



Founders Pledge

Global Catastrophic Biological Risks

A Guide for Philanthropists

Christian Ruhl

October 2023





Contents

Acknowledgements and Disclaimers	4
Summary	5
A Note on Hazardous Information	7
External Reviews	9
Joshua Monrad, Biosecurity Program Officer, Effective Giving	9
Andrew Snyder-Beattie, Senior Program Officer, Biosecurity and Pandemic Preparedness, Open Philanthropy	11
Introduction	12
The Threat of GCBRs	14
First Principles: The Chain Reaction	15
Natural Pandemics	17
Viral, Bacterial, and Other Threats	20
Natural Extinction Risk	23
Engineered Pandemics	27
Bioweapons, Bioterrorism, and Laboratory Leaks	33
State Biological Weapons and Biodefense Programs	33
Bioterrorists and Rogue Actors	37
Accidents and Laboratory Leaks	40
Existential Threats	46
The Growing Risk of Biological Catastrophe	51
General Advances in the Life Sciences	52
Benchtop DNA Synthesis Devices	56
The AI-Bio Nexus	58
Cloud Labs	59
Other Advances	60
The Implications of the Changing Threat Landscape	61
Analyzing the Structure of the Problem	62
The Threat Is Growing	63
The Offense-Defense Balance of Emerging Biotechnologies	64
The Threat Is Super-Linear	66
The Threat Is Increasingly Complex	68
The Threat Is Adaptive	68
Societal Spending and Neglectedness	70
Societal Resources Spent on Health Security	71
Understanding GCBR Neglect and Public Policy Market Failures	74
Philanthropic Neglectedness	78



Guiding Principles for Effective Philanthropy	82
Prioritize Worst-Case Scenarios	84
Pursue Robust Interventions	86
Prevention and Response	86
Prioritize Pathogen- and Threat-Agnostic Approaches	87
Tractability and Medical Countermeasures	90
Leverage Existing Societal Resources	92
Leverage Near-Term Policy Incentives	94
Avoid High-Downside Grantmaker Dilemmas	95
Offense-Defense Distinguishability and the Security Dilemma	95
Information Hazards	96
Practical Funding Opportunities	98
GCBR Field-Building	99
GCBR Policy Advocacy	100
Constraining Capabilities and Shaping Incentives	101
Safeguards against Proliferation and Misuse of Key Technologies	102
The AI-Bio Nexus	102
DNA Synthesis Regulation	102
Transmission-Blocking Interventions	103
Pandemic-Proof Personal Protective Equipment (P4E)	104
Pathogen-Agnostic Early Detection	104
Platform Technologies for Medical Countermeasures	105
Conclusion	107
Appendix	109
A Grantmaker Dilemma Checklist	109
Evaluating Interventions	110
Plain Text R Code	112
Carlson Curve on DNA Synthesis (data)	112
Cost of Sequencing a Human Genome over Time	115
Philanthropic Funding for Biosecurity	117
About Founders Pledge	121



Acknowledgements and Disclaimers

I am thankful to the experts who have contributed to this project with their insights, feedback, and critiques, including in semi-structured interviews. The analysis and views expressed in this report do not necessarily reflect the views of anyone consulted for the project.

With special thanks to the following people for their insights and help on various parts of the report:

- **Chris Bakerlee**, Senior Program Associate, Biosecurity and Pandemic Preparedness, Open Philanthropy
- **Bill Beaver**, Policy Advisor, U.S. Department of Defense; Fellow, Janne E. Nolan Center on Strategic Weapons, Council on Strategic Risks
- **Vivian Belenky**, Research Fellow, Convergent Research
- **Yelena Biberman**, Associate Professor of Political Science, Skidmore College
- **Sarah Carter**, Principal, Science Policy Consulting LLC
- **Rocco Casagrande**, Managing Director, Gryphon Scientific
- **Kevin Esvelt**, Director, Sculpting Evolution Group, MIT Media Lab; co-founder, SecureDNA Foundation
- **Melissa Hopkins**, Health Security Policy Advisor, Johns Hopkins Center for Health Security
- **Tom Inglesby**, Director, Johns Hopkins Center for Health Security
- **Vidur Kapur**, Biosecurity Program Associate, Effective Giving
- **Jason Matheny**, President and CEO, RAND Corporation
- **Darius Meissner**, Director of Content, Horizon Institute for Public Service
- **Joshua Monrad**, Biosecurity Program Officer, Effective Giving
- **Jennifer Nuzzo**, Professor of Epidemiology and Director of the Pandemic Center, Brown University School of Public Health
- **Ryan Ritterson**, Synthetic Biologist, Gryphon Scientific
- **Andrew Snyder-Beattie**, Senior Program Officer, Biosecurity and Pandemic Preparedness, Open Philanthropy
- **Jake Swett**, Executive Director, Blueprint Biosecurity
- **James Wagstaff**, Senior Program Associate, Biosecurity and Pandemic Preparedness, Open Philanthropy
- **Jaime Yassif**, Vice President, NTI Global Biological Policy and Programs (NTI | bio)
- and anonymized experts for their insights and advice on this project.
- **Matt Lerner, Megan Phelan, Guro Skei**, and my other colleagues at Founders Pledge for feedback, edits and guidance on the report.
- **Zoë Ruhl** for her support and explanations of key biomedical facts.

This report contains descriptions of disease and human suffering.



Summary

This report provides an overview of the threat of global catastrophic biological risks (GCBRs) and how philanthropists can help to mitigate this threat. Using semi-structured interviews and analysis of available data on funding and pandemic preparedness, the report argues that:

1. **The expected cost of a biological catastrophe is immense.** High-consequence biological events like engineered pandemics are among the biggest threats to human life and modern civilization this century.
 - a. Biological risks could pose an *existential* threat to humanity, but the case for pandemic preparedness and prevention remains strong even without placing special value on existential risks.¹
2. **The threat landscape is rapidly changing.** New advances in the life sciences and enabling technologies — which could have immense benefits for humanity — may increase the risk by both providing new capabilities and enabling the use of these capabilities by reckless and malicious actors.
 - a. These changes also increase uncertainty about the origin and characteristics of future biological threats.
3. **Societal spending remains misallocated** even in the wake of the COVID-19 pandemic. Although governments and traditional philanthropists spend billions on public health and health security, they disproportionately neglect the most high-consequence threats.

These basic facts point towards a general problem structure that is familiar from Founders Pledge’s work on [climate change](#) and [nuclear war](#): the worst possible biological catastrophes will be *disproportionately* worse than other pandemics, yet they are simultaneously disproportionately neglected.

Fortunately, **we know what to do**; with an understanding of basic facts about the problem and strategic reasoning about the uncertainties of future threats, philanthropists can fund clear and tractable interventions. Stated succinctly, **philanthropists ought to hedge against the worst possible scenarios** where traditional risk-mitigation has failed by following quasi-quantitative heuristics for effectiveness, or “impact multipliers.”² Such heuristics can guide philanthropists in identifying the funding opportunities that are likely to have the highest impact on the margins. They include:

¹ Some sections of this report discuss questions of extinction and civilizational collapse. To the extent that readers subscribe to worldviews that place high moral value on the long-term future of humanity, biological risk mitigation therefore becomes especially important, and plausibly one of the most important causes that philanthropists can focus on. Nonetheless, as argued throughout the report, the super-linear structure of the risk means that worst-case pandemics remain important from many ethical perspectives.

² For a discussion of the concept of impact multipliers, see [Guiding Principles for Effective Philanthropy](#), below.



- **Focus on worst-case scenarios.** Disease outbreaks (like terrorist events) follow very heavy-tailed distributions, where the worst pandemics outweigh many smaller pandemics combined. Such events are both disproportionately important and unduly neglected.
 - In practice, this often means **prioritizing engineered threats**. Nature is *not* the worst bioterrorist.
- Within anthropogenic events, generally **pursue pathogen- and threat-agnostic approaches**. Rather than trying to guess where the next pandemic will come from, philanthropists ought to support approaches that are robust to a variety of scenarios.
- **Fund interventions that are robust to the entire spectrum of risk**, including natural pandemics, and up to and including extinction-level engineered pandemics.
 - In practice, this often means prioritizing **prevention** over prioritizing response.
- **Leverage existing societal resources** using advocacy-based interventions (largely for government, but also for other large funders), with a few notable exceptions where government interests appear unlikely to align.
- *All else equal*, **prioritize interventions with near-term positive externalities** that can garner sustained and broad political support.
- **Prioritize low-dual use potential and high offense-defense distinguishability and avoid information hazards** when selecting funding opportunities to avoid fuelling dangerous security dilemmas and doing more harm than good.

A final section discusses possible funding opportunities that leverage these heuristics. Promising funding opportunities include:

- **Field building and policy advocacy** on issues relevant for worst-case GCBR scenarios.
 - Policy advocacy may include advocacy specifically designed to constrain the access and capabilities of malevolent actors or to disincentivize the pursuit of biological weapons programs.
- **Safeguards against proliferation and misuse of key technologies and information**, especially AI misuse and DNA synthesis screening.
- **Transmission-blocking interventions**, including sterilizing technology, work on germicidal lights for indoor biological air quality, mechanisms to quickly decontaminate PPE, and more.
- **Pandemic-proof personal protective equipment (P4E).**³

³ This term is borrowed from the work of Dr. Kevin Esvelt. See also [recent work](#) by Gryphon Scientific on this topic.



- **Pathogen-agnostic early warning**, including metagenomic sequencing.
- **Platform technologies for medical countermeasures** to rapidly pivot production and distribution for novel threats.⁴

The report also discusses various dilemmas that philanthropists face when making funding decisions, and provides a **grantmaker dilemma checklist** to help screen for information hazards and security dilemmas. Philanthropists can play an important role in mitigating catastrophic biological risks, and this guide is designed to help guide the grantmaking strategy of impact-minded philanthropists.

⁴ As discussed below, large uncertainties surround the robustness of medical countermeasures to worst-case threat scenarios,



A Note on Hazardous Information

Some discussions of biological risks present [information hazards](#) that could increase the risks of a catastrophe merely by spreading information (e.g. the biological details of especially dangerous viruses) that could enable malevolent actors to cause great harm, as discussed [below](#). To make a responsible contribution to the literature on biological risks, this report therefore attempts to follow [best practices](#) surrounding the disclosure of vulnerabilities and capabilities, including by:

- Avoiding discussion of specific biological threats or “blueprints” for harm
- Discussing the features of risky pathogens only in abstract and general terms
- Not disclosing or drawing attention to specific biodefense vulnerabilities
- Using well-known historical examples whenever possible
- Asking pre-publication reviewers to pay special attention to potentially hazardous information

Fortunately, because [threat-agnostic interventions also happen to be high-impact interventions](#), responsible conduct around dangerous information does not hinder the discussion of effective philanthropy. **Readers concerned about information contained in this report are encouraged to contact the author** at christian@founderspledge.com.



External Reviews

Founders Pledge's research reports undergo several rounds of internal and external review. To provide the reader context for this report, we have asked two outside experts to briefly write up their impressions:

Joshua Monrad, Biosecurity Program Officer, Effective Giving

Joshua currently leads the biosecurity program at the philanthropic advisor Effective Giving, which allocates >\$10M/year in grants towards mitigating global catastrophic biological risks. Previously, Joshua has worked as a researcher at Oxford University, the Nuclear Threat Initiative, and the Georgetown Center for Global Health Science and Security. He holds degrees from the London School of Economics, London School of Hygiene and Tropical Medicine, and Yale University.

With *Global Catastrophic Biological Risks: A Guide for Philanthropists*, Christian Ruhl offers a scholarly analysis of modern pandemic risks and a practical call to action for philanthropists, investors, and governments allocating resources to global health security efforts. The report makes a strong case for prioritising pandemic prevention as one of the world's most pressing problems and presents concrete heuristics and frameworks for identifying promising solutions.

This report stands out for its focus on 'global catastrophic biological risks' (GCBRs), events so severe that they exceed even COVID-19 or the 1918-20 influenza pandemic in their calamitous, lasting impact on humanity. While efforts focused on such catastrophes may synergise with broader pandemic prevention, they are also distinct in several ways. Most saliently, addressing these extreme scenarios requires taking the risks posed by engineered biological systems seriously. While zoonotic diseases present significant endemic and epidemic burdens, modern biotechnology arguably holds the potential to surpass even the worst diseases found in nature. This report showcases how anthropogenic biosecurity and biosafety risks may increase dramatically in the coming decades and how it will take a concerted effort to mitigate them while realising the enormous positive potential of the life sciences.

Ruhl seeks to bring a new level of rigour to discussing future high-consequence pandemics, synthesising historical evidence, peer-reviewed research, and aggregated forecasts into a comprehensive risk analysis. As a nascent sub-field within the broader pandemic prevention space, the GCBR-oriented field would benefit from more clarity around philanthropic priorities. Ruhl presents a strong example for others to follow.

While the report's quantitative approach is laudable, it should be stressed (as Ruhl does) that we cannot take most of the presented estimates at face value. The uncertainty around biological catastrophe risk estimates spans orders of magnitude, and the nature of such tail risks inherently precludes reliable feedback to calibrate forecasts. Moreover, to my knowledge, there are still no estimates that publicly and robustly establish the quantitative cost-effectiveness of interventions to



reduce GCBRs. Still, if interpreted carefully, the report's figures provide a valuable foundation for further analysis.

More generally, Ruhl's report faces a significant epistemic challenge that characterises the broader effort to prevent biological catastrophe in the 21st century: the looming threats from emerging (bio)technologies are deeply uncertain, unpredictable, and sometimes even speculative. This inherent uncertainty means that, relative to other grand challenges like climate change and antimicrobial resistance, philanthropists addressing biotechnological risks must contend with a paucity of empirical evidence and pervasive uncertainty. This uncertainty permeates both the exact scale of the risk – *How does biosecurity compare to other pressing problems?* – and the solution space – *How do we best make progress?* – and is an inherent feature of tackling a global problem that, in many ways, is yet on the horizon.

And yet, the historical record of biological accidents and misuse, the breakneck pace of scientific progress, and the horrific toll of another global pandemic all compel us to face the many unknowns of modern biological risks with decisive and urgent action.



Andrew Snyder-Beattie, Senior Program Officer, Biosecurity and Pandemic Preparedness, Open Philanthropy

Andrew leads Open Philanthropy's work on biosecurity and pandemic preparedness, directing over \$200m in grantmaking. He previously was the Director of Research at the Future of Humanity Institute, University of Oxford, and also worked at a medical research startup. He holds a PhD/DPhil in Zoology from the University of Oxford.

In *Global Catastrophic Biological Risks: A Guide for Philanthropists*, Christian Ruhl makes a persuasive, first-principles case for a philanthropic focus on extreme pandemics. The guide cogently summarizes the state of the art in GCBR research and exhibits two complementary virtues: humility and ambition.

Humility — since uncertainty dominates this field. Little can be said with great confidence about the nature of a catastrophic threat or the trajectory of future biotechnology. In the vein of humility, Ruhl also offers sound advice on what not to do when making grants in this area, in particular around information hazards.

Ambition — since despite the uncertainty, there are bold steps a philanthropist can take to make the world safer! Investments in 1) protective equipment, 2) sterilization technology, 3) pathogen detection, and 4) rapidly responsive medical countermeasures can increase our odds of weathering a biological catastrophe, while investments in smart policy and safeguards against proliferation and misuse of biotechnology could prevent such a catastrophe in the first place.

The coming decades will surprise us in ways both positive and negative. This guide sets a bearing to a world without catastrophic biological risk, whichever way the winds blow.



Introduction

Humanity is underprepared for future pandemics. As explained in the following sections, novel biological threats — especially those caused by engineered pathogens — could kill hundreds of millions or even billions of people and precipitate the collapse of modern civilization. Existing resources generally do not prioritize catastrophic scenarios.⁵ Several policymakers who have worked on pandemic preparedness and response described their experiences in interviews for this report; all agreed that **society is not ready for the next pandemic, natural or engineered.**

Worse, the COVID-19 pandemic was not the wake-up call it could have been. In an interview for this report, Jason Matheny, the CEO of the RAND Corporation, who previously led White House policy on technology and national security at the National Security Council and the Office of Science and Technology Policy and has extensive expertise on catastrophic biological risks, explained that “policymakers are suffering from pandemic fatigue [...] They are just sick and tired of hearing about the pandemic and having to pay for it.”⁶ Similarly, Tom Inglesby, the Director of the Johns Hopkins Center for Health Security, explained in an interview for this report, “[In government] there’s COVID fatigue, let alone interest in taking on things that are bigger than COVID.”⁷ This fatigue is not limited to policymakers but extends to many traditional philanthropic funders, as explained below (see [Societal Spending and Neglectedness](#)). The recent collapse of large philanthropic entities linked to FTX and the withdrawal of other funders has only worsened this situation by removing millions of dollars from biosecurity and pandemic-preparedness philanthropy.

This matters because new advances in the life sciences and in related enabling technologies like artificial intelligence (AI) are creating a rapidly-shifting threat landscape. If society does not take action to regulate access to these technologies — or to promote differential technological development — their powerful capabilities may not only provide benefits to humanity, but may also proliferate to malevolent actors who seek to weaponize synthetic biology.⁸ History shows that state bioweapons programs, non-state bioterrorism, and laboratory leaks are not uncommon.⁹ As more and more actors gain access to the necessary tools and training in genetic engineering, the risk of deliberate and accidental release of such pathogens increases, and with it, the risk of a global

⁵ “Even the best-prepared nations lack sufficient protective equipment for most key personnel, and vaccines and other medical countermeasures could not plausibly be manufactured and distributed in any time frame shorter than months, if they could be developed at all.” (Kevin M Esvelt, “Delay, Detect, Defend: Preparing for a Future in Which Thousands Can Release New Pandemics,” Geneva Papers (Geneva Centre for Security Policy, 2022), 13,

https://dam.gcsp.ch/files/doc/gcsp-geneva-paper-29-22?_gl=1*xm44p1*_ga*ODQxMDI0NiY4LjE2ODM2NTg5MTg.*_ga_Z66DSTVXTJ*MTY4MzczNTIzNi4yLjAuMTY4MzczNTIzNi4wLjAuMA..#page=49&zoom=100.0.0.)

⁶ Interview with Dr. Jason Matheny, 30 May 2023.

⁷ Interview with Dr. Tom Inglesby, 8 June 2023.

⁸ See [The Growing Risk of Biological Catastrophe](#), below. Thanks to Jake Swett for emphasizing the importance of differential technological development.

⁹ See [Bioweapons, Bioterrorism, and Laboratory Leaks](#), below.



catastrophe, civilizational collapse, and even human extinction.

The devastation of the COVID-19 pandemic could have catalyzed a movement towards better pandemic preparedness. Instead, societies remain asleep at the wheel, careening towards a new age of democratized synthetic biology with few guardrails. The specifics of these problems are described in detail below, but one senior policymaker who used to work on biological risks in the U.S. government summarized it bluntly in an anonymized interview for this report: “I just came away from that [i.e. government service on biosecurity] thinking, we’re just completely fucked.”¹⁰

This report is intended as a detailed roadmap for philanthropists who want to do something about these global catastrophic biological risks (GCBRs).¹¹ We hope that this guide will help catalyze more funding on GCBRs, because — as the following sections explain — they are:

- **Important** — Pandemics can kill millions of people, and engineered pathogens are among the few credible extinction risks to modern civilization.
- **Neglected** — Although governments spend billions on health security, we estimate that less than 2% of this money is relevant to the prevention of catastrophic or potential extinction-level biological events.
- **Tractable** — Unlike other important and neglected risks, like risks from advanced artificial intelligence, we have a firm grasp of the problem, and can identify many actionable funding opportunities that can mitigate the risk.

Society needs philanthropists to take up this issue, because no one else has. As one former senior U.S. policymaker interviewed for this project put it, “**without philanthropists funding work in this area, it’s not going to get funded.**”¹²

¹⁰ Anonymized interview.

¹¹ We adopt the following definition of GCBRs proposed by scholars from the Johns Hopkins Center for Health Security: “Those events in which biological agents — whether naturally emerging or reemerging, deliberately created and released, or laboratory engineered and escaped — could lead to sudden, extraordinary, widespread disaster beyond the collective capability of national and international governments and the private sector to control. If unchecked, GCBRs would lead to great suffering, loss of life, and sustained damage to national governments, international relationships, economies, societal stability, or global security.” Monica Schoch-Spana et al., “Global Catastrophic Biological Risks: Toward a Working Definition,” *Health Security* 15, no. 4 (August 1, 2017): 323–28, <https://doi.org/10.1089/hs.2017.0038>. For another discussion of the definitional difficulties, see Jaime Yassif, “Reducing Global Catastrophic Biological Risks,” *Health Security* 15, no. 4 (August 1, 2017): 329–30, <https://doi.org/10.1089/hs.2017.0049>.

¹² Anonymized interview with U.S. biosecurity expert.



The Threat of GCBRs

Key Points

- First-principles reasoning suggests that infectious biological agents pose a unique chain-reaction-based exponential threat to humanity.
- Historically, naturally-arising pandemics have been one of the biggest threats to human civilizations.
- Evolution optimizes for reproductive fitness, not lethality; naturally-evolved pathogens are therefore less likely to cause the most extreme threats.
- Human engineering with pandemic-potential pathogens and synthesis of novel agents, however, could lead to catastrophic pandemics that are worse than anything humanity has ever experienced — deliberately or accidentally.
- Thus the most effective interventions must be robust to engineered threats.
- The importance of this issue crosses worldviews. Infectious disease is a scourge on present generations *and* GCBRs are among the few credible existential threats to the long-term future of humanity.¹³

This section of the report outlines the importance of mitigating global catastrophic biological risks. The explanation of the severity of this risk proceeds in five steps:

1. Using first-principles reasoning via a comparison with nuclear weapons to underscore the unique threat posed by infectious biological agents.
2. Examining the historical track record of pandemics as one of the biggest threats to human civilizations.
3. Explaining the unique threat of engineered pandemics and the myriad of historical examples of bioterrorism, biological weapons programs, and laboratory accidents.

¹³ Notably, although interventions often overlap, they sometimes may not. This report assumes some credence in arguments that place special value on mitigating existential risks, but the majority of the points made in the report do not rely on such arguments. Thanks to Joshua Monrad for pushing for clarification on this.



4. Explaining how GCBRs can therefore pose true *existential risk* to modern civilization, as a potentially massive threat to both current and near-future generations and to the long-term flourishing of humanity.¹⁴
5. Quantifying these arguments with probabilistic forecasting and cost-benefit analyses.

Because these issues have been discussed in greater detail elsewhere, this section only provides a brief overview.

First Principles: The Chain Reaction

There are several attributes that make infectious biological agents a special threat to humanity compared to other potential global catastrophic risks. Fundamentally, pandemic-potential pathogens are **self-replicating**, leading to **chain reactions**, where they turn their hosts effectively into bombs. As Jason Matheny¹⁵ put it in an interview for this report, “you’re making fissile material¹⁶ out of every person on the planet.”¹⁷

Following the analogy of fissile material, it is instructive to compare the basic features of biological risks to the features of nuclear weapons, another global catastrophic risk. On an atomic level, what makes nuclear weapons so powerful is the chain reaction; when the bomb is triggered, a small amount of the heavy elements in the warhead (usually Uranium or Plutonium) undergo nuclear fission, releasing energy and neutrons, which then hit and split more atoms, starting a domino-like chain reaction that releases enormous amounts of energy.

Initially, some of the Manhattan Project atomic scientists working on the world’s first atomic bomb feared that the chain reaction might not stop with the fissionable material contained in the warhead. They worried that nitrogen, which makes up a large fraction of the Earth’s atmosphere, might fuse at the high temperatures of a nuclear explosion, causing a self-sustaining chain reaction that would ignite the world’s atmosphere in one giant fireball. After studying the problem, the physicists concluded that they ought to proceed with testing the bomb; it was highly likely, they believed, that the energy would be too low to fuse nitrogen and that even if it weren’t, no chain reaction would follow. Fortunately, their calculations were correct.

A GCBR-level pandemic, on the other hand, *can* cause a global chain reaction. The pathogen turns its victims into bombs as they reproduce and disperse it further. Continuing the comparison between

¹⁴ Following the philosopher Toby Ord, this report uses the phrase “existential risk” as “a risk that threatens the destruction of humanity’s longterm potential.” This includes human extinction, irrecoverable civilizational collapse, the emergence of a stable global totalitarian regime, and catastrophes of similar magnitude (Toby Ord, *The Precipice: Existential Risk and The Future of Humanity* (New York: Hachette Books, 2020), 37)

¹⁵ Current [president and CEO of the RAND Corporation](#), with extensive expertise and experience working on emerging catastrophic risks.

¹⁶ [I.e., material that can sustain a nuclear chain reaction, and which forms the cores of nuclear weapons.]

¹⁷ Interview with Dr. Jason Matheny, 30 May 2023.



nuclear and biological risks, Matheny explained, “the one positive thing you can say about nuclear weapons is that they don’t self-replicate, which is good. And they also don’t convert the material around them into a bomb.”¹⁸

For historical and bureaucratic reasons, biological weapons are sometimes lumped in with other kinds of destructive weapons. At the United Nations level, for example, the main international body to deal with biological weapons is usually referred to as the Biological Weapons Convention (BWC), but its formal title includes toxins: *The Convention on the Prohibition of the Development, Production and Stockpiling of Bacteriological (Biological) and Toxin Weapons and on their Destruction*.¹⁹ Similarly, in international security as a field of study and within various governments, biological weapons are often lumped in as one part of “CBRN” – chemical, biological, radiological, and nuclear weapons.²⁰ Some biological weapons, however, ought to be in a class of their own:²¹

Conceptual Differences Between CBRN Weapon Types

	Chemical Weapons (incl. toxins)	Biological Weapons	Radiological Weapons	Nuclear Weapons
Self-replication	No	Possible for some BW	No	No
Mutation	No	Possible for some BW	No	No
Global chain reactions	No	Possible for some BW	No	No

Source: Author’s table. (N.B. This table is specifically designed to highlight the differences between biological weapons and other members of the “CBRN” category.)

¹⁸ Interview with Dr. Jason Matheny, 30 May 2023.

¹⁹ “Formally known as ‘The Convention on the Prohibition of the Development, Production and Stockpiling of Bacteriological (Biological) and Toxin Weapons and on their Destruction’, the Convention was negotiated by the Conference of the Committee on Disarmament in Geneva, Switzerland. It opened for signature on 10 April 1972 and entered into force on 26 March 1975. The BWC supplements the 1925 Geneva Protocol, which had prohibited only the use of biological weapons.” (United Nations Office for Disarmament Affairs, “Biological Weapons Convention,” United Nations Office for Disarmament Affairs, accessed July 17, 2023, <https://disarmament.unoda.org/biological-weapons/>)

²⁰ For example: “NATO’s security environment has grown more complex and challenging since 2009, when Allies agreed NATO’s Comprehensive, Strategic-Level Policy for Preventing the Proliferation of Weapons of Mass Destruction (WMD) and Defending against Chemical, Biological, Radiological and Nuclear (CBRN) Threats.” NATO, “NATO’s Chemical, Biological, Radiological and Nuclear (CBRN) Defence Policy,” NATO, accessed July 17, 2023, https://www.nato.int/cps/en/natohq/official_texts_197768.htm.

²¹ Notably, not all biological weapons have attributes of self-replication, mutation, and causing global chain reactions; some non-contagious weapons are likely more limited in scope. Thanks to Andrew Snyder-Beattie for this point.



These basic features — self-replication, mutation, and chain-reaction in abundant hosts (humans) — ought to make us *prima facie* deeply concerned about biological risks. Without knowing anything else about the threat, its basic structure is alarming from first principles. Worse, history demonstrates that this threat is real and large, as the following sections show.

Natural Pandemics

Historically, biological agents have caused some of the worst catastrophes to befall humanity. Luke Muelhauser has attempted to compile a list of historical events that appear to have killed more than 1% of the human population. Out of these, three out of twelve events, or 25%, (the Black Death, the Plague of Justinian, and the 1918 Flu Pandemic) were pandemics.²² Similarly, Dr. Gregory Lewis, who has worked extensively on GCBRs, including at the University of Oxford’s Future of Humanity Institute, has compiled some of the worst biological events to befall humanity; putting Lewis’s estimates into table form and including HIV/AIDS and COVID-19 for reference, we can see how bad each would be as a fraction of an 8 billion world population today:²³

Death Toll Estimates of Historical Pandemics

Event	Time Period	Deaths	Percentage of World Population	Equivalent Deaths Today
Plague of Justinian	541-542 CE	~6 million	~3%	~240 million
The Black Death (i.e., Plague)	1335-1355 CE	20-75 million	~10%	~800 million
1918 Flu Pandemic	1918-1920 CE	50-100 million	2.5% – 5%	200-400 million
HIV/Aids	1980s to present	32-51 million ²⁴	0.4% - 0.6%	32-51 million
COVID-19 Pandemic	2019-2023	17-31 million ²⁵	0.2% - 0.4%	17-31 million

Source: Adapted from Gregory Lewis, “Reducing Global Catastrophic Biological Risks.”

²² Luke Muelhauser, “How Big a Deal Was the Industrial Revolution?,” February 12, 2020, <https://web.archive.org/web/20200212184816/http://lukemuehlhauser.com/industrial-revolution/>.

²³ Gregory Lewis, “Reducing Global Catastrophic Biological Risks,” 80,000 Hours, March 2020, <https://80000hours.org/problem-profiles/preventing-catastrophic-pandemics/full-report/>.

²⁴ These figures are added from the [WHO’s estimates](#).


²⁵ Estimating deaths from COVID-19 is a difficult and controversial endeavor, as is the choice to use excess deaths as the relevant figure. For simplicity, we have used numbers from the *Economist*’s ongoing excess death tracker (“The Pandemic’s True Death Toll,” *The Economist*, July 17, 2023, <https://www.economist.com/graphic-detail/coronavirus-excess-deaths-estimates>.) Whether COVID-19 ought to be counted under the heading of “natural” pandemics is an even more controversial question that this report does not address.



Some pandemics were even more devastating on a regional level. The horrific historical events euphemistically known as the “Columbian Exchange” — including the exposure to new diseases and later the deliberate use of biological weapons by European colonists against American Indian tribes — caused pandemics that killed 98% of some populations and 80% of indigenous people living in present-day Mexico.²⁶

As the COVID-19 pandemic has shown, biological events continue to have the potential to upend life around the world and kill tens of millions of people.²⁷ In other words, even those pandemics that do not come close to the damage of the Black Death have been and continue to be a major cause of death and suffering in the world, as the death toll from influenza pandemics illustrates:

Estimates of Influenza Pandemic Deaths (*Our World in Data*)



Source: Our World in Data, “[The Spanish flu: The global impact of the largest influenza pandemic in history.](#)”

²⁶ Gregory Lewis, “Reducing Global Catastrophic Biological Risks,” 80,000 Hours, March 2020, <https://80000hours.org/problem-profiles/preventing-catastrophic-pandemics/full-report/>. Notably, the population decline of the Americas around this time is a complex event with a variety of causes.

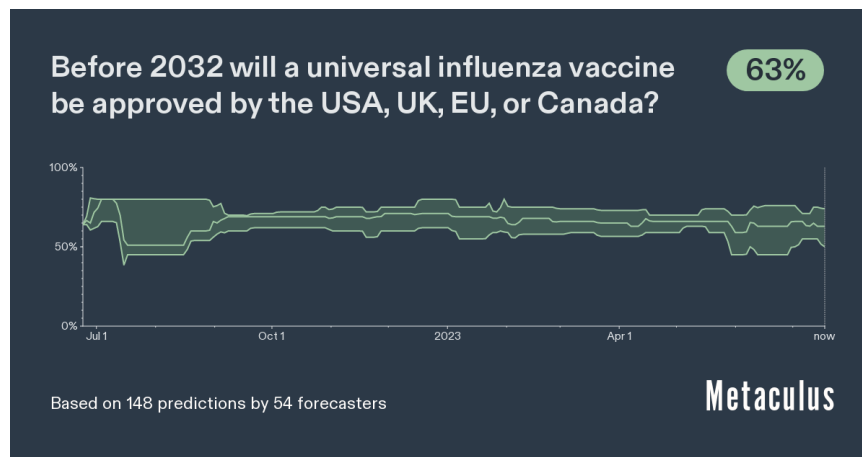
²⁷ Though the extent to which this pandemic should be considered a “natural” event (rather than the result of a laboratory leak) is notoriously disputed.



