenarios.[184] We would also benefit from more research on technical solutions that could reduce global catastrophic risks, such as clean energy sources, or how to how to align the actions of above-human-level AI with our values.[185] Since catastrophic risk reduction is a public good, it is unlikely that this research will be carried out by companies under competitive pressure. Instead, it will most likely require public or charitable funding. In fact, several academic institutions that focus on global catastrophic risk have already been set up through public and philanthropic funds. Among them are Oxford's Future of Humanity Institute and Cambridge's Centre for the Study of Existential Risk.

Beyond this, the research community should contribute to global catastrophic risk reduction by promoting a culture of safety within areas of research that could have the potential to cause a catastrophe through accident or misuse. This is especially relevant for emerging technologies, where it is not always clear in advance whether there are any risks.[186]

4.2.4. INDUSTRY

Competitive pressures mean that it is often hard for companies to make large moves on issues which do not improve their profits. They can, how-

ever, show leadership on reducing the risk of catastrophic climate change by choosing low carbon options at the margin, or by developing clean energy solutions. Just like the wider research community, they can and should promote a culture of safety in biotechnology and AI research as, for example, Google DeepMind has done by setting up an AI ethics board.[187]

4.2.5. GENERAL PUBLIC

Some global catastrophic risks can be addressed by individual action. For instance, individuals can decrease the risk of catastrophic climate change (if by ever so little) by making low-carbon consumption choices. Perhaps the more promising route to reduce global catastrophic risk for individuals is, however, by exerting political pressure on policymakers. For instance, voters could try to influence politicians to agree to carbon emissions cuts, and then to actually implement the agreements once they are in place.

4.2.6. THE NON-PROFIT SECTOR

The non-profit sector contributes significantly to global catastrophic risk reduction. Non-profits are often less constrained than companies and national governments, which means that they are free to work on intergenerational global public goods such as global catastrophic risk reduction -- if they can find donors that support that cause. There are many charities working on individual global catastrophic risks. In recent years, donors have been increasingly interested in global catastrophic risk reduction as a general category. This has led to the creation of groups such as the Future of Life Institute in Boston.

Besides supporting research, charities can also support direct interventions, e.g. to increase pandemic preparedness in developing countries.[188] They can also exert political pressure, as International Physicians for the Prevention of Nuclear War did to reduce the risk of a nuclear exchange during the cold war, winning the Nobel Peace Prize in 1985.[189]

The research community should contribute to global catastrophic risk reduction by promoting a culture of safety within areas of research that could have the potential to cause a catastrophe through accident or misuse.

Chapter 5

What can the world do to reduce global catastrophic risk?

In Chapter 2, we looked at different global catastrophic risks and their mechanisms. There was also some preliminary discussion of actions available to limit the risks. In Chapter 3 we explored the different factors which affect these risks. In Chapter 4 we considered why global catastrophic risk is probably neglected, and how different actors can help with it. Finally in this chapter we draw these strands together and outline a few of the most promising steps that existing communities can take or are already taking in order to reduce global catastrophic risk. Some of these steps pertain to individual risks, whereas others are cross-cutting opportunities which may reduce the chance or impact of several different risks at once.

Our aim here is to offer some starting points for considering action on risks of global catastrophe and to demonstrate that there are real avenues to making progress. For more detailed discussion of the actions available, there exists a rich literature on most of the specific risks.

Our aim here is to offer some starting points for considering action on risks of global catastrophe and to demonstrate that there are real avenues to making progress.

CATASTROPHIC CLIMATE CHANGE

- Research communities should increase their focus on understanding the pathways to and likelihood of catastrophic climate change, and possible ways to respond.[190]
- 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.[191]
- Developed nations could commit to the goal of building no new coal-fired power stations without carbon capture and sequestration.[192]

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.[193]
- Nuclear weapon states can adopt military strategies that reduce the chance of nuclear winter.[194]

