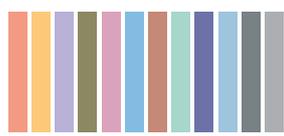


Global Challenges

 12 Risks that threaten human civilisation

Executive Summary

The case for a new risk category





Global Challenges Foundation

The Global Challenges Foundation works to raise awareness of the greatest threats facing humanity. In particular climate change, other environmental damage and political violence, and how these threats are linked to poverty and the rapid growth in global population. These problems appear insurmountable without an international body with decision-making mandate. The Foundation is therefore working to identify possible solutions and models as to how the United Nations can develop, and initiate new ideas on working global governance.

<i>Contents</i>	<i>page</i>
- <i>Introduction</i>	<i>1</i>
- <i>12 Global risks</i>	<i>2</i>
- <i>The goals of the report</i>	<i>4</i>
- <i>Report structure</i>	<i>5</i>
- <i>A new category of global risk</i>	<i>6</i>
- <i>Infinite impacts and thresholds</i>	<i>7</i>
- <i>Methodology</i>	<i>8</i>
- <i>Quick overview of each risk</i>	<i>9</i>
- <i>Relations between global risks</i>	<i>16</i>
- <i>The technical difficulty of reducing the risk and the difficulty of collaboration</i>	<i>17</i>
- <i>Uncertainties</i>	<i>18</i>
- <i>Probability</i>	<i>19</i>
- <i>Possible ways forward</i>	<i>20</i>

Introduction

This is the executive summary of a report about a limited number of global risks that pose a threat to human civilisation, or even possibly to all human life.

With such a focus it may surprise some readers to find that the report's essential aim is to inspire action and dialogue as well as an increased use of the methodologies used for risk assessment. The real focus is not on the almost unimaginable impacts of the risks the report outlines. Its fundamental purpose is to encourage global collaboration and to use this new category of risk as a driver for innovation.

The idea that we face a number of global challenges threatening the very basis of our civilisation at the beginning of the 21st century is well accepted in the scientific community, and is studied at a number of leading universities.¹ However, there is still no coordinated approach to address this group of challenges and turn them into opportunities.



Full report can be downloaded at globalchallenges.org/globalrisks

History: the LA-602 document

It is only 70 years ago that Edward Teller, one of the greatest physicists of his time, with his back-of-the-envelope calculations, produced results that differed drastically from all that had gone before. His calculations showed that the explosion of a nuclear bomb – a creation of some of the brightest minds on the planet, including Teller himself – could result in a chain reaction so powerful that it would ignite the world's atmosphere, thereby ending human life on Earth.

Robert Oppenheimer, who led the Manhattan Project to develop the nuclear bomb, halted the project to see whether Teller's calculations were correct. The resulting document, LA-602: Ignition of the Atmosphere with Nuclear Bombs, concluded that Teller was wrong. But the sheer complexity drove the assessors to end their study by writing that "further work on the subject [is] highly desirable". The LA-602 document can be seen as the first global challenge report addressing a category of risks where the worst possible impact in all practical senses is infinite.

¹ http://en.wikipedia.org/wiki/Global_catastrophic_risk

12 Global risks

This report has, to the best of the authors' knowledge, created the first list of global risks with impacts that for all practical purposes can be called infinite. It is also the first structured overview of key events related to such challenges and has tried to provide initial rough quantifications for the probabilities of these impacts. In the next phase of the project, these placeholder estimates will be improved and refined by a variety of methods (expert elicitation, fault trees, simulations, etc.) appropriate to each specific risk.

The report conducts its exploration within carefully defined bounds, resulting in a list of twelve risks with potentially infinite outcomes.

There were many challenges which might have been included on the list because of their ability to pose severe damage to humanity. They were excluded for one or more of three reasons:

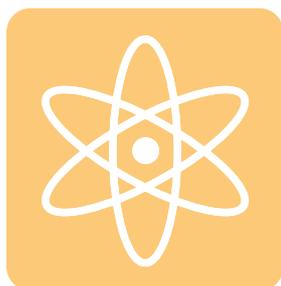
1. **Limited impact** – tsunamis, for example, and chemical pollution.
2. **No effective countermeasures** – the report focuses on promoting effective interventions and so ignores challenges where nothing useful can be done to prevent or mitigate the impact, as with nearby gamma-ray bursts.
3. **Included in other challenges** – many challenges are already covered by others, or are very similar to them. Population growth, for one, is significant for climate change and ecosystem catastrophe, but without direct large-scale impacts of its own.

It is worth noting that complex systems are often stable only within certain boundaries outside which the system can collapse and rapidly change to a new stable state. Such a collapse can trigger a process where change continues for a long time until a new stable state is found. None of the risks in this report are likely to result directly in an infinite impact, and some cannot do so physically. All the risks however are big enough to reach a threshold where the social and ecological systems become so unstable that an infinite impact could ensue.

This is a report about two extremes, not one. It is about how a better understanding of the magnitude of the challenges can help the world to address the risks it faces, and can help to create a path towards more sustainable development. It is a scientific assessment about the possibility of oblivion, certainly, but more than that it is a call for action based on the assumption that humanity is able to rise to challenges and turn them into opportunities. We are confronted with possibly the greatest challenge ever and our response needs to match this through global collaboration in new and innovative ways.



Extreme Climate Change



Nuclear War



Global Pandemic



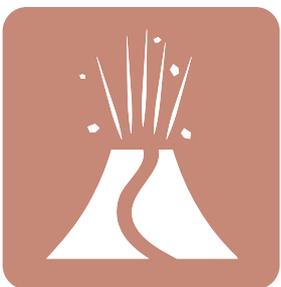
Ecological Catastrophe



Global System Collapse



Major Asteroid Impact



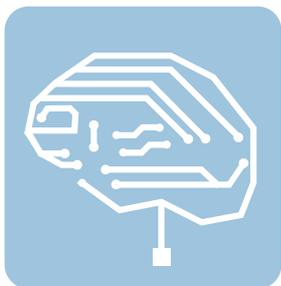
Super-volcano



Synthetic Biology



Nanotechnology



Artificial Intelligence



Unknown Consequences



Future Bad Global Governance

The goals of the report

The four main goals of this report are to acknowledge, inspire, connect and deliver.