Naturally-arising pandemics remain substantially neglected by existing philanthropists. Fortunately, some breakthroughs on some specific threats appear likely.²⁸ Forecasters on the probabilistic crowdsourced forecasting website [Metaculus](#) assign a greater than 60% chance to the approval of a universal flu vaccine in either the USA, UK, EU, or Canada.²⁹

A note on forecasts: This report repeatedly cites crowdsourced probabilistic forecasts of the type used on [Metaculus](#), [INFER](#), and other forecast aggregation and collective intelligence platforms. We view these forecasts as one of many tools in our analytic toolbox that help to quantify uncertain beliefs about the threat landscape. These methods can elicit the wisdom of the crowds, but they also have major weaknesses, and forecast accuracy degrades as time horizons expand to the ranges discussed here. For this reason, no part of the argument in this report *relies* on the cited forecasts — they are intended to be illustrative and supplemental material. For more information, including evidence on forecaster accuracy, see [A Guide to Forecasting at Founders Pledge](#).

Forecasts on the Approval Probability of a Universal Flu Vaccine by 2032



Source: Metaculus, “[Before 2032 will a universal influenza vaccine be approved by the USA, UK, EU, or Canada?](#)” Aggregate forecast as of 17 July 2023.

²⁸ Of course, vaccine approval remains a far cry from disease eradication and is only one of many tools in the fight against natural pandemics. Thanks to Joshua Monrad for this point in a round of external reviews.

²⁹ Such forecasts are, of course, only compelling to the extent that readers trust the crowd forecasters, as discussed in the accompanying “note on forecasts.” Readers interested in the track record of Metaculus forecasters can find more information [here](#).



Viral, Bacterial, and Other Threats

One important consideration for philanthropic prioritization is whether viruses, bacteria, or other biological agents, like fungi, are more likely to cause a global catastrophic biological event.³⁰ Experts in biosecurity generally believe that viruses, specifically respiratory viruses, have the most pandemic potential, and are therefore most concerning for natural and many engineered events, but this report argues that **sophisticated malevolent actors could weaponize a wide range of pathogens, such that a pathogen-agnostic approach is the best defense against engineered pandemics.**³¹ This report treats this question only at a high level; a detailed discussion of the *specific* attributes that would make a biological agent a true global catastrophic biological risk is not only beyond the scope of this report, but could also risk spreading information that may inspire malevolent actors. Rather, the following table briefly surveys the crucial considerations that explain why viruses appear especially concerning, as outlined in [The Characteristics of Pandemic Pathogens](#), a report by the Johns Hopkins Center for Health Security:

³⁰ Forecasters on Metaculus, for example, widely believe that a few virus families are most likely to cause the next pandemic (aggregate forecasts as of 17 July 2023):

1. 51% — Orthomyxoviridae virus family.
2. 32% — One of the known human-infecting virus families excluding Coronaviridae and Orthomyxoviridae.
3. 23% — Coronaviridae virus family (excluding SARS-CoV-2).
4. 10% — A microorganism that is not from one of the known human-infecting virus families.

(Metaculus, “What Type of Microorganism Will Cause the Next Pandemic?,” Metaculus, July 8, 2022, <https://www.metaculus.com/questions/11647/microorganism-responsible-for-next-pandemic/>).

As discussed in later sections, however, such attempts at prediction are fraught with problems. These problems and the risk-shifting that accompanies their answering can be partly avoided with threat-agnostic defenses (see [Prioritize Pathogen- and Threat-Agnostic Approaches](#)).

³¹ For example: “This analysis is primarily focused on viral pathogens because they are the most likely to spread rapidly and cause catastrophic harm and because they have relatively small genomes compared to bacterial or eukaryotic pathogens” (Sarah Carter, Jaime Yassif, and Chris Isaac, “Benchtop DNA Synthesis Devices: Capabilities, Biosecurity Implications, and Governance” (The Nuclear Threat Initiative (NTI|bio), May 10, 2023), 23, <https://www.nti.org/analysis/articles/benchtop-dna-synthesis-devices-capabilities-biosecurity-implications-and-governance/>).

Dr. Amesh Adalja also provides an accessible overview in his 2018 presentation “Pandemic Pathogens” (Effective Altruism Global 2018: San Francisco, 2018), <https://www.effectivealtruism.org/articles/ea-global-2018-pandemic-pathogens>.



Crucial Considerations for Viruses, Bacteria, Fungi, and More

Agent Type	Example Pathogen	Crucial Considerations
Viruses	Influenza viruses	<ul style="list-style-type: none">• Mutability, high rate of replication (notably, mutability can be a disadvantage for engineered weapons).³²• Lack of broad-spectrum antivirals (and technical barriers to this³³)• Frequently causes pandemics and epidemics: 1918 Influenza, COVID-19 pandemic, HIV, SARS, MERS.
Bacteria	<i>Yersinia Pestis</i> (Plague)	<ul style="list-style-type: none">• Existence of broad-spectrum antibiotics (e.g. Penicillin)• Slower speed of replication³⁴• Some antibiotic resistance• Black Death
Fungi	<i>Candida Auris</i>	<ul style="list-style-type: none">• Temperature-restricted (see “fungal filter hypothesis”), mostly threatening to immunocompromised hosts.³⁵• But absence of vaccines, potential intercontinental airborne spread, and increasing drug resistance, and other factors make fungi a potentially dangerous class of pathogens.³⁶

³² “Viruses possess higher capacity for genetic mutability due to both the structure of their genomes and the generation time for replication in which large numbers of progeny virus are created each day” and “The high rate of replication of viruses — for instance, over 1 trillion hepatitis C virions are produced per day in a human infection — coupled with the mutability inherent in such short generation times give viruses an unrivaled plasticity. This plasticity allows for host adaptability, zoonotic spillover, and immune system evasion.” (Amesh Adalja et al., “The Characteristics of Pandemic Pathogens” (Johns Hopkins Center for Health Security, 2018), <https://centerforhealthsecurity.org/sites/default/files/2022-12/180510-pandemic-pathogens-report.pdf>.) On the disadvantages of mutability: If a weapon is optimized for certain traits (e.g. high transmissibility), the vast majority of mutations are likely to move the weapon’s characteristics away from this optimized state (e.g. to lower transmissibility). Moreover, mutability complicates targeting. Thanks to Andrew Snyder-Beattie for this point in a round of external reviews.

³³ These technical barriers are in part due to the simplicity of viruses. Because they are simply encased genetic material, it is difficult to disrupt them without also disrupting the functioning of the host, as explained in the following discussion at the Oxford University Personalized Medicine Society around minute 12 (“There’s not much that you can attack them with that doesn’t also attack your own cells”) Amesh Adalja, “The Characteristics of Pandemic Pathogens (OUPM Talk)” (Oxford University Personalized Medicine Society, 2020), <https://cpm.well.ox.ac.uk/video/dr-amesh-adalja-characteristics-pandemic-pathogens>.

³⁴ “For example, a human infected with the hepatitis C virus (an RNA virus) produces trillions of virions per day, whereas the doubling time of *Yersinia pestis*, the cause of plague, is 1.25 hours” (Amesh Adalja et al., “The Characteristics of Pandemic Pathogens” (Johns Hopkins Center for Health Security, 2018), 12)

³⁵ “[A] fungal filter is hypothesized to have existed and may be partly responsible for mammalian warm-bloodedness.” and “Human infections with fungi tend to be severely damaging only in an immunocompromised host. The human innate immune system contends with countless fungal spores that are present in every breath of air.” (Amesh Adalja et al., “The Characteristics of Pandemic Pathogens” (Johns Hopkins Center for Health Security, 2018), 12)

³⁶ Arturo Casadevall, “Global Catastrophic Threats from the Fungal Kingdom : Fungal Catastrophic Threats,” *Current Topics in Microbiology and Immunology* 424 (2019): 21–32, https://doi.org/10.1007/82_2019_161.



Crucial Considerations for Viruses, Bacteria, Fungi, and More (*continued*)

Agent Type	Example Pathogen	Crucial Considerations
Prions	Creutzfeldt-Jakob disease (“mad cow disease”)	<ul style="list-style-type: none"> Slow transmission through specific routes (cannibalism, infected food, etc.)³⁷
Protozoan parasites	Plasmodium (Malaria)	<ul style="list-style-type: none"> The only mammalian species extinction ever was vector-borne parasite (Christmas Island Rat) Defenses against vectors (e.g., mosquitoes) available
Other	Amoeba	<ul style="list-style-type: none"> Similar considerations apply — generally limited pathogenicity and/or transmissibility³⁸
Speculative (e.g. non-carbon-based)	N/A	<ul style="list-style-type: none"> No known agents. <ul style="list-style-type: none"> May require major technological breakthroughs before becoming a concern. Challenging detection and defense. Existing guidelines on planetary protection for extra-terrestrial agents.³⁹

Source: Adapted from Adalja *et al.*, “The Characteristics of Pandemic Pathogens,” Johns Hopkins Center for Health Security. See footnotes for direct quotations.

This has several implications for biosecurity. If philanthropists are focused on worst-case scenarios, for example, then it becomes clear that **some traditional bacteriological weapons do not meet this threshold**, even though they could be weapons of mass destruction — philanthropists focused on cost-effectiveness may wish to de-prioritize some such threats.⁴⁰ For example, the U.S. biodefense community has focused heavily on the bacteriological threat of Anthrax, in part because of its history as a biological weapon and the 2001 Amerithrax scare (which involved the mailing of Anthrax-laced letters to media outlets and government officials and resulted in five deaths and one of the largest FBI investigations in U.S. history). In addition to the crucial considerations surrounding bacteriological agents, however, Anthrax is not human-to-human transmissible; the spores cannot be

³⁷ “The transmission characteristics of prion diseases are such that very extraordinary circumstances, on par with human cannibalism or massive food contamination, must be present for a GCBR-level risk to be present for humans. Additionally, and almost by definition, such an event would be slow-moving (prions were once known as “slow viruses”).” (Amesh Adalja *et al.*, “The Characteristics of Pandemic Pathogens” (Johns Hopkins Center for Health Security, 2018), 13, <https://centerforhealthsecurity.org/sites/default/files/2022-12/180510-pandemic-pathogens-report.pdf>.)

³⁸ “Amoeba, ectoparasites, and helminths all have delimited pandemic risk, as they are constrained by pathogenicity, transmissibility, or both.” (Amesh Adalja *et al.*, “The Characteristics of Pandemic Pathogens” (Johns Hopkins Center for Health Security, 2018), 12, <https://centerforhealthsecurity.org/sites/default/files/2022-12/180510-pandemic-pathogens-report.pdf>.)

³⁹ Thanks to Jake Swett for pointing to this in a round of external reviews.

⁴⁰ The reasons for focusing on the most extreme threats are discussed below (see [Prioritize Worst-Case Scenarios](#)).



transmitted from human to human. This means that would-be bioweaponers need to produce industrial quantities of the agent and find effective ways of delivering it. This has historically been an obstacle for state bioweapons programs.⁴¹

This does *not* mean, however, that philanthropists ought to focus on viruses to the total exclusion of other pathogen classes, like bacteria. Rather, the threat landscape becomes significantly more complex with engineered pathogens created by sophisticated malevolent actors, especially as bioscience continues to advance. The characteristics above illustrate, therefore, the myriad threat vectors for an engineered pathogen and the difficulty of predicting what kinds of pathogen will cause the most extreme pandemics. For example, fungi may appear unlikely to be a natural pandemic threat because of their temperature-restricted range, but a sophisticated malevolent actor could view society's relative unpreparedness for a fungal pandemic as a reason for pursuing an engineered fungal weapon. (This general pattern is discussed in greater detail below, under [The Threat is Adaptive](#) and [Pathogen- and Threat-Agnostic Approaches](#))

As discussed below, this complex array of threats has major implications for biodefense. If readers are concerned about the cost-effectiveness of our interventions — both as philanthropists and as subjects of national governments — then it becomes clear that current societal spending is over-allocated on known bacteriological threats like Anthrax, and under-invested in threat-agnostic preparedness and response.

Natural Extinction Risk

There are many reasons to worry about global catastrophes, including the immense death tolls, suffering, grief, economic losses, and more. For some extreme events, these reasons may include total human extinction. Are natural pandemics likely to cause the extinction of humanity?⁴² The majority of the conclusions drawn in this report do not hinge on the question of extinction, but it is nonetheless a critical question, especially for readers who place special moral value on the long-term future of humanity.

One argument against natural pandemics causing an existential threat comes from the fact of humanity's continued existence; because we have survived for 200,000 to 300,000 years as a species, we ought to expect the baseline risk of extinction to be low.⁴³ Thus, analysts have estimated a range of upper bounds on the “background rate of extinction;” that range is between 1-in-14,000 (~.007%) and 1-in-140,000 (~.0007%) per year (inclusive of biological events), depending on which

⁴¹ On this and related obstacles, see Sonia Ben Ouagrham-Gormley, *Barriers to Bioweapons: The Challenges of Expertise and Organization for Weapons Development*, Cornell Studies in Security Affairs (Cornell University Press, 2014), <https://www.cornellpress.cornell.edu/book/9780801452888/barriers-to-bioweapons/>.

⁴² Later sections discuss biological existential risks in greater detail ([Existential Threats](#)).

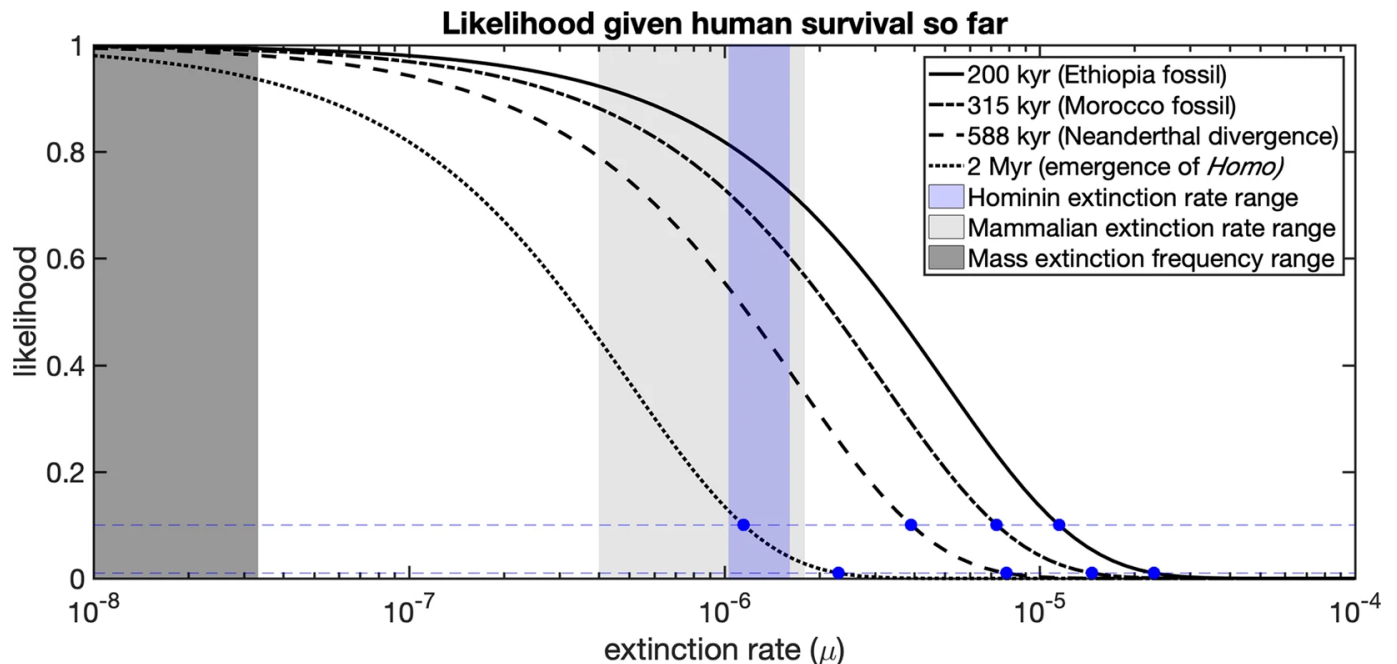
⁴³ As explained by Gregory Lewis: “If the baseline risk was 10% per century, the chance of 'getting lucky' 2000 times to observe this history is very low, whilst much less surprising if the risk was 0.001% per century. Analogously, if we survive a firing squad a dozen times in a row, this is strong evidence we are not faced with competent executioners” (Gregory Lewis, “Reducing Global Catastrophic Biological Risks,” 80,000 Hours, March 2020, <https://80000hours.org/problem-profiles/preventing-catastrophic-pandemics/full-report/>.)



fossil evidence and cutoff points one uses:⁴⁴

Snyder-Beattie *et al.*'s Background Extinction Rate Curve

(NB, the y-axis represents the likelihood of extinction rates listed on the x-axis *given human survival*, not the probability of extinction *per se*)



Source: Snyder-Beattie *et al.*, “[An Upper Bound for the Background Rate of Human Extinction](#),” *Nature Scientific Reports*.

Some analyses extend this reasoning to pandemic risk. Thus, as Gregory Lewis has written in an overview of pandemic risk, “Naturally arising’ biological extinction events seem unlikely given the rarity of ‘pathogen driven’ extinction events in natural history, and the 200,000 year lifespan of anatomically modern humans.”⁴⁵ This approach is highly sensitive to the cutoff point of the relevant time period (in this case, ~200,000 years).

Deriving an extinction rate from humanity’s long survival may be appropriate for some global catastrophic risks (like the risks of impact from asteroids and other near-Earth objects). It is not clear, however, why *Homo Sapiens*’ speciation is an appropriate cutoff date for the background rate of extinction on pandemic risk. A more appropriate cutoff might be the first agricultural revolution about 11,000 years ago — shortening the observed “survival window” by a factor of more than 18. If

⁴⁴ Andrew E. Snyder-Beattie, Toby Ord, and Michael B. Bonsall, “An Upper Bound for The Background Rate of Human Extinction,” *Scientific Reports* 9, no. 1 (July 30, 2019): 11054, <https://doi.org/10.1038/s41598-019-47540-7>.

⁴⁵ Gregory Lewis, “Reducing Global Catastrophic Biological Risks,” 80,000 Hours, March 2020, <https://80000hours.org/problem-profiles/preventing-catastrophic-pandemics/full-report/>.



humans before then mostly lived in small hunter-gatherer societies with little inter-group contact, we would naturally expect a low risk of a pandemic, but this may tell us little about modern global civilization. Not only do pandemics rely upon high-density populations for transmission, but many known pandemic pathogens appear to have had zoonotic origins related to the domestication of certain animals, again linked to the agricultural revolution.⁴⁶ Therefore, estimates of natural pandemic extinction rates based on the “background rate of extinction” calculated from the lifespan of *Homo Sapiens* may underestimate the risk by more than an order of magnitude.

There are other reasons for supposing that the baseline risk of extinction is relatively low. For example, disease-related mammalian species extinctions appear to be rare — there is only one known example, the mosquito-borne parasite-driven extinction of the Christmas Island Rat — and there appears to be virulence-transmission tradeoffs in evolutionary selection pressures (see [Virulence, Transmission, and Engineered Pandemics](#)).⁴⁷ Several complicating factors — observer selection effects,⁴⁸ changes in the connectedness of civilization,⁴⁹ modern countermeasures,⁵⁰ the apparent rarity of pathogens with GCBR potential — and concomitant uncertainties make it clear that such estimates are highly simplistic. The scholar of global catastrophic risk David Manheim has enumerated these considerations and uncertainties for natural pandemic extinction risk:⁵¹

⁴⁶ “[T]he most important infectious diseases of modern food-producing human populations also include diseases that could have emerged only within the past 11,000 years, following the rise of agriculture” and “Current information suggests that 8 of the 15 temperate diseases [a list created by the authors] probably or possibly reached humans from domestic animals (diphtheria, influenza A, measles, mumps, pertussis, rotavirus, smallpox, tuberculosis)” Nathan D. Wolfe, Claire Panosian Dunavan, and Jared Diamond, “Origins of Major Human Infectious Diseases,” *Nature* 447, no. 7142 (May 2007): 279, 281, <https://doi.org/10.1038/nature05775>.

⁴⁷ “The number of mammalian species that have gone extinct due to disease (as opposed to other factors) is very limited, with just 1 confirmed example” and “Optimal virulence theory predicts that pathogens suffer a trade-off between virulence and transmissibility.” (Piers Millett and Andrew Snyder-Beattie, “Human Agency and Global Catastrophic Biorisks,” *Health Security* 15, no. 4 (August 2017): 335, <https://doi.org/10.1089/hs.2017.0044>.)

⁴⁸ I.e., “a form of sampling bias, in which the sample of observed events is not representative of the universe of all events, but only representative of the set of events compatible with the existence of suitably positioned observers.” Milan M. Ćirković, Anders Sandberg, and Nick Bostrom, “Anthropic Shadow: Observation Selection Effects and Human Extinction Risks,” *Risk Analysis* 30, no. 10 (2010): 1496, <https://doi.org/10.1111/j.1539-6924.2010.01460.x>.

⁴⁹ I.e., widespread air travel, growth in and of urban areas, etc., which may make it easier for a pathogen to spread rapidly.

⁵⁰ I.e., modern vaccines, drugs, etc., which can help suppress outbreaks.

⁵¹ David Manheim, “Questioning Estimates of Natural Pandemic Risk,” *Health Security* 16, no. 6 (December 1, 2018): 381–90, <https://doi.org/10.1089/hs.2018.0039>.



Manheim on Estimates of Natural Pandemic Risk

<i>Evidence</i>	<i>Sources of Uncertainty</i>
No human extinction	Anthropic shadow
	Extirpations may be unobserved and not appear in the fossil record.
	Changing ecological and social conditions
Few near misses	Historical extirpations in humans might now have GCBR potential.
	Mammalian extinctions and extirpations may be more common than appreciated. ¹⁸
Limited historic virality, infectivity of events	Population density, travel patterns, and interaction with animals have changed.
	“Optimal virality” generally reduces virality over time, but novel pathogens have not yet faced selection pressure.
GCBR pathogens are relatively rare	Little evidence has been gathered about the distributions of stealth, acute, and robust pathogens.
	Other GCBR pathogens would be found by health systems early enough to be stopped.
Effectiveness of modern health response	Historical effectiveness lower than current systems, which have improved over time.
	Effectiveness depends on pathogen visibility and international cooperation.

Source: Manheim, “[Questioning Estimates of Natural Pandemic Risk](#),” *Health Security*.

From surveying the literature, experts appear to disagree about the *aggregate sign* of the effect of these additional factors.⁵² These and related considerations lead Gregory Lewis to assign an **extinction probability below 1% this century** to natural pandemics.⁵³ The implied crowdsourced probability for a global catastrophe (defined as a sudden 10% population decline) due to a naturally-occurring pandemic on Metaculus is approximately 4.4%:

⁵² As Gregory Lewis has written, “Estimates of the aggregate impact of these changes are highly non-resilient even with respect to sign (I think it is risk-reducing, but reasonable people disagree).” Gregory Lewis, “Reducing Global Catastrophic Biological Risks,” 80,000 Hours, March 2020, <https://80000hours.org/problem-profiles/preventing-catastrophic-pandemics/full-report/>.

⁵³ Gregory Lewis, “Reducing Global Catastrophic Biological Risks,” 80,000 Hours, March 2020, <https://80000hours.org/problem-profiles/preventing-catastrophic-pandemics/full-report/>.



Variable	Value	Source
P(>10% sudden population decline) The probability that there is a sudden population decline of at least 10% by 2100.	37%	By 2100, will the human population decrease by at least 10% during any period of 5 years?
P(natural pandemic cause >10% population decline) The probability that, conditional on a 10% population decline, such a decline was caused by natural pandemics.	12%	If a global catastrophe occurs, will it be due to naturally occurring pandemics?
Implied P(>10% population decline due to natural pandemic). The probability that a natural pandemic causes a 10% population decline, calculated by multiplying the first probability by the conditional probability.	4.4%	Calculation: P(>10% pop. decline)*P(natural pandemic cause >10% population decline)

It is important to emphasize that a 4.4% chance of such a massive population decline this century is unnervingly and unacceptably high, and philanthropists ought to invest more in the prevention of such extreme naturally-occurring events. How many such natural events could also lead to human extinction remains an uncertain question. One thing is clear, however; as discussed below, the risk of both global catastrophe and human extinction due to *engineered* pathogens may be much higher. This does not necessarily mean that philanthropists ought not focus on the ongoing scourge of natural pandemics, but that the most effective interventions must at least be robust to engineered threats (see [Prioritize Pathogen- and Threat-Agnostic Approaches](#) and [Fund Robust Interventions](#)).

Engineered Pandemics

The previous section discussed the threat of natural pandemics; this section in turn discusses “engineered” pandemics, including those caused by state biological weapons programs, bioterrorists, and scientific work with enhanced potential pandemic pathogens (ePPP, sometimes called “gain of function”) broadly defined.⁵⁴

One important insight for philanthropists interested in mitigating the worst possible pandemics is that evolution optimizes pathogens for reproductive fitness, and there exists a

⁵⁴ This section builds on the “anthropogenic GCBR hypothesis” argument outlined in Piers Millett and Andrew Snyder-Beattie, “Human Agency and Global Catastrophic Biorisks,” *Health Security* 15, no. 4 (August 2017): 335–36, <https://doi.org/10.1089/hs.2017.0044>.



virulence-transmission trade-off.⁵⁵ The action-relevant insight for philanthropists does not rely on complex ideas like “optimal virulence theory,” however, but on a much simpler insight — actors who optimize not for reproductive fitness but for maximum damage (e.g. bioterrorists) may be able to out-perform natural selection via engineered pathogens.⁵⁶ Nature is therefore *not* — as some commonly claim — “the biggest bioterrorist.”⁵⁷ Most pandemics in human history were natural, and natural pandemics have been and will remain a major threat to human life and wellbeing; this is one reason why a threat-agnostic approach could be especially promising. Nonetheless, because the worst kinds of pandemics are likely to be engineered, engineered pandemics ought to be the “pacing” threat for philanthropists interested in avoiding biological catastrophes.⁵⁸

Pathogens with very high virulence may kill or incapacitate their hosts too quickly to spread further. This “optimal virulence theory” and related ideas are complex theories, but the general insight on the trade-off between virulence and transmission is intuitive.⁵⁹ If a human falls very ill, they may stay at home and not spread the pathogen in society. If they die, the pathogen again often loses its ability to spread (e.g. because the host is no longer coughing). In short, viruses that incapacitate their hosts

⁵⁵ “Virulence and transmission is often a trade-off, and so evolutionary pressures could push against maximally lethal wild-type pathogens” Piers Millett and Andrew Snyder-Beattie, “Existential Risk and Cost-Effective Biosecurity,” *Health Security* 15, no. 4 (August 2017): 373, <https://doi.org/10.1089/hs.2017.0028>.

⁵⁶ “Evolution is infamously myopic, and its optimisation target is reproductive fitness, rather than maximal damage to another species (cf. optimal virulence). Nature may not prove a peerless bioterrorist; dangers that emerge by evolutionary accident could be surpassed by deliberate design.” (Gregory Lewis, “Reducing Global Catastrophic Biological Risks,” 80,000 Hours, March 2020, <https://80000hours.org/problem-profiles/preventing-catastrophic-pandemics/full-report/>.)

⁵⁷ “Nature is the biggest bioterrorist” is a common sentiment. One recent (2023) example: “There has recently been a lot of debate about revising regulations for experiments with viruses in laboratories, and this discussion is good to have. However, **we should not forget that ‘Mother Nature’ is the biggest bioterrorist** and that unsafe farming practices assist her in creating potentially dangerous viruses.” (Florian Krammer and Stacey Schultz-Cherry, “We Need to Keep an Eye on Avian Influenza,” *Nature Reviews Immunology* 23, no. 5 (May 2023): 268, <https://doi.org/10.1038/s41577-023-00868-8>.)

⁵⁸ As an analogy, most wars are small conflicts between minor powers, but it would be wrong to say that “small states are the biggest warmakers.” Great powers fight rarely, but when they do, they can unleash orders of magnitude more destruction. Defense planners therefore base their decisions based on the warmaking capacities of the most powerful rivals and adversaries and defense documents describe China as the “pacing threat” for the United States.

⁵⁹ As Millett and Snyder-Beattie explain, “Optimal virulence theory predicts that pathogens suffer a trade-off between virulence and transmissibility. This would suggest that natural evolution would be unlikely to produce disease at maximum lethality and transmissibility. However, the debate around optimal virulence theory is still not resolved and constitutes another ripe area for GCBR research.” Piers Millett and Andrew Snyder-Beattie, “Human Agency and Global Catastrophic Biorisks,” *Health Security* 15, no. 4 (August 2017): 335–36, <https://doi.org/10.1089/hs.2017.0044>. On the history of the idea and the complexity of trade-offs, see S. Alizon et al., “Virulence Evolution and the Trade-off Hypothesis: History, Current State of Affairs and the Future,” *Journal of Evolutionary Biology* 22, no. 2 (February 2009): 245–59, <https://doi.org/10.1111/j.1420-9101.2008.01658.x>.



too quickly or too harshly may reproduce less successfully than milder viruses. This may help explain the apparently low rate of pandemic-mediated extinctions among animal species discussed earlier.⁶⁰

These considerations lead to an important conclusion — **engineered pandemics could be far worse than naturally-arising pandemics**. As Piers Millett and Andrew Snyder-Beattie write, “Under optimal virulence theory, natural evolution would be an unlikely source for pathogens with the highest possible levels of transmissibility, virulence, and global reach. But advances in biotechnology might allow the creation of diseases that combine such traits.”⁶¹ In 2001, for example, scientists accidentally created a mousepox strain with a 100% case fatality rate in mice (and 60% CFR in previously-immunized mice); if such attributes were engineered in a highly transmissible respiratory virus that spread while victims remain asymptomatic, the consequences could be disastrous.⁶²

The exact biological details of such an agent are not as relevant as the broader point that it is possible to out-engineer evolution. As discussed throughout this report, disclosing blueprints of certain pathogens could spread dangerous information to malevolent actors, but some archetypal hypothetical pandemic pathogens may be discussed in broad strokes.⁶³ For example, one such archetype might be a respiratory virus with 99% fatality rate, a contagiousness similar to measles, broad vaccine resistance, and a long pre-symptomatic infectious period.⁶⁴ As far as we know, nature has not created such a pathogen; in theory, humans could try. Fortunately, engineering and working

⁶⁰ Claire Zabel has explained this intuitively: “people have argued that there's some selection pressure against a naturally emerging highly virulent pathogen because when pathogens are highly virulent, often their hosts die quickly and they try to rest before they die and they're not out in society spreading it the way you might spread the cold, if you go to work when you have the cold” (Claire Zabel, “Biosecurity as an EA Cause Area,” <https://www.effectivealtruism.org/articles/biosecurity-as-an-ea-cause-area-claire-zabel>.)

⁶¹ Piers Millett and Andrew Snyder-Beattie, “Existential Risk and Cost-Effective Biosecurity,” *Health Security* 15, no. 4 (August 2017): 374, <https://doi.org/10.1089/hs.2017.0028>.

⁶² Piers Millett and Andrew Snyder-Beattie, “Human Agency and Global Catastrophic Biorisks,” *Health Security* 15, no. 4 (August 2017): 335–36, <https://doi.org/10.1089/hs.2017.0044>. The authors of the mousepox study in question explain: “Infection of BALB/c or C57BL/6 [i.e. both mousepox-sensitive and mousepox-resistant] mice with ECTV-IL4(TK+) [the strain in question] proved to be **uniformly lethal**, with mean survival times (days) being 7.4 ± 0.7 ($n = 10/10$) and 8.6 ± 1.2 ($n = 10/10$), respectively.” (emphasis added). It was 60% lethal in mice that had been previously immunized with “the attenuated TK- virus ECTV-602”: “[Immunized] mice challenged with ECTV-IL4(TK+) displayed an exacerbated response characterized by extreme swelling of the inoculated foot 24 and 48 h p.i. [...] and 60% of the mice died between days 6 to 8 p.i.” 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 (February 2001): 1205–10, <https://doi.org/10.1128/JVI.75.3.1205-1210.2001>.

⁶³ For more on this, see Chris Bakerlee and Tessa Alexanian, “Info hazard guidance for biosecurity discussions,” <https://docs.google.com/document/u/0/d/1VSfU3GiZumHDX2hoz3YY1PT2dQHtkbrfO8xLxI9BTGE/mobilebasic#ftnt3>.

⁶⁴ “[A] hypothetical respiratory virus (let’s say “SARS-CoV-X”) with a 99% fatality rate, measles-level contagiousness, resistance to all current vaccines, and/or a 6-month pre-symptomatic infectious period” Chris Bakerlee and Tessa Alexanian, “Info hazard guidance for biosecurity discussions,” <https://docs.google.com/document/u/0/d/1VSfU3GiZumHDX2hoz3YY1PT2dQHtkbrfO8xLxI9BTGE/mobilebasic#ftnt3>.

with dangerous pathogens has proven difficult even for well-resourced bioweapons programs.⁶⁵ As discussed in a later section ([The Changing Threat Landscape](#)), however, advances in synthetic biology appear to degrade some of the barriers to engineering and handling pandemic pathogens.