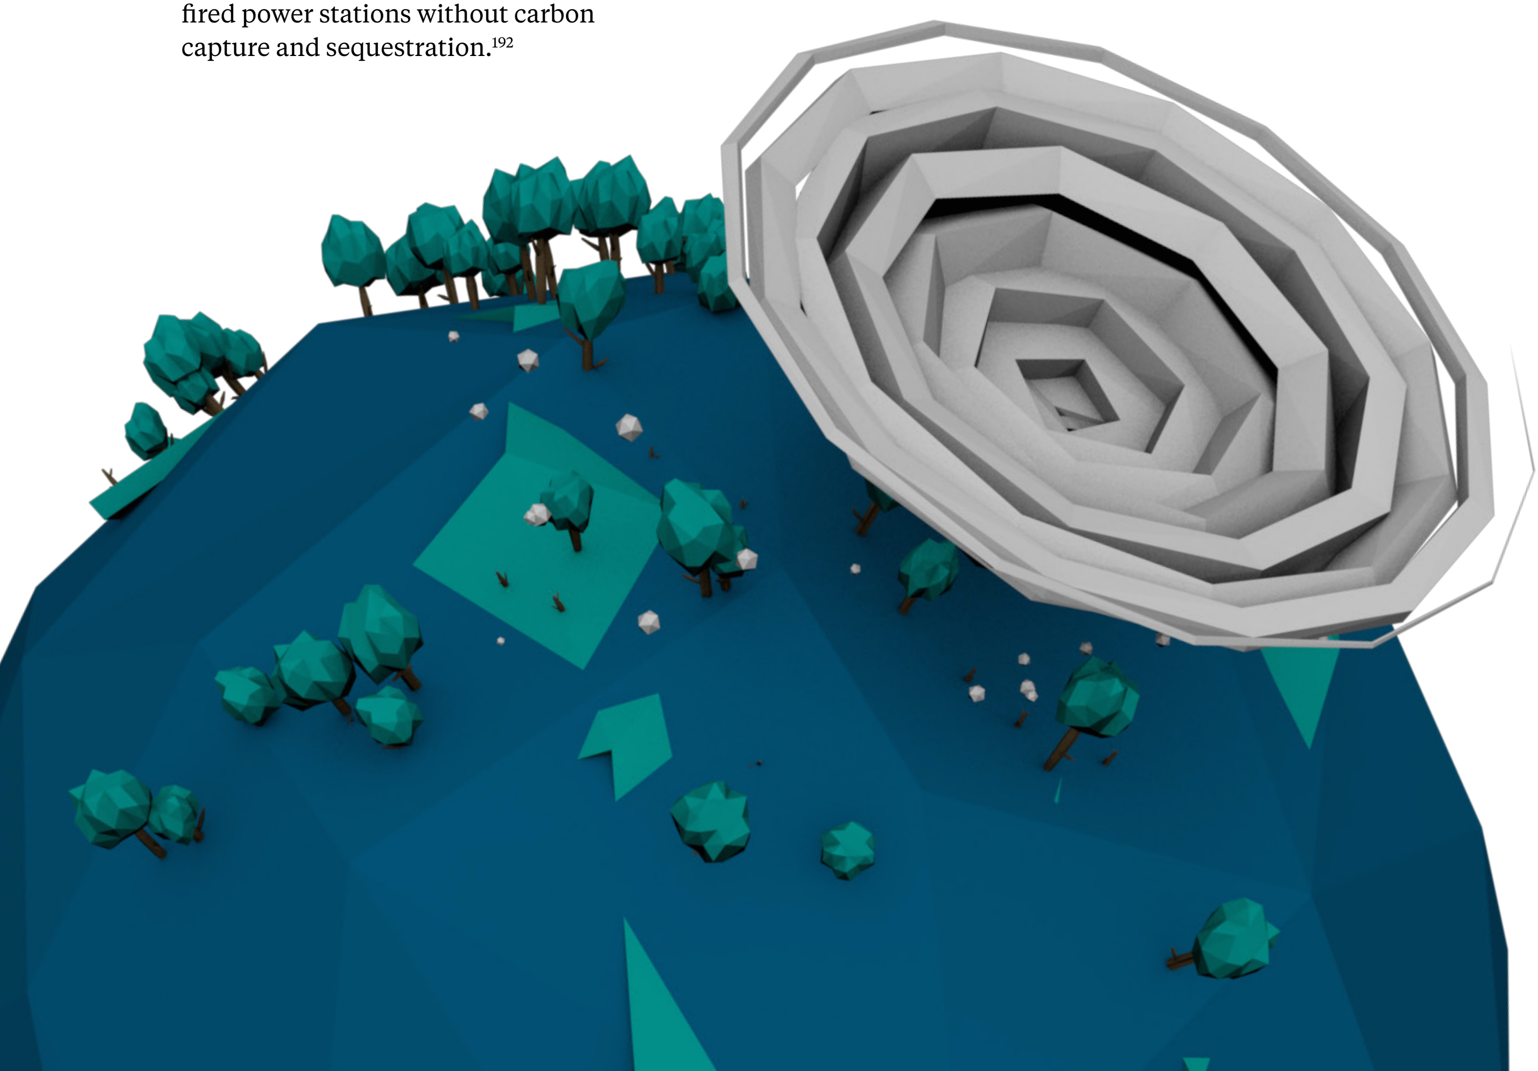

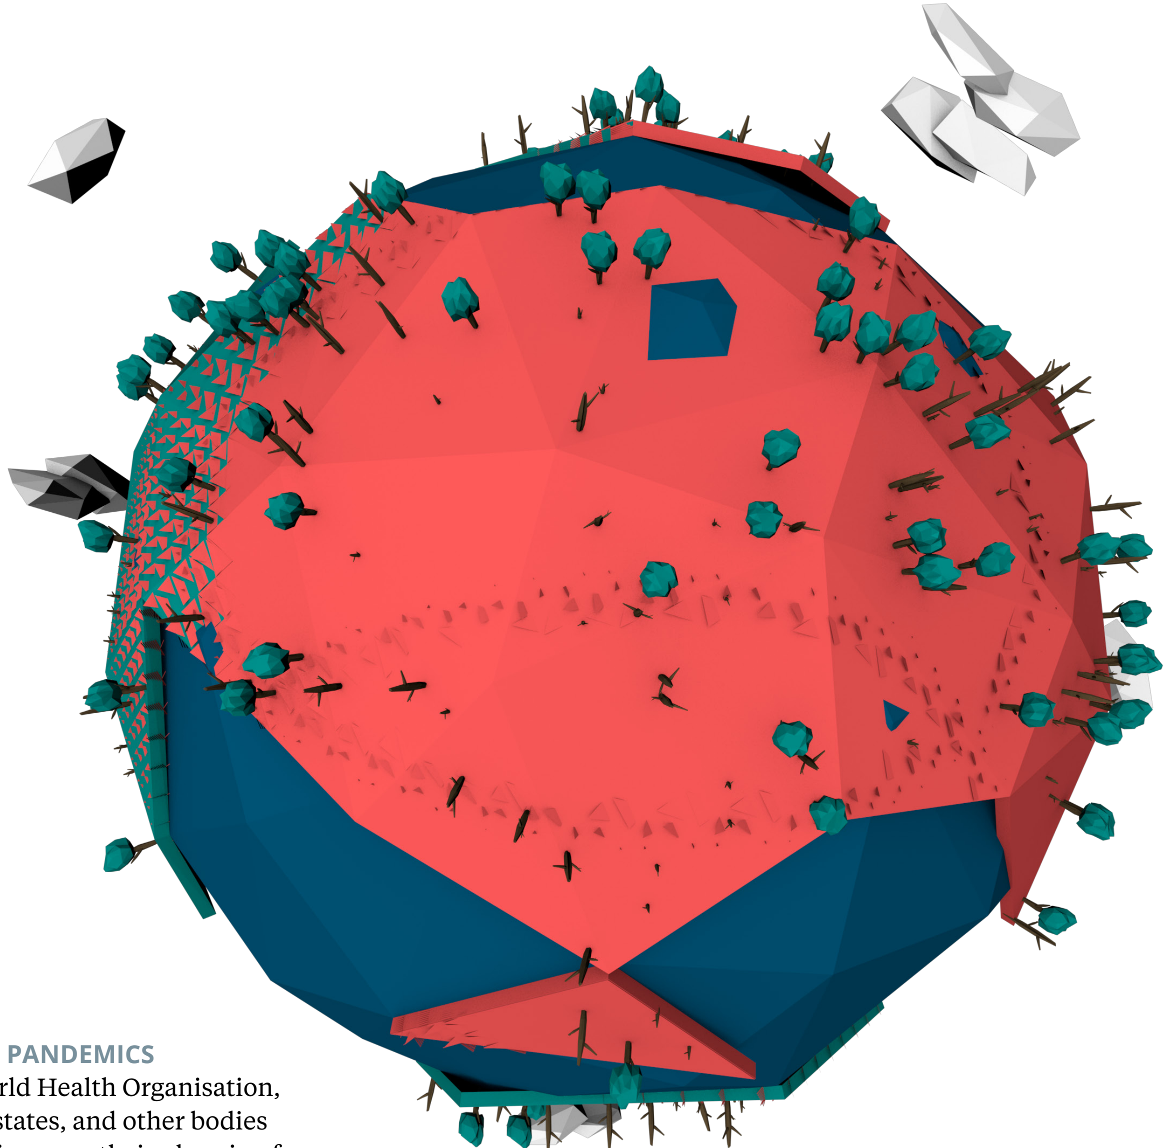

NATURAL PANDEMICS

- The World Health Organisation, nation states, and other bodies should increase their planning for extremely bad pandemics.[195]
- International and research communities could improve disease surveillance, for example by reducing the delay between scientific breakthrough and the availability of diagnostic tools.
- 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.[196]

ASTEROIDS AND COMETS

- Research communities should continue working to detect and track asteroids and commets with a diameter of 1 km or more.[197]

SUPERVOLCANIC ERUPTIONS

- Research communities should continue to work on understanding their causes, to increase predictability.[198]