The first of the report's goals – acknowledging the existence of risks with potentially infinite impact – seeks to help key stakeholders to acknowledge the existence of the category of risks that could result in infinite impact, and to show them that we can reduce or even eliminate most of them.

The second goal is to inspire by showing the practical action that is taking place today. This report seeks to show that helping to meet these global challenges is perhaps the most important contribution anyone can make today, and highlights concrete examples to inspire a new generation of leaders.

The third goal is to connect different groups at every level, so that leaders in different sectors connect with each other to encourage collaboration. This will need a specific focus on financial and security policy, where significant risks combine to demand action beyond the incremental.

The fourth goal is to deliver actual strategies and initiatives that produce actual results. The report is a first step and its success will ultimately be measured only on how it contributes to concrete results.

The report will have achieved its goals when key decision-makers recognise the magnitude of the possible risks and our ability to reduce or even eliminate most of them.

The goals

- 1. to acknowledge the existence of risks with potentially infinite impact.*
- 2. to inspire by showing the practical action that is taking place today.*
- 3. to connect different groups at every level.*
- 4. to deliver actual strategies and initiatives that produce actual results.*

Report structure

The first part of the report introduces and defines the global challenges and includes the methodology for selecting them.

The second part is an overview of the twelve challenges and key events that illustrate strategic work to address them. It also lists for each challenge five important factors that influence its probability or impact. The challenges are divided into four different categories:

- **current challenges** includes those which currently threaten humanity because of its economic and technological development;
- **exogenic challenges** are those where the basic probability of an event is beyond human control, but where the probability and magnitude of the impact can be influenced;
- **emerging challenges** could both help reduce the risks associated with current challenges and also result in infinite impacts;
- the last of the twelve challenges are **global policy challenges**, threats arising from future global governance as it resorts to destructive policies in response to the categories of challenge listed above.

The third part of the report discusses the relationship between the different challenges, as action to address one can increase the risk of another. Many solutions can also address multiple challenges, so there are significant benefits from understanding how they are linked.

The fourth part is an overview, the first ever to the authors' knowledge, of the probabilities of global challenges with potentially infinite impacts.

The fifth part presents some of the most important underlying trends that influence the challenges, which often build up slowly to a threshold where very rapid changes can ensue.

The sixth part presents an overview of possible ways forward.

A new category of global risk

The idea that there may be risks where the impact can be described as infinite, defined as the end of human civilisation or even human life, is not new. However, it excites relatively little political or academic interest, and the way it is treated in popular culture makes a serious discussion more difficult.

For several reasons the potentially infinite impacts of the challenges in this report are not as well known as they should be. One reason is the way that extreme impacts are often masked by most of the theories and models used by governments and business today.

Climate change is a good example, where almost all of the focus is on the most likely scenarios, and there are few public studies that include the low-probability high-impact scenarios. In most reports about climate impacts, those caused by warming beyond five or six degrees Celsius are omitted from tables and graphs. Other aspects that contribute to this relative invisibility include the fact that extreme impacts are difficult to translate into monetary terms, as they have a global scope and often require a time-horizon of a century or more. They cannot be understood simply by linear extrapolation of current trends, and they lack historical precedents. There is also the fact that the measures required to significantly reduce the probability of infinite impacts will be radical compared to a business-as-usual scenario.

A scientific approach requires us to base our decisions on the whole probability distribution.

The review of literature indicates that, under a business as usual scenario, new risks with potential infinite impact are probably inseparable from the rapid technological development in areas like synthetic biology, nanotechnology and AI.

Most risks are linked to increased knowledge, economic and technical development that has brought many benefits. E.g. climate change is a result from the industrial revolution and fossil fuel based development. The increased potential for global pandemics is one consequence of an integrated global economy where goods and services move quickly internationally. Similar challenges can be expected for synthetic biology, nanotechnology and AI.

There are remedies, including technological and institutional, for all risks. But they will require collaboration of a sort humanity has not achieved before, and the creation of systems which can deal with problems pre-emptively. It is important to understand that much of the knowledge and many tools that we have, and will develop, can be both a risk and a solution to risks depending on context.