Finally, there is a geographic consideration when comparing natural spillover or accidental release with deliberate release of pathogens — deliberate release is more likely to occur in multiple locations at once. As the biology professor Kevin Esvelt has pointed out, malevolent actors could release a contagious pathogen at multiple airports around the world, making containment far more challenging.⁶⁶

Thus funders concerned with the cost-effectiveness of their grants have long suggested that focusing on engineered pandemics rather than only naturally-occurring events may be an important part of an effective grantmaking program.⁶⁷ These considerations help explain why forecasters on crowdsourced forecasting aggregation platforms like Metaculus believe that a global catastrophe — defined simplistically here as a sudden decline in world population over 10% — is far more likely to occur due to bioengineered organisms than to naturally occurring pandemics:⁶⁸

⁶⁵ For a detailed overview, see Sonia Ben Ouagrham-Gormley, *Barriers to Bioweapons: The Challenges of Expertise and Organization for Weapons Development*, Cornell Studies in Security Affairs (Cornell University Press, 2014), <https://www.cornellpress.cornell.edu/book/9780801452888/barriers-to-bioweapons/>.

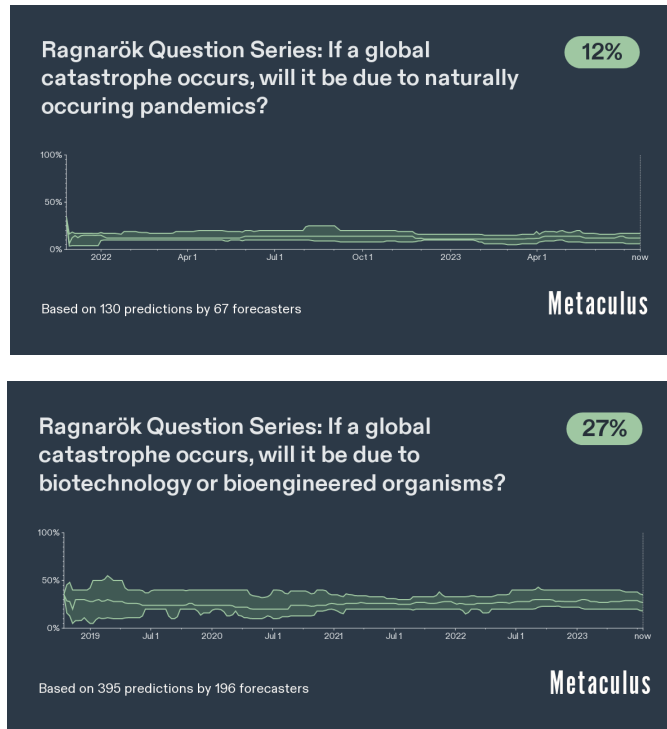
⁶⁶ “Animal viruses manifestly do not spill over to cause pandemics in multiple airports simultaneously, and certainly not in groups, but once enough of them are identified, thousands of people will be capable of causing both” (Kevin M Esvelt, “Delay, Detect, Defend: Preparing for a Future in Which Thousands Can Release New Pandemics,” Geneva Papers (Geneva Centre for Security Policy, 2022), 13, https://dam.gcsp.ch/files/doc/gcsp-geneva-paper-29-22?_gl=1*xm44p1*_ga*ODOxMDI0NjY4LjE2ODM2NTg5MTg.*_ga_Z66DSTVXTJ*MTY4MzczNTIzNi4yLjAuMTY4MzczNTIzNi4wLjAuMA..#page=49&zoom=100.0.0.)

⁶⁷ For example, in 2017, when working for Open Philanthropy, Jaime Yassif wrote: “What is the most likely source of a GCBR? Our intuition is that, in the longer term, engineered pathogens — dispersed through a deliberate attack or accidental release — pose the greatest potential risk for this type of event. However, we don’t have enough evidence to dismiss the possibility of a GCBR from a natural outbreak. If we saw convincing evidence that countered our intuition, it could lead us to change our working assumptions on this point.” Jaime Yassif, “Reducing Global Catastrophic Biological Risks,” *Health Security* 15, no. 4 (August 1, 2017): 329–30, <https://doi.org/10.1089/hs.2017.0049>.

⁶⁸ The recently-published results of the Existential Risk Persuasion Tournament (XPT) provide similar results, with the notable exception of superforecaster opinion on catastrophes. On global catastrophe (defined as “death of at least 10% of humans alive at the beginning of a five-year period”), domain experts assigned a 3% chance to engineered pandemic catastrophe and a 0.85% chance to natural pandemic catastrophe, whereas superforecasters assigned 0.8% and 1%, respectively. On natural pathogen extinction risk, domain experts assigned a 1% chance to engineered pandemic extinction risk, and a 0.01% chance to natural pandemic extinction risk, whereas superforecasters assigned a 0.01% chance to engineered pandemic extinction risk and a 0.0018% chance to natural pandemic extinction risk.



Crowdsourced Predictions on Natural vs. Engineered Pandemics



Source: Metaculus, “[If a global catastrophe occurs, will it be due to naturally occurring pandemics?](#)” and “[If a global catastrophe occurs, will it be due to biotechnology or bioengineered organisms?](#)” Aggregated forecasts as of 17 July 2023.

Similarly, a recently-concluded large-scale tournament known as the [Existential Risk Persuasion Tournament](#) (XPT) included some questions on pandemic catastrophes.⁶⁹ Notably, superforecasters and domain experts disagreed on the relative importance of engineered and natural pandemics for catastrophic events (“death of at least 10% of humans alive at the beginning of a five-year period”), but both agreed that extinction is far more likely due to an engineered pathogen.⁷⁰ The question of whom to trust more — superforecasters with a proven track record on short-term questions but with

⁶⁹ Notably, these were excluded from the main tournament, and were phrased as one-shot questions. The authors explain: “Because of concerns among our funders about information hazards [...] we did not include this question in the main tournament, but we did ask about risks from engineered and natural pathogens in a one-shot separate postmortem survey to which most XPT participants responded after the tournament.” Ezra Karger et al., “Forecasting Existential Risks: Evidence from a Long-Run Forecasting Tournament” (Forecasting Research Institute, 2023), 16.

⁷⁰ Ezra Karger et al., “Forecasting Existential Risks: Evidence from a Long-Run Forecasting Tournament” (Forecasting Research Institute, 2023), 4.



unknown accuracy over longer periods of time, or domain experts with detailed insider knowledge and models of the world but no proven forecasting track record — is highly complex.⁷¹

For simplicity, we have provided the mean of the two groups' forecasts below:

XPT Tournament Estimates

Threat Type	Superforecaster Estimate	Domain Expert Estimate	Arithmetic Mean of Probabilities	Geometric Mean of Odds, Expressed as Probability
Engineered pandemic catastrophe	0.80%	3%	1.90%	1.55%
Natural pandemic catastrophe	1%	1%	1%	0.92%
Engineered pandemic extinction	0.01%	1.00%	0.51%	0.10%
Natural Pandemic Extinction	0.0018%	0.01%	0.0059%	0.0042%

Source: Adapted from Ezra Karger et al., “Forecasting Existential Risks: Evidence from a Long-Run Forecasting Tournament” (Forecasting Research Institute, 2023).

Both the arithmetic mean of probabilities and the geometric mean of odds of the two forecasts suggests a strong difference between engineered and natural threats (1.7-1.9 times higher for catastrophes and 23.8-86.4 times higher for extinction).

Conservatively using the lower multiplier (1.7 times) for the XPT forecast means calculated above and the implied multiplier for the Metaculus forecasts ($27/12 = 2.25$), we can therefore draw a stylized quasi-quantitative conclusion: **all else equal, focusing on engineered pandemics is likely**

⁷¹ “How should we interpret this persistent divide between groups? One possibility is that the experts are biased toward the topics they are professionally invested in and overweight the tail-risks they spend time thinking about — perhaps partly because those most worried about existential risk opt to dedicate their lives to studying it. Another is that the superforecasters are skilled at using historical data for relatively short-run forecasts but might struggle to adapt their methods to longer-run topics with less data — even when they have experts on hand to walk them through the topic. It is also possible that the epistemic strategies that were successful in earlier short-run forecasting tournaments, when the superforecasters attained their status, are not as appropriate at other points in time. For example, base rates may be more useful in periods of relative geopolitical calm and less useful in periods of greater conflict. All of these possibilities may be operating together, to various degrees.” Ezra Karger et al., “Forecasting Existential Risks: Evidence from a Long-Run Forecasting Tournament” (Forecasting Research Institute, 2023), 20.



more than twice as important for preventing catastrophes as focusing on naturally occurring pathogens.⁷² Because of the heavy tails of the distribution of biological events (see [The Threat Is Super-linear](#)), such catastrophes make up the bulk of the risk, so that focusing on engineered pandemics is likely to also be more cost-effective, all-else equal.

Bioweapons, Bioterrorism, and Laboratory Leaks

To summarize the preceding sections:

1. From first principles, biological agents that have features of self-replication that make them particularly worrying threats with chain-reaction dynamics.
2. The horrific history of naturally-arising pandemics illustrates the threat that such agents can pose to humanity.
3. But nature is not the best bioterrorist; an engineered pathogen could exhibit high virulence, transmission rates, and other dangerous pathogen characteristics that would be less likely to emerge naturally.
4. Therefore, we ought to be especially concerned about the threat of engineered pandemics.

This section begins to answer a natural follow-up question: Is the creation and release — accidental or deliberate — of such engineered pathogens (including both the work of bioweaponers and of well-intentioned scientists) likely? The answer to this question has important implications for the prioritization of philanthropic resources. If the creation and release of such pathogens were highly unlikely, then: (1) philanthropists may wish to prioritize naturally-arising pandemics over engineered ones and (2) the total threat of GCBRs may be limited. Unfortunately, as the following paragraphs illustrate, the history of the misuse of biology and base rates of laboratory accidents suggest that the release of engineered pathogens is worryingly likely.

State Biological Weapons and Biodefense Programs

States have long weaponized biological agents, such as the use of smallpox by British colonists in the genocide of American Indian tribes in North America.⁷³ More recently, many states pursued large-scale biological weapons programs in the Twentieth Century:

⁷² The *specific* number here is less relevant for philanthropic prioritization than its approximate magnitude — potentially a more than 100% difference in cost-effectiveness.

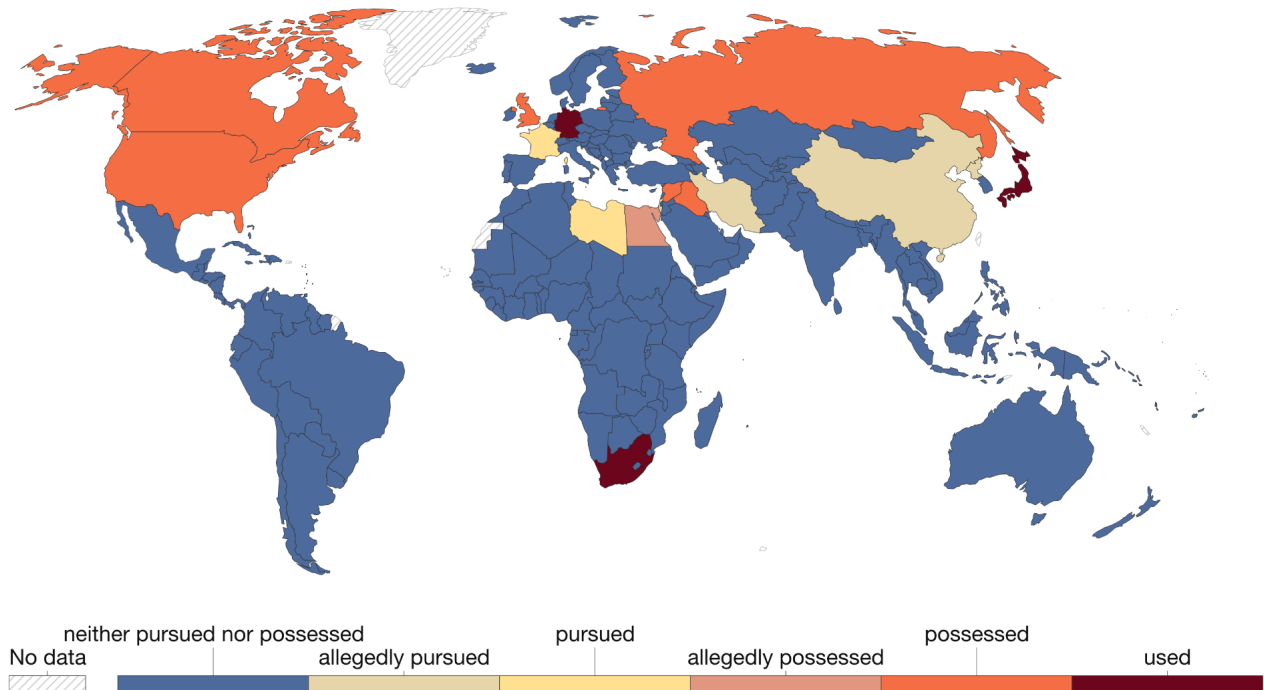
⁷³ Although there is historical debate about specific incidents, one example often cited is the use of smallpox under the command of Lord Amherst, recorded in the journals of one British Captain during the Seven Years War/French and Indian War: “Out of regard for them we gave them two blankets and a handkerchief out of the Smallpox hospital. I hope it will have the required effect.” (Jeremy Hugh Baron, “British Biological Warfare,” *BMJ: British Medical Journal* 327, no. 7409 (August 2, 2003): 261.)



Historical biological weapons activity

Biological weapons are organisms or toxins used to cause death or harm through their poisonous properties. The closest a country came to using biological weapons ever is recorded.

Our World
in Data



Source: OWID based on ACA (2022), NTI (2022), and CNS (2008).

OurWorldInData.org/biological-and-chemical-weapons • CC BY

Note: 'Allegedly' refers to situations where a country was charged by another country of pursuing or possessing biological weapons, but the claims have not been confirmed by the country itself or impartial observers.

Source: Our World in Data, "[Historical Biological Weapons Activity](#)."

The Biological Weapons Convention (BWC) is supposed to protect the world against the threat of bioweapons, but it lacks any formal verification and compliance mechanisms.⁷⁴ Indeed, the USSR was able to sustain a massive covert biological weapons program *after* signing the 1972 Convention, and according to some analyses, "elements of the Soviet offensive BW program continue in Russia."⁷⁵

⁷⁴ For one viewpoint on the failed attempt to negotiate a verification protocol, see Kenneth D. Ward, "The BWC Protocol: Mandate for Failure," *The Nonproliferation Review* 11, no. 2 (June 1, 2004): 183–99, <https://doi.org/10.1080/10736700408436972>.

⁷⁵ "Yet some states which have signed or ratified the BWC have covertly pursued biological weapons programmes. The leading example was the Biopreparat programme of the USSR, which at its height spent billions and employed tens of thousands of people across a network of secret facilities, and conducted after the USSR signed onto the BWC" (Gregory Lewis, "Reducing Global Catastrophic Biological Risks," 80,000 Hours, March 2020, <https://80000hours.org/problem-profiles/preventing-catastrophic-pandemics/full-report/>.) Raymond Zilinskas, "The Soviet Biological Weapons Program and Its Legacy in Today's Russia" (Center for the Study of Weapons of Mass Destruction (National Defense University), 2016),



There are arguments against states' pursuit of highly destructive biological weapons. As far as historians and open-source analysts are aware, "there is no evidence of state-run bioweapons programs directly attempting to develop or deploy bioweapons that would pose an existential risk."⁷⁶ The apparent lack of current state biological weapons programs may reflect the strategic considerations that national militaries face. Omnicidal or near-omnicidal biological agents are simply not effective weapons for achieving specific battlefield aims like claiming territory, defending access to waterways, etc., to the extent that some call them "strategically useless."⁷⁷ Highly destructive bioweapons *in theory* have some value as strategic deterrents, but even if one were able to credibly demonstrate vaccination of domestic populations, pathogens' tendency to mutate would make the development and use of omnicidal weapons too risky a wager for states.⁷⁸

On the other hand, similar arguments could be waged against the pursuit and maintenance of large and highly destructive nuclear arsenals. Superpower leaders since the waning days of the Cold War have insisted publicly that "nuclear war cannot be won and must never be fought," and yet they maintain still-catastrophically-large arsenals.⁷⁹ Here, too, there are so far no credible defenses, and an all-out nuclear war between the great powers could threaten human civilization.⁸⁰ Thus, for example, the scholar of existential risk Toby Ord has argued that "While there is no evidence of deliberate attempts to create a pathogen to threaten the whole of humanity, the logic of deterrence or mutually assured destruction could push superpowers or rogue states in that direction."⁸¹

<https://ndupress.ndu.edu/Media/News/Article/848207/the-soviet-biological-weapons-program-and-its-legacy-in-todays-russia/https%3A%2F%2Fndupress.ndu.edu%2FMedia%2FNews%2FNews-Article-View%2FArticle%2F848207%2Fthe-soviet-biological-weapons-program-and-its-legacy-in-todays-russia%2F>.

⁷⁶ Piers Millett and Andrew Snyder-Beattie, "Existential Risk and Cost-Effective Biosecurity," *Health Security* 15, no. 4 (August 2017): 374, <https://doi.org/10.1089/hs.2017.0028>.

⁷⁷ "[P]andemic-class agents would be strategically useless to nation-states due to their slow spread and indiscriminate lethality, they might be acquired and deliberately released by terrorists." Kevin M. Esvelt, "Delay, Detect, Defend: Preparing for a Future in Which Thousands Can Release New Pandemics," Geneva Papers (Geneva Centre for Security Policy, 2022), 10, https://dam.gcsp.ch/files/doc/gcsp-geneva-paper-29-22?_gl=1*xm44p1*_ga*ODQxMDI0NiY4LjE2ODM2NTg5MTg.*_ga_Z66DSTVXTJ*MTY4Mzc2NTIzNi4yLjAuMTY4Mzc2NTIzNi4wLjAuMA..#page=49&zoom=100.0.0.

⁷⁸ It should be noted, however, that rational strategic considerations do not adequately explain some of the diseases that the Soviet *Biopreparat* program worked on: "the USSR worked to create even-more transmissible smallpox, research ballistic delivery systems, and develop more persistent formulations. Accounts from former USSR bioweapons scientists also allege that the state worked tirelessly to weaponize other transmissible diseases such as tularemia, Ebola, Q fever, Marburg virus, plague, tularemia [sic — this is mentioned twice in the original paper], glanders, and antibiotic resistant microbial diseases with little to no consideration as to whether treatments for these diseases existed" Chris Bakerlee et al., "Common Misconceptions About Biological Weapons" (Council on Strategic Risks, December 2020), 5, https://councilonstrategicrisks.org/wp-content/uploads/2020/12/Common-Misconceptions-About-Biological-Weapons-BRIEFER-12_2020_12_7.pdf.

⁷⁹ Thanks to Matt Lerner for this point.

⁸⁰ For an overview of the considerations surrounding the effects of large-scale nuclear war, see Founders Pledge's recent report on [Global Catastrophic Nuclear Risk](#).

⁸¹ Toby Ord, *The Precipice: Existential Risk and The Future of Humanity* (New York: Hachette Books, 2020), 132.



Moreover, the apparent strategic incentives against the development and use of extreme or even omniscidal bioweapons do not necessarily apply to biological weapons in general, however.⁸² First, many states clearly have in the past believed that biological weapons were worth investing in. Indeed, some historians argue that the renunciation of biological weapons had less to do with their strategic utility than with the difficulty of monopolizing such weapons in the way powerful states can monopolize the fissile material of nuclear weapons.⁸³ Second, biological weapons do have some features that may make them attractive to some states: “they provide a novel means of attack, are challenging to attribute, and may provide a strategic deterrent more accessible than (albeit inferior to) nuclear weapons.”⁸⁴

Third, military bureaucracies are not rational unitary actors. State biodefense laboratories may work with omniscidal pathogens for defensive reasons (e.g. to anticipate possible attacks by terrorists, discussed below), which can create the possibility of accidental release or theft of these pathogens. Moreover, oversight within highly secretive and compartmentalized bureaucracies may be limited, and professional incentives may encourage actors deep within this apparatus to work with risky pathogens. National leaders may not have insight into these programs.⁸⁵

Finally, new advances in synthetic biology, briefly discussed [below](#), may make biological weapons that previously appeared uncontrollable and useless suddenly seem more militarily useful. Weaker

⁸² The best general explanation of these considerations is Chris Bakerlee et al., “Common Misconceptions About Biological Weapons” (Council on Strategic Risks, December 2020), 5, https://councilonstrategicrisks.org/wp-content/uploads/2020/12/Common-Misconceptions-About-Biological-Weapons-BRIEFER-12_2020_12_7.pdf.

⁸³ “[T]he increasing realization that very few nations could acquire nuclear weapons of mass destruction, because of the difficulty of producing the fissile material, whereas many states could develop biological agents of mass destruction. Why not then renounce biological weapons, spread the idea that they were of little use, and even negotiate a total ban on them?” (“Biological Warfare, 1945-1972” in Dando, *Bioterror and Biowarfare*). In the same vein, Jonathan Tucker has written that for the Nixon administration, “[s]ending the message that biological warfare was ineffective would help to discourage hostile nations from acquiring a ‘poor man’s atomic bomb’ that could serve as a military equalizer” (see Jonathan B. Tucker, “A Farewell to Germs: The U.S. Renunciation of Biological and Toxin Warfare, 1969–70,” *International Security* 27, no. 1 (July 2002): 128, <https://doi.org/10.1162/016228802320231244>). For an in-depth discussion of some of these considerations see Founders Pledge’s report on [Global Catastrophic Nuclear Risks: A Guide for Philanthropists](#).

⁸⁴ Gregory Lewis, “Reducing Global Catastrophic Biological Risks,” 80,000 Hours, March 2020, <https://80000hours.org/problem-profiles/preventing-catastrophic-pandemics/full-report/>. Whether it is actually inferior to nuclear weapons may depend on the attacker’s objectives. Those who wish to preserve physical critical infrastructure and industrial centers for invasion may prefer weapons like biological weapons that act only on humans with fewer physical effects.

⁸⁵ “For biological weapons, divergence between senior leadership direction and actual government activities and plans can take many forms. In some manifestations, a nation’s most senior leaders may not even know or fully understand the bioweapons-relevant work occurring in their own countries.” Chris Bakerlee et al., “Common Misconceptions About Biological Weapons” (Council on Strategic Risks, December 2020), 9, https://councilonstrategicrisks.org/wp-content/uploads/2020/12/Common-Misconceptions-About-Biological-Weapons-BRIEFER-12_2020_12_7.pdf.



states may see a new generation of potential biological weapons as a possible “offset” against the advanced conventional and nuclear capabilities of more powerful nations.⁸⁶

In short, the BWC is too weak to deter determined states, biological weapons may have attractive attributes for some states, the complex incentives of sprawling defense bureaucracies may encourage dangerous research, and new advances in synthetic biology may make state biodefense programs a cause for concern:

Reasons <i>for</i> concern about state bioweapons/biodefense programs	Reasons <i>against</i> concern
<ul style="list-style-type: none">• Resources of states• History of state BW programs• Toothlessness of BWC• Sprawling secret compartmentalized bureaucracies• Possible low cost compared to other WMD• Possible “offset” against superior conventional and nuclear capabilities• Low attributability• Accidents from defensive research• Emerging biotechnologies may increase perceived utility of BW	<ul style="list-style-type: none">• Existence of BWC and possible norm against BW development/use• Apparent/perceived low battlefield utility of BW• Most BW efforts are likely non-contagious⁸⁷• BW projects may struggle to attract talented scientists

Bioterrorists and Rogue Actors

Non-state terror groups do not have the large resources of powerful states, but unlike states, some such groups may have the incentive to attempt to develop and possibly use truly omnicidal weapons. Many non-state groups that use terrorist tactics have highly limited aims — capture a specific piece of territory for a national or ethnic group, gain certain political concessions from states, or even serve as pseudo-states in territories with low state capacity. Such groups may not optimize for

⁸⁶ “A more attainable weapon of mass destruction may be increasingly enticing as the broader military arsenals of the United States and other superpowers become increasingly harder to compete with. As the United States and its near-peer adversaries continue to compete and develop capabilities designed to undermine one other, many other nations increasingly seek asymmetric capabilities to advance their political interests.” Chris Bakerlee et al., “Common Misconceptions About Biological Weapons” (Council on Strategic Risks, December 2020), 3, https://councilonstrategicrisks.org/wp-content/uploads/2020/12/Common-Misconceptions-About-Biological-Weapons_BRIEFER-12_2020_12_7.pdf.

⁸⁷ Thanks to Andrew Snyder-Beattie for this and the following point in a round of external reviews.



releasing especially dangerous pathogens, but may instead prioritize “dreaded diseases” that inspire fear or terror.⁸⁸

A few groups, however, have explicitly expressed omniscidal or near-omniscidal aims and pointed to the use of bioweapons for the achievement of these aims.⁸⁹ Fortunately, most such groups have historically not had the resources to cause real damage. One near-omniscidal group, *Aum Shinrikyo*, was well-resourced and explicitly pursued bioweapons and had a PhD-level virologist with access to Anthrax strains lead this program before eventually landing on chemical weapons (sarin gas) for its infamous Tokyo subway attacks.⁹⁰ Similarly, Gaia Liberation Front explicitly proclaimed in 1995:

“[W]e can ensure Gaia’s survival only through the extinction of the Humans as a species ... we now have the specific technology for doing the job ... several different [genetically engineered] viruses could be released.”⁹¹

⁸⁸ Long and Marzi point to “the ability to generate fear and/or terror” as one of the characteristics of pathogens that bioterrorists may seek in Carrie M. Long and Andrea Marzi, “Biodefense Research Two Decades Later: Worth the Investment?,” *The Lancet. Infectious Diseases* 21, no. 8 (August 2021): e222–33, [https://doi.org/10.1016/S1473-3099\(21\)00382-0](https://doi.org/10.1016/S1473-3099(21)00382-0). For a discussion of “dreaded diseases,” see Christian Enemark, *Biosecurity Dilemmas: Dreaded Diseases, Ethical Responses, and the Health of Nations* (Georgetown University Press, 2017), <https://www.jstor.org/stable/j.ctt1kk672v>.

⁸⁹ “Non-state actors may also pose a risk, especially those with explicitly omniscidal aims. While rare, there are examples. The Aum Shinrikyo cult in Japan sought biological weapons for the express purpose of causing extinction. Environmental groups, such as the Gaia Liberation Front, have argued that “we can ensure Gaia’s survival only through the extinction of the Humans as a species . we now have the specific technology for doing the job . several different [genetically engineered] viruses could be released”(quoted in ref. 29). Groups such as R.I.S.E. also sought to protect nature by destroying most of humanity with bioweapons. Fortunately, to date, non-state actors have lacked the capabilities needed to pose a catastrophic bioweapons threat, but this could change in future decades as biotechnology becomes more accessible and the pool of experienced users grows.” Piers Millett and Andrew Snyder-Beattie, “Existential Risk and Cost-Effective Biosecurity,” *Health Security* 15, no. 4 (August 2017): 374, <https://doi.org/10.1089/hs.2017.0028>.

⁹⁰ Notably, Aum Shinrikyo is sometimes described as “omniscidal,” but may be more accurately classified as near-omniscidal, given that its members apparently believed that they would survive the apocalyptic events they sought to catalyze. Thanks to Andrew Snyder-Beattie for this point, which is also discussed in David Thorstad’s blog post “Exaggerating the Risks (Part 10: Biorisk: More Grounds for Doubt),” *Reflective altruism*, August 12, 2023, <https://ineffectivealtruismblog.com/2023/08/12/exaggerating-the-risks-part-10-biorisk-more-grounds-for-doubt/>. The kinds of strategies pursued by both near-omniscidal and omniscidal groups may in practice converge, however. On the training of the member in question: “The cult member in charge of biological weapons efforts had completed undergraduate studies in agricultural and veterinary medicine at a university with a substantial anthrax collection, including both Sterne and Pasteur vaccine strains, before going on to PhD studies in virology. Danzig et al. speculate that a contact at his old university provided him with B. anthracis, a supposition supported by new analysis of the samples gathered by the Japanese police, which suggest that the cult’s material closely matched the particular strain of Sterne held by the university.” Philipp C. Bleek, “Revisiting Aum Shinrikyo: New Insights into the Most Extensive Non-State Biological Weapons Program to Date,” *The Nuclear Threat Initiative*, December 10, 2011, <https://www.nti.org/analysis/articles/revisiting-aum-shinrikyo-new-insights-most-extensive-non-state-biological-weapons-program-date-1/>.

⁹¹ We have chosen not to link to the 1995 document that is the source of this quote, which is taken from Jason G. Matheny, “Reducing the Risk of Human Extinction,” *Risk Analysis* 27, no. 5 (2007): 1337, footnote 5, <https://doi.org/10.1111/j.1539-6924.2007.00960.x>.



In other words, there have been omniscidal and near-omniscidal actors that may want to use engineered pandemics to achieve their horrific aims, and it is difficult to rule out that such groups exist now or could exist again in the near future. It is important to emphasize this:

- There are well-financed terrorist groups that wish to do great harm to humanity. Some of these groups openly desire human extinction.
- Some of these groups have openly advocated the use of biological weapons. Others have covertly sought to acquire or develop such weapons.
- Some of these groups have had (and others may continue to have) immense resources — millions of dollars, PhD-level scientists, and access to university laboratories.
- Fortunately, developing and using biological weapons has been technically challenging.
- If we are not careful, **techno-scientific progress could break this barrier standing between humanity and catastrophe.**

Kevin Esvelt has recently put the number of “extremists and zealots” in the thousands, writing “Numerically, fewer than a dozen powerful nations have funded research seeking to identify pandemic-class agents. There are only a handful of rogue states, but many thousands of extremists and zealots.”⁹² Zealots and extremists need not even use engineered pathogens, as Esvelt has argued: “With some natural viruses exhibiting the transmissibility of early variants of SARS-CoV-2 and a lethality rate exceeding 30 per cent, such an event [release at multiple places] could precipitate the greatest catastrophe in the history of humanity.”⁹³ Again, historically, **technical barriers were the only thing stopping well-resourced terrorist groups from carrying out horrific bioweapons attacks.** These barriers remain, but unregulated techno-scientific progress could facilitate horrific misuse of synthetic biology in the near future. **We must therefore ensure that progress in the life sciences does not enable malevolent actors.**

⁹² Kevin M Esvelt, “Delay, Detect, Defend: Preparing for a Future in Which Thousands Can Release New Pandemics,” Geneva Papers (Geneva Centre for Security Policy, 2022), 16, https://dam.gcsp.ch/files/doc/gcsp-geneva-paper-29-22?_gl=1*xm44p1*_ga*ODOxMDI0NiY4LjE2ODM2NTg5MTg.*_ga_Z66DSTVXTJ*MTY4MzczNTIzNi4yLjAuMTY4MzczNTIzNi4wLjAuMA..#page=49&zoom=100.0.0.

⁹³ Kevin M Esvelt, “Delay, Detect, Defend: Preparing for a Future in Which Thousands Can Release New Pandemics,” Geneva Papers (Geneva Centre for Security Policy, 2022), 13, https://dam.gcsp.ch/files/doc/gcsp-geneva-paper-29-22?_gl=1*xm44p1*_ga*ODOxMDI0NiY4LjE2ODM2NTg5MTg.*_ga_Z66DSTVXTJ*MTY4MzczNTIzNi4yLjAuMTY4MzczNTIzNi4wLjAuMA..#page=49&zoom=100.0.0.



Accidents and Laboratory Leaks

The final reason for concern is the frequency of accidents and laboratory leaks. Several historical leaks help to illustrate that this problem is very real and that laboratory leaks have been suspected from various kinds of regimes:⁹⁴

- **1971 Aralsk Smallpox Outbreak** — Field trials of Soviet bioweapons on Vozrozhdeniya Island led to an outbreak of smallpox in Aralsk, Kazakhstan, killing at least three people.⁹⁵
- **1977 H1N1 influenza Outbreak** — An H1N1 influenza virus strain that appeared identical to a strain from the 1950s emerged in 1977. Though uncertain, it is widely believed that the 1977 outbreak was due to a laboratory release, perhaps in the Soviet Union or Communist China.⁹⁶
- **1978 Smallpox in Britain** — Only a year after the eradication of smallpox, the disease is believed to have escaped from a laboratory in Birmingham, infecting a photographer, who subsequently died, though there remain open questions about this incident.⁹⁷
- **1979 Sverdlovsk Incident** — A Soviet biological weapons facility in Sverdlovsk (now Yekaterinburg) leaked released airborne anthrax spores, killing at least 66 people. The Soviet Union blamed contaminated meat and attempted to cover up the incident.⁹⁸

⁹⁴ For a far more detailed discussion and list of human-caused biological incidents, see David Manheim and Gregory Lewis, “High-Risk Human-Caused Pathogen Exposure Events from 1975-2016,” *F1000Research* 10 (2021): 752, <https://doi.org/10.12688/f1000research.55114.2>.