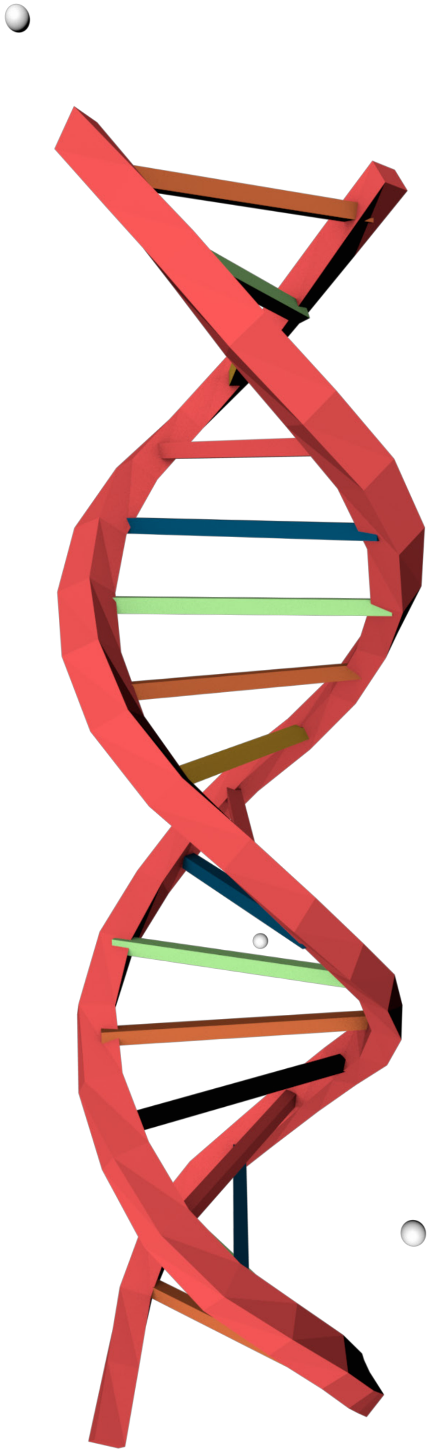

BIOTECHNOLOGY

- Research communities should further investigate the possible risks from emerging capabilities in biotechnology, and develop countermeasures.[199]
- Policymakers should continue to work with researchers to understand the biosafety and biosecurity issues that are likely to arise, and build planned adaptive risk regulation.[200]

ARTIFICIAL INTELLIGENCE

- Research communities should further investigate the possible risks from artificial intelligence, and work on developing possible solutions.[201]
- Policymakers can work with researchers to understand the implications of advanced artificial intelligence.

GEO-ENGINEERING

- The international community should continue work to stabilise and eventually reduce concentrations of greenhouse gases in the atmosphere in order to reduce the incentives to use solar radiation management.
- Research communities can further investigate the potential impacts of solar radiation management.[202]

UNKNOWN RISKS

- Research communities can continue to develop methods and tools for horizon scanning and reduction of unknown unknowns.
- Research communities should identify and carefully investigate speculative threats.

CROSS-CUTTING OPPORTUNITIES

- Research communities should further investigate 'particulate winter' scenarios, both in forecasting and mitigation strategies.[203]
- Nations and local communities can continue to take steps to build their resilience to catastrophe.[204]
- 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.[205]
- Nations should work to incorporate the interests of future generations into their decision-making frameworks.[206]

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.

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. We discuss this in more detail in chapter 4.
15. http://www.atomicarchive.com/Docs/Begin/Einstein.shtml
16. Ali Nouri and Christopher F. Chyba, "Biotechnology and Biosecurity," in *Global Catastrophic Risks*, ed. Nick Bostrom and Milan M. Ćirković (Oxford University Press, 2008).
17. Future of Life Institute, "AI Open Letter," accessed February 3, 2016, http://futureoflife.org/ai-open-letter/.
18. 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/.
19. Gernot Wagner and Martin L. Weitzman, *Climate Shock : The Economic Consequences of a Hotter Planet* (Princeton: Princeton University Press, 2015).
20. 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.
21. NASA, "The Relentless Rise of Carbon Dioxide,".
22. Sergey A. Zimov, Edward AG Schuur, and F. Stuart Chapin III, "Permafrost and the Global Carbon Budget," accessed January 18, 2016, http://www.imedea.uib-csic.es/master/cambioglobal/Modulo_V_cod101619/Permafrost%20response.pdf.
23. Christopher B. Field et al., Climate Change 2014: Impacts, Adaptation, and Vulnerability: Summary for Policymakers (Contribution of Working Group II to the Fifth Assessment Report of the Intergovernmental Panel on Climate Change, 2014), 14, http://epic.awi.de/37531/.
24. David King et al., "Climate Change–a Risk Assessment" (Centre for Science Policy, University of Cambridge, 2015), 45–46, www.csap.cam.ac.uk/projects/climate-change-risk-assessment/.
25. Gernot Wagner and Martin L. Weitzman, Climate Shock : The Economic Consequences of a Hotter Planet (Princeton: Princeton University Press, 2015), 53–56.
26. King et al., "Climate Change–a Risk Assessment," chap. 10.
27. King et al., "Climate Change–a Risk Assessment," pt. 1.
28. "Green Light," The Economist, December 19, 2015, http://www.economist.com/news/international/21684144-what-expect-after-deal-exceeded-expectations-green-light.
29. Rajendra K. Pachauri et al., *Climate Change 2014: Synthesis Report. Contribution of Working Groups I, II and III to the Fifth Assessment Report of the Intergovernmental Panel on Climate Change,* 2014, 10-13
30. Wagner and Weitzman, Climate Shock, 53–56.
31. Field et al., Climate Change 2014: Impacts, Adaptation, and Vulnerability: Summary for Policymakers, chap. 16.
32. Nicholas Stern, "The Economics of Climate Change," The American Economic Review 98, no. 2 (May 1, 2008): 1–37; William D. Nordhaus, "A Review of the 'Stern Review on the Economics of Climate Change,'" Journal of Economic Literature 45, no. 3 (September 1, 2007): 686–702; Wagner and Weitzman, Climate Shock.

33. Wagner and Weitzman, Climate Shock, chap. 3.

34. H. Damon Matthews et al., "The Proportionality of Global Warming to Cumulative Carbon Emissions," Nature 459, no. 7248 (June 11, 2009): 829–32.

35. Malte Meinshausen et al., "Greenhouse-Gas Emission Targets for Limiting Global Warming to 2 °C," *Nature* 458, no. 7242 (April 30, 2009): 1158–62.

36. Ibid.

37. J. G. Shepherd, Geoengineering the Climate: Science, Governance and Uncertainty (Royal Society, 2009), chap. 2.

38. Ibid., chap. 3.

39. *Gernot Wagner and Martin L. Weitzman, Climate Shock : The Economic Consequences of a Hotter Planet (Princeton: Princeton University Press, 2015), 54.*

40. See Martin E. Hellman, "How Risky Is Nuclear Optimism?," *Bulletin of the Atomic Scientists 67*, no. 2 (March 1, 2011): 47–56.

41. Bulletin of the Atomic Scientists, "Nuclear Notebook," Bulletin of the Atomic Scientists, accessed February 3, 2016, http://thebulletin.org/nuclear-notebook-multimedia.

42. Hans M. Kristensen and Robert S. Norris, "US Nuclear Forces, 2015," Bulletin of the Atomic Scientists 71, no. 2 (2015): 108.