$$\text{Risk} = \text{Probability} \times \text{Impact}$$

Infinite impacts and thresholds

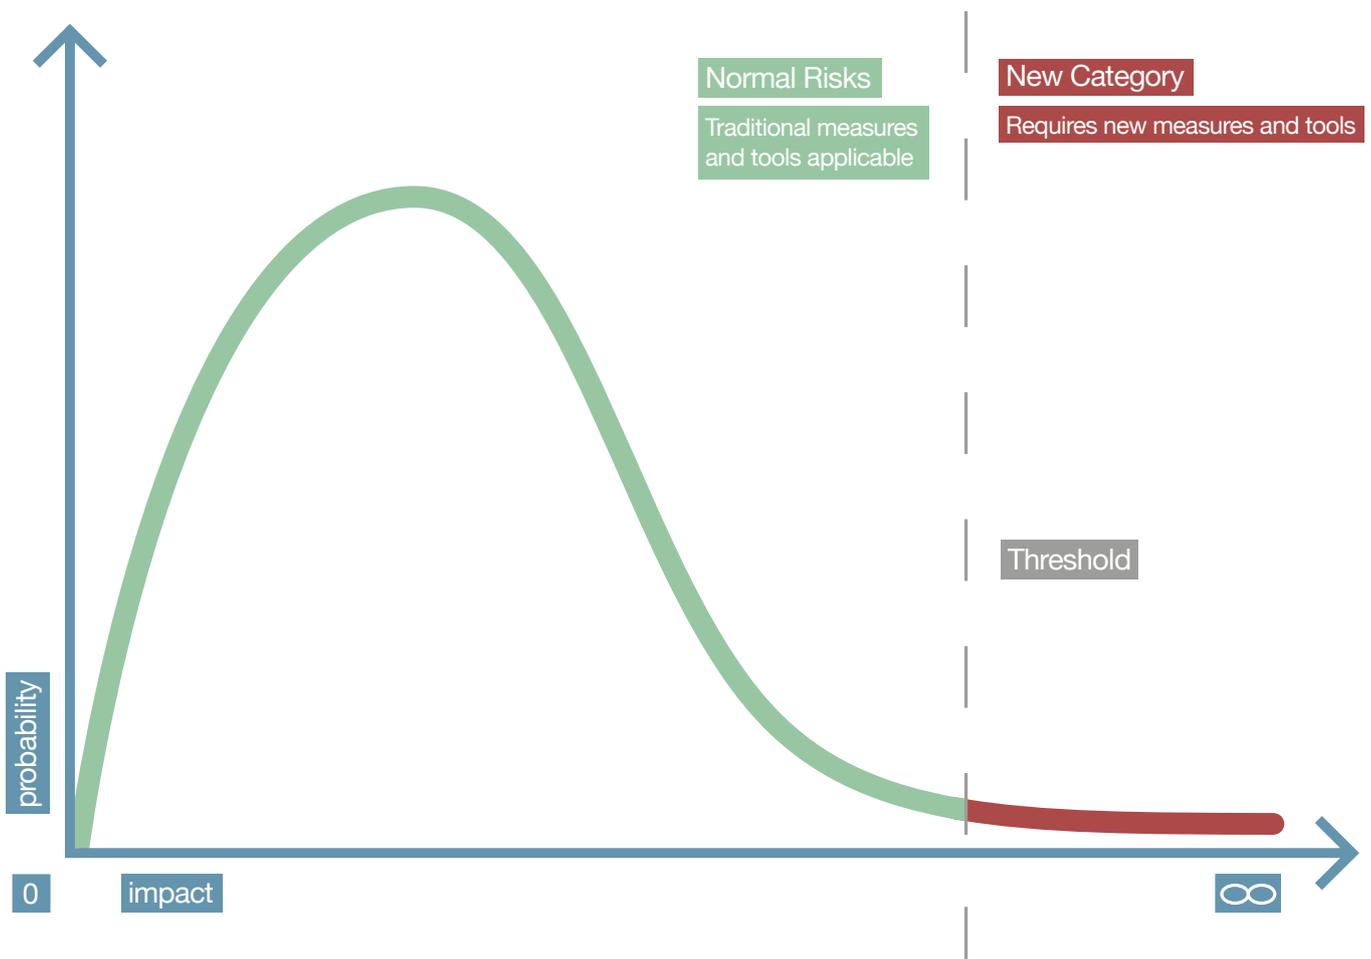
There is a clear ethical dimension to the concept of infinite impact, because a very small group alive today can take decisions that will fundamentally affect all future generations.

Using traditional economic tools is problematic and can generate disagreement over issues such as discounting, which the report examines in some detail, considering for example the role of tipping points.

The report distinguishes between the concepts of infinite impact – where civilisation collapses to a state of great suffering and does not recover, or a situation where all human life ends – and infinite impact threshold – an impact that can trigger a chain of

events that could result first in a civilisation collapse, and then later result in an infinite impact. Such thresholds are especially important to recognise in a complex and interconnected society where resilience is decreasing.

A collapse of civilisation is defined as a drastic decrease in human population size and political/economic/social complexity, globally and for an extended time.



Methodology

In order to establish a list of global challenges with potentially infinite impact, a methodological triangulation was used, consisting of:

- A quantitative assessment of relevant literature.
- A strategic selection of relevant organisations and their priorities.
- A qualitative assessment with the help of expert workshops.

Two workshops were arranged where the selection of challenges was discussed, one with risk experts in Oxford at the Future of Humanity Institute and the other in London with experts from the financial sector. No challenge was excluded at the workshops, but one was added: the participants agreed to include Global System Collapse as a category.

Relevant literature

Identification of credible sources: search relevant literature in academic literature included in World of Knowledge and Google Scholar.

1

Estimations of impact

Only literature where there is some estimation of impact that indicates the possibility of an infinite impact is included.

2

Leading organisations' priorities

In order to increase the probability of covering all relevant risks an overview of leading organisations' work was conducted. This list was then compared with the initial list and subjected to the same filter regarding the possibility to affect the probability or impact.

3

Possibility of addressing the risk

From the risks gathered from literature and organisations, only those where the probability or impact can be affected by human actions are included.

4

Expert review

Qualitative assessment: Expert review in order to increase the probability of covering all relevant global risks.

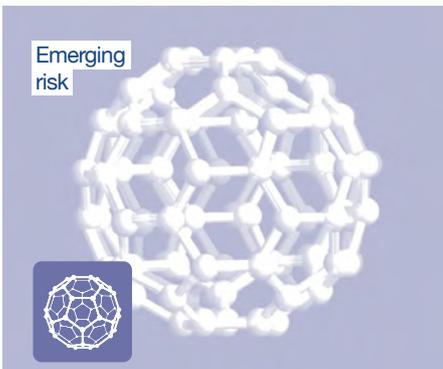
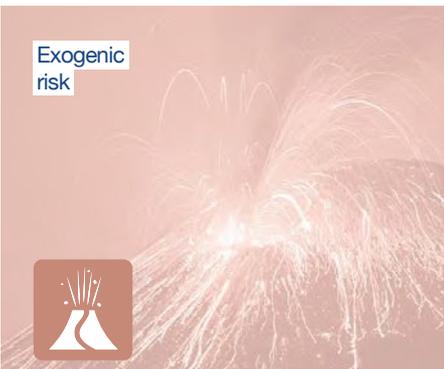
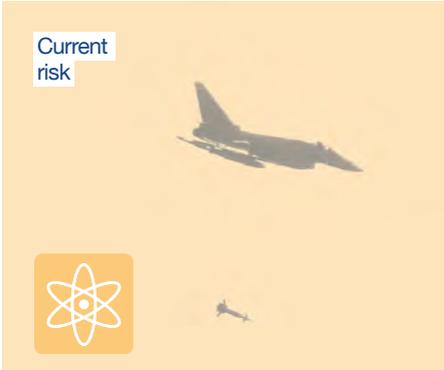
5

List of risks

Result: List of risks with potentially infinite impacts.