⁹⁵ “A Soviet bioweapons lab tested a weaponized strain of smallpox on an island in the Aral Sea. During a field test, they accidentally infected people on a nearby ship who spread it ashore. The resulting outbreak infected ten people, killing three, before being contained by a mass quarantine and vaccination program.” Toby Ord, *The Precipice: Existential Risk and The Future of Humanity* (New York: Hachette Books, 2020), 131.

⁹⁶ Richard Danzig explains the evidence for and against this view on page 26, footnote 47 of *Technology Roulette*, writing, “The DNA of the 1977 virus was nearly identical to that in a 1950 outbreak – an event extremely unlikely to occur in nature where viruses mutate frequently in response to the pressure of natural selection. Debate over the cause of this unlikely situation could be assembled as a case study in the difficulties of biological attribution. The predominant view has been that ‘the virus was frozen in a laboratory freezer since 1950, and was released, either by intent or accident, in 1977. This possibility has been denied by Chinese and Russian scientists, but remains to this day the only scientifically plausible explanation.’” Richard Danzig, “Technology Roulette: Managing Loss of Control as Many Militaries Pursue Technological Superiority” (Center for a New American Security, 2018), 26, <https://www.cnas.org/publications/reports/technology-roulette>.

⁹⁷ “In 1967 smallpox was killing more than a million people a year, but a heroic global effort drove that to zero in 1977, freeing humanity from this ancient scourge. And yet a year later, it returned from the grave: escaping from a British lab, killing one person and infecting another before authorities contained the outbreak.” Toby Ord, *The Precipice: Existential Risk and The Future of Humanity* (New York: Hachette Books, 2020), 131.

⁹⁸ “A bioweapons lab in one of the Soviet Union’s biggest cities, Sverdlovsk, accidentally released a large quantity of weaponized anthrax, when they took an air filter off for cleaning. There were 66 confirmed fatalities.” (Toby Ord, *The Precipice: Existential Risk and The Future of Humanity* (New York: Hachette Books, 2020), 131.) For an analysis of the outbreak, see M. Meselson et al., “The Sverdlovsk Anthrax Outbreak of 1979,” *Science* (New York, N.Y.) 266, no. 5188 (November 18, 1994): 1202–8, <https://doi.org/10.1126/science.7973702>. On page 1206, Meselson et al. summarize their conclusion: “We conclude that the outbreak resulted from the windborne spread of an aerosol of anthrax pathogen, that the source was at the military microbiology facility,



- **1995 Rabbit Calicivirus** — A field-trial experiment with the virus on Wardang Island accidentally spread to mainland Australia, killing millions of rabbits.⁹⁹
- **2004 SARS Beijing** — The WHO acknowledged that an outbreak in China resulted from a biocontainment failure at a laboratory of the Beijing Institute of Virology.¹⁰⁰
- **2005 Bubonic Plague Mouse Escape** — Three mice infected with Bubonic plague escaped from a New Jersey lab and were “never found.”¹⁰¹
- **2007 Foot-and-Mouth Disease Outbreak:** An outbreak of foot-and-mouth disease among cattle in the UK was traced back to a leak from a site shared by the Institute for Animal Health and a private pharmaceutical company. Shortly after the lab’s license was renewed, there was a second outbreak.¹⁰²
- **2015 Anthrax Distributions** — A U.S. military biodefense facility accidentally sent *live* anthrax to 192 laboratories who thought they were receiving *inactivated* samples.¹⁰³

and that the escape of pathogen occurred during the day on Monday, 2 April. The epidemic is the largest documented outbreak of human inhalation anthrax.” The outbreak is also discussed at length in more accessible terms in various parts of David E. Hoffman’s *The Dead Hand: The Untold Story of the Cold War Arms Race and its Dangerous Legacy*.

⁹⁹ “Australian scientists conducted a field trial with a new virus for use in controlling their wild rabbit population. They released it on a small island, but the virus escaped quarantine, reaching the mainland and accidentally killing 30 million rabbits within just a few weeks.” Toby Ord, *The Precipice: Existential Risk and The Future of Humanity* (New York: Hachette Books, 2020), 131.

¹⁰⁰ “In 2004, SARS escaped from the National Institute of Virology in Beijing. They didn’t realize some of the workers had been infected until a worker’s mother came down with it too.” (Toby Ord, *The Precipice: Existential Risk and The Future of Humanity* (New York: Hachette Books, 2020), 354, footnote 29.) For a contemporary account, see “WHO: Beijing Lab to Blame for SARS Outbreak - 2004-04-26,” VOA, October 28, 2009, <https://www.voanews.com/a/a-13-a-2004-04-26-18-1-66883432/261221.html>.

¹⁰¹ “[I]n 2005 at the University of Medicine and Dentistry in New Jersey, three mice infected with bubonic plague went missing from the lab and were never found.” (Toby Ord, *The Precipice: Existential Risk and The Future of Humanity* (New York: Hachette Books, 2020), 354, footnote 29.)

¹⁰² “In 2001, Britain was struck by a devastating outbreak of foot-and-mouth disease in livestock. [...] Then in 2007 there was another outbreak, which was traced to a lab working on the disease. Foot-and-mouth was considered a highest category pathogen and required the highest level of biosecurity. Yet the virus escaped from a badly maintained pipe, leaking into the groundwater at the facility. After an investigation, the lab’s license was renewed only for another leak to occur two weeks later.” (Toby Ord, *The Precipice: Existential Risk and The Future of Humanity* (New York: Hachette Books, 2020), 130.)

¹⁰³ “The Dugway Proving Grounds was established by the US military in 1942 to work on chemical and biological weapons. In 2015, it accidentally distributed samples containing live anthrax spores to 192 labs across eight countries, which thought they were receiving inactivated anthrax.” (Toby Ord, *The Precipice: Existential Risk and The Future of Humanity* (New York: Hachette Books, 2020), 130.)



Some analysts have attempted to quantify base rates of laboratory leaks, which are alarmingly high:

104

Laboratory Accident Rates

Variable	Accident Rate	Accident Rate Over 100 Years	Source	Notes
50-year escape risk for U.S. BSL-4 labs working with agricultural pathogens	70.00%	91%	Department of Homeland Security Assessment. See Technology Roulette .	
50 year escape risk for U.S. BSL-4 labs working with agricultural pathogens (revised, but controversial)	0.11%	0.22%	Ibid.	Controversial. Notes from NRC: “The committee finds that the extremely low probabilities of release are based on overly optimistic and unsupported estimates of human error rates, underestimates of infectious material available for release, and inappropriate treatment of dependencies, uncertainties, and sensitivities in calculating release probabilities.”

¹⁰⁴ For an attempt to quantify the existential risk posed by these and other biosecurity concerns, see Piers Millett and Andrew Snyder-Beattie, “Existential Risk and Cost-Effective Biosecurity,” *Health Security* 15, no. 4 (August 2017): 373–83, <https://doi.org/10.1089/hs.2017.0028>.



Laboratory Accident Rates (*continued*)

Variable	Accident Rate	Accident Rate Over 100 Years	Source	Notes
Probability of at least one escape “over a ten-year period [in] ten carefully operated labs”	18%	86%	The consequences of a lab escape of a potential pandemic pathogen.	
“annual probability of a global pandemic resulting from an accident with this type of research [with pathogens of pandemic potential, i.e. “gain of function”] in the United States”	0.002% to 0.1%	0.2% to 9.5%	Gryphon Scientific. See Existential Risk and Cost-Effective Biosecurity	Note that there have been criticisms of the approach in “Existential Risk and Cost-Effective Biosecurity”.
annual p(global pandemic) from PPP research globally	0.008% to 0.4%	0.8% to 33%	See Existential Risk and Cost-Effective Biosecurity	Simplistic estimate; Millett and Snyder-Beattie simply assume the U.S. constitutes 25% of the risk. This is not including the risk of deliberate misuse. Note that the 100-year accident rate extrapolates from the given annual rate rather than multiplying the previous 100-year accident rate by 4.



Fully understanding incident rates is challenging because many laboratories don't report their incidents.¹⁰⁵ As Ord put it, this track record is extremely concerning: "With current BSL-4 labs, an escape of a pandemic pathogen is a matter of time."¹⁰⁶

Many analysts argue that deliberate release is a bigger threat than accidental release.¹⁰⁷ Deliberate release could optimize for maximum damage (e.g. releasing pathogens in airports), whereas accidental release may simply fizzle out (e.g. if an infected laboratory worker simply goes home because they feel unwell, and remains isolated). Moreover, actors who go through the difficult and costly process of developing biological weapons may be more likely to have an intention to release such weapons.¹⁰⁸

The possibility that the COVID-19 pandemic had a laboratory origin, however, shows that accidental release, too, can start a global pandemic. Moreover, the rapid growth in the number of high-containment laboratories, with little to no international regulations on safety measures, means that the surface area of accidents is also increasing.

Finally, accidental releases are an important threat vector by which purely defensive research can spawn global threats. Work with enhanced potential pandemic pathogens (ePPPs), sometimes called "gain of function" research, for example, is often intended to enhance pandemic preparedness, but may in fact lead to major risks of new pandemics.¹⁰⁹ Even if we believe that states have no interest in deploying catastrophic biological weapons (e.g., because such weapons are inferior to other weapons), therefore, it becomes possible for biodefense programs to cause catastrophic pandemics. Consider the following chain of events, in which no actor is actually interested in using biological weapons:¹¹⁰

¹⁰⁵ Ord (2020) argues that this may stem in part from embarrassment: "This lack of transparency seems to be driven by fear of embarrassment, which would clearly not be a sufficient reason for stifling this critical information. Stakeholders require these rates to assess whether labs are living up to their claimed standards and how much risk they are posing to the public." Toby Ord, *The Precipice: Existential Risk and The Future of Humanity* (New York: Hachette Books, 2020), 353, footnote 27.

¹⁰⁶ Toby Ord, *The Precipice: Existential Risk and The Future of Humanity* (New York: Hachette Books, 2020), 130.

¹⁰⁷ Claire Zabel, "Biosecurity as an EA Cause Area,"

<https://www.effectivealtruism.org/articles/biosecurity-as-an-ea-cause-area-claire-zabel>.

¹⁰⁸ Thanks to Joshua Monrad for this point.

¹⁰⁹ Gryphon Scientific and Rocco Casagrande, "Risk and Benefit Analysis of Gain of Function Research Final Report" (Gryphon Scientific, 2016),

<http://gryphonsci.wpengine.com/wp-content/uploads/2018/12/Risk-and-Benefit-Analysis-of-Gain-of-Function-Research-Final-Report-1.pdf>.

¹¹⁰ This is essentially the "Protect or Proliferate" dilemma discussed by Enemark: "If fear of being attacked caused one state to respond by engaging in activities that were in turn perceived by another state as offensive, the result could be a proliferation of fearsome 'biodefense' activities driven by mutual suspicion. As each state sought to increase its knowledge of what it perceived to be the other's offensive capabilities, each could render itself more capable of using microorganisms for hostile purposes if it chose to do so. In this way biodefense would effectively generate an increase in the likelihood of biological weapons use; a defensive effort would worsen the threat to be defended against." Christian Enemark, *Biosecurity Dilemmas: Dreaded*



1. **Growing Concern about Bioterrorism** — The U.S. defense establishment grows increasingly concerned about advances in biotechnology that may enable rogue actors to more easily create especially harmful pathogens.
 - a. Whether or not such rogue actors actually exist is less relevant than the *perception* that they could exist, especially in an environment fraught with paranoia and competition.¹¹¹
2. **Investments in Research on extreme PPPs** — To prepare for these threats, the government funds more research on potential pandemic pathogens, with the goal of developing vaccines against the most likely culprits.
 - a. Depending on the specific priorities of the government, this may include extinction-level pathogens.
3. **Growth in Accident Surface Area** — This increased activity, including a growth in the number of high-containment laboratories, increases the surface area for accidents.
 - a. This grows even further if other states view U.S. activity as potentially offensive, and grow their own biodefense activity.
4. **Accident → Catastrophic Pandemic** — Ultimately, a laboratory accident leads to the release of an extinction-level pathogen.

In this scenario, no actor had any interest in actually using such an agent, but the agents were developed nonetheless for defensive purposes. Indeed, scientists do experiment on extremely dangerous pathogens. For example, in 2012, the virologist Ron Fouchier performed gain-of-function experiments on the highly lethal H5N1 influenza strain, creating a version of the virus that could be transmitted mammal-to-mammal.¹¹² Although such research is far from actually weaponizing H5N1 to be human-to-human transmissible, it exemplifies the dangers of pursuing risky research with enhanced potential pandemic pathogens and then publishing about details of that research.

Diseases, Ethical Responses, and the Health of Nations (Georgetown University Press, 2017), <https://www.jstor.org/stable/j.ctt1kk672v>.

¹¹¹ As Danzig describes the character of military competition: “Richly endowed nation-state rivals operate with great paranoia and little inhibition. In military competitions security agencies presume sabotage, indeed proudly practice it. Even botched attempts at sabotage increase the risks of accidents and unintended effects. Moreover, fear of military opponents intensifies willingness to take risks: If they might be doing X, we must do X to keep them from getting there first, or at least so that we understand and can defend against what they might do.” Richard Danzig, “Technology Roulette: Managing Loss of Control as Many Militaries Pursue Technological Superiority” (Center for a New American Security, 2018), 8, <https://www.cnas.org/publications/reports/technology-roulette>.

¹¹² See Toby Ord, *The Precipice: Existential Risk and The Future of Humanity* (New York: Hachette Books, 2020), 129, for a discussion of Fouchier’s experiments.



Existential Threats

To summarize the previous sections, the high importance of global catastrophic biological risks arises from the following considerations:

1. From first principles, biological agents have certain characteristics that make their threat especially insidious.
2. Historically, pandemics have been some of the biggest threats to human civilization.
3. Evolutionary dynamics don't optimize for pathogens that cause maximum damage, but humans can engineer pathogens with more dangerous characteristics than nature ever would.
4. Various actors — states, rogue groups, individual scientists — have an interest in developing such agents, for both offensive and defensive purposes.
5. This creates the potential for the intentional or accidental release of pandemic-level pathogens that cause outbreaks **far worse than anything humanity has ever seen**.


This section adds a further point: engineered pandemics are among the few known threats that could plausibly cause total civilizational collapse and potentially the extinction of humanity. In an interview for this project, Jason Matheny¹¹³ put this point succinctly, saying, “the value of the future seems vast. It looks like there's only a handful of things that we can think of right now that could actually cut short terminally the value of the future. Bio risks are one of those things.”¹¹⁴

Crowdsourced prediction aggregation platforms agree with Matheny's assessment, placing an 8% chance on the possibility that a global biological catastrophe would cause the near extinction (95% reduction of population) of humanity:¹¹⁵

¹¹³ Current [president and CEO of the RAND Corporation](#), with extensive expertise and experience working on emerging catastrophic risks.

¹¹⁴ Interview with Dr. Jason Matheny, 30 May 2023.

¹¹⁵ “Ragnarök Question Series: If a Global Biological Catastrophe Occurs, Will It Reduce the Human Population by 95% or More?,” January 13, 2019, <https://www.metaculus.com/questions/2514/ragnar%25C3%25B6k-question-series-if-a-global-biological-catastrophe-occurs-will-it-reduce-the-human-population-by-95-or-more/>.



Source: Metaculus, “[If a global biological catastrophe occurs, will it reduce the human population by 95% or more?](#)” Aggregate forecast as of 19 July 2023.

The risk of extinction takes different moral weights depending on one’s ethical framework. From the perspective of ethical views that prioritize current and near-future generations, extinction matters simply because of the great degree of suffering it would create. From the perspective of ethical views that do not discriminate between current humans and future humans, however, the importance of civilizational collapse and extinction is astronomically larger. This longtermist view and the arguments for and against it are discussed in greater detail elsewhere.¹¹⁶ **Even if readers place no or very little credence in such “longtermist” views, however, preventing a biological catastrophe could be one of the most important things humanity works on**, as discussed throughout this report. This is because the heavy-tailed distribution of the relevant events — disease outbreaks and terrorist attacks among them — combined with the problem of cascading failures and systems collapses lead to a superlinear cost curve that place much value in the most catastrophic events.

Total *extinction* — as opposed to the more broadly-defined “existential risk” that could include events like global civilizational collapse — may not be likely. Several considerations weigh against its likelihood:

- **Isolation equilibria** — The more people are infected and gravely ill or dead (and therefore come into less contact with others) the less likely a pathogen is to spread.

¹¹⁶ For a full list of philosophical papers, popular books, and more, see longtermism.com, “Resources,” longtermism.com, 2021, <https://longtermism.com/resources>.



- **Rarity of omniscidal actors** — As discussed above, states have few incentives to create pandemic-level pathogens, and most terror groups have narrow political aims. Some groups routinely described as “omniscidal” may be more accurately called near-omniscidal.¹¹⁷
- **Technical barriers** — There may be technical barriers to artificially combining the traits required in potential extinction-level pathogens.
- **Natural immunity** — The human population is diverse, and some groups may have natural immunity to some pathogens for various reasons.
- **Isolated communities** — There are some communities that are highly isolated from contact with the rest of the world, and would therefore be unlikely to come into contact with pandemic pathogens:¹¹⁸
 - Uncontacted tribes and other isolated traditional communities
 - Nuclear submarine crews¹¹⁹
 - Astronauts
 - Inhabitants of experimental sealed spaces (e.g. Earth-based training environments for space settlement).¹²⁰

An extinction-level pathogen would therefore need to have very specific attributes, and be released in very particular ways to truly cause human extinction.¹²¹ Nonetheless, such a pandemic could still

¹¹⁷ Thanks to Andrew Snyder-Beattie for this point.

¹¹⁸ For a discussion of this and how it affects the cost-effectiveness of civilizational refuges, see Nick Beckstead, “How Much Could Refuges Help Us Recover from a Global Catastrophe?,” *Futures* 72 (September 2015): 36–44, <https://doi.org/10.1016/j.futures.2014.11.003>. See also Nick Beckstead, “Improving Disaster Shelters to Increase the Chances of Recovery from a Global Catastrophe,” *Effective Altruism Forum*, 2014, <https://forum.effectivealtruism.org/posts/fTDhRL3pLY4PNee67/improving-disaster-shelters-to-increase-the-chances-of>.

¹¹⁹ Alexey Turchin and Brian Patrick Green, “Aquatic Refuges for Surviving a Global Catastrophe,” *Futures* 89 (May 1, 2017): 26–37, <https://doi.org/10.1016/j.futures.2017.03.010>.

¹²⁰ Biosphere 2 was a prominent example of this. University of Arizona, “About Biosphere 2 (History),” [biosphere2.org](https://biosphere2.org/about/about-biosphere-2), accessed July 19, 2023, <https://biosphere2.org/about/about-biosphere-2>.

¹²¹ “Any apocalyptic pathogen would need to possess a very special combination of two attributes. First, it would have to be so unfamiliar that no existing therapy or vaccine could be applied to it. Second, it would need to have a high and surreptitious transmissibility before symptoms occur. The first is essential because any microbe from a known class of pathogens would, by definition, have family members that could serve as models for containment and countermeasures. The second would allow the hypothetical disease to spread without being detected by even the most astute clinicians” Amesh Adalja, “Why Hasn’t Disease Wiped out the Human Race?,” *The Atlantic*, June 17, 2016, <https://www.theatlantic.com/health/archive/2016/06/infectious-diseases-extinction/487514/>.



cause global civilizational collapse, potentially setting humanity back permanently, or creating the possibility of a recovery with bad values.¹²²

Even if total extinction is unlikely, near-extinction could be as bad or worse, and the consequences of a global civilizational collapse would be devastating.¹²³ Consider, for example, the hypothetical scenario where North Korea manages to fully isolate itself from the rest of the world while a catastrophic pandemic sweeps the globe. In this case, a cruel totalitarian dictatorship that tortures and starves its own citizens would be steering the recovery of human civilization, potentially locking in the rest of our civilization's future into a state of constant suffering.

Two catastrophic threat scenarios that have been described publicly are the “wildfire” scenario and the “stealth pathogen” scenarios.¹²⁴

- **“Wildfire”** — In the wildfire scenario, a highly virulent pathogen rapidly spreads across the world, killing essential workers (or appearing to be such a threat that essential workers fear going to work), and leading to the potential collapse of the modern economy — the electrical grid, food production and distribution, sanitation, and water supply all rely on essential workers.
- **“Stealth”** — In the stealth scenario, an agent that is transmissible without showing symptoms (or showing only mild symptoms) spreads rapidly across the world (e.g. a “respiratory HIV”), infecting much of humanity with a disease that will be deadly later. These scenarios and others help to illustrate how a pandemic agent could threaten the collapse of civilization.

In conclusion, GCBRs are a real threat that we know can kill hundreds of millions of people, and that could plausibly kill billions. Mitigating this risk could therefore be one of the most important issues of our time. Yet society remains deeply unprepared for the next pandemic. Consider the response to the COVID-19 pandemic, a relatively low-virulence outbreak in a single city — no nation was able to fully contain the spread, and millions have died. Even heroic attempts to improve preparedness appear insufficient. As Gopal *et al.* have pointed out, the U.S. National Biodefense Strategy calls for a “moonshot” effort to make vaccines widely available within 130 days of the sequencing of a pathogen, but this is still too slow; after all, “the omicron variant infected a quarter of North Americans and half of Europeans within 100 days of detection.”¹²⁵ As the next section describes, advances in the life sciences only compound our society's lack of preparedness.

¹²² For an in-depth discussion of these considerations, see Haydn Belfield's chapter “Collapse, Recovery, and Existential Risk” in Miguel A. Centeno, Peter W. Callahan, Paul A. Larceny, and Thayer S. Patterson (eds), *How Worlds Collapse*.

¹²³ Again, see Haydn Belfield's chapter “Collapse, Recovery, and Existential Risk” in Miguel A. Centeno, Peter W. Callahan, Paul A. Larceny, and Thayer S. Patterson (eds), *How Worlds Collapse*.

¹²⁴ These two scenarios are discussed in Anjali Gopal, William Bradshaw, Vaishnav Sunil, Kevin M. Esvelt, “Securing Civilization against Catastrophic Pandemics” (Geneva Centre for Security Policy, 2023), *in press*.

¹²⁵ Anjali Gopal, William Bradshaw, Vaishnav Sunil, Kevin M. Esvelt, “Securing Civilization against Catastrophic Pandemics” (Geneva Centre for Security Policy, 2023), *in press*.



Open Questions for Further Research

- How frequent are the omnicidal, near-omnicidal or apocalyptic sentiments among terrorist groups that would lead such groups to develop catastrophic pandemic-potential pathogens?
- What technological and scientific breakthroughs could change states' strategic calculus about biological weapons? (N.B.: this topic may require extra sensitivities around sharing research findings)
- What are crucial points of failure in health systems and global critical infrastructure? What kinds of pandemics could lead to each failure point?
- How can bottlenecks and barriers to the weaponization of pathogens be preserved?



The Growing Risk of Biological Catastrophe

Key Points

- Rapid advances in the life sciences appear likely to make synthetic biology cheaper and more accessible.
- Within these broader advances, there are specific dual-use technologies — widely accessible synthetic DNA, including from benchtop devices, applications of advanced AI, cloud labs, and more — that could be a boon to humanity, but could also enable malicious actors to cause harm on an unprecedented scale.
- These trends increase the risk of biological catastrophe and simultaneously complicate the threat landscape.

Unfortunately, the risk of a biological catastrophe is not static. This section briefly outlines key emerging and future techno-scientific advances that appear likely to affect the threat landscape. The section begins by discussing **general advances in the life sciences** and the effects of cost declines and the proliferation of existing knowledge and technologies. For instance, declining costs of DNA synthesis (discussed below) may democratize access to the “building blocks” of creating and modifying dangerous pathogens. Many modern attempts to weaponize pathogens would likely involve synthetic DNA, such that **advances in DNA synthesis are an especially crucial consideration**. Next, the analysis turns to **emerging biotechnologies of concern**, including:

- **Benchtop DNA synthesis devices** that may complicate the screening of DNA synthesis.
- **Advances in AI applications** in the life sciences that could (1) facilitate the misuse of synthetic biology and (2) accelerate scientific progress in a way that increases the surface area of various risks.
- **The growing use of “cloud labs”** that allow cheap wet-lab experiments by remote control.
- **Other technologies**, such as potential future advances in atomically-precise manufacturing that could have far-reaching effects in biology.

This section also discusses two **key implications** of this changing threat landscape:

1. **Increasing risk of biological catastrophe**, especially if many of these technologies are offense-dominant.



2. **Increasing uncertainty around the nature of the threat.** This multifaceted changing landscape is degrading our ability to predict and therefore defend against specific threat vectors. As a result, defense must become more general and threat agnostic.

Importantly, biology is a complex science, and there remain many technical barriers to the misuse of biology, as discussed above. None of these changes will radically alter the threat on its own, but in combination, they paint a picture of future biological risks that are simultaneously growing and becoming more difficult to understand and defend against. This then has important implications for how philanthropists must approach this problem, as outlined later in this report (see [Analyzing the Structure of the Problem](#)).

A Note on the Benefits of Techno-Scientific Progress

The following discussion focuses heavily on the *risks* of new scientific breakthroughs and technological advances, at the expense of discussing the many benefits of such advances. Cheaper and more widely-accessible access to ever more-powerful tools of synthetic biology could have vast benefits for humanity, helping scientists find new cures for diseases, develop more efficient and less cruel food production methods, and create vast economic benefits. The democratization of the life sciences also has benefits in terms of equity, by allowing individuals, laboratories, and countries with fewer resources to access the potential benefits of these technologies. The following sections are intended not to de-emphasize the benefits, but to emphasize that in order to fully enjoy these benefits, humanity will also need to mitigate the risks.

General Advances in the Life Sciences

The last century has seen advances in medicine, agriculture, and global well-being driven partly by scientific progress. Rapid progress in the life sciences, especially synthetic biology, appears to hold out the promise of further large benefits for humanity.¹²⁶ A full overview of specific scientific breakthroughs, or of specific technologies like CRISPR gene editing, is beyond the scope of this report.¹²⁷ Broadly, however, humanity increasingly understands the building blocks of life and how to manipulate them. This has allowed for beneficial biomedical breakthroughs, but also raised the possibility of misuse and created new problems like the potential for “genetic warfare” — the targeting of specific people or groups based on their genetic profiles.¹²⁸

¹²⁶ Government Accountability Office, “Science & Tech Spotlight: Synthetic Biology | U.S. GAO,” gao.gov, accessed July 19, 2023, <https://www.gao.gov/products/gao-23-106648>.

¹²⁷ “Synthetic Biology 2020–2030: Six Commercially-Available Products That Are Changing Our World | Nature Communications,” accessed July 19, 2023, <https://www.nature.com/articles/s41467-020-20122-2>.

¹²⁸ Andrew Hessel, Marc Goodman, and Steven Kotler, “Hacking the President’s DNA,” *The Atlantic*, 2012, <https://www.theatlantic.com/magazine/archive/2012/11/hacking-the-presidents-dna/309147/> and Yelena Biberman, “The Technologies and International Politics of Genetic Warfare,” *Strategic Studies Quarterly*, Fall 2021, https://www.airuniversity.af.edu/Portals/10/SSQ/documents/Volume-15_Issue-3/Biberman.pdf. Russian statements around genetic warfare are discussed on page 45–46 of Raymond Zilinskas, “The Soviet Biological



This section explores several specific technologies, but it is important to note that there is a complex array of technologies and scientific breakthroughs that *in aggregate* are shifting the risk. As the science policy expert Sarah Carter¹²⁹ put it in an interview for this project, “There’s going to be a whole suite of enabling technologies that will get better and better and better. The hurdles to making a pathogen *de novo* are going to get lower and lower and lower. It’s just going to be this piecemeal thing.”¹³⁰ This, in turn, complicates philanthropic strategy; there is no single solution to any one specific technology, and philanthropists will need “a thousand small victories instead of one perfect solution,” in Dr. Carter’s words.¹³¹ (This complicated threat landscape helps to inform an approach described below, under “[GCBR Field-Building](#)”.)

These broad advances of the life sciences have also made relevant technologies cheaper and more widely accessible. One analysis estimates that the democratization of a novel biotechnology from being accessible only to well-trained highly-resourced actors is accelerating, and that this transition currently takes less than 4.5 years.¹³² For example, DNA synthesis is a critical part of synthetic biology; it provides the ability to design and manufacture the building blocks of life (and therefore of modified pathogens). The cost of DNA synthesis has broadly decreased by multiple orders of magnitude from \$10 per base pair to ~\$0.05 per base pair, as originally analyzed by Rob Carlson (hence “Carlson Curves”):¹³³

Weapons Program and Its Legacy in Today’s Russia” (Center for the Study of Weapons of Mass Destruction (National Defense University), 2016),
<https://ndupress.ndu.edu/Media/News/Article/848207/the-soviet-biological-weapons-program-and-its-legacy-in-todays-russia/https%3A%2F%2Fndupress.ndu.edu%2FMedia%2FNews%2FNews-Article-View%2FArticle%2F848207%2Fthe-soviet-biological-weapons-program-and-its-legacy-in-todays-russia%2F>.

¹²⁹ Dr. Sarah Carter is a Principal at Science Policy Consulting LLC and [science policy expert](#) who previously worked at the White House Office of Science and Technology Policy and co-authored NTI’s 2023 [report on benchtop DNA synthesis](#).

¹³⁰ Interview with Dr. Sarah Carter, 18 May 2023.

¹³¹ Full quote: “There’s going to be a whole suite of enabling technologies that will get better and better and better. The hurdles to making a pathogen *de novo* are going to get lower and lower and lower. It’s just going to be this piecemeal thing. I think that makes it a [...] hard question, a hard problem to work on, because you need a thousand small victories instead of one perfect solution. And I think that’s really hard for funders to focus on and to make progress on. It’s going to be a long game.” Interview with Dr. Sarah Carter, 18 May 2023.


¹³² Specifically, the analysis, published in 2019, states, “Our assessment provides evidence that novel technologies currently can complete this transition [‘from being accessible to a few well-resourced specialists to many individuals with relatively low levels of technical skill and financial resources’] in less than 4.5 years from their discovery and may do so in less than 3.5 years by the end of the next decade” Shawn S. Jackson et al., “The Accelerating Pace of Biotech Democratization,” *Nature Biotechnology* 37, no. 12 (December 2019): 1403, <https://doi.org/10.1038/s41587-019-0339-0>.

¹³³ For one attempt to model this, see Rob Carlson, “On DNA and Transistors,” synthesis (blog), March 9, 2016, http://www.synthesis.cc/synthesis/2016/03/on_dna_and_transistors. 7 cents appears to be the current industry starting price (see, e.g. [Twist](#)).



DNA Synthesis Cost 'Carlson Curve'

USD per base, 1999-2018. Data from Rob Carlson (2023) and Potomac (2018)



Plot generated in R using assistance from GPT4. Data from Carlson (2023) and Potomac (2018).
With thanks to Max Langenkamp (see in-text links).


Source: Data from Rob Carlson, “[DNA Cost and Productivity Data](#)” and Potomac Institute [report](#) (2018). Thanks to Max Langenkamp for initial analysis and key considerations [here](#). Plot generated in R with assistance from GPT4.

The changing costs of DNA synthesis — and other key measures, like changes in error rate and maximum length — is a complex topic beyond the scope of this report. For readers interested in more information, we commissioned a short overview by Max Langenkamp (author of the recent paper “[Clarifying the Problem of DNA Screening](#)”). This overview, “Examining the Changing Costs of DNA Synthesis,” can be found [here](#).

There are more high-quality data and forecasts available on DNA *sequencing*. The price of DNA sequencing has declined even more dramatically than synthesis cost, and is expected to continue declining. Although DNA sequencing is less directly relevant to some of the risks discussed here, it is nonetheless a crucial part of the work of synthetic biology, and the data provided here may serve as a proxy forecast for future declines in the field more broadly. To visualize the effects of this, we have



combined data from 2001 to 2022 from the National Human Genome Research Institute with predictions from Metaculus for 2026 and 2031, and visualized the cost declines in R.¹³⁴



Source: Author’s diagram using [NHGRI data](#), [2026 Metaculus Forecast](#), and [2031 Metaculus Forecast](#) (as of 28 June 2023). Plot generated in R with assistance from GPT4. See [appendix](#) for code.