43. Hans M. Kristensen and Robert S. Norris, "Russian Nuclear Forces, 2015," Bulletin of the Atomic Scientists 71, no. 3 (2015): 85.

44. "A Nuke by Any Other Name," Bulletin of the Atomic Scientists, May 17, 2012, http://thebulletin.org/nuke-any-other-name.

45. US Office of Technology Assessment, "The Effects of Nuclear War," 1979, 8–9, http://ota.fas.org/reports/7906.pdf; Joseph Cirincione, "The Continuing Threat of Nuclear War," in Global Catastrophic Risks, ed. Nick Bostrom and Milan M. Ćirković (Oxford University Press, 2008), 386–390.

46. This area has received less attention since the end of the Cold War. Other estimates give death tolls within the same order of magnitude. Ira Helfand et al., "Projected US Casualties and Destruction of US Medical Services from Attacks by Russian Nuclear Forces," Medicine & Global Survival 7, no. 2 (2002): 68–76; WHO, "Health Effects of the Chernobyl Accident and Special Health Care Programmes" (WHO, 2006), http://www.who.int/ionizing_radiation/chernobyl/who_chernobyl_report_2006.pdf; Owen B. Toon et al., "Atmospheric Effects and Societal Consequences of Regional Scale Nuclear Conflicts and Acts of Individual Nuclear Terrorism," Atmospheric Chemistry and Physics 7, no. 8 (2007): 1973–2002.

47. For a good summary see Seth D. Baum, "Winter-Safe Deterrence: The Risk of Nuclear Winter and Its Challenge to Deterrence," Contemporary Security Policy 36, no. 1 (January 2, 2015): 126, doi:10.1080/13523260.2015.1012346.

48. Alan Robock and Owen Brian Toon, "Local Nuclear War, Global Suffering," Scientific American 302, no. 1 (2010): 74–81. See the discussion in Givewell, "Nuclear Weapons Policy," September 2015, http://www.givewell.org/labs/causes/nuclear-weapons-policy.

49. Martin E. Hellman, "How Risky Is Nuclear Optimism?"

50. Anthony M. Barrett, Seth D. Baum, and Kelly Hostetler, "Analyzing and Reducing the Risks of Inadvertent Nuclear War Between the United States and Russia," Science & Global Security 21, no. 2 (May 1, 2013): 106–33, doi:10.1080/08929882.2013.798984.

51. Graham Allison, "At 50, the Cuban Missile Crisis as Guide," The New York Times, June 15, 2012.

52. Zachary Keck, "Russia Threatens Nuclear Strikes Over Crimea," July 11, 2014; National Institute for Public Policy, "Russia's Nuclear Posture" (Fairfax, VA, 2015).

53. Hellman, "How Risky Is Nuclear Optimism?"; Tom Hussain, "Are India and Pakistan Heading for a Nuclear Showdown?," Al Jazeera, March 2016, http://www.aljazeera.com/indepth/opinion/2016/03/india-pakistan-heading-nuclear-showdown-160303053541342.html.

54. Michael Krepon, "Pakistan's Nuclear Strategy and Deterrence Stability," Stimson Center 10 (2012), http://indianstrategicknowledgeonline.com/web/Krepon_-_Pakistan_Nuclear_Strategy_and_Deterrence_Stability.pdf.

55. See Cirincione, "The Continuing Threat of Nuclear War"; Hellman, "How Risky Is Nuclear Optimism?"

56. Cirincione, "The Continuing Threat of Nuclear War," 396.

57. Matthew R. Costlow, "Do More Nukes Really Mean More Nuclear Crises? Not Necessarily," Bulletin of the Atomic Scientists, December 8, 2015, http://thebulletin.org/do-more-nukes-really-mean-more-nuclear-crises-not-necessarily8965.

58. For other suggestions for weapons systems see Seth D. Baum, "Confronting the Threat of Nuclear Winter," Futures, Confronting Future Catastrophic Threats To Humanity, 72 (September 2015): 69–79, doi:10.1016/j.futures.2015.03.004.

59. http://thebulletin.org/nuclear-notebook-multimedia

60. See the work of the Nuclear Threat Initiative at http://www.nti.org/.

61. Alan Robock, Luke Oman, and Georgiy L. Stenchikov, "Nuclear Winter Revisited with a Modern Climate Model and Current Nuclear Arsenals: Still Catastrophic Consequences," Journal of Geophysical Research: Atmospheres 112, no. D13 (July 16, 2007): 10, doi:10.1029/2006JD008235.

62. Ole J. Benedictow, "The Black Death: The Greatest Catastrophe Ever," History Today, 2005, http://www.historytoday.com/ole-j-benedictow/black-death-greatest-catastrophe-ever; Edwin Dennis Kilbourne, "Plagues and Pandemics: Past, Present, and Future," in Global Catastrophic Risks, ed. Nick Bostrom and Milan M. Ćirković (Oxford: Oxford University Press, 2008), 295. Percentages deduced using population data 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.

63. 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.

64. WHO, "Current WHO Phase of Pandemic Alert for Avian Influenza H5N1," WHO, accessed February 15, 2016, http://apps.who.int/csr/disease/avian_influenza/phase/en/index.html; Declan Butler, "Death-Rate Row Blurs

Mutant Flu Debate," Nature 482, no. 7385 (February 13, 2012): 289–289, doi:10.1038/482289a.

65. Marc Lipsitch and Thomas V. Inglesby, "Moratorium on Research Intended To Create Novel Potential Pandemic Pathogens," mBio 5, no. 6 (December 31, 2014): e02366–14, doi:10.1128/mBio.02366-14.

66. Butler, "Death-Rate Row Blurs Mutant Flu Debate."

67. Christopher JL Murray et al., "Estimation of Potential Global Pandemic Influenza Mortality on the Basis of Vital Registry Data from the 1918–20 Pandemic: A Quantitative Analysis," The Lancet 368, no. 9554 (December 23, 2006): 2211–18, doi:10.1016/S0140-6736(06)69895-4.

68. Michael T. Osterholm, "Preparing for the Next Pandemic," New England Journal of Medicine 352, no. 18 (May 5, 2005): 1839–42, doi:10.1056/NEJMp058068; Neil M. Ferguson et al., "Strategies for Mitigating an Influenza Pandemic," Nature 442, no. 7101 (July 27, 2006): 448–52, doi:10.1038/nature04795.

69. Bruine de Bruin, Wändi, et al. "Expert judgments of pandemic influenza risks." Global Public Health 1.2 (2006): 179-194.

70. For two recent overviews see Peter Sands, Carmen Mundaca-Shah, and Victor J. Dzau, "The Neglected Dimension of Global Security — A Framework for Countering Infectious-Disease Crises," New England Journal of Medicine 0, no. 0 (January 13, 2016): null, doi:10.1056/NEJMsr1600236; WHO and Report of the High-level Panel on the Global Response to Health Crises, "Protecting Humanity from Future Health Crises" (WHO, January 25, 2016), http://www.un.org/News/dh/infocus/HLP/2016-02-05_Final_Report_Global_Response_to_Health_Crises.pdf.