6

Quick overview of each risk



Current risk



5 key factors:

- 1 The uncertainties in climate sensitivity models, including the tail.
- 2 The likelihood - or not - of global coordination on controlling emissions.
- 3 The future uptake of low carbon economies, including energy, mobility and food systems.
- 4 Whether technological innovations will improve or worsen the situation, and by how much.
- 5 The long-term climate impact caused by global warming.

Extreme Climate Change

As for all risks there are uncertainties in the estimates, and warming could be much more extreme than the middle estimates suggest. Feedback loops could mean global average temperatures increase by 4°C or even 6°C over pre-industrial levels. Feedbacks could be the release of methane from permafrost or the dieback of the Amazon rainforest. The impact of global warming would be strongest in poorer countries, which could become completely uninhabitable for the highest range of warming.

Mass deaths and famines, social collapse and mass migration are certainly possible in this scenario. Combined with shocks to the agriculture and biosphere-dependent industries of the more developed countries, this could lead to global conflict and possibly civilisation collapse. Further evidence of the risk comes from signs that past civilisation collapses have been driven by climate change.

Current risk



5 key factors:

- 1 How relations between current and future nuclear powers develop.
- 2 The probability of accidental war.
- 3 Whether disarmament efforts will succeed in reducing the number of nuclear warheads.
- 4 The likelihood of a nuclear winter.
- 5 The long-term effects of a nuclear war on climate, infrastructure and technology. A new category of global risk.

Nuclear War

The likelihood of a full-scale nuclear war between the USA and Russia has probably decreased. Still, the potential for deliberate or accidental nuclear conflict has not been removed, with some estimates putting the risk in the next century or so at around 10%. A larger impact would depend on whether or not the war triggered what is often called a nuclear winter or something similar – the creation of a pall of smoke high in the stratosphere that would

plunge temperatures below freezing around the globe and possibly also destroy most of the ozone layer. The detonations would need to start firestorms in the targeted cities, which could lift the soot up into the stratosphere. The risks are severe and recent models have confirmed the earlier analysis. The disintegration of the global food supply would make mass starvation and state collapse likely.

Current risk



5 key factors:

- 1 What the true probability distribution for pandemics is, especially at the tail.
- 2 The capacity of international health systems to deal with an extreme pandemic.
- 3 How fast medical research can proceed in an emergency.
- 4 How mobility of goods and people, as well as population density, will affect pandemic transmission.
- 5 Whether humans can develop novel and effective anti-pandemic solutions.

Global Pandemic

An epidemic of infectious disease that has spread through human populations across a large region or even worldwide. There are grounds for suspecting that such a high-impact epidemic is more probable than usually assumed. All the features of an extremely devastating disease already exist in nature: essentially incurable (Ebola), nearly always fatal (rabies), extremely infectious (common cold), and long incubation periods (HIV). If a pathogen were to emerge that somehow combined these features

(and influenza has demonstrated antigenic shift, the ability to combine features from different viruses), its death toll would be extreme. The world has changed considerably, making comparisons with the past problematic. Today it has better sanitation and medical research, as well as national and supra-national institutions dedicated to combating diseases. But modern transport and dense human population allow infections to spread much more rapidly, and slums can be breeding grounds for disease.

Current risk



5 key factors:

- 1 The extent to which humans are dependent on the ecosystem.
- 2 Whether there will be effective political measures taken to protect the ecosystem on a large scale.
- 3 The likelihood of the emergence of sustainable economies.
- 4 The positive and negative impacts on the ecosystems of both wealth and poverty.
- 5 The long-term effects of an ecological collapse on ecosystems.

Ecological Collapse

This is where an ecosystem suffers a drastic, possibly permanent, reduction in carrying capacity for all organisms, often resulting in mass extinction. Humans are part of the global ecosystem and so fundamentally depend on it. Species extinction is now far faster than the historic rate, and attempts to quantify a safe ecological operating space place humanity well outside it. Many of the problems of ecological degradation interact to multiply

the damage and (unlike previous, localised collapses) the whole world is potentially at risk. It seems plausible that some human lifestyles could be sustained in a relatively ecosystem independent way, at relatively low costs. Whether this can be achieved on a large scale in practice, especially during a collapse, will be a technological challenge and whether it is something we want is an ethical question.

Current risk



5 key factors:

- 1 Whether global system collapse will trigger subsequent collapses or fragility in other areas.
- 2 What the true trade-off is between efficiency and resilience.
- 3 Whether effective regulation and resilience can be developed.
- 4 Whether an external disruption will trigger a collapse.
- 5 Whether an internal event will trigger a collapse.

Global System Collapse

An economic or societal collapse on the global scale. The term has been used to describe a broad range of conditions. Often economic collapse is accompanied by social chaos, civil unrest and sometimes a breakdown of law and order. Societal collapse usually refers to the fall or disintegration of human societies, often along with their life support systems. The world economic and political system is made up of many actors with many objectives and many links between them. Such intricate, interconnected systems are subject to unexpected system-wide failures caused by the

structure of the network – even if each component of the network is reliable. This gives rise to systemic risk, when parts that individually may function well become vulnerable when connected as a system to a self-reinforcing joint risk that can spread from part to part, potentially affecting the entire system and possibly spilling over to related outside systems. Such effects have been observed in ecology, finance and critical infrastructure such as power grids. The possibility of collapse becomes more acute when several independent networks depend on each other.