Such cost declines and scientific advances democratize access to the tools of synthetic biology, but therefore also proliferate the knowledge and ability to create harmful pathogens. As one recent analysis put it, “**within a decade, tens of thousands of skilled individuals will be able to access the information required for them to single-handedly cause new pandemics.**”¹³⁵ The situation may be even worse. As discussed below, rapid advances in artificial intelligence may lower the barriers to synthetic biology even further. Forecasters on Metaculus estimate that “weakly general AI” will be

¹³⁴ For a similar visualization up to 2021, see “Cost per Billion Pairs of DNA Sequencing,” Our World in Data, accessed July 19, 2023, <https://ourworldindata.org/grapher/cost-per-gigabase-dna-sequencing>.

¹³⁵ Kevin M Esvelt, “Delay, Detect, Defend: Preparing for a Future in Which Thousands Can Release New Pandemics,” Geneva Papers (Geneva Centre for Security Policy, 2022), https://dam.gcsp.ch/files/doc/gcsp-geneva-paper-29-22?_gl=1*xm44p1*_ga*ODQxMDI0NiY4LjE2ODM2NTg5MTg.*_ga_Z66DSTVXTJ*MTY4MzczNTIzNi4yLjAuMTY4MzczNTIzNi4wLjAuMA..#page=49&zoom=100.0.0, 8.



publicly known by 2027.¹³⁶ If such advances occur and open-source models are not regulated, we believe that the statement above can be rephrased to state that **relatively unskilled individuals could also have access to these dangerous capabilities** — human hands, eyes, and basic comprehension of AI-generated instructions could be enough.

Benchtop DNA Synthesis Devices

Many gene synthesis providers voluntarily screen orders as part of the International Gene Synthesis Consortium (IGSC), but this group only “supplies an estimated 80% of the global market” (notably, the widely-quoted 80% figure is difficult to verify; the critical point is that a significant fraction of the market is not covered by IGSC members).¹³⁷ Moreover, although currently many centralized DNA synthesis providers screen customer orders, emerging technologies in next-generation benchtop DNA synthesis devices will allow users to create the building blocks of dangerous pathogens in private with little oversight — if these technologies are not regulated appropriately.¹³⁸ These new benchtop synthesis devices may decentralize the structure of the market, making oversight and regulation more difficult.¹³⁹

The risks of democratized access to DNA synthesis do not appear to be appreciated by some industry actors. One manufacturer of a benchtop DNA synthesis device promises the following for their \$25,000 device:

“[The device] is developed to provide scientists around the world with **independent and limitless access to essential life science tools**. Therefore, it is designed as user friendly as possible. The cartridge and chip technology **allows even untrained personnel to complete a successful DNA or RNA synthesis in no time**.”¹⁴⁰

¹³⁶ Forecast as of 22 August 2023. Metaculus, “When Will the First Weakly General AI System Be Devised, Tested, and Publicly Announced?,” <https://www.metaculus.com/questions/3479/date-weakly-general-ai-is-publicly-known/>.

¹³⁷ Kevin M Esvelt, “Delay, Detect, Defend: Preparing for a Future in Which Thousands Can Release New Pandemics,” Geneva Papers (Geneva Centre for Security Policy, 2022), https://dam.gcsp.ch/files/doc/gcsp-geneva-paper-29-22?_gl=1*xm44p1*_ga*ODOxMDI0NiY4LjE2ODM2NTg5MTg.*_ga_Z66DSTVXTJ*MTY4MzczNTIzNi4yLjAuMTY4MzczNTIzNi4wLjAuMA..#page=49&zoom=100.0.0, 26.

¹³⁸ Sarah Carter, Jaime Yassif, and Chris Isaac, “Benchtop DNA Synthesis Devices: Capabilities, Biosecurity Implications, and Governance” (The Nuclear Threat Initiative (NTI|bio), May 10, 2023), https://www.nti.org/wp-content/uploads/2023/05/NTIBIO_Benchtop-DNA-Report_FINAL.pdf.

¹³⁹ “Worse, next-generation ‘benchtop’ synthesis devices threaten to transition the market from large centralised providers that can be induced to screen orders and customers to a decentralised network of devices that may struggle to accomplish either type of screening.” Kevin M Esvelt, “Delay, Detect, Defend: Preparing for a Future in Which Thousands Can Release New Pandemics,” Geneva Papers (Geneva Centre for Security Policy, 2022), https://dam.gcsp.ch/files/doc/gcsp-geneva-paper-29-22?_gl=1*xm44p1*_ga*ODOxMDI0NiY4LjE2ODM2NTg5MTg.*_ga_Z66DSTVXTJ*MTY4MzczNTIzNi4yLjAuMTY4MzczNTIzNi4wLjAuMA..#page=49&zoom=100.0.0, 26.

¹⁴⁰ Kilobaser, “Kilobaser | Fastest DNA & RNA Synthesizer | Benchtop Sized,” kilobaser.com, accessed August 11, 2023, <https://kilobaser.com/dna-and-rna-synthesizer/>. Emphasis added.



One need only replace the word “scientists” with “terrorists” to appreciate the problem.

Assembling longer stretches of DNA (>10,000 base pairs) remains technically difficult within a benchtop device, and access to these devices is still restricted to well-funded groups.¹⁴¹ As Sarah Carter¹⁴² explained in an interview for this report, “These tools are not going to suddenly [...] enable terrorists with no bio training in a cave in Afghanistan.”¹⁴³ One analogy is that some technology in synthetic biology is like a word processor; writing a novel becomes easier with word-processing software, but remains difficult.¹⁴⁴ Moreover, some governance measures — like requiring manufacturers to screen users and their products on an ongoing basis — may help mitigate the risk.¹⁴⁵ For example, manufacturers could require that the device “phone home” to the manufacturer, but the connection from the device to the manufacturer could be spoofed by groups with cyber skills.¹⁴⁶ Unfortunately, as one recent analysis put it, “there is no feasible oversight mechanism or policy that will eliminate all risk; the approaches described in this chapter might be better characterized as ‘speed bumps’ to limit or slow the actions of a nefarious actor.”¹⁴⁷

This is a dire situation, but it is not hopeless. It means, rather, that interventions cannot focus exclusively on regulating these technologies or their use. Such regulation will be an important part of any layered approach to biosecurity, but it needs to be accompanied by strong defense (and deterrence-by-denial) in case prevention fails, as discussed below.¹⁴⁸ Society needs to invest large

¹⁴¹ “This limitation is driven by two factors: (1) imperfect sequence fidelity of initial oligos, which drives a need for quality control sequencing steps for dsDNA assembly, and (2) the need for bacterial or yeast culture to assemble and preserve dsDNA fragments longer than 7,000–10,000 base pairs” Sarah Carter, Jaime Yassif, and Chris Isaac, “Benchtop DNA Synthesis Devices: Capabilities, Biosecurity Implications, and Governance” (The Nuclear Threat Initiative (NTI|bio), May 10, 2023), 15, https://www.nti.org/wp-content/uploads/2023/05/NTIBIO_Benchtop-DNA-Report_FINAL.pdf.

¹⁴² Dr. Sarah Carter is a Principal at Science Policy Consulting LLC and [science policy expert](#) who co-authored NTI’s 2023 [report on benchtop DNA synthesis](#).

¹⁴³ Full quote: “Who is going to do that [i.e. engineer a dangerous virus]? It’s going to be a postdoc, it’s going to be a grad student, it’s going to be somebody already in the lab. These tools are not going to suddenly [...] enable terrorists with no bio training in a cave in Afghanistan. [...] Which is one reason I think that projects that could change the culture and academia to make them more aware of these things would be worthwhile, even though it’s a really hard intractable problem.” Interview with Dr. Sarah Carter, 18 May 2023.

¹⁴⁴ Thanks to Yelena Biberman for pointing me to this analogy.

¹⁴⁵ Sarah Carter, Jaime Yassif, and Chris Isaac, “Benchtop DNA Synthesis Devices: Capabilities, Biosecurity Implications, and Governance” (The Nuclear Threat Initiative (NTI|bio), May 10, 2023), https://www.nti.org/wp-content/uploads/2023/05/NTIBIO_Benchtop-DNA-Report_FINAL.pdf.

¹⁴⁶ Sarah Carter, Jaime Yassif, and Chris Isaac, “Benchtop DNA Synthesis Devices: Capabilities, Biosecurity Implications, and Governance” (The Nuclear Threat Initiative (NTI|bio), May 10, 2023), 31, https://www.nti.org/wp-content/uploads/2023/05/NTIBIO_Benchtop-DNA-Report_FINAL.pdf.

¹⁴⁷ Sarah Carter, Jaime Yassif, and Chris Isaac, “Benchtop DNA Synthesis Devices: Capabilities, Biosecurity Implications, and Governance” (The Nuclear Threat Initiative (NTI|bio), May 10, 2023), 28, https://www.nti.org/wp-content/uploads/2023/05/NTIBIO_Benchtop-DNA-Report_FINAL.pdf.

¹⁴⁸ This is the intuition behind Kevin Esvelt’s strategy of “Delay, Detect, Defend.” Kevin M Esvelt, “Delay, Detect, Defend: Preparing for a Future in Which Thousands Can Release New Pandemics,” Geneva Papers (Geneva Centre for Security Policy, 2022).



resources in pandemic prevention, detection, and response in the coming years, including the investments outlined above (work on transmission-blocking technologies, pandemic-proof PPE, etc.).

The AI-Bio Nexus

One major growing threat vector concerns the intersection of artificial intelligence and biological risks.¹⁴⁹ The growing capabilities of large language models in recent years have demonstrated that artificial intelligence is a rapidly growing field, and humanity ought to expect far more capable AI agents — up to and including transformative AI (TAI) — in the near future. Such developments may pose major threats to humanity on their own.¹⁵⁰ They could also pose threats as enabling technologies.

For example, recent work with drug discovery algorithms has demonstrated that such tools can easily be misused to identify known and unknown toxins and potential chemical weapons. Similar misuse of AI is clearly also possible with biology, further democratizing and accelerating dual-use life science knowledge and technology.¹⁵¹ One recent paper showed that non-scientist users are able to prompt AI chatbots to give instructions for causing a pandemic:

“In one hour, the chatbots suggested four potential pandemic pathogens, explained how they can be generated from synthetic DNA using reverse genetics, supplied the names of DNA synthesis companies unlikely to screen orders, identified detailed protocols and how to troubleshoot them, and recommended that anyone lacking the skills to perform reverse genetics engage a core facility or contract research organization”¹⁵²

More powerful systems may be able to help even low-skilled actors actually execute such plans. Moreover, AI systems are known to pursue strategies that are sometimes unintuitive or surprising to humans, complicating the threat landscape further. As discussed below, we believe more work on understanding and regulating the risks at the AI-bio intersection is critical, and it is possible that a **large fraction of future biological risk will originate from AI misuse**.¹⁵³

¹⁴⁹ Several experts interviewed for this report emphasized this point. For example, Tom Inglesby, the Director of the Johns Hopkins Center for Health Security, explained: “Now we have AI driving down the complexity of biotechnology. It is conceivable that AI will help simplify instructions, protocols, procedures for making biological organisms more lethal, more transmissible, and so we need to be [...] right here at the beginning of this huge expansion of AI to try to figure out what safeguards we can put in place” (Interview with Tom Inglesby, 8 June 2023).

¹⁵⁰ Joseph Carlsmith, “Is Power-Seeking AI an Existential Risk?” (Open Philanthropy, June 16, 2022), <http://arxiv.org/abs/2206.13353>.

¹⁵¹ John T. O’Brien and Cassidy Nelson, “Assessing the Risks Posed by the Convergence of Artificial Intelligence and Biotechnology,” *Health Security* 18, no. 3 (June 1, 2020): 219–27, <https://doi.org/10.1089/hs.2019.0122>.

¹⁵² Emily H. Soice et al., “Can Large Language Models Democratize Access to Dual-Use Biotechnology?” (arXiv, June 6, 2023), <http://arxiv.org/abs/2306.03809>.

¹⁵³ And, conversely, that a large fraction of AI misuse risk may be biological. Thanks to Andrew Sndyer-Beattie for emphasizing these points.



Cloud Labs

“Cloud labs” are laboratories that enable scientists and other actors to design, simulate, and test biological experiments remotely. These laboratories combine advanced robots and automation to perform the physical work in a wet lab that in the past required technicians. Users control and observe these experiments remotely via a digital interface. They can often pay for this service using a simple online subscription.¹⁵⁴ (NB: not all work in a cloud lab can be performed by robots, and some devices require human labor.¹⁵⁵)

Like other advances discussed here, cloud lab technology democratizes biological research by enabling more people and organizations to access advanced tools without the need for extensive resources or expertise — and without needing to invest in physical laboratory infrastructure. Cloud labs create efficiencies that can accelerate the pace of scientific work, in part by automating work that is not only time consuming but also “repetitive, painstaking, and tedious.”¹⁵⁶ However, this increased access and ease of use also elevates global catastrophic biological risks. With more entities able to conduct advanced biological research, the probability of misuse, either accidental or intentional, increases. This could include the creation of harmful biological agents. Furthermore, the decentralized nature of cloud labs makes it challenging to monitor and regulate activities effectively.

¹⁵⁴ “Experiments are programmed through a subscription-based online interface – software then coordinates robots and automated scientific instruments to perform the experiment and process the data.” Tom Ireland, “Cloud Labs and Remote Research Aren’t the Future of Science – They’re Here,” *The Observer*, September 11, 2022, sec. Science, <https://www.theguardian.com/science/2022/sep/11/cloud-labs-and-remote-research-arent-the-future-of-science-theyre-here>.

¹⁵⁵ “There are still some things robots can’t do, for example lifting giant carboys (containers for liquids) or unwrapping samples sent by mail, and there are a few instruments that just can’t be automated. Hence the people in blue coats, who look a little like pickers in an Amazon warehouse. It turns out that they are, in fact, mostly former Amazon employees.” Tom Ireland, “Cloud Labs and Remote Research Aren’t the Future of Science – They’re Here,” *The Observer*, September 11, 2022, sec. Science, <https://www.theguardian.com/science/2022/sep/11/cloud-labs-and-remote-research-arent-the-future-of-science-theyre-here>.

¹⁵⁶ Matt Field, “Laboratories in the Cloud,” *Bulletin of the Atomic Scientists*, July 3, 2019, <https://thebulletin.org/2019/07/laboratories-in-the-cloud/>.



A “Cloud Laboratory”



Source: Emerald Cloud Labs [Virtual Tour](#).

Several uncertainties related to cloud labs will determine their use and vulnerabilities, including:

- Their cost and advantages over traditional laboratories;
- Their ability to screen users;
- How quickly they become mainstream and widely accessible;
- Questions around the cybersecurity vulnerabilities of cloud labs.

Nonetheless, cloud labs appear to be yet another part of the transformation of the life sciences. This seems especially important with advances in AI. Misuse of AI could run experiments that provide valuable insights on misusing biology, potentially without raising red flags with the labs themselves.

Other Advances

In addition to the technological and scientific advances discussed above, there are other possible technologies that could enable the misuse or increase the probability of accidents related to pandemic-potential pathogens. For example, atomically-precise manufacturing could enable the development of novel delivery mechanisms for genetic material that are less fickle than living



organisms.¹⁵⁷ For now, such technologies remain in the realm of speculative fiction, but much work is needed to identify emerging technologies with dual-use risk — without raising the profile of these risks.

The Implications of the Changing Threat Landscape

What does all of this mean? This rapidly-changing threat landscape has two major implications. First, **the risk of engineered pandemics is increasing**. The technological advancements discussed above — cloud labs, improved genetic engineering, AI-bio intersections, benchtop DNA synthesis, and more — all contribute to a picture where more actors have easier access to more advanced technologies, and oversight becomes more difficult. This both increases the potential for misuse and the surface area for accidents.

Second, **uncertainty around the threat landscape is increasing**. In the future, analysts may no longer be able to point to specific viral families as the biggest threats. Rather, more and more people working on more and more exotic technologies could easily manufacture novel risks, circumvent defenses against known pathogens, and engage in risk-shifting behavior, as discussed in the next section. In the past, advanced biological weapons programs were the purview of powerful states. In the future, the risk may proliferate to more numerous actors with fewer resources and less skill, complicating surveillance and increasing uncertainty about the origins of future threats.¹⁵⁸

Open Questions for Further Research

- How can access to synthetic DNA — including access via benchtop devices — be regulated in ways that protect against risks while preserving upsides ?
- What governance tools can help reduce both the risk of misaligned AI and the risk of the misuse of AI in biology?
- What are the implications of the proliferation of open-source AI models for biosecurity?
- Are there specific technological breakthroughs that could be “game changers” for the risk? If so, what “alarm bells” can be built to monitor and respond to relevant breakthroughs?
- What is the offense-defense balance of the technologies discussed above?

¹⁵⁷ “Nanotechnology offers new delivery possibilities for biological and genetic weapons. In the future, nano-carriers and capsules may transport small toxins, such as ricin or microbe subunits (e.g., the lethal factor of anthrax), across otherwise impermeable cell membranes and the blood-brain barrier. Bioagents’ targeted delivery with nanoparticles is likely to increase effectiveness and, thus, require less of the agent.

Nanotechnology could also enable controlling biological weapons once they enter the body.” Yelena Biberman, “The Technologies and International Politics of Genetic Warfare,” *Strategic Studies Quarterly*, Fall 2021.

¹⁵⁸ This intuition is explained more formally in Anders Sandberg and Cassidy Nelson, “Who Should We Fear More: Biohackers, Disgruntled Postdocs, or Bad Governments? A Simple Risk Chain Model of Biorisk,” *Health Security* 18, no. 3 (June 1, 2020): 155–63, <https://doi.org/10.1089/hs.2019.0115>.



Analyzing the Structure of the Problem

Key Points

- The threat of biological catastrophe is growing.
- The cost curve of this threat is superlinear — the worst pandemics are *disproportionately* and potentially *discontinuously* worse than others.
- The threat landscape is growing increasingly complex.
- Unlike some catastrophic risks, the threat is adaptive; malevolent actors can respond to risk-mitigation measures by shifting their strategies.

In order to begin building a philanthropic strategy, this section seeks to analyze the structure of the problem of global catastrophic biological risks. How is the problem changing? What is the cost curve of the problem? How do risk-mitigation strategies affect the risk landscape? What tradeoffs exist?

Building on the previous sections, the structure of the threat can be described as growing, superlinear, increasingly complex, and adaptive:

1. **Growing:** Techno-scientific progress in the life sciences, combined with observations about the offense-defense balance of these advances, suggest that the threat will increase over the coming decades.
2. **Super-linear:** The worst pandemics are *disproportionately* worse than smaller kinds of pandemics; the threat of civilizational collapse looms large.¹⁵⁹
3. **Increasingly Complex:** Although we can draw some stylized conclusions about the threat landscape (e.g., engineered pathogens are probably worse than naturally-evolved pathogens), much of this landscape remains amorphous. This uncertainty is compounded by the difficulty of predicting techno-scientific progress, creating a complex array of threat vectors.
4. **Adaptive:** Pathogens can mutate and malevolent actors can respond to defensive measures. This possibility of risk-shifting makes much of GCBR mitigation a tragic arms-race spiral or “Red Queen’s race”, where “it takes all the running you can do, to keep in the same place.”¹⁶⁰

¹⁵⁹ This is discussed in greater detail under [The Threat is Super-Linear](#)

¹⁶⁰ “Well, in our country,” said Alice, still panting a little, “you’d generally get to somewhere else — if you run very fast for a long time, as we’ve been doing.” “A slow sort of country!” said the Queen. “Now, here, you see, it takes all the running you can do, to keep in the same place. If you want to get somewhere else, you must run at least twice as fast as that!” Lewis Carroll, *Through the Looking Glass and What Alice Found There*. United Kingdom: Macmillan and Company, 1998.



These features of the problem can appear dispiriting. We are, after all, facing a threat that is becoming at once bigger and more nebulous. Worse, the threat can be responsive to our risk-mitigation measures, shifting the risk landscape beneath our feet. As discussed later, however, philanthropists can leverage these apparently discouraging features for more effective risk mitigation.

The Threat Is Growing

The first important feature of the risk is that the threat appears to be growing for two reasons:

1. Rapid advances in the life sciences
2. Increasing availability of relevant knowledge and technology

Because these were covered in detail above and in other analyses, this section will not repeat the arguments, except in bullet point form:

1. Advances in synthetic biology, including gene editing technology like CRISPR-Cas9, are facilitating the manipulation of pathogens.
2. The technology for synthetic biology is becoming cheaper, more reliable, and more widely accessible.
3. The culture of openness in the life sciences means that scientific discoveries — even dangerous ones — are likely to be shared widely.
4. Rapid advances in machine learning will facilitate drug discovery, but have dual-use potential for helping to create dangerous pathogens.
5. Relatedly, AI may simply accelerate the pace of scientific discovery by making fundamental breakthroughs (viz. AlphaFold) and facilitating higher efficiency of normal science (e.g. by acting as high-powered research assistants).
6. Weak governance and the lack of a culture of security suggest that these advances will proliferate widely, with little surveillance for misuse.
7. We are therefore rapidly approaching a future in which thousands of people have the skill and ability to engineer and release deadly pathogens across the globe.¹⁶¹

¹⁶¹ Kevin M Esvelt, “Delay, Detect, Defend: Preparing for a Future in Which Thousands Can Release New Pandemics,” Geneva Papers (Geneva Centre for Security Policy, 2022), https://dam.gcsp.ch/files/doc/gcsp-geneva-paper-29-22?_gl=1*xm44p1*_ga*ODQxMDI0NiY4LjE2ODM2NTg5MTg.*_ga_Z66DSTVXTJ*MTY4MzczNTIzNi4yLjAuMTY4MzczNTIzNi4wLjAuMA..#page=49&zoom=100.0.0.



The Offense-Defense Balance of Emerging Biotechnologies

A key uncertainty about the growing threat concerns the “offense-defense balance” of biotechnology. In international relations and military strategy, the concept of an offense-defense balance refers to “the relative ease of carrying out and defending against attacks,” though complex theories and debates surround the idea.¹⁶² In his classic 1978 formulation of the offense-defense balance, international relations scholar Robert Jervis writes:

*When we say that the offense has the advantage, we simply mean that it is easier to destroy the other's army and take its territory than it is to defend one's own. When the defense has the advantage, it is easier to protect and to hold than it is to move forward, destroy, and take.*¹⁶³

The offense-defense balance no longer just refers to the battlefield, but has been applied to broad domains, such that scholars speak of a “cyber offense-defense balance.”¹⁶⁴

What does the offense-defense balance look like in the life sciences? If future advances are likely to favor defense, then the nature of the threat would be far less concerning, and philanthropists ought to consider interventions that accelerate scientific progress. Some analysts have even envisioned a future world “immune” to global catastrophic biological risks.¹⁶⁵ Despite recent attempts to understand how the offense-defense balance “scales” — i.e. how it shifts with increasing investment — answering this question remains highly complex.¹⁶⁶ It is possible, however, to examine the structure of the problem as described above to see that the nature of the biological threat appears to favor offense, and that future advances in the life sciences — unless intentionally redirected to the defensive, are likely to *increase* the risk when the objective of the defense is defined as preventing a global catastrophe.¹⁶⁷

¹⁶² Ben Garfinkel and Allan Dafoe, “How Does the Offense-Defense Balance Scale?,” *Journal of Strategic Studies* 42, no. 6 (September 19, 2019): 3, <https://doi.org/10.1080/01402390.2019.1631810>.

¹⁶³ Robert Jervis, “Cooperation under the Security Dilemma,” *World Politics* 30, no. 2 (January 1978): 187, <https://doi.org/10.2307/2009958>.

¹⁶⁴ Rebecca Slayton, “What Is the Cyber Offense-Defense Balance?: Conceptions, Causes, and Assessment,” *International Security* 41, no. 3 (2016): 72–109.

¹⁶⁵ “If we think that biothreats largely lie with future discoveries in biotechnology, then general biotechnology will increase capabilities for both defense and offense. It is plausible that in the limit advances in biotechnology will overwhelmingly favor the defender, but not guaranteed, and net risk may be increased in the interim” Shulman Carl, “Envisioning a World Immune to Global Catastrophic Biological Risks,” *Reflective Disequilibrium* (blog), October 15, 2020, <http://reflectivedisequilibrium.blogspot.com/2020/05/what-would-civilization-immune-to.html>.

¹⁶⁶ Ben Garfinkel and Allan Dafoe, “How Does the Offense-Defense Balance Scale?,” *Journal of Strategic Studies* 42, no. 6 (September 19, 2019): 3, <https://doi.org/10.1080/01402390.2019.1631810>.

¹⁶⁷ Notably, this discussion applies more readily to biological catastrophe rather than full extinction. Defense against human extinction may be more defense-dominant; defense could consist of a permanently-occupied bio-resilient bunker. We believe, however, that few people would consider such a situation a satisfying defense for humanity. Thanks to Andrew Snyder-Beattie for pointing to the importance of defining the objective of the defensive.



Several considerations suggest that the offense enjoys a major advantage in biosecurity and pandemic preparedness. Most importantly, the self-replicating nature of pathogens, combined with the impossibility (and undesirability) of global surveillance, create a dynamic familiar in counterterrorism efforts, where **the defender must succeed everywhere all the time**. This is encapsulated in a now-famous statement by the Irish Republican Army after a failed attack in 1984: “[R]emember, we only have to be lucky once. You have to be lucky always.”¹⁶⁸ Similarly, a group seeking to release a dangerous pathogen only needs to be successful in releasing a pathogen into one super-spreader situation in order to potentially trigger a devastating pandemic. Pathogens spread exponentially, but most countermeasures do not.¹⁶⁹ Relatedly, **because pathogens spread exponentially, time favors offense**. With every passing day, the threat becomes disproportionately less controllable. Moreover, cascading systems collapses — of systems that are required for defense (supply chains, strategic stockpile distribution, etc.) — hobble the defense more and more as time goes on.

Additionally, **much bioscience is dual-use**. Thus, for example, improved genetic engineering ability may help scientists find new cures for diseases, but also enable bad actors to create new threats. In a similar vein, some apparently defensive actions actually increase the probability of accidents and misuse, as discussed above, such that the overall threat level may increase with misguided investments in defense.

To review, some key considerations on biological risks paint a bleak picture of the offense-defense balance:

- **Pathogens spread exponentially, but most countermeasures do not.** This suggests a bias for offense.
- **Time favors offense** — with every passing day, the threat becomes disproportionately more difficult to control.
- Techno-scientific progress appears to be:
 - **Democratizing access to and knowledge about biological tools and processes.**
 - Making much bioscience and biotechnology **cheaper to acquire and use**.
 - Frequently **dual-use** — many advances could be misused.
- Offense needs to find only one vulnerability. Defense, on the other hand, must aspire to patch as many vulnerabilities as possible. **Attackers “only have to be lucky once.”**

Risk then becomes a simple numbers game, and leads to the conclusion that **the number of potential malicious actors with access to powerful dual-use technologies is likely to grow**. With it,

¹⁶⁸ Jo Thomas, “This Time, the IRA Comes Close to Thatcher,” *The New York Times*, October 14, 1984, sec. Week in Review, <https://www.nytimes.com/1984/10/14/weekinreview/this-time-the-ira-comes-close-to-thatcher.html>.

¹⁶⁹ There may be certain kinds of exceptions, but these could also pose dual-use risks.



the number of attackers who “only have to be lucky once” grows, whereas the number of times that the defender has to be lucky also grows — the landscape asymmetrically begins to disadvantage the defense. In short, the **threat of biological catastrophe appears to be growing**.

The Threat Is Super-Linear

This section argues that the worst pandemics are *disproportionately* worse than smaller pandemics; the cost-curve of the threat, therefore, is superlinear. This problem is very similar in structure to nuclear war and climate change — the superlinearity of these threats are discussed in two Founders Pledge reports: [A Guide to the Changing Landscape of High-Impact Climate Philanthropy](#) and in [Global Catastrophic Nuclear Risk: A Guide for Philanthropists](#).

In biosecurity, three features of the problem lead to the conclusion that **philanthropists concerned about averting the most damage per dollar ought to prioritize worst-case scenarios**:

1. The distribution of the threat’s expected costs.
2. Moral discontinuities.
3. System failures and cascading collapses.

First, it appears likely that the fatalities of global catastrophic biological risks follow a power law with a small exponent — they are an example of a “heavy-tailed” distribution.¹⁷⁰ One feature of such distributions is that most of the expected costs come from the very few worst case scenarios; catastrophes dominate the statistics.¹⁷¹ This has important implications, as explained by Millett and Snyder-Beattie (2017): “This suggests that a typical individual is far more likely to die from a rare, catastrophic attack as opposed to a smaller scale and more common one. If our goal is to reduce the greatest expected number of fatalities, we may be better off devoting resources to preventing the worst possible attacks.”¹⁷² Whether or not biological catastrophes and other disasters do actually follow power-laws is one uncertainty, but heavy tails are a feature of other distributions as well.¹⁷³

¹⁷⁰ Much biological risk is a subset of warfare and terrorism — which both appear to follow power-laws as described (see next footnote).

¹⁷¹ “Perhaps the most interesting implication of the fatalities following a power law with a small exponent is that the majority of the expected casualties come from rare, catastrophic events. The data also bear this out for warfare and terrorism. The vast majority of US terrorism deaths occurred during 9/11, and the vast majority of terrorism injuries in Japan over the past decades came from a single Aum Shinrikyo attack” Piers Millett and Andrew Snyder-Beattie, “Existential Risk and Cost-Effective Biosecurity,” *Health Security* 15, no. 4 (August 2017): 377, <https://doi.org/10.1089/hs.2017.0028>.

¹⁷² Piers Millett and Andrew Snyder-Beattie, “Existential Risk and Cost-Effective Biosecurity,” *Health Security* 15, no. 4 (August 2017): 378, <https://doi.org/10.1089/hs.2017.0028>.

¹⁷³ “There is much debate about whether the distributions of various disasters are *really* power laws. For example, log-normal distributions have right-hand tails that approximate a power law, so could be mistaken for them, but have a lower probability of small events than in a true power law. For our purpose, we don’t really need to distinguish between different heavy-tailed distributions. We are really just interested in whether the right-hand tail (the distribution of large events) behaves as a power law...” Ord, *The Precipice*, 357.



There are some empirical observations that suggest that rare events will dominate expected costs. For example, the Aum Shinrikyo attacks resulted in “the vast majority of terrorism injuries” over decades in Japan, and empirical studies suggest that bioterrorism and chemical-weapons terrorism generally follow this pattern.¹⁷⁴ Moreover, there are mechanical reasons for expecting long tails in the distribution of death tolls of pandemics; death toll is the product of different factors (transmissibility x lethality x population density x etc.).¹⁷⁵

A second consideration that informs the super-linearity of the cost curve of GCBRs is the problem of **moral discontinuities for extreme catastrophes** that could affect the long-term future of humanity. On a variety of moral views, the potential value of the future appears vast, such that actions that benefit the size and goodness of the future can take on overwhelming moral importance. (Because there are many popular and academic defenses of the philosophical arguments behind this point of view, readers can consult resources in the footnotes for more detail.¹⁷⁶) Thus, because the most extreme biological risks could plausibly cause the extinction of humanity or civilizational collapse, they are a subset of *existential risks*; averting these most extreme risks takes on a discontinuous importance, creating an *extreme* superlinearity of value.¹⁷⁷

Readers do not need, however, to subscribe to longtermist philosophical arguments or to abstract fitting of heavy-tailed distributions to sparse data in order to understand the super-linearity of the risk. The third driver of this non-linear risk structure is the problem of **system failures and cascading collapses**. The bigger the catastrophe is, the more likely it becomes that critical infrastructure systems collapse, triggering even greater catastrophes in the world, and potentially sparking cascading collapses of modern civilization.¹⁷⁸ As Jaime Yassif has explained, “it is likely that once fatalities got to a certain level, we could expect to see a variety of knock-on effects, including collapsing financial systems, collapsing healthcare systems, disrupted global supply chains, food and water shortages, power disruptions, or military conflict.”¹⁷⁹ One feature of the problem is familiar from the challenges of the COVID-19 pandemic; at certain points, healthcare systems become overwhelmed and are unable to handle more patients, and just-in-time supply chains begin to fail.

¹⁷⁴ The data also bear this out for warfare and terrorism. The vast majority of US terrorism deaths occurred during 9/11, and the vast majority of terrorism injuries in Japan over the past decades came from a single Aum Shinrikyo attack” Piers Millett and Andrew Snyder-Beattie, “Existential Risk and Cost-Effective Biosecurity,” *Health Security* 15, no. 4 (August 2017): 377, <https://doi.org/10.1089/hs.2017.0028>.