71. Stanley A. Plotkin, Adel A.F. Mahmoud, and Jeremy Farrar, "Establishing a Global Vaccine-Development Fund," *New England Journal of Medicine 373*, no. 4 (July 23, 2015): 297–300, doi:10.1056/NEJMp1506820.

72. Catherine Yen et al., "The Development of Global Vaccine Stockpiles," The Lancet Infectious Diseases 15, no. 3 (March 2015): 340–47, doi:10.1016/S1473-3099(14)70999-5.

73. WHO, "International Health Regulations: Support to Global Outbreak Alert and Response, and Building and Maintaining National Capacities," 2015, http://apps.who.int/iris/bitstream/10665/199747/1/WHO_HSE_GCR_2015.7_eng.pdf; WHO, "Pandemic Influenza Preparedness Framework," 2011, http://www.who.int/influenza/resources/pip_framework/en/.

74. Emily H. Chan et al., "Global Capacity for Emerging Infectious Disease Detection," Proceedings of the National Academy of Sciences 107, no. 50 (December 14, 2010): 21701–6, doi:10.1073/pnas.1006219107.

75. Sands, Mundaca-Shah, and Dzau, "The Neglected Dimension of Global Security — A Framework for Countering Infectious-Disease Crises."

76. Edwin Dennis Kilbourne, "Plagues and Pandemics: Past, Present, and Future," in Global Catastrophic Risks, ed. Nick Bostrom and Milan M. Ćirković (Oxford: Oxford University Press, 2008); William Rosen, Justinian's Flea : Plague, Empire, and the Birth of Europe (New York: Viking, 2007); 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; Ole J. Benedictow, "The Black Death: The Greatest Catastrophe Ever," History Today, 2005, http://www.historytoday.com/ole-j-benedictow/black-death-greatest-catastrophe-ever; WHO and others, Bugs, Drugs and Smoke: Stories from Public Health, 2011, chap. 1, http://apps.who.int/iris/handle/10665/44700; WHO, "HIV/AIDS Fact Sheet," WHO, November 2015, http://www.who.int/mediacentre/factsheets/fs360/en/.

77. Arnon Dar, "Influence of Supernovae, Gamma-Ray Bursts, Solar Flares, and Cosmic Rays on the Terrestrial Environment," in Global Catastrophic Risks, ed. Nick Bostrom and Milan M. Ćirković (Oxford: Oxford University Press, 2008).

78. See the draft manuscript Toby Ord, "Will We Cause Our Own Extinction? Natural versus Anthropogenic Extinction Risks" 2014. Ord's argument is about extinction risks and not the broader category of global catastrophic risks, but it is very plausible that similar arguments would also show that exogenous non-extinction global catastrophic risks are less likely than anthropogenic non-extinction global catastrophic risks.

79. Michael Rampino, "Super-Volcanism and Other Geophysical Processes of Catastrophic Import," in Global Catastrophic Risks, ed. Nick Bostrom and Milan M. Ćirković (Oxford: Oxford University Press, 2008).

80. G. A. Zielinski et al., "Potential Atmospheric Impact of the Toba Mega-Eruption 71,000 Years Ago," Geophysical Research Letters 23, no. 8 (April 15, 1996): 837–40, doi:10.1029/96GL00706; Rampino, "Super-Volcanism and Other Geophysical Processes of Catastrophic Import," 209–210.

81. Rampino, "Super-Volcanism and Other Geophysical Processes of Catastrophic Import," 211–212; Christine S. Lane, Ben T. Chorn, and Thomas C. Johnson, "Ash from the Toba Supereruption in Lake Malawi Shows No Volcanic Winter in East Africa at 75 Ka," Proceedings of the National Academy of Sciences 110, no. 20 (May 14, 2013): 8025–29, doi:10.1073/pnas.1301474110.

82. S. Sparks et al., Super-Eruptions: Global Effects and Future Threats. London, Report of a Geological Society of London Working Group, 2005, 6.

83. W. Aspinall et al., "Volcano Hazard and Exposure in GFDRR Priority Countries and Risk Mitigation Measures," Volcano Risk Study 0100806-00-1-R, 2011, 15; Rampino, "Super-Volcanism and Other Geophysical Processes of Catastrophic Import," 212–213; Susan Loughlin et al., Global Volcanic Hazards and Risk (Cambridge University Press, 2015), 97.

84. Sparks et al., *Super-Eruptions*, 20.

85. Peter Schulte et al., "The Chicxulub Asteroid Impact and Mass Extinction at the Cretaceous-Paleogene Boundary," Science 327, no. 5970 (March 5, 2010): 1214–18, doi:10.1126/science.1177265.

86. National Research Council (U. S.). Committee to Review Near-Earth-Object Surveys and Hazard Mitigation Strategies, Defending Planet Earth: Near-Earth Object Surveys and Hazard Mitigation Strategies (Washington, DC: National Academies Press, 2010), 21.

87. National Research Council (U. S.). Committee to Review Near-Earth-Object Surveys and Hazard Mitigation Strategies, Defending Planet Earth, 21.