Exogenic risk



5 key factors:

- 1 Whether detection and tracking of asteroids and other dangerous space objects is sufficiently exhaustive.
- 2 How feasible it is to deflect an asteroid.
- 3 Whether measures such as evacuation could reduce the damage of an impact.
- 4 The short- and long-term climate consequences of a collision.
- 5 Whether our current civilisation could adapt to a post-impact world.

Major Asteroid Impact

Large asteroid collisions – with objects 5 km or more in size – happen about once every twenty million years and would have an energy a hundred thousand times greater than the largest bomb ever detonated. A land impact would destroy an area the size of a nation like Holland. Larger asteroids could be extinction-level events. Asteroid impacts are probably one of the best understood of all risks in this report.

There has been some discussion about possible methods for deflecting asteroids found on a collision course with the planet. Should an impact occur the main destruction will not be from the initial impact, but from the clouds of dust projected into the upper atmosphere. The damage from such an “impact winter” could affect the climate, damage the biosphere, affect food supplies, and create political instability.

Exogenic risk



5 key factors:

- 1 Whether countries will coordinate globally against super-volcano risk and damage.
- 2 The predictability of super-volcanic eruptions.
- 3 How directly destructive an eruption would be.
- 4 The effectiveness of general mitigation efforts.
- 5 How severe the long-term climate effects would be.

Super-volcano

Any volcano capable of producing an eruption with an ejecta volume greater than 1,000 km³. This is thousands of times larger than normal eruptions. The danger from super-volcanoes is the amount of aerosols and dust projected into the upper atmosphere. This dust would absorb the Sun's rays and cause a global volcanic winter. The Mt Pinatubo eruption of 1991 caused an average global cooling of surface temperatures by 0.5°C over three years,

while the Toba eruption around 70,000 years ago is thought by some to have cooled global temperatures for over two centuries.

The effect of these eruptions could be best compared with that of a nuclear war. The eruption would be more violent than the nuclear explosions, but would be less likely to ignite firestorms and other secondary effects.

Emerging risk



5 key factors:

- 1 The true destructive potential of synthetic biology, especially the tail risk.
- 2 Whether the field will be successfully regulated, or successfully manage to regulate itself.
- 3 Whether the field will usher in a new era of bio-warfare.
- 4 Whether the tools of synthetic biology can be used defensively to create effective counter measures.
- 5 The dangers of relying on synthetic biologists to estimate the danger of synthetic biology.

Synthetic Biology

The design and construction of biological devices and systems for useful purposes, but adding human intentionality to traditional pandemic risks. Attempts at regulation or self-regulation are currently in their infancy, and may not develop as fast as research does. One of the most damaging impacts from synthetic biology would come from an engineered pathogen targeting humans or a crucial component of the ecosystem.

This could emerge through military or commercial bio-warfare, bio-terrorism (possibly using dual-use products developed by legitimate researchers, and currently unprotected by international legal regimes), or dangerous pathogens leaked from a lab. Of relevance is whether synthetic biology products become integrated into the global economy or biosphere. This could lead to additional vulnerabilities (a benign but widespread synthetic biology product could be specifically targeted as an entry point through which to cause damage).

Emerging risk



5 key factors:

- 1 The timeline for nanotech development.
- 2 Which aspects of nanotech research will progress in what order.
- 3 Whether small groups can assemble a weapons arsenal quickly.
- 4 Whether nanotech tools can be used defensively or for surveillance.
- 5 Whether nanotech tools or weaponry are made to be outside human control.

Nanotechnology

Atomically precise manufacturing, the creation of effective, high-throughput manufacturing processes that operate at the atomic or molecular level. It could create new products – such as smart or extremely resilient materials – and would allow many different groups or even individuals to manufacture a wide range of things. This could lead to the easy construction of large arsenals of conventional or more novel weapons made possible by atomically precise manufacturing.

Of particular relevance is whether nanotechnology allows the construction of nuclear bombs. But many of the world's current problems may be solvable with the manufacturing possibilities that nanotechnology would offer, such as depletion of natural resources, pollution, climate change, clean water and even poverty. Some have conjectured special self-replicating nanomachines which would be engineered to consume the entire environment. The misuse of medical nanotechnology is another risk scenario.

Emerging risk



5 key factors:

- 1 The reliability of AI predictions.
- 2 Whether there will be a single dominant AI or a plethora of entities.
- 3 How intelligent AIs will become.
- 4 Whether extremely intelligent AIs can be controlled, and if so, how.
- 5 Whether whole brain emulations (human minds in computer form) will arrive before true AIs.

Artificial Intelligence

AI is the intelligence exhibited by machines or software, and the branch of computer science that develops machines and software with human-level intelligence. The field is often defined as “the study and design of intelligent agents”, systems that perceive their environment and act to maximise their chances of success. Such extreme intelligences could not easily be controlled (either by the groups creating them, or by some international regulatory regime), and would probably act to boost their own intelligence and acquire maximal resources for almost all initial AI motivations.

And if these motivations do not detail the survival and value of humanity, the intelligence will be driven to construct a world without humans. This makes extremely intelligent AIs a unique risk, in that extinction is more likely than lesser impacts. On a more positive note, an intelligence of such power could easily combat most other risks in this report, making extremely intelligent AI into a tool of great potential. There is also the possibility of AI-enabled warfare and all the risks of the technologies that AIs would make possible. An interesting version of this scenario is the possible creation of “whole brain emulations”: human brains scanned and physically represented in a machine. This would make the AIs into properly human minds, possibly alleviating a lot of problems.

Emerging risk



5 key factors:

- 1 Whether there will be extensive research into unknown risks and their probabilities.
- 2 The capacity to develop methods for limiting the combined probability of all uncertain risks.
- 3 The capacity for estimating “out of-model” risks.
- 4 The culture of risk assessment in potentially risky areas.
- 5 Whether general, non-risk-specific mitigation or resilience measures are implemented.