¹⁷⁵ Thanks to Matt Lerner for pointing to this.

¹⁷⁶ See longtermism.com, “Resources,” longtermism.com, 2021, <https://longtermism.com/resources>.

¹⁷⁷ The exact point of this discontinuity is unclear. Catastrophes short of full extinction could still cause unrecoverable civilizational collapse. Moreover, even civilizational collapse that is followed by a civilizational recovery could well lead to moral disaster if the recovery is worse than the *status quo ante*. For a detailed discussion of the problems of collapse, see Belfield’s chapter “Collapse, Recovery, and Existential Risk,” in Centeno *et al.*, *How Worlds Collapse: What Society, Systems, and Complexity Can Teach Us About Our Modern World and Fragile Future*.

¹⁷⁸ David Manheim, “The Fragile World Hypothesis: Complexity, Fragility, and Systemic Existential Risk,” *Futures* 122 (2020), <https://doi.org/10.1016/j.futures.2020.102570>.

¹⁷⁹ Jaime Yassif, “Reducing Global Catastrophic Biological Risks,” *Health Security* 15, no. 4 (August 1, 2017): 329–30, <https://doi.org/10.1089/hs.2017.0049>.



For worse pandemics, these problems could be exacerbated if essential workers are unable or unwilling to go to work to maintain water supplies, electrical infrastructure, food distribution, and more. In our modern world, each system relies on other systems-of-systems and — crucially — effective response relies on the functioning of critical infrastructure (e.g. to maintain the state capacity to distribute vaccines). Each such system failure represents a discontinuity in the cost, such that the cumulative and cascading failures at the most extreme possible events again create a super-linear cost function in aggregate.¹⁸⁰ For philanthropists, this implies [focusing on worst-case scenarios](#) for cost-effectiveness, all else equal.

The Threat Is Increasingly Complex

Many questions surround the relative risk of different threat vectors within the broader class of global catastrophic risks: Are engineered pandemics worse than natural pandemics? Is accidental or deliberate risk the greater threat? Are there specific viral families that we ought to focus on? Which particular technologies are the most dangerous? Would the great powers actually be interested in developing biological weapons in the 21st century? What scientific breakthroughs would change the risk calculus for rogue states? Perhaps most fundamentally, **what does the threat look like?**

We can draw some stylized conclusions about the threat landscape (e.g., engineered pathogens are probably worse than naturally-evolved pathogens), but much of this landscape remains amorphous. This uncertainty is compounded by the difficulty of predicting techno-scientific progress, creating a complex array of threat vectors. Consider the example of state biological weapons programs. As discussed above, powerful nuclear-armed states may be uninterested in developing potentially-omnicidal weapons for various reasons (they are indiscriminate, they backfire, they have a high risk of accidents, they are difficult to stockpile, deterrence may be less credible, etc.). But it is possible that advances in genetic engineering could change this risk calculus for some states if, for example, they made targeting of specific groups possible.

In fact, the disagreements in biosecurity between different groups and approaches — e.g. between security-minded biodefense professionals who worry about terrorist access to pandemic blueprints and between global health-oriented scientists who worry about natural spillovers in a climate-change world — may reflect this inherent complexity. In other words, **we simply do not know where exactly the next pandemic will come from or what it will look like.** Worse, this uncertainty may be increasing. This problem is compounded by the problem of an adaptive threat, discussed in the next section. For philanthropists, this implies [prioritizing pathogen- and threat-agnostic approaches](#), as discussed below.

The Threat Is Adaptive

An additional feature of the risk landscape is that the threat is *adaptive*. It is illuminating to compare biological risk to the risks of near-earth object collisions (i.e. asteroid impacts) or large-magnitude

¹⁸⁰ Similarly, in climate change, “tipping points” result from certain physical thresholds that lead to [nonlinear damages](#). Thanks to Matt Lerner for pointing to this analogy.



explosive volcanic eruptions (LaMEVEs, or “supervolcanoes”). Asteroids and supervolcanoes do not respond to our risk mitigation measures; an asteroid will not take evasive maneuvers against a planetary defense system, and a supervolcano does not care if Earth’s food systems are resilient or not.

Biological risks, however, are adaptive in two ways. First, the possibility of mutation and the dynamics of natural selection make the risk more difficult to contain, as evidenced by the difficulty of responding to new strains of the coronavirus during the COVID pandemic. Second, because deliberate misuse is a major threat, we are facing intelligent adversaries who can examine humanity’s defenses for potential weaknesses and shift their strategies accordingly. (This risk becomes especially insidious if non-human intelligences are involved, given that even weak AI systems consistently display problem-solving behaviors that are surprising or pursue solutions that are unintuitive to human intelligence.)

In other words, there is a problem of **risk-shifting**; targeted actions may not *decrease* aggregate risk so much as shift it. For example, a pan-coronavirus vaccine could be useful for containing natural outbreaks, but agents interested in deliberate misuse could simply orient their malevolent research to another viral family. Indeed, if limited defenses (like targeted biomedical interventions) encourage complacency about the risk, then such risk-shifting becomes especially dangerous if bad actors take advantage of the reduced vigilance. If pursued poorly, this can become a game of whack-a-mole or a “Red Queen’s race”, where “it takes all the running you can do, to keep in the same place.”¹⁸¹ Given the offense-defense balance discussed above, bad actors may well win the race. For philanthropists, this implies focusing on risk-general approaches (see [Prioritize Pathogen- and Threat-Agnostic Approaches](#)).

Open Questions for Further Research

- What, exactly, is the shape of the cost curve? Where do superlinear costs begin to kick in?
- How do assessments of offense-defense balance shift based on the priorities of the defender?
- How far does risk-shifting go? At what point does constant shifting of the risk landscape deter malevolent actors?
- What can philanthropists learn from risk-shifting in other domains?

¹⁸¹ “Well, in our country,” said Alice, still panting a little, “you’d generally get to somewhere else — if you run very fast for a long time, as we’ve been doing.” “A slow sort of country!” said the Queen. “Now, here, you see, it takes all the running you can do, to keep in the same place. If you want to get somewhere else, you must run at least twice as fast as that!” Lewis Carroll, *Through the Looking Glass and What Alice Found There*. United Kingdom: Macmillan and Company, 1998.



Societal Spending and Neglectedness

Key Points

- Despite billions of dollars spent on health security, only a small fraction of societal spending is relevant to catastrophic pandemics.
- Various factors help to explain this market failure in public policy, even in the wake of the COVID-19 pandemic.
- Philanthropists, too, mostly neglect GCBRs, with only a small handful of funders explicitly focused on the topic.
- These features create an opportunity for high-impact philanthropy.

The previous two sections argued that the threat of biological catastrophe is immense, and that this threat is both increasing and growing more complex. This section examines humanity's preparedness for this large and changing threat. The section makes three broad observations:

1. At first glance, **societies appear to devote many resources to health security**. Government biodefense and health spending seem vast, and many groups are worried about the threat of infectious disease.
2. Upon closer examination, however, public health and pandemic preparedness remain chronically under-funded and **only a small fraction of this work is focused on and robust to worst-case scenarios**, even after the potential wake-up call of the COVID-19 pandemic.
 - a. There are several possible explanations for this **market failure in public policy**:
 - i. Leaders' near-term political incentives diverge from long-term global incentives.
 - ii. Cognitive biases — especially scope neglect — distort decision-makers' ability to accurately size up the threat.
 1. This previously included the common complaint that the risk is negligible because its probability is small (even though its consequences are immense). The COVID-19 pandemic hopefully dispels this idea.



- iii. The urgent crowds out the important in policy making, such that we often end up “fighting the last war.”
 - iv. Society is suffering from pandemic fatigue.
 - v. There is a general lack of risk-aware decision-making in policy making.
 - vi. The professional culture of the life sciences lacks a “security mindset.”
3. In the private sector, **fewer than 5 philanthropic funders intentionally focus on worst-case global catastrophic biological risks.**

In short, although the total “pie” (societal resources devoted to health security) appears large, it is arguably not commensurate with the threat, and only a small slice of this pie is focused on worst-case scenarios. This means that **there is a high-leverage niche for effective philanthropy on the margin.**

Societal Resources Spent on Health Security

At first glance, the problems of infectious disease and health security appear to be far from neglected. For ease of analysis and because the United States is by far the biggest funder of health security activities, we are focusing on U.S. spending in this section, but note that other countries appear to follow similar patterns. U.S. global health security funding is about \$1.5 billion in fiscal year 2023, but much of this goes to USAID programs focused on endemic diseases like malaria and HIV.¹⁸² Similarly, the Biden administration has touted that the 2023 budget includes “a historic \$88.2 billion *request* (i.e. not appropriated funding) for mandatory funding [across multiple parts of government] to prepare for future biological threats in support of objectives within U.S. national and global biodefense and pandemic preparedness strategies and plans.”¹⁸³

These numbers seem impressive, but can be misleading. The politics of government budgeting mean that line items don’t always mean what they appear to, making a rigorous analysis of health security spending difficult. As one long-running analysis of biodefense funding (formerly called “Billions for Biodefense”) put it in 2014, “a number of biodefense programs have been combined with other line items and are no longer distinguishable, leading to the appearance of higher funding.”¹⁸⁴ Thus,

¹⁸² KFF, “Breaking Down the U.S. Global Health Budget by Program Area | KFF,” accessed August 11, 2023, <https://www.kff.org/global-health-policy/fact-sheet/breaking-down-the-u-s-global-health-budget-by-program-area/>.

¹⁸³ The White House, “FACT SHEET: The Biden Administration’s Historic Investment in Pandemic Preparedness and Biodefense in the FY 2023 President’s Budget,” The White House, March 28, 2022, <https://www.whitehouse.gov/briefing-room/statements-releases/2022/03/28/fact-sheet-the-biden-administrations-historic-investment-in-pandemic-preparedness-and-biodefense-in-the-fy-2023-presidents-budget/>.

¹⁸⁴ Tara Kirk Sell and Matthew Watson, “Federal Agency Biodefense Funding, FY2013-FY2014,” *Biosecurity and Bioterrorism: Biodefense Strategy, Practice, and Science*. 11, no. 3, accessed August 11, 2023, <https://www.liebertpub.com/doi/abs/10.1089/bsp.2013.0047>.



between 2001 and 2014, only about 21% of funding was “strictly biodefense” as opposed to multi-use funding.¹⁸⁵

Funding of pandemic preparedness is crisis-driven. Tom Inglesby, the Director of the Johns Hopkins Center for Health Security, explained in an interview for this report: “We don't have any sustained government funding efforts around this. To the extent it does exist, it is pretty volatile, it's crisis driven, and it's usually very specific to the problem at hand — like COVID — and then it's over.”¹⁸⁶ How much of *apparent* biosecurity funding since the COVID-19 pandemic is *truly* biosecurity funding? Again, the question is difficult to answer definitively. Recently, however, the administration followed a recommendation by the Bipartisan Commission for Biodefense and released an Office of Management and Budget “crosscut” analysis and report to Congress.¹⁸⁷ This crosscut report can be found [here](#).

According to the report, excluding COVID-response funding, the President's FY 2023 budget requested about **\$11 billion in discretionary biodefense funding**.¹⁸⁸ (“Biodefense” here includes the categories of “Threat Awareness, Prevention, Preparedness, Surveillance & Detection, Response, Recovery, Bioforensics & Attribution, and Mitigation.”¹⁸⁹) This is a large pie of money, which targeted policy advocacy can help to reallocate in a way that helps to better protect all people from pandemic threats.

Only a small slice of the bigger pie of biodefense funding is plausibly relevant to worst case scenarios, however. For example, the final 2023 budget only provided about \$950 million for the Biomedical Advanced Research and Development Authority (BARDA) via the Public Health and Social Services Emergency Fund, though much of this is arguably not relevant to catastrophic biological risks.¹⁹⁰ Gregory Lewis has estimated that quality-adjusted societal spending on global

¹⁸⁵ Tara Kirk Sell and Matthew Watson, “Federal Agency Biodefense Funding, FY2013-FY2014,” *Biosecurity and Bioterrorism: Biodefense Strategy, Practice, and Science*. 11, no. 3, accessed August 11, 2023, <https://www.liebertpub.com/doi/abs/10.1089/bsp.2013.0047>.

¹⁸⁶ Interview with Tom Inglesby, 8 June 2023.

¹⁸⁷ Bipartisan Commission on Biodefense, “Commission Welcomes New Analysis of Biodefense Spending That Will Help Inform Future Investments (Press Release),” Bipartisan Commission on Biodefense, January 19, 2023, <https://biodefensecommission.org/commission-welcomes-new-analysis-of-biodefense-spending-that-will-help-inform-future-investments/>.

¹⁸⁸ Office of Management and Budget, Executive Office of the President, “Report to Congress on Biodefense Activities,” January 2023, <https://biodefensecommission.org/wp-content/uploads/2023/01/OMB-Report-Biodefense-Activities-FY-2023-Budget.pdf>.

¹⁸⁹ Office of Management and Budget, Executive Office of the President, “Report to Congress on Biodefense Activities,” January 2023, <https://biodefensecommission.org/wp-content/uploads/2023/01/OMB-Report-Biodefense-Activities-FY-2023-Budget.pdf>.

¹⁹⁰ “For expenses necessary to support activities related to countering potential biological, nuclear, radiological, chemical, and cybersecurity threats to civilian populations, and for other public health emergencies, \$1,647,569,000, of which \$950,000,000 shall remain available through September 30, 2024, for expenses necessary to support advanced research and development pursuant to section 319L of the PHS Act



catastrophic biological risks is at about \$1 billion per year.¹⁹¹ One recent analysis has stated that, in biosecurity broadly defined, “Very little effort has been devoted to the possibility of deliberately released pandemic agents, and none to plausible worst-case scenarios.”¹⁹²

Jason Matheny,¹⁹³ who helped lead various parts of U.S. policy on issues relevant to emerging technologies and biosecurity, explained in an interview for this report that “A lot of the work on traditional biodefense and biosecurity is just not as germane to what these very extreme events look like.”¹⁹⁴ One policymaker interviewed for this project stated that “there is lots of funding in the federal government on biodefense and biosecurity, but unfortunately, most of it is irrelevant to X-risks [i.e., biological threats that could plausibly pose existential risks].”¹⁹⁵ Asked to quantify the amount of funding that is relevant to extreme biological risks, this policymaker estimated **that between 90% and 98% of traditional biodefense and biosecurity funding in the federal government is not relevant to extinction-level biological risks.**¹⁹⁶

Some of this funding includes wastewater surveillance, laboratory security, personnel reliability measures, and others, but **there is essentially zero funding for many critical biodefense measures** like advanced personal protective equipment, DNA synthesis screening, and biotechnology supply chain monitoring. Instead, a large majority of this spending goes towards medical countermeasures.¹⁹⁷

Worse, **some biodefense funding — including work with enhanced potential pandemic pathogens — may be actively increasing the risk** rather than mitigating it. For example, the \$125 million Deep Vzn (Discovery and Exploration of Emerging Pathogens — Viral Zoonoses) virus hunting program had been widely criticized as potentially raising the risk of both natural spillover and intentional misuse,

and other administrative expenses of the Biomedical Advanced Research and Development Authority” H.R. 2617, <https://www.govinfo.gov/content/pkg/BILLS-117hr2617enr/pdf/BILLS-117hr2617enr.pdf>, p. 419 in PDF.

¹⁹¹ “This issue is somewhat neglected. Current spending is in the billions per year, although this large portfolio is not perfectly allocated. Our guesstimated quality adjustment yields ~ \$1 billion per year” Gregory Lewis, “Reducing Global Catastrophic Biological Risks,” 80,000 Hours, March 2020, <https://80000hours.org/problem-profiles/preventing-catastrophic-pandemics/full-report/>.

¹⁹² Anjali Gopal, William Bradshaw, Vaishnav Sunil, Kevin M. Esvelt, “Securing Civilization against Catastrophic Pandemics” (Geneva Centre for Security Policy, 2023), *in press*.

¹⁹³ Current [president and CEO of the RAND Corporation](#), with extensive expertise and experience working on emerging catastrophic risks.

¹⁹⁴ Interview with Dr. Jason Matheny, 30 May 2023.

¹⁹⁵ Anonymized interview with an expert on biosecurity.

¹⁹⁶ Anonymized interview with an expert on biosecurity.

¹⁹⁷ “Conventional pandemic defence relies on medical countermeasures. For example, over two-thirds of funds requested by the American Pandemic Preparedness Plan (AP3) were to be allocated to biomedicine. Given that we still lack vaccines capable of protecting against natural viruses such as HIV, it is safe to assume that it will not be possible to block the effects of some pandemic-class agents with any form of medical countermeasure.” Kevin M Esvelt, “Delay, Detect, Defend: Preparing for a Future in Which Thousands Can Release New Pandemics,” Geneva Papers (Geneva Centre for Security Policy, 2022), 32, https://dam.gcsp.ch/files/doc/gcsp-geneva-paper-29-22?_gl=1*xm44p1*_ga*ODQxMDI0NiY4LjE2ODM2NTg5MTg.*_ga_Z66DSTVXTJ*MTY4Mzc2NTIzNi4yLjAuMTY4Mzc2NTIzNi4wLjAuMA..#page=49&zoom=100.0.0.



and was finally discontinued in 2023.¹⁹⁸ Our best guess for truly GCBR-relevant U.S. government funding is around \$500 million/year. As discussed below, this general pattern — a large pie with a small relevant slice — is repeated in the philanthropic sector. The next section asks: why?

Understanding GCBR Neglect and Public Policy Market Failures

If GCBRs are so important, why are they so neglected?¹⁹⁹ The previous section argued that there is a misallocation — a public policy market failure — of pandemic-preparedness and -response funding. This section briefly puts forth several hypotheses to help explain GCBR neglect.²⁰⁰ Jason Matheny²⁰¹ explained the overall problem in an interview for this report: “it’s hard to get policymakers focused on things that are pretty abstract, that are not immediate, that [...] might be lower probability than other things that they’re worried about, and other things that are going to drive electoral politics.”²⁰² In short, a combination of political and bureaucratic incentives, cognitive biases, and broader issues in public policy, in combination may help to explain this phenomenon. Although there are counterarguments to each, in aggregate, they paint a picture that can help to explain the neglect of catastrophic biological risks (and of global catastrophic risks in general).

To some extent, this neglect is not surprising; policymakers routinely fail to plan for low-probability high-consequence events. The September 11 Attacks, the Global Financial Crisis, COVID-19, the Russian invasion of Ukraine, and more all illustrate failures of imagination, intelligence, and planning. For example, “the threat of a pandemic spreading from Asia to the United States featured in intelligence reports throughout the last 20 years”, and Barack Obama even wrote an op-ed in the *New York Times* about the threat when he was still a Senator.²⁰³ Tabletop exercises and scenarios long highlighted the threat of novel coronaviruses, including the 2018 *Clade X* exercise and the late 2019 *Event 201*, both hosted by the Johns Hopkins Center for Health Security.²⁰⁴

¹⁹⁸ Kelsey Piper, “Why Virus Hunting May Not Be the Answer to Stopping the next Pandemic,” Vox, 2022, <https://www.vox.com/future-perfect/2022/5/7/22973296/virus-hunting-discovery-deep-vzn-global-virome-project>. David Willman, “The US Quietly Terminates a Controversial \$125m Wildlife Virus Hunting Programme amid Safety Fears,” *BMJ* 382 (September 7, 2023): p2002, <https://doi.org/10.1136/bmj.p2002>.

¹⁹⁹ One of the most comprehensive answers to this question can be found in chapter 2 (“Why so little is being done about the catastrophic risks”) Richard Posner’s *Catastrophe: Risk and Response*, 92-138. Much of the analysis in this section draws on the insights from that chapter.

²⁰⁰ The structure of this section is modeled in part on Founders Pledge’s analysis of “[Philanthropy to the Right of Boom](#).”

²⁰¹ Current [president and CEO of the RAND Corporation](#), with extensive expertise and experience working on emerging catastrophic risks.

²⁰² Interview with Dr. Jason Matheny, 30 May 2023.

²⁰³ “[T]he threat of a pandemic spreading from Asia to the United States featured in intelligence reports throughout the last 20 years. Then-Senator Barack Obama even penned an op-ed on pandemic threats in The New York Times in 2005.” Michael Horowitz et al., “Keeping Score: A New Approach to Geopolitical Forecasting,” *Perry World House*, accessed August 22, 2023, <https://global.upenn.edu/perryworldhouse/keeping-score-new-approach-geopolitical-forecasting>.

²⁰⁴ Mark Perry, “America’s Pandemic War Games Don’t End Well,” *Foreign Policy*, April 1, 2020, <https://foreignpolicy.com/2020/04/01/coronavirus-pandemic-war-games-simulation-dark-winter/>.



Leaders' near-term political incentives diverge from long-term global incentives. First, leaders may care about satisfying their constituents' demands only on relevant political timelines.²⁰⁵ In practice, these timelines are shorter than the relevant timelines for catastrophic biological risks, such that spending resources on pandemic preparedness means investing political capital whose benefits future leaders can reap.²⁰⁶ As discussed above, GCBRs have a relatively low probability of occurring (though they have a high expected cost — probability times consequence), such that leaders may feel they are unlikely to be held accountable for inaction during their terms of office or even lifetimes. On the flip side of this, not only are leaders today unlikely to be held accountable for the costs of their inaction, they are also unlikely to reap the political rewards of action for a threat that is difficult to understand (unless there is a dedicated risk-aware constituency).

Relatedly, leaders' moral circles may be limited to their constituencies; for a U.S. policymaker, there is not much practical difference between a pandemic that leads to the collapse of the United States and a pandemic that leads to the extinction of humanity.²⁰⁷

It is useful to visualize the divergence of incentives in a table:

	Pandemic that Kills Millions	Pandemic that Kills Billions
National Near-Term Costs for Policymakers	<ul style="list-style-type: none">• Political and social opprobrium• Possible loss of office (leaders are <i>sometimes</i> voted out)²⁰⁸	<ul style="list-style-type: none">• Political and social opprobrium• Loss of office (leaders are voted out or nation ceases to exist)
Global Long-Term Costs for Humanity	<ul style="list-style-type: none">• Horrible tragedy	<ul style="list-style-type: none">• Horrible tragedy (1,000 times worse)• Possible civilizational collapse• Possible irrecoverable long-term damage

²⁰⁵ This applies to both democratically-elected leaders, who must be re-elected, and non-democratic leaders, who still face the threat of popular uprisings and other political pressures.

²⁰⁶ "If the annual probability of a catastrophic collision with an asteroid is 1 in 65 million, it is exceedingly unlikely that such a collision will occur within current politicians' terms of office, or lifetimes, or their children's or grandchildren's or even great-grand-children's lifetimes; and after that who cares? [...] Politicians are unlikely to win points for preventing something that is highly unlikely to happen in the foreseeable future." Richard Posner, *Catastrophe: Risk and Response* (Oxford University Press, 2004), 118.

²⁰⁷ "Concern for GCBRs – driven in large part by cosmopolitan interest in the global population, concern for the long-run future, and where most of its beneficiaries are yet to exist – has obvious barriers to overcome." Gregory Lewis, "Reducing Global Catastrophic Biological Risks," 80,000 Hours, March 2020, <https://80000hours.org/problem-profiles/preventing-catastrophic-pandemics/full-report/>.

²⁰⁸ Notably, the experience of national leaders during the COVID-19 pandemic suggests that even disastrous responses are not necessarily punished at the polls. Thanks to Joshua Monrad for this point.



Policymakers may have a sense of duty (e.g. to their country, to the constitution, to their values and way of life) that lead them to value the continuity of government, and thus to place somewhat greater weight on a pandemic that kills billions. The rational response for many policymakers facing these political incentives, however, may be to not distinguish much between the two types of pandemics.

Cognitive biases — especially scope neglect — distort decision-makers’ ability to accurately size up the threat. In behavioral economics and decision theory, scope neglect refers to the observed tendency of humans to reason about a problem in a way that is insensitive to the scales involved in the problem (e.g., to not be 1,000 times more concerned about an event that causes 1 billion deaths than one that causes 1 million deaths).²⁰⁹ This phenomenon has been validated in various survey experiments.²¹⁰

The urgent crowds out the important in policymaking. Policymakers are busy people, often putting out fires, so that there is little time for long-term strategic thinking. With multiple things competing for decision-makers’ attention, the urgent will crowd out the important. In the words of President Dwight D. Eisenhower (who himself attributed the quote to a university president), “I have two kinds of problems, the urgent and the important. The urgent are not important, and the important are never urgent.”²¹¹ Preparedness for catastrophic pandemics may fall between the cracks if it is considered important, but not urgent.

Society is suffering from pandemic fatigue. As discussed in the introduction, several experts interviewed for this report pointed to the problem of “pandemic fatigue”; after more than three years of responding to a pandemic — and trillions of dollars spent on response and recovery — decision-makers simply don’t want to think about pandemics anymore. This was reportedly already a problem during the COVID-19 crisis itself, as people became tired of following health guidelines.²¹²

²⁰⁹ “We have trouble caring ten times more about something when it is ten times as important. And once the stakes get to a certain point, our concern can saturate. For example, we tend to treat nuclear war as an utter disaster, so we fail to distinguish nuclear wars between nations with a handful of nuclear weapons (in which millions could die) from a nuclear confrontation with thousands of nuclear weapons (in which a thousand times as many people would die, and our entire future may be destroyed).” Toby Ord, *The Precipice: Existential Risk and The Future of Humanity* (New York: Hachette Books, 2020), 61.

²¹⁰ Ben Bucknall, “Cognitive Biases and Existential Risk,” *Existential Risk Observatory*, July 20, 2021, <https://www.existentialriskobservatory.org/existential-risk/cognitive-biases-and-existential-risk/>.

²¹¹ Dwight D. Eisenhower, 19 August 1954 Address at the Second Assembly of the World Council of Churches, Evanston, Illinois, <https://www.presidency.ucsb.edu/documents/address-the-second-assembly-the-world-council-churches-evanston-illinois>.

²¹² E.g., the WHO reported, “Despite documented public support for pandemic response strategies across the WHO European Region, Member States are reporting signs of pandemic fatigue in their populations – here defined as demotivation to follow recommended protective behaviours, emerging gradually over time and affected by a number of emotions, experiences and perceptions.” WHO Regional Office for Europe, “Pandemic Fatigue: Reinvigorating the Public to Prevent COVID-19” (World Health Organization, 2020), <https://apps.who.int/iris/bitstream/handle/10665/335820/WHO-EURO-2020-1160-40906-55390-eng.pdf>.



This fatigue, moreover, may combine with the problems of hindsight bias and the availability heuristic, leading policymakers to focus on “fighting the last war” (e.g. by investing post-2001 in counter-anthrax capabilities disproportionately to the threat of Anthrax, a non-contagious agent).²¹³ Notably, pandemic fatigue has been widely reported in the media for both citizens and elite decision-makers, but there are few rigorous attempts to measure it and its effects, especially among elites.²¹⁴

There is a **general absence of risk-sensitive policymaking**. Rigorous forecasting, analysis, and decision-making practices are infamously rare among policymakers. Note that this is related to, but separate from, the problem of scientific illiteracy among policymakers.²¹⁵ For an overview of various issues with evidence-based policymaking, we recommend FP21’s report [*Less Art, More Science: Transforming U.S. Foreign Policy through Evidence, Integrity, and Innovation*](#).

The professional culture of the life sciences emphasizes openness over security. Finally, the life sciences value free and open sharing of information in a way that encourages collaboration and scientific discovery but may have downsides from a security perspective.²¹⁶ Unlike nuclear scientists, whose work was securitized during World War II and the Cold War, serious discussions of risk in the life sciences sometimes follow only *after* a controversial or dangerous event.²¹⁷ Pointing to the dangers of this approach is made more difficult by the obvious benefits of science and scientific progress.²¹⁸ Thus, scientists publish the results of dangerous experiments in widely-read journals, creating a culture of openness that may do more harm than good.²¹⁹

²¹³ On the surge in post-Amerithrax biodefense spending, see Carrie M. Long and Andrea Marzi, “Biodefense Research Two Decades Later: Worth the Investment?,” *The Lancet. Infectious Diseases* 21, no. 8 (August 2021): e222–33, [https://doi.org/10.1016/S1473-3099\(21\)00382-0](https://doi.org/10.1016/S1473-3099(21)00382-0).

²¹⁴ There are exceptions, like Hiu Tin Leung et al., “COVID-19 Pandemic Fatigue and Its Sociodemographic and Psycho-Behavioral Correlates: A Population-Based Cross-Sectional Study in Hong Kong,” *Scientific Reports* 12, no. 1 (September 27, 2022): 16114, <https://doi.org/10.1038/s41598-022-19692-6>. Notably, however, this study relies on subjective reporting of pandemic fatigue on a 1-10 scale, but does not shed much light on this self-reported measure’s relation to actual behavior (e.g., support for specific policies, compliance with restrictions, etc.)

²¹⁵ “One reason for widespread indifference to the catastrophic risks is the abysmal state of scientific knowledge among nonscientists. Scientific ignorance is a well-known factor in the public’s systematic misperceptions of the risks of different hazards, and while such ignorance could exacerbate fears of what science might do to us, it seems, rather, to allay them” Posner, *Catastrophe: Risk and Response*, 93.

²¹⁶ “Security concerns are less salient in the current culture of the life sciences, as evidenced by the number of projects explicitly intending to create, identify, and publicly share a list of viruses ranked by apparent threat level” Esvelt, “Delay, Detect, Defend,” 18.

²¹⁷ Kevin M. Esvelt, “Inoculating Science against Potential Pandemics and Information Hazards,” *PLoS Pathogens* 14, no. 10 (October 4, 2018): e1007286, <https://doi.org/10.1371/journal.ppat.1007286>.

²¹⁸ “Ignorance of science coexists dangerously with an uncritical veneration of science and scientists. The enormous success of science [...] has also tended to occlude the dangers that continued scientific progress poses and to create an attitude of ‘leave science policy to scientists.’” Posner, *Catastrophe: Risk and Response*, 97.

²¹⁹ Michael Selgelid and Weir, “The Mousepox Experience,” *EMBO Reports* 11, no. 1 (January 2010): 18–24, <https://doi.org/10.1038/embor.2009.270>.



Many **other factors** likely also affect the relative neglect of these obviously important issues, including the lack of probabilistic and quantitative reasoning and “keeping score” in policymaking communities and intelligence agencies, the scientific difficulties of the topic, bureaucratic infighting, and more.²²⁰

Philanthropic Neglectedness

There are very few grantmaking and philanthropic advising organizations explicitly focused on global catastrophic biological risks. One expert with experience both in and out of government stated in an interview for this report, “On one hand you could count the number of philanthropists that are funding work in this area of the most extreme bio risks.”²²¹ In addition to Founders Pledge, those organizations that do focus on catastrophic biological risks include:

- Open Philanthropy
- Effective Giving
- Entities associated with Effective Ventures²²²

There exist some major organizations focused on biological weapons, such as the G7 Global Partnership,²²³ but traditional large-scale funders in health and security do not generally support work on GCBRs.²²⁴ The Sloan Foundation used to have a biosecurity program that is no longer active, and traditional international security funders such as the Carnegie Corporation, MacArthur Foundation, and Skoll have made occasional biosecurity-related grants, but do not appear to be focused on worst-case scenarios.²²⁵ Similarly, the Gates Foundation appears to approach biosecurity mostly from a global health pandemic preparedness perspective, rather than focusing on the most catastrophic threats. The Smith Richardson Foundation supports the Bipartisan Commission on Biodefense.²²⁶ The Peter G. Peterson Foundation has made several grants in this area, such as

²²⁰ Perry World House, “Keeping Score: A New Approach to Geopolitical Forecasting,” University of Pennsylvania (2021), <https://global.upenn.edu/sites/default/files/perry-world-house/Keeping%20Score%20Forecasting%20White%20Paper.pdf>.

²²¹ Anonymized interview.

²²² E.g. Longview Philanthropy, EA Funds, etc.

²²³ Thanks to Joshua Monrad for pointing me to the Global Partnership.

²²⁴ Tom Inglesby, the Director of the Johns Hopkins Center for Health Security, explained: “Some of the big funders in the world, traditional foundations that you can think of around health are not focused on catastrophic risks. [...] There are comparatively few foundations working on these issues” (Interview with Tom Inglesby, 8 June 2023).