88. Clark R. Chapman, "The Hazard of near-Earth Asteroid Impacts on Earth," Earth and Planetary Science Letters 222, no. 1 (May 15, 2004): 11, doi:10.1016/j.epsl.2004.03.004.
89. National Research Council (U. S.). Committee to Review Near-Earth-Object Surveys and Hazard Mitigation Strategies, Defending Planet Earth, 23.
90. Alan Harris, "What Spaceguard Did," Nature 453, no. 7199 (June 26, 2008): 1178–79, doi:10.1038/4531178a; NASA, "NASA Space Telescope Finds Fewer Asteroids near Earth," September 2011, http://www.nasa.gov/mission_pages/WISE/news/wise20110929.html.
91. Jason C. Reinhardt et al., "Asteroid Risk Assessment: A Probabilistic Approach," Risk Analysis, July 1, 2015, 1-18, doi:10.1111/risa.12453.
92. Source: Alan Harris, "What Spaceguard Did," Nature 453, no. 7199 (June 26, 2008): 1178–79, doi:10.1038/4531178a.
93. National Research Council (U. S.). Committee to Review Near-Earth-Object Surveys and Hazard Mitigation Strategies, Defending Planet Earth, 23.
94. Ibid., 22.
95. Harris, "What Spaceguard Did."
96. A. Harris et al., "Asteroid Impacts and Modern Civilization: Can We Prevent a Catastrophe," Asteroids IV. University of Arizona Press, Tucson, AZ, 2015, http://elib.dlr.de/100120/1/3004_Harris_reprint.pdf.
97. Robert Carlson, "The Changing Economics of DNA Synthesis," Nature Biotechnology 27, no. 12 (December 2009): 1091–94, doi:10.1038/nbt1209-1091.
98. Terrence M. Tumpey et al., "Characterization of the Reconstructed 1918 Spanish Influenza Pandemic Virus," Science 310, no. 5745 (2005): 77–80.
99. Markus Schmidt, "Diffusion of Synthetic Biology: A Challenge to Biosafety," Systems and Synthetic Biology 2, no. 1–2 (June 2008): 1–6, doi:10.1007/s11693-008-9018-z.
100. Ronald J. Jackson et al., "Expression of Mouse Interleukin-4 by a Recombinant Ectromelia Virus Suppresses Cytolytic Lymphocyte Responses and Overcomes Genetic Resistance to Mousepox," Journal of Virology 75, no. 3 (2001): 1205–10.
101. Christopher F. Chyba and Alex L. Greninger, "Biotechnology and Bioterrorism: An Unprecedented World," Survival 46, no. 2 (2004): 148–149.
102. Sander Herfst et al., "Airborne Transmission of Influenza A/H5N1 Virus between Ferrets," Science 336, no. 6088 (2012): 1534–41; Masaki Imai et al., "Experimental Adaptation of an Influenza H5 HA Confers Respiratory Droplet Transmission to a Reassortant H5 HA/H1N1 Virus in Ferrets," Nature 486, no. 7403 (June 21, 2012): 420–28, doi:10.1038/nature10831.
103. Nicholas Greig Evans, Marc Lipsitch, and Meira Levinson, "The Ethics of Biosafety Considerations in Gain-of-Function Research Resulting in the Creation of Potential Pandemic Pathogens," Journal of Medical Ethics 41, no. 11 (November 1, 2015): 901–8, doi:10.1136/medethics-2014-102619.
104. Richard A. Posner, Catastrophe : Risk and Response (Oxford: Oxford University Press, 2004), 78–79.
105. George Church, "A Synthetic Biohazard Non-Proliferation Proposal," 2004, http://arep.med.harvard.edu/SBP/Church_Biohazard04c.htm.
106. Nathan Myhrvold, "Strategic Terrorism: A Call to Action," SSRN Scholarly Paper (Rochester, NY: Social Science Research Network, July 3, 2013), http://papers.ssrn.com/abstract=2290382.
107. Martin J. Rees, Our Final Century : Will Civilisation Survive the Twenty-First Century? (London: Arrow, 2004).
108. Marc Lipsitch and Alison P. Galvani, "Ethical Alternatives to Experiments with Novel Potential Pandemic Pathogens," PLoS Med 11, no. 5 (May 20, 2014): e1001646, doi:10.1371/journal.pmed.1001646.
109. Lynn C. Klotz and Edward J. Sylvester, "The Unacceptable Risks of a Man-Made Pandemic," Bulletin of the Atomic Scientists, August 7, 2012, http://thebulletin.org/unacceptable-risks-man-made-pandemic. http://www.the-scientist.com/?articles.view/articleNo/41263/title/Moratorium-on-Gain-of-Function-Research/
110. Lipsitch and Inglesby, "Moratorium on Research Intended To Create Novel Potential Pandemic Pathogens," 2.
111. Ibid.
112. Gautam Mukunda, Kenneth A. Oye, and Scott C. Mohr, "What Rough Beast?," Politics and the Life Sciences 28, no. 2 (September 1, 2009): 8, doi:10.2990/28_2_2.
113. Church, "A Synthetic Biohazard Non-Proliferation Proposal."
114. Owen Cotton-Barratt, Sebastian Farquhar, and Andrew Snyder-Beattie, "Beyond Risk-Benefit Analysis: Pricing Externalities for Gain-of-Function Research of Concern" (Future of Humanity Institute, Global Piorities Project, March 2016).
115. Chyba and Greninger, "Biotechnology and Bioterrorism," 151–153.
116. "DNA Sequencing Costs," accessed March 31, 2016, http://www.genome.gov/sequencingcosts/.
117. Future of Life Institute, "AI Open Letter," accessed February 3, 2016, http://futureoflife.org/ai-open-letter/.
118. Eliezer Yudkowsky, "Artificial Intelligence as a Positive and Negative Factor in Global Risk," in *Global Catastrophic Risks,* ed. Nick Bostrom and Milan M. Ćirković, vol. 1 (Oxford University Press, 2008), 303.
119. Carl Benedikt Frey and Michael A. Osborne, "The Future of Employment" 2013; Nick Bostrom, Superintelligence : Paths, Dangers, Strategies (Oxford: Oxford University Press, 2014), chap. 11.
120. Bostrom, Superintelligence, chap. 4.
121. For extensive discussion of the 'AI control problem' see Bostrom, Superintelligence.
122. Bostrom, Superintelligence, 18–21. For more detail about the surveys see Vincent C. Müller and Nick Bostrom, "Future Progress in Artificial Intelligence: A Survey of Expert Opinion," in Fundamental Issues of Artificial Intelligence, ed. Vincent C. Müller (Berlin: Synthese Library; Springer, 2014).
123. Bostrom, Superintelligence, 18–21
124. Bostrom, Superintelligence, 20–21.
125. Shepherd, Geoengineering the Climate, chap. 2.
126. Shepherd, Geoengineering the Climate.
127. Alan Robock, Kirsten Jerch, and Martin Bunzl, "20 Reasons

Why Geoengineering May Be a Bad Idea," Bulletin of the Atomic Scientists 64, no. 2 (2008): 14–59.

128. Shepherd, Geoengineering the Climate, chap. 3. For criticism of the 'climate emergency' argument for SRM see Oliver Morton, The Planet Remade : How Geoengineering Could Change the World (London: Granta, 2015), chap. 6; Jana Sillmann et al., "Climate Emergencies Do Not Justify Engineering the Climate," Nature Climate Change 5, no. 4 (2015): 290–92.

129. Scott Barrett, "The Incredible Economics of Geoengineering," Environmental and Resource Economics 39, no. 1 (2008): 45–54.

130. Alan Robock, Luke Oman, and Georgiy L. Stenchikov, "Regional Climate Responses to Geoengineering with Tropical and Arctic SO2 Injections," Journal of Geophysical Research: Atmospheres 113, no. D16 (August 27, 2008): D16101, doi:10.1029/2008JD010050.

131. Morton, The Planet Remade, chap. 4..

132. Shepherd, Geoengineering the Climate, chap. 4.

133. Ibid., chap. 3.

134. Morton, The Planet Remade, chap. 4.

135. Seth D. Baum, Timothy M. Maher Jr, and Jacob Haqq-Misra, "Double Catastrophe: Intermittent Stratospheric Geoengineering Induced by Societal Collapse," Environment Systems & Decisions 33, no. 1 (January 8, 2013): 168–80, doi:10.1007/s10669-012-9429-y.

136. Shepherd, Geoengineering the Climate, chap. 4.

137. See for example Shepherd, *Geoengineering the Climate*; National Academy of Sciences, *Climate Intervention*: Reflecting Sunlight to Cool Earth (Washington, D.C.: National Academies Press, 2015), http://www.nap.edu/catalog/18988; Solar Radiation Management Governance Initiative et al., *Solar Radiation Management: The Governance of Research* (SRMGI, 2012).

138. Shepherd, *Geoengineering the Climate*; Barrett, "The Incredible Economics of Geoengineering."

139. Arnon Dar, "Influence of Supernovae, Gamma-Ray Bursts, Solar Flares, and Cosmic Rays on the Terrestrial Environment," in Global Catastrophic Risks, ed. Nick Bostrom and Milan M. Ćirković (Oxford: Oxford University Press, 2008).