Unknown Consequences

These represent the unknown unknowns in the family of global catastrophic challenges. They constitute an amalgamation of all the risks that can appear extremely unlikely in isolation, but can combine to represent a not insignificant proportion of the risk exposure. One resolution to the Fermi paradox – the apparent absence of alien life in the galaxy – is that intelligent life destroys itself before beginning to expand into the galaxy. Results that increase or decrease the probability of this explanation modify the

generic probability of intelligent life (self-)destruction, which includes uncertain risks. Anthropoc reasoning can also bound the total risk of human extinction, and hence estimate the unknown component. Non risk-specific resilience and post-disaster rebuilding efforts will also reduce the damage from uncertain risks, as would appropriate national and international regulatory regimes. Most of these methods would also help with the more conventional, known risks, which badly need more investment.

Global Policy risk



5 key factors:

- 1 How the severity of non-deadly policy failures can be compared with potential casualties.
- 2 Whether poor governance will result in a collapse of the world system.
- 3 How mass surveillance and other technological innovations will affect governance.
- 4 Whether there will be new systems of governance in the future.
- 5 Whether a world dictatorship may end up being constructed.

Future Bad Global Governance

There are two main divisions in governance disasters: failing to solve major solvable problems, and actively causing worse outcomes. An example of the first would be failing to alleviate absolute poverty; of the second, constructing a global totalitarian state. Technology, political and social change may enable the construction of new forms of governance, which may be either much better or much worse.

Two issues with governance disasters are first, the difficulty of estimating their probability, and second, the dependence of the impact of these disasters on subjective comparative evaluations: it is not impartially obvious how to rank continued poverty and global totalitarianism against billions of casualties or civilisation collapse.

Relations between global risks

Two things make the understanding of the relation between the global risks particularly important.

1. Impacts: The risks are interconnected in different ways. Often the situation resembles a set of dominoes: if one falls, many follow. Even small impacts can start a process where different risks interact.
2. Specific measures to address a risk: Global risks often require significant changes, which will result in situations where measures to reduce the risk in one area affect the probability and/or the impact in other areas, for better or worse.



The technical difficulty of reducing the risk and the difficulty of collaboration

In order to better understand the relations between different global risks, work could start to analyse similarities and differences.

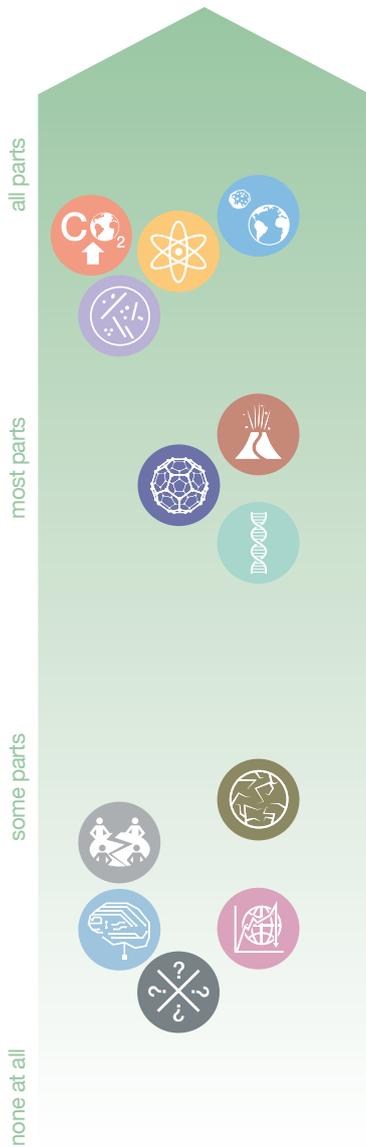
Below is an example of an overview of how different global risks can be plotted depending on the technical difficulty of reducing the risk and the difficulty of collaborating to reduce it.



Uncertainties

As the different challenges are very different and the status of probability estimates varies significantly, the initial probability numbers are provided together with estimates regarding:

1. Understanding of sequence



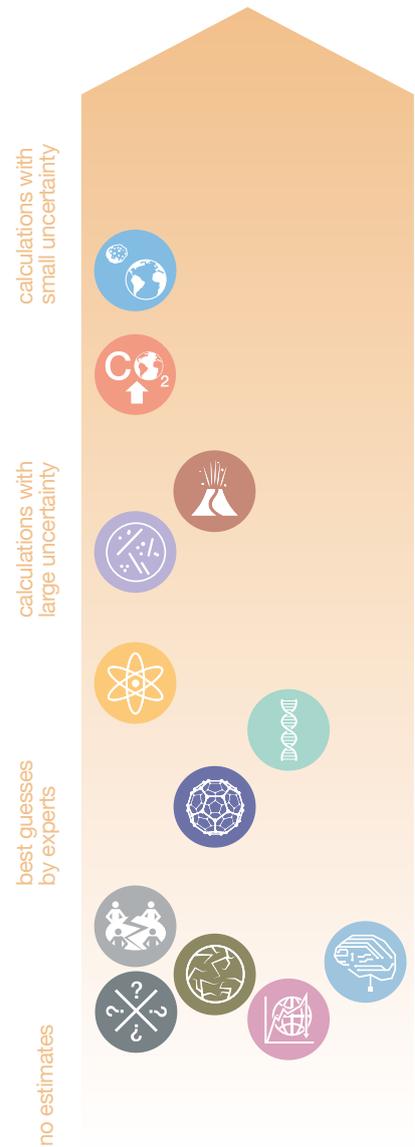
degree of events from today's actions to infinite impact

2. Data availability



amount of data to make probability assessment on all relevant steps of the sequence