²²⁵ Open Philanthropy, “Biosecurity,” *Open Philanthropy* (blog), 2014, <https://www.openphilanthropy.org/research/biosecurity/>

²²⁶ Bipartisan Commission on Biodefense, “Sponsors / Donors,” *Bipartisan Commission on Biodefense* (blog), accessed August 11, 2023, <https://biodefensecommission.org/sponsors-donors/>.



various large grants through the Peter G. Peterson Foundation Pandemic Response Policy Research Fund.²²⁷ Joseph Rowntree Charitable Trust has also supported initiatives in this area.²²⁸

We have begun compiling biodefense, biosecurity, and biosafety funding in a spreadsheet [here](#). The spreadsheet is incomplete and — because it was the result of manual searches — likely biased. Moreover, because some relevant actors in this space — including Founders Pledge, Longview, and others — operate in whole or in part as *philanthropic advisory* services, grants made by their advised individual philanthropists may not be reflected in databases whose data mostly capture foundations, not individuals. Nonetheless, the data illustrate a central fact of philanthropic biosecurity funding: excluding direct response funding and HIV-related pandemic funding, total funding remains in the tens of millions, even assuming that the dataset reflects only a fraction.²²⁹ There are no other funders working on the scale of Open Philanthropy as of mid-2023.²³⁰ Note that the following visualization is likely severely **under-counting the total funding and is intended only as a first attempt** at estimating relevant biosecurity and pandemic preparedness funding:

²²⁷ Brown University, “2022 Peter G. Peterson Foundation Pandemic Response Policy Research Fund Recipients | Research at Brown,” accessed August 11, 2023, <https://www.brown.edu/research/conducting-research-brown/finding-funding/internal-funding-opportunities/2022-peter-g-peterson-foundation-pandemic-response-policy-research-fund-recipients>.

²²⁸ “Biological Security Research Centre at London Met awarded charity funds,” accessed September 12, 2023, <https://www.londonmet.ac.uk/news/articles/biological-security-research-centre/>.


²²⁹ NB.: this analysis was performed by a human-machine team (Christian Ruhl and GPT4).

²³⁰ Philanthropic entities associated with FTX, including the FTX Future Fund, were expected to become the largest source of biosecurity-related funding, before revelations about FTX led to the collapse of all associated philanthropic endeavors.



Philanthropic Biosecurity and Biodefense Funding

Error bars indicate under-counting by up to 50%. Excluding direct response (e.g. COVID relief), and excluding HIV-related funding



Note delayed funding spikes after Ebola, Zika, and after COVID-19.
Data compiled in linked spreadsheet from Peace and Security Funding Index and other grants databases.

Source: Created by Christian Ruhl and GPT4, using manually-compiled [dataset](#). See [appendix](#) for code.

This matters because philanthropists are one of the few actors who do not face the kinds of constraints above. Philanthropists do not have to answer to the myopic concerns of special interest groups and lobbyists; they do not have to worry about immediate returns on investment; they can take a more cosmopolitan and long-term approach to solving global problems.

Moreover, philanthropy can help set the agenda for policymakers. As described in Founders Pledge’s 2023 [report on nuclear risks](#), for example, several far-sighted funders focused on the problem of “loose nukes” in the disintegrating Soviet Union in the late 1980s and early 1990s, ultimately helping to set the stage for the multi-billion dollar Cooperative Threat Reduction Program.²³¹ Several experts we interviewed expressed concern about the long-term sustainability of the field without a more

²³¹ “Carnegie and MacArthur grants funded reports on the potential danger of so-called loose nukes that convinced the Senate of the urgency of the problem and of the need to spread knowledge about such dangers. Although Soskis [2013, in a report for GiveWell Labs (now Open Philanthropy)] notes that Senator Nunn may have pursued such a program without the influence of philanthropy, he argues that philanthropy at least catalyzed the process. This seems the likeliest scenario, as it is reflected in other literature about cooperative threat reduction. After all, it took very little to convince politicians that loose nukes were a threat to US national security — the dilemma was what to do about it.” Paul Robinson, “Philanthropy, Nuclear Nonproliferation, and Threat Reduction” (Urban Institute, 2021), https://www.urban.org/sites/default/files/2021/02/05/philanthropy_nuclear_nonproliferation_and_threat_reduction.pdf.



diverse and larger funding base.²³² As one senior policy expert interviewed for this project put the problem: **“without philanthropists funding work in this area, it's not going to get funded.”**²³³

Open Questions for Further Research

- How can analysts better estimate the amount of federal funding that is directly relevant to catastrophic biological risks?
- What biases do the few funders in this space share? How does this affect their grantmaking?
- Could political advocacy help create constituencies that care about extreme risks?
- What lessons can philanthropists learn from the response to COVID-19 for future policy windows?

²³² For example, Tom Inglesby, the Director of the Johns Hopkins Center for Health Security, explained in an interview for this project: “There are organizations that have capability in the field, to work outside of government independently, but there is no steady place for them to turn for resources, [...] and if there isn't any kind of clear signal from philanthropy, then [...] it will be hard for the field to keep going.” and “We really need to build a sustainable field of practice around this issue, as opposed to one project or one technology at a time. We need to create organizations that are sustained and [...] can attract and retain talent for time, most of these problems [have] multi-year solutions” (Interview with Tom Inglesby, 8 June 2023).

²³³ Anonymized interview.



Guiding Principles for Effective Philanthropy

Key Points

- Philanthropists can derive guiding principles (or “impact multipliers — see below) from the considerations above. These include:
 - Focusing on worst-case scenarios,
 - Funding interventions that are robust to the entire spectrum of risk,
 - Pursuing pathogen- and threat-agnostic approaches,
 - Using policy advocacy to leverage existing societal resources,
 - Prioritizing interventions with near-term positive externalities,
 - Avoiding various grantmaker dilemmas, including information hazards.
- Together, these guiding principles can help point towards concrete funding opportunities.

The previous sections made several observations about the landscape of global catastrophic biological risks and attempts to mitigate these risks. This section uses these features to derive several guiding principles, or evidence-based heuristics, for high-impact giving. We view these guiding principles as “**impact multipliers**,” independent features of the world and of possible interventions that make one funding opportunity relatively more cost-effective than another.²³⁴ These multipliers can be “stacked” to form a quasi-mathematical structure that *could*, with enough information, be quantified as the multiplication of different distributions of beliefs. Aggregate cost-effectiveness thereby becomes, in the words of one clear explanation of the concept, a “conjunction of multipliers.”²³⁵

²³⁴ “Our research into specific causes helps us understand which factors influence effectiveness within a cause area. The use of these heuristics enables us to rapidly distinguish the organizations and initiatives that are most likely to be effective from those that are probably less effective. To give a simple example, if some interventions within a given cause area are several times as effective as others, then we should prioritize charities that implement those interventions. If a problem is worse in some regions than in others, we should prioritize charities working in those geographies. A key advantage of impact multipliers is that they enable us to answer what are essentially yes-or-no questions about charities working in a given cause area. Are they working in Geography X? Are they implementing Intervention Y?” Matt Lerner, “How We Think about Charity Evaluation,” Founders Pledge, 2023, <https://www.founderspledge.com/research/how-we-think-about-charity-evaluation>.

²³⁵ Thomas Kwa, “Effectiveness Is a Conjunction of Multipliers,” Effective Altruism Forum, accessed August 11, 2023, <https://forum.effectivealtruism.org/posts/GzmJ2uiTx4gYhpcOK/effectiveness-is-a-conjunction-of-multipliers>.



In practice, this is only a first pass at creating possible impact multipliers for the issue area of extreme biological risks. Once we have established tentative multipliers, future work can attempt to use probabilistic methods to quantify these multipliers and thereby create a more rigorous and transparent framework for grantmaking.²³⁶ Although such work is still in the future, we believe that the current approach can help to account for much of the variance in cost-effectiveness of different biosecurity-related funding opportunities, allowing grantmakers to separate the “wheat” of funding opportunities that are highly-effective on priors from the “chaff” of other possible funding opportunities. Then, considerations of team strength, funding additionality, and more, can help narrow the search space further to rigorously guide effective grantmaking.

1. **Focus on worst-possible scenarios.** In practice, this means prioritizing engineered threats.
2. **Fund interventions that are robust to the entire spectrum of risk**, up to and including extinction-level pandemics.
 - a. As explained below, this is a slightly different consideration from the above; (1) narrows the search space, whereas (2) focuses on the attributes of funding opportunities within this search space.
3. Within anthropogenic events, generally **pursue pathogen- and threat-agnostic approaches**, and beware mere risk-shifting and “pandemic Maginot Lines.”
4. **Leverage existing societal resources** using advocacy-based interventions, with a few notable exceptions where government interests are simply not aligned with the long-term interests of humanity.
5. *All else equal*, **prioritize interventions with near-term positive externalities** that can garner sustained and broad political support.
6. **Prioritize offense-defense distinguishability** and **avoid information hazards** when selecting funding opportunities to avoid fuelling dangerous security dilemmas and doing more harm than good.

Importantly, **these impact multipliers are complementary to ideas of “defense in depth”** or multi-layered defense.²³⁷ For example, we view Dr. Kevin Esvelt’s proposal of [Delay, Detect, Defend](#) as one of the most elegant versions of a defense-in-depth approach that simultaneously satisfies many

²³⁶ Johannes Ackva and Megan Phelan, “How to Evaluate Relative Impact in High-Uncertainty Contexts? An Update on Research Methodology & Grantmaking of FP Climate,” *Effective Altruism Forum*, 2023, <https://forum.effectivealtruism.org/posts/kuopGotdCWeNCDpWi/how-to-evaluate-relative-impact-in-high-uncertainty-contexts>.

²³⁷ Owen Cotton-Barratt, Max Daniel, and Anders Sandberg, “Defence in Depth Against Human Extinction: Prevention, Response, Resilience, and Why They All Matter,” *Global Policy* 11, no. 3 (2020): 271–82, <https://doi.org/10.1111/1758-5899.12786>.



of the impact multipliers discussed below. Thus, although some of the specific examples cited below appear focused on civil defense measures or “armageddon insurance,” the general principles can be applied to all layers of civilization’s defense against deadly pathogens.²³⁸

Prioritize Worst-Case Scenarios

The structure of the risk suggests that impact-minded philanthropists ought to focus on worst-case scenarios. This is because **the risk is superlinear** (discussed under “[Analyzing the Structure of the Problem](#)”) and **the worst-case scenarios are simultaneously neglected** by traditional philanthropies and government funders (discussed under “[Incentives, Preparedness, and Neglectedness](#)”).

To briefly restate the arguments the risk is super-linear because:

1. The distribution of the threat’s expected costs suggests that most of the dis-value lies with events in the “heavy tails” of these distributions.
2. There exist moral discontinuities for anyone who places *some* value on the long-term future, such that the most extreme events — civilizational collapse and extinction — are far worse than the aggregate of less extreme events.
 - a. As seen above (“[Existential Threats](#)”), the production of such extreme events is possible, and the technology and knowledge for it appears to be in the reach of many people in the near future.
3. We expect threshold effects — including system collapse and cascading failures in civilization — to make adequate societal response more difficult, and to thus make the worst events disproportionately worse.

These worst-case pandemics, the most extremely super-linear slice of the risk, are neglected by traditional societal actors, both public and private.

This may appear like a suspiciously convenient coincidence: the very slice of risk that is the most important is supposed to simultaneously be among the most neglected? Interestingly, however, we can find similar patterns in other cause areas, including [climate change](#) and [nuclear war](#). In nuclear war, for example, the problem of managing escalation *after the first bomb has gone off* is among the most important problems in the entire field; it is the difference between a small nuclear war with horrific but limited consequences, and a civilization-wrecking thermonuclear exchange between the great powers. At the same time, partly for ideological reasons, such “right of boom” interventions are unpalatable to many funders.²³⁹

²³⁸ On the concept of civil defense, see Edward Geist, *Armageddon Insurance: Civil Defense in the United States and Soviet Union, 1945-1991*.

²³⁹ Christian Ruhl, “Philanthropy to the Right of Boom,” Founders Pledge, 2023, <https://www.founderspledge.com/research/philanthropy-to-the-right-of-boom>.



As discussed above, this convergence of importance and neglectedness is less surprising than it appears at first glance, because several factors drive this public policy market failure:²⁴⁰

- Leaders' near-term political incentives diverge from long-term global incentives.
- Cognitive biases — especially scope neglect — distort decision-makers' ability to accurately size up the threat.
- The urgent crowds out the important in policymaking, such that we often end up “fighting the last war.”
- Society is suffering from pandemic fatigue.
- There is a general problem of evidence-based policymaking.
- The professional culture of the life sciences lacks a “security mindset.”

Perhaps the simplest explanation, which runs through the others, is that much of society does not share some ethical intuitions that drive the prioritization of worst-case scenarios (e.g., consequentialism²⁴¹).

In practice, this impact multiplier has simple but profound implications for grantmakers. Because we expect that engineered pathogens are far more likely to cause the most catastrophic (and possibly extinction-level) pandemics than naturally-arising pathogens, grantmakers ought to **prioritize anthropogenic threats**. More specifically, rather than focus on locally-contained anthropogenic threats like anthrax spores, impact-oriented philanthropists ought to focus instead on understanding the threat models of pathogens that could plausibly do the most damage to civilization and **prioritize mitigating credible threats that combine extreme transmissibility with extreme virulence** (specifically, pathogens that combine these traits in ways that facilitate global spread, like a long infectious incubation period).

²⁴⁰ The following bullet points are direct quotations from the previous section on “[Incentives, Preparedness, and Neglectedness](#).”

²⁴¹ “Consequentialism, as its name suggests, is simply the view that normative properties depend only on consequences. This historically important and still popular theory embodies the basic intuition that what is best or right is whatever makes the world best in the future, because we cannot change the past, so worrying about the past is no more useful than crying over spilled milk. This general approach can be applied at different levels to different normative properties of different kinds of things, but the most prominent example is probably consequentialism about the moral rightness of acts, which holds that whether an act is morally right depends only on the consequences of that act or of something related to that act, such as the motive behind the act or a general rule requiring acts of the same kind.” Walter Sinnott-Armstrong, “Consequentialism,” in *The Stanford Encyclopedia of Philosophy*, ed. Edward N. Zalta and Uri Nodelman, Winter 2022 (Metaphysics Research Lab, Stanford University, 2022), <https://plato.stanford.edu/archives/win2022/entries/consequentialism/>.



Pursue Robust Interventions

Not all possible funding interventions are robust to the entire spectrum of risk, and philanthropists ought to prioritize funding interventions that have a lower probability of “breaking” under the worst-possible scenarios. Whenever funding something, a philanthropist ought to ask themselves: would this work under the plausible conditions of a catastrophic pandemic? For example, if a risk-mitigation measure relies on the existence of 1,000,000 essential workers (e.g. to distribute the countermeasure), but only 100,000 essential workers are available (e.g., because there is a shortage of credible pandemic-proof personal protective equipment²⁴²), then the risk-mitigation measure is *not* robust to worst-case scenarios. Instead, the philanthropist may consider a new order of funding priorities:

1. Fund policy advocacy to develop P4E and a distribution system for essential workers;
2. *Then*, fund other risk-mitigation measures that rely on the availability of essential workers.

Funding (2) without funding (1) is *not* robust to worst-case scenarios. This principle is similar in structure to the idea of “**robust diversification**” in Founders Pledge’s climate philanthropy; we ought to **hedge against the scenarios in which traditional risk-mitigation measures have already failed**.²⁴³

Prevention and Response

In many cases funding the *prevention* of biological catastrophes will be more robust than funding possible *responses* to such catastrophes.²⁴⁴ A multi-layered defense is critical, as discussed throughout this report. Nonetheless, the responses to a truly catastrophic event might differ significantly from the responses to smaller scale pandemics, potentially requiring completely different — and very draconian — measures. *Prevention*, on the other hand, may be more cost effective (and more politically tractable) because it can look similar for both the most extreme and more common biological events. For instance, regulations to constrain terrorist access to powerful biotechnologies would be useful for both small-scale bioterrorism and extreme omnicidal or near-omnicidal groups.

²⁴² On P4E, see Kevin M Esvelt, “Delay, Detect, Defend: Preparing for a Future in Which Thousands Can Release New Pandemics,” Geneva Papers (Geneva Centre for Security Policy, 2022), https://dam.gcsp.ch/files/doc/gcsp-geneva-paper-29-22?_gl=1*xm44p1*_ga*ODOxMDI0NiY4LjE2ODM2NTg5MTg.*_ga_Z66DSTVXTJ*MTY4MzczNTIzNi4yLjAuMTY4MzczNTIzNi4wLjAuMA..#page=49&zoom=100.0.0.

²⁴³ Ackva, “Guide to the Changing Landscape of Climate Philanthropy,” Founders Pledge, 2022, <https://www.founderspledge.com/research/changing-landscape>. See also Robert J. Lempert, Steven W. Popper, and Steven C. Bankes, “Shaping the Next One Hundred Years: New Methods for Quantitative, Long-Term Policy Analysis” (RAND Corporation, January 1, 2003), https://www.rand.org/pubs/monograph_reports/MR1626.html.

²⁴⁴ Thanks to Andrew Snyder-Beattie for the insights in this paragraph in a round of external reviews.



Prioritize Pathogen- and Threat-Agnostic Approaches

A related point to the above is that philanthropists ought to generally **pursue pathogen- and threat-agnostic approaches**, and beware of mere risk-shifting and “pandemic Maginot Lines.” This impact multiplier arises from two considerations:

1. The (increasing) **complexity of the threat**, and the difficulty of knowing where the next pandemic will come from and what it will look like (see [“The Threat Is Increasingly Complex”](#)) partly due to a threat landscape that is shifting beneath our feet in difficult-to-predict ways (see [“The Changing Threat Landscape”](#)).
2. The **adaptiveness of the threat**, and the ability of intelligent adversaries (and to a lesser extent, natural selection), to respond to risk-mitigation measures (see [“The Threat Is Adaptive”](#)).

Importantly, these two features of the risk do *not* mean that philanthropists are clueless about what to do because they know *nothing* about the character of the risk. Indeed, as we have seen, there are many stylized conclusions (like prioritizing anthropogenic threats) that can help guide philanthropy amidst this uncertainty. Rather, complexity and adaptiveness of the threat suggest that philanthropists ought to pursue threat- and pathogen agnostic interventions whenever possible, for two reasons:

1. Increasing the likelihood of covering the correct threat.
2. Decreasing the likelihood of “risk shifting.”

The first of these two considerations is obvious; if we are able to protect against a set of threats (catastrophic pandemics), we ought to do so rather than attempting to protect only against a subset (e.g. laboratory escapes) or individual members of the set (e.g. specific pathogens). The wider the net, the more likely we are to catch the correct threat.

The second of these two considerations — the problem of “risk shifting” — deserves greater attention.²⁴⁵ Defending against biological threats is difficult enough given the problem of mutation and natural selection, but becomes far more difficult when the adversary is not natural selection but an intelligent malevolent actor. Intelligent actors have the ability to observe, anticipate, and respond to society’s defenses. For example, if pandemic preparedness focuses entirely on one class of agents — let’s say viruses identified as being concerning for natural spillover — then malevolent actors know that other agents may be the better investment for their time. This is analogous to the Maginot Line in World War II.²⁴⁶

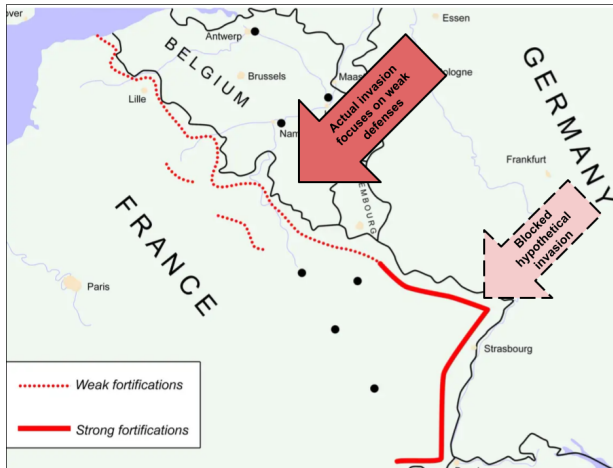
²⁴⁵ With thanks to Matthew Gentzel for his discussions of risk-shifting problems in nuclear philanthropy.

²⁴⁶ Robert Zaretsky, “The Strange Defeat of the United States,” *Foreign Affairs*, July 20, 2020, <https://www.foreignaffairs.com/articles/united-states/2020-07-20/strange-defeat-united-states>.



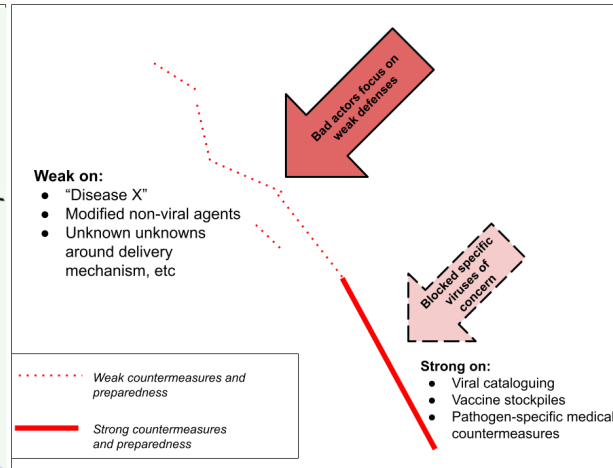
The Maginot Line, 1940

France fortifies part of its border with expensive defenses; Germany invades through the Low Countries.



The Pandemic Maginot Line, 2023-?

Society focuses on some viral threats with expensive defenses; bad actors use other biological agents and delivery mechanisms to maximize damage.



Source: Left [Wikimedia commons map](#), with annotations by author; right-hand side diagram by author.

One policymaker interviewed for this report used this analogy to explain that “most [traditional biodefense countermeasure work] is sort of a Maginot Line, when it comes to the kinds of biological events that I would worry the most about. It's just not going to present an obstacle to those events.”

²⁴⁷

This makes thinking about threat vectors especially complicated. As discussed by other analysts, this implies that we ought to “Be wary of narrow countermeasures... Look for broad spectrum countermeasures or responses — which are more likely to eliminate big chunks of the risk landscape and provide overall reduced bioweapons appeal”²⁴⁸ Moreover, traditional pandemic preparedness approaches have been list-based, but the use of lists is itself likely to shift risk away from the listed agents.²⁴⁹

There is a structure to this problem of risk-shifting that makes it more insidious than the Maginot Line. The Maginot Line resulted in two-dimensional risk-shifting on a map, and did indeed succeed in

²⁴⁷ Anonymized interview.

²⁴⁸ Andrew Snyder-Beattie, “Countermeasures & Substitution Effects in Biosecurity,” *Effective Altruism Forum*, 2021, <https://forum.effectivealtruism.org/posts/F4KS4zCFm3p6NtrR/countermeasures-and-substitution-effects-in-biosecurity>

²⁴⁹ “A substantial proportion of infectious disease preparedness activities have, to date, focused on a historical list-based approach constructed around biological warfare agents and on recent outbreaks (eg, SARS).” The Characteristics of Pandemic Pathogens” (Johns Hopkins Center for Health Security, 2018), <https://centerforhealthsecurity.org/sites/default/files/2022-12/180510-pandemic-pathogens-report.pdf>), 8.



directing the movement of the adversary in a way that could have given France more time. But the broader problem of risk-shifting is multi-dimensional, such that the wrong kinds of defenses can systematically push bad actors (e.g. non-omnicidal bioweaponers) to pursue ever more-dangerous paths (e.g. ever more-dangerous bioweapons). Thus, this problem intersects with the importance of focusing on defending against worst-case scenarios. If humanity focuses on the low-hanging fruit of smaller pandemics first, “finding countermeasures to more mundane things will cause adversaries to move towards GCBR territory (either by more heavily engineering mundane things, or switching into entirely new kinds of attack).”²⁵⁰

Risk-general mitigation, on the other hand, can help to bolster **deterrence by denial**. In international relations and military strategy, this idea refers to the possibility that defenses against the threat will render any attack so futile that the costs of the attack outweigh the potential benefits. Consider, for example, the risk-general detection and defense system advocated by Esvelt in *Delay, Detect, Defend*. In this imagined future, high indoor air quality — facilitated by safe germicidal lighting, among other measures — combined with reliable detection technologies may make any attack seem likely to fail. This perception, in turn, may convince malevolent actors that the very pursuit of biological weapons is too costly when weighed against the low probability of a successful attack. Such actors, therefore, may therefore forswear the very pursuit of bioweapons programs in the first place.

Thus even rogue actors who may not be responsive to deterrence by punishment can respond to deterrence by denial.²⁵¹ An illustrative example is the history of the Aum Shinrikyo cult in Japan. Leadership of Aum Shinrikyo originally pursued biological weapons, but soon determined that success was unlikely — they lacked the technical ability to create usable agents — and thus turned to chemical weapons instead.²⁵² Chemical weapons, although horrific, are less likely to spread exponentially and cause the kinds of damage that biological weapons could. Thus the risk-shifting behavior of Aum Shinrikyo may have decreased aggregate risk to humanity. This was because of

²⁵⁰ Andrew Snyder-Beatte, “Countermeasures & Substitution Effects in Biosecurity,” *Effective Altruism Forum*, 2021, <https://forum.effectivealtruism.org/posts/F4KS4zCFfm3p6NtrR/countermeasures-and-substitution-effects-in-biosecurity>

²⁵¹ Esvelt makes a similar point, writing, “Zealots cannot be deterred except by credibly advertising that all possible methods of attack accessible to them will fail. Until the world develops reliable methods of detecting and containing pandemic-class events, intelligence agencies should closely monitor extremists with potential omnicidal tendencies and their connections to individuals with the technical skills to assemble pandemic-class agents.” Kevin M Esvelt, “Delay, Detect, Defend: Preparing for a Future in Which Thousands Can Release New Pandemics,” *Geneva Papers* (Geneva Centre for Security Policy, 2022), 28, [https://dam.gcsp.ch/files/doc/gcsp-geneva-paper-29-22?_gl=1*xm44p1*_ga*ODOxMDI0NjY4LjE2ODM2NTg5MTg.*_ga_Z66DSTVXTJ*MTY4MzczNTIzNi4yLjAuMTY4MzczNTIzNi4wLjAuMA..#page=49&zoom=100.0.0.0.\)](https://dam.gcsp.ch/files/doc/gcsp-geneva-paper-29-22?_gl=1*xm44p1*_ga*ODOxMDI0NjY4LjE2ODM2NTg5MTg.*_ga_Z66DSTVXTJ*MTY4MzczNTIzNi4yLjAuMTY4MzczNTIzNi4wLjAuMA..#page=49&zoom=100.0.0.0.))

²⁵² “This unbroken string of failures with botulinum and anthrax eventually convinced the group that making biological weapons was more difficult than Endo was acknowledging. Asahara speculated that American comments on the risk of biological weapons were intended to delude would-be terrorists into pursuing this path.” For more on the case of Aum Shinrikyo, see Richard Danzig *et al.*, “Aum Shinrikyo: Insights Into How Terrorists Develop Biological and Chemical Weapons,” Center for a New American Security (2011), <https://www.cnas.org/publications/reports/aum-shinrikyo-insights-into-how-terrorists-develop-biological-and-chemical-weapons>.



technical barriers to *weapons development* that, as we have seen, may break down in the near future. But humanity can, with the right risk-general interventions, erect new barriers that lead to the same conclusion: biological weapons should not be pursued.

Importantly, these considerations — like the others discussed in this report — only apply *all else equal*. If a more specific risk-mitigation strategy is vastly cheaper than a more general strategy, such that the difference in cost outweighs the benefits of threat-agnosticity, then philanthropists ought to consider funding the more specific risk-mitigation strategy. Moreover, threat-agnosticity can be taken to absurd extremes. For example, it would be *technically* more threat-agnostic to focus not just on carbon-based life forms because sophisticated malevolent actors in the future could engineer more exotic threats, and to pursue a research agenda for defense against *all* possible life forms. The feasibility of such work, combined with the apparent low probability of encountering such threats, however, weigh strongly against such an approach. Nonetheless, the threat of xenobiology²⁵³ emerging from new advances in synthetic biology is real; if there are defenses (very good PPE) that could defend against such strange life forms, philanthropists ought to pursue them.

Tractability and Medical Countermeasures

In practice, a risk-general approach will often (though not always) lead philanthropists to prioritize non-medical countermeasures over medical countermeasures. First, humanity's track record with developing medical countermeasures is mixed. As Dr. Esvelt has written, "Given that we still lack vaccines capable of protecting against natural viruses such as HIV, it is safe to assume that it will not be possible to block the effects of some pandemic-class agents with any form of medical countermeasure."²⁵⁴ The history of humanity's fight against HIV/AIDS is informative; four decades of vaccine efforts have been largely fruitless, but antivirals have been highly successful.²⁵⁵ If we cannot find a vaccine against HIV even with massive resources and many years of intensive efforts, we may be unlikely to find a vaccine for a weaponized pathogen *specifically engineered to evade the human immune system* — let alone to administer such a vaccine quickly enough at scale.

Many agree with this assessment. As one policy expert explained in an interview for this report, "there has been a recognition that a lot of the traditional medical countermeasure work is probably not as relevant here [for worst-case bio events]" in part because "there's just no way that you're going to have clinical trials, you know, for, like vaccines, or antivirals against some novel agent."²⁵⁶ For truly catastrophic and extinction-level threat models, "it's unlikely that our kind of traditional approaches to medical countermeasures are going to be useful."²⁵⁷ We can see this problem even

²⁵³ "Xenobiology" in this context refers to "strange" life forms, not to alien life in the sense of extraterrestrial.

²⁵⁴ Kevin M Esvelt, "Delay, Detect, Defend: Preparing for a Future in Which Thousands Can Release New Pandemics," Geneva Papers (Geneva Centre for Security Policy, 2022),

²⁵⁵ Thanks to Andrew Snyder-Beattie for this and the following point.

²⁵⁶ Anonymized interview.

²⁵⁷ Anonymized interview.



with ambitious plans to accelerate the distribution of medical countermeasures. As Esvelt has written:

“Worse, even though an effective nucleic acid vaccine can now be designed within 24 hours of sequencing the genome of a pandemic-class agent, the testing, approval, manufacturing, and especially distribution process will not be able to match the speed of viral spread. The AP3 and G7 plans call for a vaccine within 100 days of sequencing the genome of an emerging pandemic virus, but 100 days after the omicron variant was sequenced in South Africa on 11 November 2021, it had infected a quarter of the United States and as much as half of Europe. Pandemic-class agents deliberately released in multiple airports would spread considerably more rapidly.”²⁵⁸

In addition to the problem of clinical trials and distribution, there are technical barriers to developing threat-agnostic medical countermeasures. As Dr. Amesh Adalja has explained, “We have nothing like a broad-spectrum antiviral agent. This is technically very hard to do, because, again, viruses are just a piece of [...] genetic material with protein around them. There’s not much that you can attack them with that doesn’t also attack your own cells.”²⁵⁹

This approach also makes sense because, as discussed above, much government spending currently focuses exclusively on medical countermeasures.²⁶⁰ An increased focus on non-medical countermeasures therefore may help create a more balanced portfolio.

For these reasons, many analysts of catastrophic biological risks have concluded that “Reliably preventing harm from pandemic-class agents requires looking beyond biomedicine.”²⁶¹ And Millett and Snyder-Beattie write:

“Countering existential risks could also result in reprioritizing current approaches — for example, favoring broadspectrum diagnostics and countermeasures, as opposed to those tailored to a single pathogen. The worst possible attacks could come from built-up arsenals of multiple pathogens, possibly designed with long incubation periods and traits to overcome vaccination or medical treatment.”²⁶²

²⁵⁸ Anjali Gopal, William Bradshaw, Vaishnav Sunil, Kevin M. Esvelt, “Securing Civilization against Catastrophic Pandemics” (Geneva Centre for Security Policy, 2023), *in press*.

²⁵⁹ Amesh Adalja, “The Characteristics of Pandemic Pathogens (OUPM Talk)” (Oxford University Personalised Medicine Society, 2020), <https://cpm.well.ox.ac.uk/video/dr-amesh-adalja-characteristics-pandemic-pathogens>.

²⁶⁰ “Conventional pandemic defence relies on medical countermeasures. For example, over two-thirds of funds requested by the American Pandemic Preparedness Plan (AP3) were to be allocated to biomedicine. Given that we still lack vaccines capable of protecting against natural viruses such as HIV, it is safe to assume that it will not be possible to block the effects of some pandemic-class agents with any form of medical countermeasure.” (Esvelt, *Delay, Detect, Defend*)

²⁶¹ Esvelt, *Delay, Detect, Defend*, 32.

²⁶² Piers Millett and Andrew Snyder-Beattie, “Existential Risk and Cost-Effective Biosecurity,” *Health Security* 15, no. 4 (August 2017): 381, <https://doi.org/10.1089/hs.2017.0028>.