140. Bryan Caplan, "The Totalitarian Threat," in Global Catastrophic Risks, ed. Nick Bostrom and Milan M. Ćirković (Oxford: Oxford University Press, 2008).

141. Seth D. Baum, "Winter-Safe Deterrence: The Risk of Nuclear Winter and Its Challenge to Deterrence," Contemporary Security Policy 36, no. 1 (January 2, 2015): 123–48, doi:10.1080/13523260.2015.1012346.

142. Joseph Cirincione, "The Continuing Threat of Nuclear War," in Global Catastrophic Risks, ed. Nick Bostrom and Milan M. Ćirković (Oxford University Press, 2008).

143. Pavel Podvig, "Reducing the Risk of an Accidental Launch," Science & Global Security 14, no. 2–3 (December 1, 2006): 75–115, doi:10.1080/08929880600992990.

144. For an overview of the literature on 'nuclear winter' see Seth D. Baum, "Winter-Safe Deterrence: The Risk of Nuclear Winter and Its Challenge to Deterrence," Contemporary Security Policy 36, no. 1 (January 2, 2015): 123–48, doi:10.1080/13523260.2015.1012346. The resilience of the food supply is also an important risk factor, and is discussed in section 3.2.

145. Centre for Science Policy, University of Cambridge et al., "Climate Change–a Risk Assessment," 38; International Energy Agency and Nuclear Energy Agency, "Technology Roadmap: Nuclear Energy," 2015; Ottmar Edenhofer and Intergovernmental Panel on Climate Change. Working Group III, Climate Change 2014 : Mitigation of Climate Change : Working Group III Contribution to the Fifth Assessment Report of the Intergovernmental Panel on Climate Change (New York, NY: Cambridge University Press, 2015), 516–518.

146. Edenhofer and Intergovernmental Panel on Climate Change. Working Group III, Climate Change 2014, 531.

147. According to Sailor et al "development of nuclear weapons has been aided in at least three countries (India, Iraq, and Israel) by use of research reactors obtained under the cover of peaceful research programs".William C. Sailor et al., "A Nuclear Solution to Climate Change?," Science 288, no. 5469 (2000): 1177–78. For a very good overview of the technical and political aspects of the impact of nuclear power on nuclear proliferation see A. Adamantiades and I. Kessides, "Nuclear Power for Sustainable Development: Current Status and Future Prospects," Energy Policy 37, no. 12 (December 2009): 5149–66, doi:10.1016/j.enpol.2009.07.052.

148. Edwin Dennis Kilbourne, "Plagues and Pandemics: Past, Present, and Future," in Global Catastrophic Risks, ed. Nick Bostrom and Milan M. Ćirković (Oxford: Oxford University Press, 2008).

149. James M. Hughes et al., "The Origin and Prevention of Pandemics," Clinical Infectious Diseases 50, no. 12 (June 15, 2010): 1636–40, doi:10.1086/652860.

150. Kate E. Jones et al., "Global Trends in Emerging Infectious Diseases," Nature 451, no. 7181 (February 21, 2008): 990–93, doi:10.1038/nature06536.

151. Peter Sands, Carmen Mundaca-Shah, and Victor J. Dzau, "The Neglected Dimension of Global Security — A Framework for Countering Infectious-Disease Crises," New England Journal of Medicine 0, no. 0 (January 13, 2016): null, doi:10.1056/NEJMsr1600236.

152. David Charles Denkenberger and Joshua Pearce, Feeding Everyone No Matter What : Managing Food Security after Global Catastrophe (Amsterdam: Academic Press, 2015).

153. S. Sparks et al., Super-Eruptions: Global Effects and Future Threats. London, Report of a Geological Society of London Working Group, 2005; National Research Council (U. S.). Committee to Review Near-Earth-Object Surveys and Hazard Mitigation Strategies, Defending Planet Earth: Near-Earth Object Surveys and Hazard Mitigation Strategies (Washington, DC: National Academies Press, 2010), chap. 2 and 5.

154. David King et al., "Climate Change–a Risk Assessment" (Centre for Science Policy, University of Cambridge, 2015), pt. 1, www.csap.cam.ac.uk/projects/climate-change-risk-assessment/.

155. Sybren Drijfhout et al., "Catalogue of Abrupt Shifts in Intergovernmental Panel on Climate Change Climate Models," Proceedings of the National Academy of Sciences 112, no. 43 (2015): E5777–86

156. For a qualified defence of this position see Oliver Morton, The Planet Remade : How Geoengineering Could Change the World (London: Granta, 2015).

157. J. G. Shepherd, Geoengineering the Climate: Science, Governance and Uncertainty (Royal Society, 2009).

158. For discussions of worries about arms races and emerging technologies see Christopher F. Chyba, "Biotechnology and the Challenge to Arms Control," Arms Control Today 36, no. 8 (2006), http://search.proquest.com/openview/3251 1e461cdf35678bea5679269ac45b/1?pq-origsite=gscholar; Future of Life Institute, "Open Letter on Autonomous Weapons," accessed March 18, 2016, http://futureoflife.org/open-letter-autonomous-weapons/; Gernot Wagner and Martin L. Weitzman, Climate Shock : The Economic Consequences of a Hotter Planet (Princeton: Princeton University Press, 2015), chap. 3.

159. Nick Bostrom, Superintelligence : Paths, Dangers, Strategies (Oxford: Oxford University Press, 2014); Ali Nouri and Christopher F. Chyba, "Biotechnology and Biosecurity," in Global Catastrophic Risks, ed. Nick Bostrom and Milan M. Ćirković (Oxford University Press, 2008).

160. For a discussion of this issue with respect to biotechnology and geo-engineering risks respectively see W. Paul Duprex et al., "Gain-of-Function Experiments: Time for a Real Debate," Nature Reviews Microbiology 13, no. 1 (January 2015): 58–64, doi:10.1038/nrmicro3405; Oliver Morton, The Planet Remade : How Geoengineering Could Change the World (London: Granta, 2015).

161. Bostrom, Superintelligence.

162. Bostrom, N. "Analyzing Human Extinction Scenarios and Related Hazards". Journal of Evolution and Technology. Vol. 9. No. 1. (2002)

163. Philip Nel and Marjolein Righarts, "Natural Disasters and the Risk of Violent Civil Conflict," International Studies Quarterly 52, no. 1 (March 1, 2008): 159–85, doi:10.1111/j.1468-2478.2007.00495.x.

164. For an overview see Halvard Buhaug, Nils Petter Gleditsch, and Ole Magnus Theisen, "Implications of Climate Change for Armed Conflict" (World Bank, 2009); Ragnhild Nordås and Nils Petter Gleditsch, "Climate Change and Conflict," Political Geography, Climate Change and Conflict, 26, no. 6 (August 2007): 627–38, doi:10.1016/j.polgeo.2007.06.003.

165. BBC, "Working Gun Made with 3D Printer," BBC News, accessed February 22, 2016, http://www.bbc.co.uk/news/science-environment-22421185.