3. Existing probability estimation



kind of estimation and uncertainty

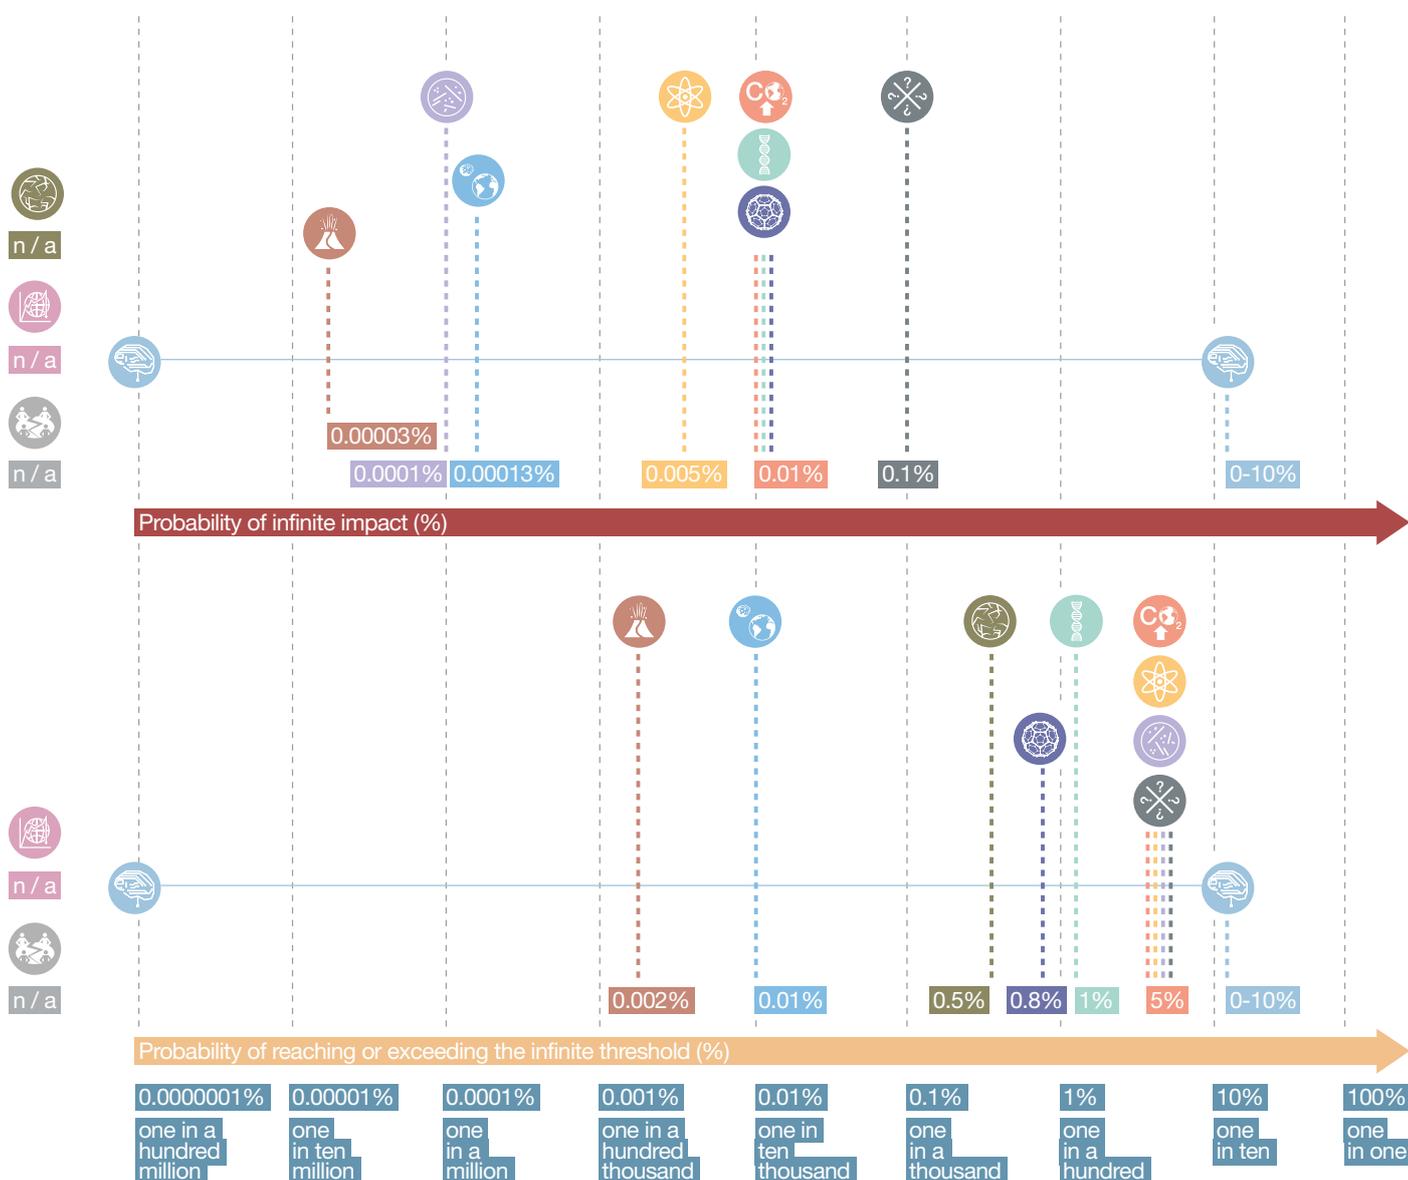
Probability

These estimates are an attempt to assemble existing estimates in order to encourage efforts to improve the numbers. They express estimates of probabilities over 100 years, except in the case of extreme climate change, where the time frame is 200 years.

Global challenges need to be seen in the light of trends which help to shape the wider society. These include:

Poverty – although it has fallen, it could increase again. This is especially relevant to climate change and pandemics.

Population growth – the UN’s estimates range from 6.8 billion people by 2100 to a high-variant projection of 16.6 bn (which would require the resources of 10 Earth-like planets to provide everyone with a modern Western lifestyle). Other trends include technological development and demographic changes.



Possible ways forward

There are ten areas that could help mitigate immediate threats while also contributing to a future global governance system capable of addressing global risks with a potential infinite impact:

1. Global challenges leadership networks
2. Better quality risk assessment for global challenges
3. Development of early warning systems
4. Encouraging visualisation of complex systems
5. Highlighting early movers
6. Including the whole probability distribution
7. Increasing the focus on the probability of extreme events
8. Encouraging appropriate language to describe extreme risks
9. Establishing a Global Risk and Opportunity Indicator to guide governance
10. Explore the possibility of establishing a Global Risk Organisation (GRO)

The main authors of this report are **Dennis Pamlin**, Executive Project Manager, Global Challenges Foundation and **Dr Stuart Armstrong**, James Martin Research Fellow, Future of Humanity Institute, Oxford Martin School & Faculty of Philosophy, University of Oxford. Stuart Armstrong wrote the chapter covering the twelve global challenges, under the direction of Dennis Pamlin who served as project manager and himself wrote and edited the rest of the report. **Seth Baum**, Executive Director of the Global Catastrophic Risk Institute and affiliate researcher at the Center for Research on Environmental Decisions, Columbia University, also played an important role as he helped develop the methodology chapter regarding the selection of the global risks with a potentially infinite impact as well as providing helpful input throughout the process. The report is the result of a collaborative approach where many people have provided invaluable contributions.

The authors would therefore like to thank a few people in particular. First and foremost **László Szombatfalvy**, Chairman of the Global Challenges Foundation, whose work is the basis for this report and whose guidance on all levels has been invaluable. The rest of the board of the Global Challenges Foundation have also contributed in many different ways, in particular, **Johan Rockström** has provided important input regarding the structure and methodology. Outside the foundation **Prof Nick Bostrom**, Professor & Director of the Future of Humanity Institute, Oxford Martin School & Faculty of Philosophy, University of Oxford, who initiated the possibility of working with the Future of Humanity Institute at the University of Oxford, played a particularly important role. **Patrick McSharry**, Head of Smith School's Catastrophe Risk Financing research area, provided invaluable input regarding complex systems and ways that the economic system can respond to risks with potentially infinite impacts. **Alex Kirby** also played a key part as he did so much more than proofread the text; the report would hardly be possible to read without his help. Various additional edits and changes were made by **Peter Brietbart**.

Others that must be mentioned, including those who participated in the workshop on 14 January 2014, at the Future of Humanity Institute (FHI), University of Oxford and the workshop at the Munich RE office in London on 15 January 2014, and helped provide input regarding the economic and finance aspects, include (in alphabetical order):

Dr Nick Beckstead, Research Fellow, Future of Humanity Institute, Oxford Martin School & Faculty of Philosophy, University of Oxford

Kennette Benedict, Executive Director and Publisher of the Bulletin of the Atomic Scientists

Oliver Bettis, Pricing Actuary, Munich RE and Fellow of the Chartered Insurance Institute and the Institute & Faculty of Actuaries

Dr Eric Drexler, Academic Visitor, Future of Humanity Institute, Oxford Martin School & Faculty of Philosophy, University of Oxford

Madeleine Enarsson, Transformative Catalyst, 21st Century Frontiers

Dr Seán Ó hÉigearthaigh, Senior Academic Manager, Future of Humanity Institute, University of Oxford and Executive Director, Centre for the Study of Existential Risk, University of Cambridge

Martin Hellman, Professor Emeritus of Electrical Engineering, Stanford University

Pan Jiahua, Director of the Institute for Urban and Environmental Studies, Chinese Academy of Social Sciences (CASS); Professor of economics at CASS; Vice-President Chinese Society for Ecological Economics; Member of the National Expert Panel on Climate Change and National Foreign Policy Advisory Committee, China

Aled Jones, Director of the Global Sustainability Institute (GSI) at Anglia Ruskin University

Nick Mabey, Chief Executive and Founding Director of E3G (Third Generation Environmentalism)

Jennifer Morgan, Founder & Co-Convener, The Finance Lab

Prof Vincent Müller, James Martin Research Fellow, Future of Humanity Institute, Oxford Martin School & Faculty of Philosophy, University of Oxford

Robert de Neufville, Professional Associate, Global Catastrophic Risk Institute

Dr Toby Ord, James Martin Research Fellow, Future of Humanity Institute, Oxford Martin School & Faculty of Philosophy, University of Oxford

Jules Peck, Founding Partner, Jericho Chambers; Trustee, New Economics Foundation

Dr Anders Sandberg, James Martin Research Fellow, Future of Humanity Institute, Oxford Martin School & Faculty of Philosophy, University of Oxford

Nick Silver, Director of Callund Consulting and founder and director of the Climate Bonds Initiative (CBI)

Andrew Simms, Author, Fellow at the New Economics Foundation and Chief Analyst at Global Witness

Andrew Snyder-Beattie, Academic Project Manager, Future of Humanity Institute, Oxford Martin School & Faculty of Philosophy, University of Oxford

James Taplan, Principal Sustainability Advisor, Forum for the Future

Raj Thamotheram, CEO, Preventable Surprises

Nathan Wolfe, Director of Global Viral and the Lorry I. Lokey Visiting Professor in Human Biology at Stanford University

Liang Yin, Investment Consultant at Towers Watson



Global Challenges Foundation

Published February 2015
by Global Challenges Foundation

globalchallenges.org

Global Challenges Foundation
Stureplan 4C
114 35 Stockholm



12 Risks that threaten human civilisation

The opinions expressed in this report are those of the authors and do not necessarily reflect those of Global Challenges Foundation.

Any inaccuracies in the report remain the responsibility of the authors.

The material and the geographical designations in this report do not imply the expression of any opinion whatsoever on the part of Global Challenges Foundation, concerning the legal status of any country, territory, or area, or concerning the delimitation of its frontiers or boundaries.

For comments and/or queries on this report, please contact the co-authors:

Dennis Pamlin

Executive Project Manager
Global Risks
Global Challenges Foundation
dennis@globalchallenges.org
globalchallenges.org

Dr Stuart Armstrong

James Martin Research Fellow
Future of Humanity Institute
Oxford Martin School & Faculty of Philosophy
University of Oxford
stuart.armstrong@philosophy.ox.ac.uk
fhi.ox.ac.uk