166. See Chapter 2 and Alan Robock, Luke Oman, and Georgiy L. Stenchikov, "Nuclear Winter Revisited with a Modern Climate Model and Current Nuclear Arsenals: Still Catastrophic Consequences," Journal of Geophysical Research: Atmospheres 112, no. D13 (July 16, 2007): D13107, doi:10.1029/2006JD008235; National Research Council (U. S.). Committee to Review Near-Earth-Object Surveys and Hazard Mitigation Strategies, Defending Planet Earth: Near-Earth Object Surveys and Hazard Mitigation Strategies (Washington, DC: National Academies Press, 2010), chap. 2; S. Sparks et al., Super-Eruptions: Global Effects and Future Threats. London, Report of a Geological Society of London Working Group, 2005.

167. Denkenberger and Pearce, Feeding Everyone No Matter What.

168. For an excellent discussion of global public goods see Scott Barrett, Why Cooperate? : The Incentive to Supply Global Public Goods (Oxford: Oxford University Press, 2007).

169. Clean air can of course also be influenced by pollution in neighbouring countries.

170. Elinor Ostrom, Governing the Commons: the evolution of institutions for collective action. 1990. Cambridge University Press.

171. It is hotly disputed how we should go about taking the interests of future generations into account, but it seems plausible that they should be given at least some consideration in public policy. Indeed, this is a standard assumption in many governmental cost-benefit analyses.

172. Mancur Olson, The Logic of Collective Action : Public Goods and the Theory of Groups, Harvard Economic Studies ; v. 124 (Cambridge, Mass: Harvard University Press, 1965).

173. This is an example of availability bias, which is discussed along with other cognitive biases that can skew perceptions of global risks in Eliezer Yudkowsky, "Cognitive Biases Potentially Affecting Judgment of Global Risks," in Global Catastrophic Risks, ed. Nick Bostrom and Milan M. Ćirković, vol. 1 (Oxford: Oxford University Press, 2008).

174. Barrett, Why Cooperate?, chap. 2.

175. WHO, "Summary of States Parties 2013 Report on IHR Core Capacity Implementation," 2014, 6, http://www.who.int/ihr/publications/WHO_HSE_GCR_2014.10/en/.

176. Lawrence O. Gostin and Devi Sridhar, "Global Health and the Law," New England Journal of Medicine 370, no. 18 (May 1, 2014): 1732–40, doi:10.1056/NEJMra1314094.

177. Barrett, Why Cooperate?

178. Treaty on the Non-Proliferation of Nuclear Weapons, Washington/Moscow/London, 1 July 1968, in force 5 March 1970, 729 UNTS 161

179. United Nations Framework Convention on Climate Change, Rio de Janeiro, 9 May 1992, in force 21 March 1994, 1771 UNTS 107. This convention provided the platform for the Kyoto and Paris climate change agreements.

180. The IAEA's mandate is limited to civil nuclear applications, but it has probably reduced the likelihood of nuclear terrorism by imposing detailed nuclear fissile accounting. This verification process as well as on-site inspection also make it very difficult for non-nuclear-weapon states to develop nuclear weapons without detection.

181. United Nations General Assembly, "Protection of Global Climate for Present and Future Generations of Mankind," December 1988, http://www.ipcc.ch/docs/UNGA43-53.pdf.

182. Gernot Wagner and Martin L. Weitzman, Climate Shock : The Economic Consequences of a Hotter Planet (Princeton: Princeton University Press, 2015).

183. See for example Alan Robock, "Nuclear Winter," Wiley Interdisciplinary Reviews: Climate Change 1, no. 3 (May 1, 2010): 418–27, doi:10.1002/wcc.45.

184. Nick Bostrom, Superintelligence : Paths, Dangers, Strategies (Oxford: Oxford University Press, 2014).

185. For example, see the discussion of the responsibility of research laboratories conducting gain of function research in W. Paul Duprex et al., "Gain-of-Function Experiments: Time for a Real Debate," Nature

Reviews Microbiology 13, no. 1 (January 2015): 58–64, doi:10.1038/nrmicro3405.

186. Charles Arthur, "DeepMind: 'Artificial Intelligence Is a Tool That Humans Can Control and Direct,'" The Guardian, June 9, 2015, sec. Technology, http://www.theguardian.com/technology/2015/jun/09/deepmind-artificial-intelligence-tool-humans-control.

187. Anja Wolz, "Face to Face with Ebola — An Emergency Care Center in Sierra Leone," New England Journal of Medicine 371, no. 12 (September 18, 2014): 1081–83, doi:10.1056/NEJMp1410179.

188. T. Frängsmyr, and I. Abrams, I., 1997. Peace, 1981-1990 (Vol. 5). World Scientific.

189. National Research Council (U. S.). Committee to Review Near-Earth-Object Surveys and Hazard Mitigation Strategies, Defending Planet Earth: Near-Earth Object Surveys and Hazard Mitigation Strategies (Washington, DC: National Academies Press, 2010), chap. 2.

190. For a defence of the importance of tail risk climate change see Martin L. Weitzman, "A Review of 'The Stern Review on the Economics of Climate Change", Journal of Economic Literature 45, no. 3 (2007): 703–24.

191. Wagner and Weitzman, Climate Shock, chap. 3.

192. See Meinshausen et al., "Greenhouse-Gas Emission Targets for Limiting Global Warming to 2 °C.".

193. On this see Pavel Podvig, "Reducing the Risk of an Accidental Launch," Science & Global Security 14, no. 2–3 (December 1, 2006): 75–115, doi:10.1080/08929880600992990.

194. On this see Seth D. Baum, "Confronting the Threat of Nuclear Winter," Futures, Confronting Future Catastrophic Threats To Humanity, 72 (September 2015): 69–79, doi:10.1016/j.futures.2015.03.004.

195. The WHO's International Health Regulations were an important development in this area, but the rules and their implementation could be improved. See Rebecca Katz and Scott F Dowell, "Revising the International Health Regulations: Call for a 2017 Review Conference," The Lancet Global Health 3, no. 7 (July 2015): e352–53, doi:10.1016/S2214-109X(15)00025-X.

196. Ibid.

197. See National Research Council (U. S.). Committee to Review Near-Earth-Object Surveys and Hazard Mitigation Strategies, Defending Planet Earth, chap. 2.

198. Sparks et al., Super-Eruptions..

199. See the discussion in Duprex et al., "Gain-of-Function Experiments.".

200. Nouri and Chyba, "Biotechnology and Biosecurity."

201. Bostrom, Superintelligence.

202. Shepherd, Geoengineering the Climate.

203. For discussion of different particulate winter scenarios see Timothy M. Maher and Seth D. Baum, "Adaptation to and Recovery from Global Catastrophe," Sustainability 5, no. 4 (March 28, 2013): 1461–79, doi:10.3390/su5041461.

204. Ibid.

205. See for example David Charles Denkenberger and Joshua Pearce, Feeding Everyone No Matter What : Managing Food Security after Global Catastrophe (Amsterdam: Academic Press, 2015).

206. Nick Bostrom, "Existential Risk Prevention as Global Priority," Global Policy 4, no. 1 (February 1, 2013): 15–31, doi:10.1111/1758-5899.12002.

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