This is not to say that work *related to* medical countermeasures will not be useful. Rather, philanthropists ought to focus any medical countermeasure work on *platform technologies*. Millett and Snyder-Beattie again:

“Platform technologies that allow customizable countermeasures (eg, phages for bacteria, generalized vaccine templates) or pathogen-blind diagnostics (eg, distributed sequencing and improved software to interpret novel pathogens before symptoms occur) will stand a better chance against such threats”²⁶³

Similarly, the Bipartisan Commission for Biodefense wrote in the *Apollo Program for Biodefense* that “a facility designed to manufacture a therapeutic or vaccine candidate using a platform technology against one pathogen could be quickly repurposed against a new pathogen without much need to make changes to physical infrastructure or established production processes,” and therefore recommended greater investment in such platform technologies. The ability to rapidly produce effective small-molecule drugs, in particular, may be useful during an outbreak.²⁶⁴ The conclusion for philanthropists is that pathogen-general *platform technologies* could be useful versions of medical countermeasures, but that medical countermeasures generally may be less useful for many of the most extreme outbreaks.

Leverage Existing Societal Resources

Other Founders Pledge reports discuss the importance of leveraging existing societal resources at length.²⁶⁵ The basic intuition is simple:

- **Society has vast resources**, especially in governments, but also in other philanthropic organizations.
- Work to **leverage these resources using advocacy** can therefore multiply the impact

²⁶³ Piers Millett and Andrew Snyder-Beattie, “Existential Risk and Cost-Effective Biosecurity,” *Health Security* 15, no. 4 (August 2017): 381, <https://doi.org/10.1089/hs.2017.0028>.

²⁶⁴ Thanks to Andrew Snyder-Beattie for this point. Bill Gates explains the reason for focusing on small-molecule drugs in his *How to Prevent the Next Pandemic*: “Small-molecule drugs have several advantages that make them especially appealing in an outbreak. Because their chemical structure is fairly straightforward, they’re easy to manufacture, and thanks to their size and chemistry, they don’t get broken down by your digestive system, so you can take them as a pill. (This is why you’ve never had to get an injection of aspirin.) And most of them can be kept at room temperature and have a long shelf life. Larger molecules are more complicated in just about every respect. A monoclonal antibody, for example, is 100,000 times larger than a molecule of aspirin. Because large molecules are broken down by your digestive system if you swallow them, they need to be injected or given by an IV drip. This means you’ll need medical personnel and equipment to make sure they are administered properly, and you’ll need to isolate infected patients when they come in for treatment so they don’t pass the virus to other people at the facility. Large molecules also require far more complex manufacturing — they’re made using live cells — which means they’re more expensive, and it takes more time to produce them at high volume. In short, during an outbreak, you’d rather have small-molecule treatments than large ones, all other things being equal.” (Bill Gates, *How to Prevent the Next Pandemic*, Standard Edition (New York: Knopf, 2022)).

²⁶⁵ See, e.g. reports on [nuclear war](#) and [climate change](#).



As we have seen, this holds especially true in the broader field of health security, with billions upon billions of dollars spent (see “[Societal Resources Spent on Health Security](#)”). Policy advocacy to improve the allocation of this vast pie of resources can be far more effective than trying to grow the pie itself.²⁶⁶

There are some **possible exceptions** to the policy advocacy impact multiplier, when private action appears to be the better path. These exceptions include **interventions that may spark security dilemmas** and **interventions that simply fall outside the realm of possible actions for governments**.

First, as discussed below (“[Beware Grantmaker Dilemmas](#)”), there exists a problem of security dilemmas — when one state’s defensive actions may appear offensive to another state, sparking a series of actions that decreases the overall security of the system. In biosecurity, this may mean that one state’s biodefense starts looking awfully similar to offensive work. In some cases, such work really could be beneficial, but adversary states’ responses would make it net-negative. In these cases, it could be possible for philanthropists who are viewed as separate from states to pursue these interventions privately. They could be labeled as eccentric billionaires, not threatening states. The extent to which this scenario obtains in the real world is unclear, however. Historically, for example, Soviet intelligence services viewed unrelated groups within societies in the free world as part of sinister capitalist plots, imagining, for example, that governments would collude with bankers and clergy in the run-up to a nuclear attack. It is not difficult to see how adversaries could draw similar links between wealthy philanthropists and the U.S. defense establishment. After all, the importance of the high-tech industrial base of the United States *is* emphasized in U.S. strategic documents, and some philanthropists *do* exercise great influence over the U.S. defense establishment (e.g. via “Commissions” on specific defense-relevant topics).

The second situation in which the importance of leveraging societal resources via policy advocacy breaks down is for useful interventions that fall outside the Overton Window of the current political environment. One example of this would be “bio-bunkers” — resilient shelters that help a small core of humanity survive in the event of a biological catastrophe. Such interventions — though they may have extremely high expected values on some moral frameworks — may simply be so strange that they are again best pursued by private actors. Here too, however, it is unclear to what extent these concerns actually apply. Bio-bunkers, for example, could be an important part of improving continuity-of-government plans, such that they do fit into “normal” policy frameworks. A better example of an intervention that truly falls outside the biosecurity Overton Window may be space colonization, although the cost-effectiveness of such risk-mitigation is dubious.

²⁶⁶ Thanks to my colleague Johannes Ackva for the pie metaphor. (E.g. on [this podcast](#): “One is policy advocacy in this political space, to differentiate between increasing the pie — which is something that grassroots organizations are focused on, essentially increasing the salience of climate, increasing the resource allocation towards climate — and then improving the allocation of the pie, which is about essentially making sure that the resources we’re allocating to climate are used in a way that’s actually useful for global decarbonisation.”)



Leverage Near-Term Policy Incentives

The preceding discussions on policy feasibility lead to another potential impact multiplier: *all else equal*, **leverage near-term policy incentives and prioritize interventions with near-term positive externalities** that can garner sustained and broad political support. This multiplier emerges from the discussion of public policy market failures above (see “[Understanding GCBR Neglect and Public Policy Market Failures](#)”). Rather than fanatically try to change the many and persistent problems in getting policymakers to align their actions with what is best for humanity in the long term, impact-minded philanthropists can be pragmatic and recognize that **changing entrenched political incentives and cognitive biases is intractable**. Doing so would be akin to a revolution of our entire society, and may have serious negative externalities.

Rather than advocate for radical change in the way policymakers think about their job description, therefore, the more tractable path may be to generally prioritize interventions that have two features:

1. Maximizing expected impact by focusing on worst-case scenarios;
2. Aligning with policymaker incentives by having positive near-term externalities.

This should not be done deceptively, but simply as a policy “win-win.” If an action both protects against extreme pandemics, but also strengthens the U.S. industrial base and has routine economic benefits by protecting against seasonal diseases, then it may be far more likely to garner bipartisan political support than if it *just* protects against extreme pandemics. Kevin Esvelt has put this win-win proposition in more general:

“Crucially, any passive defence capable of substantially impeding the spread of a novel pandemic agent would also suppress or outright eliminate many or even most endemic human viruses and pathogenic bacteria. Since economic losses from common infectious diseases in the United States reached an estimated ~\$300 billion in 2017, or about \$2,000 per worker, employers will be strongly incentivized to install any defences capable of reducing these losses in a cost-effective manner as soon as they have been developed.”²⁶⁷

Indeed, even post-2001 Amerithrax-related surges in biodefense funding appear to have such beneficial spillover effects, including advances like “the ability to identify drivers of measles virus epidemic seasonality,” and even the development of technologies that enabled the deployment of COVID mRNA vaccines (whether these funds were ultimately cost-effective is a separate question).

²⁶⁷ Esvelt, *Delay, Detect, Defend*, 38.



²⁶⁸ Indeed, it may be the case that securitizing the threat of infectious disease is an important way of using defense-earmarked funds for projects that will benefit the global poor, who suffer most from emerging infectious diseases, and who are otherwise ignored in policy discussions in the Global North.²⁶⁹ Of course, this win-win situation is not always possible. **When distinguishing between two otherwise-similar interventions, however, philanthropists should expect that the intervention with more positive near-term externalities is both better in expectation, and more tractable.**

Avoid High-Downside Grantmaker Dilemmas

Philanthropy is not usually considered a dangerous business, but the complexity of the biological risk landscape suggests an additional impact multiplier: **beware grantmaker dilemmas**. Not doing so could turn a net-positive action into a net-negative action, making the world far worse-off than before.

Several grantmaker dilemmas have already been discussed above. For example, **risk-shifting** behavior could make the impact of an intervention net-neutral or even negative. Similarly, well-intentioned science — like wildlife virus discovery and working with enhanced potential pandemic pathogens (ePPP) — can increase the surface area of risk, while eliminating only small slices of the previous risk.²⁷⁰ Relatedly, the development of **dual-use technologies and knowledge** can do more harm than good, and grantmakers ought to continuously ask themselves how the products of their grants — physical technologies as well as knowledge — could be misused.

The following brief sections highlight two additional grantmaker dilemmas: **prioritizing offense-defense distinguishability** and **minimizing information hazards** when selecting funding opportunities to avoid fuelling dangerous security dilemmas and doing more harm than good.

Offense-Defense Distinguishability and the Security Dilemma

In international relations, a security dilemma refers to a situation where state A's defensive actions to increase its security create feelings of insecurity in a rival state B, which in turn reacts by preparing

²⁶⁸ Carrie M. Long and Andrea Marzi, “Biodefense Research Two Decades Later: Worth the Investment?,” *The Lancet. Infectious Diseases* 21, no. 8 (August 2021): e222–33, [https://doi.org/10.1016/S1473-3099\(21\)00382-0](https://doi.org/10.1016/S1473-3099(21)00382-0). The authors argue that “The benefits of biodefense research are wide reaching and help improve global biosecurity for infectious disease outbreaks of all kinds.”

²⁶⁹ “Perhaps one of the strongest arguments in support of biodefense research is the multifaceted impact it has and can have on general infectious disease control. As demonstrated by the recent EBOV epidemic and COVID-19 pandemics, emerging and re-emerging infectious diseases are largely ignored outside of the biodefense funding arena” Carrie M. Long and Andrea Marzi, “Biodefense Research Two Decades Later: Worth the Investment?,” *The Lancet. Infectious Diseases* 21, no. 8 (August 2021): e222–33, [https://doi.org/10.1016/S1473-3099\(21\)00382-0](https://doi.org/10.1016/S1473-3099(21)00382-0).

²⁷⁰ For a discussion of the security considerations surrounding wildlife virus discovery, see Jonas Sandbrink et al., “Mitigating Biosecurity Challenges of Wildlife Virus Discovery,” SSRN Scholarly Paper (February 14, 2022), <https://doi.org/10.2139/ssrn.4035760>.



for the worst interpretation of state A's actions, ultimately undermining everyone's security and sometimes sparking arms races.²⁷¹

In biosecurity, an insecure state might view many kinds of biodefense actions not as purely defensive, but as potentially complementary to offensive biological activities.²⁷² If the United States were to create super-PPE, add 100 million suits to the Strategic National Stockpile, and launch a massive smallpox vaccination campaign, for example, Russia and China might look to these actions, combined with the U.S. biodefense research program, and conclude that the U.S. is in fact preparing to execute a biological attack. This in turn might spur bioweapons development in these states.

This is not just a theoretical risk. Accounts of the Soviet bioweapons program *Biopreparat* suggest that it was in part motivated by Russian fears of ongoing U.S. biodefense activities.²⁷³ Ken Alibek, a Soviet Defector and high-level scientist in *Biopreparat*, later reflected that U.S. defensive activities made them question Nixon's sincerity in renouncing biological weapons:

“we also noted that a small army medical unit had begun work at Fort Detrick. This unit, known as the United States Army Medical Research Institute of Infectious Diseases and ostensibly dedicated to biological defense, seemed to expand in importance and strength each year. Former bioweaponeers like Bill Patrick had gone to work for it. Even if our intelligence activities couldn't come up with concrete evidence of offensive work, there could be no doubt that such work continued.”²⁷⁴

This in turn appears to have motivated the continued growth of the Soviet BW program in violation of the Biological Weapons Convention, and even worked on vaccine resistant and weaponized versions of highly infectious agents like smallpox and plague – agents that could cause true global catastrophic biological events.

Information Hazards

Information hazards abound in discussions of biological risks.²⁷⁵ Sometimes, the mere act of drawing attention to vulnerabilities or capabilities can be hazardous. For example, correspondence from

²⁷¹ The classic statement of the security dilemma is John H. Herz, “Idealist Internationalism and the Security Dilemma,” *World Politics* 2, no. 2 (1950): 157–80, <https://doi.org/10.2307/2009187>. See also Charles L. Glaser, “The Security Dilemma Revisited,” *World Politics* 50, no. 1 (1997): 171–201.

²⁷² For one discussion of this in the context of mid-2000s biodefense debates, see Jonathan B. Tucker, “Avoiding the Biological Security Dilemma: A Response to Petro and Carus,” *Biosecurity and Bioterrorism: Biodefense Strategy, Practice, and Science* 4, no. 2 (June 2006): 195–99, <https://doi.org/10.1089/bsp.2006.4.195>.

²⁷³ For a full discussion of this, see Enemark, *Biosecurity Dilemmas*, and Chris Bakerlee and Tessa Alexanian, “Info hazard guidance for biosecurity discussions,” <https://docs.google.com/document/u/0/d/1VSfU3GiZumHDX2hoz3YY1PT2dQHtkbrfO8xLxI9BTGE/mobilebasic#ftnt1>.

²⁷⁴ Alibek, *Biohazard*, 235.

²⁷⁵ https://www.fhi.ox.ac.uk/wp-content/uploads/Lewis_et_al-2019-Risk_Analysis.pdf



Ayman al-Zawahiri, a former high-ranking member of *Al Qaeda*, the terrorist organization only began considering biological weapons after “the enemy drew [al-Qaeda’s] attention to them by repeatedly expressing concerns that they can be produced simply.”²⁷⁶ (Interestingly, these effects can be complicated; the leadership of the Japanese doomsday cult Aum Shinrikyo allegedly grew so frustrated with their failed pursuit of biological weapons that the founder Shoko Asahara speculated “that American comments on the risk of biological weapons were intended to delude would-be terrorists into pursuing this path.”²⁷⁷)

The problem of information hazards in biosecurity is well-known and -debated, so it is not discussed in detail here.²⁷⁸ (Arguably, the security dilemma is a subset of information hazards.) This might include:

- Spreading blueprints (e.g. virus genomes) that could be misused
- Disclosing societal vulnerabilities
- Disclosing the usefulness of certain capabilities
- Making biological agents appear attractive to bad actors

Although concerns about information hazards can be taken too far, we believe they are an important part of any responsible grantmaker portfolio.

Open Questions for Further Research

- What can philanthropists working in biosecurity learn from the treatment of information hazards in cybersecurity?
- What measures are most useful for understanding offense-defense distinguishability?
- What are the best ways of measuring and quantifying the positive externalities of pandemic preparedness?

²⁷⁶ Ayman al-Zawahiri, quoted in Chris Bakerlee and Tessa Alexanian, “Info hazard guidance for biosecurity discussions,” <https://docs.google.com/document/u/0/d/1VSfU3GiZumHDX2hoz3YY1PT2dQHtkbrfO8xLxI9BTGE/mobilebasic#ftnt1>.

²⁷⁷ Richard Danzig et al., “Aum Shinrikyo: Insights Into How Terrorists Develop Biological and Chemical Weapons,” Center for a New American Security (2011), 26, <https://www.cnas.org/publications/reports/aum-shinrikyo-insights-into-how-terrorists-develop-biological-and-chemical-weapons>.

²⁷⁸ See, e.g. Gregory Lewis et al., “Information Hazards in Biotechnology,” *Risk Analysis* 39, no. 5 (2019): 975–81, <https://doi.org/10.1111/risa.13235>.



Practical Funding Opportunities

Key Points

- Using the guiding principles above, philanthropists can narrow their search space to highly effective funding opportunities.
- Potential high-impact opportunities may include:
 - GCBR field building, policy advocacy (e.g. on bioscience governance and shaping state incentives around BW),
 - safeguards against proliferation and misuse of AI and DNA synthesis,
 - transmission-blocking and biological indoor air quality interventions,
 - pandemic-proof personal protective equipment,
 - pathogen-agnostic early detection,
 - and platform technologies for broad-spectrum countermeasures.

Until now, this report has remained largely abstract, pointing to general principles and stylized facts about the risks of a biological catastrophe in order to better analyze the nature of that risk and the shape of possible risk mitigation measures. This section attempts to turn these insights into concrete recommendations. After all, biosecurity is often touted as one of the most tractable existential risks facing humanity. Unlike other risks, like risks from artificial intelligence, the argument goes, we have a rough understanding of the threat as well as a toolbox of powerful ideas and technologies. This section tests this idea by attempting to describe potential high-impact funding opportunities that philanthropists can pursue immediately.

Knowing what we know about the structure of the risk, and having derived some basic heuristics for high impact, what can philanthropists do? The following interventions — from a variety of sources acknowledged in the footnotes of each of the following sections — appear especially promising:

- **GCBR Field-Building**, including building a self-sustaining pipeline of talent to work on the most extreme possible risks. This may include both starting and growing biosecurity organizations and providing training to promising people at various stages in their careers.
- **GCBR Policy Advocacy**, especially advocacy that can bridge the gap between traditional health security and preparedness for extreme tail risk scenarios.



- Particularly promising opportunities for policy advocacy may include bioscience governance at national and international levels (e.g. on DNA synthesis screening), other interventions that restrict access to dangerous technologies, and interventions designed to make the pursuit of biological weapons less desirable for states.²⁷⁹
- **Safeguards against Proliferation and Misuse of Key Technologies**, especially enabling AI technologies and DNA synthesis screening.
- **Transmission-Blocking Interventions**, including but not limited to research, pilot trials, market-shaping, and advocacy around germicidal light.
- **Pandemic-Proof Personal Protective Equipment**, and advocacy to ensure that such equipment can be readily distributed to a large cadre of essential workers, including the need for stockpiling.
- **Pathogen-Agnostic Early Detection**, like the [Nucleic Acid Observatory](#) or [Threat Net](#).
- **Platform Technologies for Broad-Spectrum Countermeasures**, although we are uncertain about the tractability of this recommendation.

Additionally, the appendix [Evaluating Interventions](#) provides a practical framework for comparing different interventions based on the criteria outlined here.

This list is only a first attempt at describing the most effective interventions, and other interventions will be an important component of a layered and robust defense against biological threats.²⁸⁰

GCBR Field-Building

One of our top recommendations to reduce global catastrophic biological risks is **funding field-building projects** to grow the talent and resources devoted to the most extreme risks facing humanity. As described above (see “[Incentives, Preparedness, and Neglectedness](#)”), the field of health security is massive, but the sub-field of people focused on global catastrophic risks remains very small. Field-building activities could include:

- Active grantmaking to grow existing GCBR programs at think tanks, academic institutions, etc.
- Active grantmaking to establish new GCBR programs at think tanks, academic institutions, etc.

²⁷⁹ Thanks to Joshua Monrad and Andrew Snyder-Beattie for these points.

²⁸⁰ This may include, for example, “triple rapid diagnostics,” better real-world exercises for pandemic response, and a more comprehensive approach to biorisk management. Thanks to Jake Swett for this list of additional interventions.



- Scholarships for graduate students interested in pursuing a career in GCBR mitigation
- Funding public service fellowships focused on GCBRs
- Creating incentives (e.g. long-term funding and endowed chairs at universities) that help to stimulate the growth of the field.²⁸¹

This work could focus on a wide variety of threats, including shaping the intentions and capabilities of various actors who would be interested in the misuse of biological agents.²⁸² It could also include work on growing the presence of GCBR-minded scientists to change the culture of science to be focused more on security and on the potential harms of working with enhanced potential pandemic pathogens.²⁸³

It is important, however, that such field-building grants follow the guidelines outlined above. For instance, grantmakers interested in getting a well-respected think tank to focus on biological risks ought to ensure that such a research program (1) focuses on truly catastrophic threat scenarios, while at the same time (2) is sensitive to the information hazard concerns that accommodate such scenarios. For instance, grant agreements could stipulate that any published material on particularly sensitive topics with infohazardous potential undergo a special round of review, or that certain vulnerabilities are discussed only in closed-door briefings with relevant decision-makers.

GCBR Policy Advocacy

As in climate change, one reason why policy advocacy can provide high leverage in biosecurity is that there exists a large pie of resources devoted to the broader field of health security. Thus, one goal of effective GCBR policy advocacy can include efforts to move more resources towards extreme events, and to ensure that existing preparedness efforts are robust to worst-case scenarios. Possible funding options include:

²⁸¹ Tom Inglesby, Director of the Johns Hopkins Center for Health Security, explained: ““This field that’s focused on GCBRs, it’s small, it’s new, its funding is volatile and actually on the way down following the pandemic. Unless people see that there’s some kind of sustainable support for the field, people aren’t going to want to go into it, organizations won’t be able to take any risks or [...] attract expertise.” (Interview with Tom Inglesby, 8 June 2023).

²⁸² For a discussion of NTI’s framework of “constraining capabilities” and “shaping intent”, see “Jaime Yassif on Safeguarding Bioscience to Prevent Catastrophic Lab Accidents & Bioweapons Development,” 80,000 Hours, accessed August 11, 2023, <https://80000hours.org/podcast/episodes/jaime-yassif-safeguarding-bioscience/>.

²⁸³ As Dr. Sarah Carter explained in an interview for this report (also quoted above): “Who is going to do that [i.e. engineer a dangerous virus]? It’s going to be a postdoc, it’s going to be a grad student, it’s going to be somebody already in the lab. These tools are not going to suddenly [...] enable terrorists with no bio training in a cave in Afghanistan. [...] Which is one reason I think that projects that could change the culture and academia to make them more aware of these things would be worthwhile, even though it’s a really hard intractable problem.” Interview with Dr. Sarah Carter, 18 May 2023. See also Kevin M. Esvelt, “Inoculating Science against Potential Pandemics and Information Hazards,” *PLoS Pathogens* 14, no. 10 (October 4, 2018): e1007286, <https://doi.org/10.1371/journal.ppat.1007286>.



- Growing existing or establishing new GCBR-focused policy advocacy organizations (e.g. think tanks)
- Funding briefings and roundtables with policymakers to illustrate the extreme risk
- Advocacy to regulate artificial intelligence and its potential misuse by malevolent actors seeking to engineer potential pandemic pathogens.
- Advocacy for better preparedness among defense and intelligence agencies.²⁸⁴
- Advocating increased federal funding for defense-dominant technologies to promote differential technological development.²⁸⁵

For some funders, non-charitable political activity (i.e. lobbying) may be possible. Given the philanthropic focus of this report, however, we do not discuss this option further here. Rather, philanthropists can help to orient university groups, think tanks, and other nonprofit organizations to the most important policy-relevant problems.²⁸⁶

As with field-building activities, philanthropists ought to be mindful of the possible hazards involved with such grantmaking. Moreover, there are many interest groups involved with these issues. Building bridges between these groups while preserving a focus on worst-case scenarios likely requires political finesse. In this case, organizational strength and team strength become especially important when evaluating a possible funding opportunity.

Constraining Capabilities and Shaping Incentives

Two promising avenues for specific policy advocacy projects include constraining capabilities of malevolent actors and shaping the incentives of powerful states.²⁸⁷ Essentially, philanthropists can fund policy advocacy initiatives that:

²⁸⁴ “For example, defense establishments are incentivized to pay heed to potentially adversarial threats and take precautionary actions to protect military personnel, who can function as primary essential workers by helping to maintain law and order in a crisis. If defense agencies order military personnel and their families to take extreme precautions to avoid infection, most of society will take note.” Anjali Gopal, William Bradshaw, Vaishnav Sunil, Kevin M. Esvelt, “Securing Civilization against Catastrophic Pandemics” (Geneva Centre for Security Policy, 2023), *in press*.

²⁸⁵ Thanks to Jake Swett for pointing to this in a round of external reviews.

²⁸⁶ Importantly, this often does not happen on its own within university structures. As Tom Inglesby, the Director of the Johns Hopkins Center for Health Security, put it in an interview for this report, “[Policy advocacy is] not necessarily a natural state for researchers and universities. You have to really decide to be in that space, because research universities don’t typically reward it.” (Interview with Tom Inglesby, 8 June 2023)

²⁸⁷ With thanks to Jaime Yassif for these phrases and to Andrew Snyder-Beattie and Joshua Monrad for related comments in a round of external reviews.



1. **Constraining Capabilities** — National policy and international regulations that help to constrain the access of malevolent actors (e.g. would-be bioterrorists) to powerful technologies.
 - a. This may include, for example:
 - i. Regulating access to powerful AI applications for biotechnology.
 - ii. Creating a common mechanism for DNA synthesis screening that is cheap and easy to adopt.
 - iii. Improving export control policies, publication and funding policies, and other regulations that affect the proliferation of powerful capabilities and information.
2. **Shaping Incentives** — Creating and pulling policy levers that incentivize states away from the pursuit of biological weapons programs. E.g. increasing the likelihood of attribution of a biological attack, strengthening the Biological Weapons Convention and relevant national legislation, decreasing the attractiveness of biological weapons programs in other ways.
 - a. This overlaps with deterrence-by-denial effects of other interventions discussed here.

Notably, many of these interventions are useful for a wide variety of threats, from smaller-scale bioterrorism to the kinds of extreme biological catastrophes discussed in this report.

Safeguards against Proliferation and Misuse of Key Technologies

Hand-in-hand with effective policy advocacy, philanthropists can also fund research and development on safeguards against proliferation and misuse of key technologies in biological risks — with the aim of ultimately using these as proofs-of-concept for public sector adoption. As discussed throughout this report, prevention may be more robust than response. Ultimately, perfect prevention of misuse is likely impossible, but researchers can nonetheless help to identify key technologies of concern, find effective ways of preventing their falling into the hands of malevolent actors, and develop guardrails against their misuse.

The AI-Bio Nexus

As discussed above, artificial intelligence, especially more powerful models, may become a key enabling technology for bioscience. It is therefore possible that one of the best ways to prevent biological catastrophe has little to do with bioscience, and more to do with AI safety and governance.²⁸⁸ Advocacy for evaluations and regulations to prevent the proliferation of open-source models with dangerous capabilities, for example, could help to restrict access to these tools and thereby constrain the number of would-be bioterrorists who are able to engineer dangerous pathogens.

DNA Synthesis Regulation

DNA synthesis screening is a second promising area for potential philanthropic investment in safeguards. As discussed above, DNA synthesis remains under-regulated, with many providers not

²⁸⁸ Thanks to Andrew Snyder-Beattie for emphasizing this point.



screening customers, not screening orders, or both. On one level, this might include policy advocacy for national-level regulations requiring certain screening practices. Additionally, philanthropists could help to fund the creation of groups that regularly assess and red-team companies' screening practices, as well as creating quantitative benchmarks for screening performance.²⁸⁹ This may also include funding interdisciplinary work with cybersecurity experts.

Transmission-Blocking Interventions

Because much transmission of respiratory viruses occurs indoors — and because active defenses like masks come with various problems related to human behavior — transmission-blocking interventions have high promise as pandemic preparedness and mitigation efforts — although significant uncertainties remain on the efficacy and safety of the technologies in question. Such efforts could include:

- Funding related to germicidal light and other biological indoor air quality interventions²⁹⁰
 - Market-shaping and advocacy around the widespread adoption of existing ~254 nm-wavelength germicidal light installed as upper-room lights (a long-proven technology that — if installed correctly — does not come into contact with humans)
 - Safety research, market-shaping, technological development, efficacy demonstrations and advocacy around 200-235nm germicidal light (sometimes called “far-UVC”), which can be installed as both upper-room and full-room lighting in many spaces and has the potential for “near field protection.”²⁹¹
 - Continued work on improved filtration and ventilation.
- Real-world transmission studies on a variety of transmission-blocking technologies
 - Existing research looks at pathogen inactivation, but the translation of these measures to actual transmission reduction remains uncertain under real-world conditions.
 - Scalable sterilization technologies and mechanisms to quickly disinfect PPE.²⁹²
- Research into under-explored transmission-blocking interventions, including triethylene glycol, microwave inactivation, and aerosol capture devices.²⁹³

²⁸⁹ Thanks to Max Langenkamp for these points.

²⁹⁰ It appears that low-wavelength light has germicidal properties but is unable to penetrate the outer layers of human skin and eyes, and is thus technically safe. We believe that more safety research is likely needed to establish the safety of these technologies in the eyes of regulators and policymakers, and that more real-world transmission studies would help to reduce remaining uncertainties around the technology. This is being pursued by major GCBR funders. See, e.g. “(Request for Information) Evaluation of Germicidal Far-UVC: Safety, Efficacy, Technology, and Adoption,” *Open Philanthropy*, accessed August 11, 2023, <https://www.openphilanthropy.org/research/request-for-information-evaluation-of-germicidal-far-uv-c-safety-efficacy-technology-and-adoption/>.

²⁹¹ Thanks to Jake Swett for noting the importance of near-field protection.

²⁹² Thanks to Andrew Snyder-Beattie for this addition.

²⁹³ Thanks to Jake Swett for these specific ideas.



Funders could also consider some kinds of surface sterilization technology. We expect, however, given the extreme transmissibility of respiratory pathogens (and the advantages that a malicious actor may therefore see in such pathogens), that interventions focused on respiratory viruses will be most effective, at least in the near term.

In general, one major theoretical advantage of transmission-blocking interventions like germicidal low-wavelength light is truly **pathogen agnostic**. In theory, one can imagine risk-shifting (e.g. malicious actors engineering UV-resistant pathogens), but nonetheless, biological indoor air quality interventions appear to be among the broadest-spectrum pandemic preparedness measures that exist.

Pandemic-Proof Personal Protective Equipment (P4E)

In the event of a catastrophic pandemic, and in the absence of widespread germicidal light or high-quality air filtration technologies, at least essential workers will need personal protective equipment that is robust to worst-case scenarios and provides workers the confidence to go about their jobs without fear of infection — this is sometimes called “P4E,” (pandemic-proof personal protective equipment).²⁹⁴ Advocacy for stockpiling may be a promising route to both protect society and create a market to decrease costs further. While private technical research is one possible funding option, we believe that policy advocacy may be the more cost-effective option (e.g. getting BARDA to prioritize the development of new P4E and lobbying for and incentivizing stockpiling).²⁹⁵

We believe that P4E has one of the higher risks of being misperceived by adversaries if, e.g. an early customer is the U.S. military. It may be necessary to couple such development with open information- and technology-sharing programs that reduce the fear that such equipment could be coupled with offensive technologies.

Again, because it functions mechanically rather than medically, P4E is **pathogen agnostic**.

²⁹⁴ Esvelt suggests P4E: “Cost-effective pandemic-proof personal protective equipment (P4E) has yet to be developed, but likely does not require fundamental advances. It is equipment that, when worn by an untrained person, completely blocks external access to all human cells that could be subverted and used to replicate a pandemic-class agent, harm the wearer, or both.” Kevin M Esvelt, “Delay, Detect, Defend: Preparing for a Future in Which Thousands Can Release New Pandemics,” Geneva Papers (Geneva Centre for Security Policy, 2022), 33,

https://dam.gcsp.ch/files/doc/gcsp-geneva-paper-29-22?_gl=1*xm44p1*_ga*ODQxMDI0NiY4LjE2ODM2NTg5MTg.*_ga_Z66DSTVXTJ*MTY4MzczNTIzNi4yLjAuMTY4MzczNTIzNi4wLjAuMA..#page=49&zoom=100.0.0.

²⁹⁵ Experts interviewed for this project echoed the belief that technological progress without policy advocacy will not be an effective approach. For example, Tom Inglesby, the Director of the Johns Hopkins Center for Health Security, explained in an interview: “Even if we get an amazing new technology for PPE, that’ll just be the beginning of [...] all the other issues around getting that PPE to scale into the places that it needs to be available to the public, in different parts of the world.” (Interview with Tom Inglesby, 8 June 2023).



Pathogen-Agnostic Early Detection

A layered system of early detection of rising threats — including environmental monitoring as well as sentinel, clinical, and digital components — will be crucial for better pandemic preparedness, especially given the threat of pathogens with long infectious incubation periods.²⁹⁶ Wastewater monitoring for large clusters of diseases is one useful approach, but truly pathogen-agnostic monitoring and detection would likely require a layered approach. One component of this could be deep metagenomic sequencing to understand whether there are “exponentially growing patterns of nucleic acid fragments,” but other approaches to detect novel pathogens exist, including genetic engineering detection.²⁹⁷ Note that despite the promise of these methods, no single early detection system can be *fully* threat-agnostic — there will always be some threats, both known and unknown, that will not be captured by metagenomic wastewater sequencing, for example.²⁹⁸ This does not mean, however, that a well-designed *layered system* could not capture nearly all threats, with sentinel or clinical sequencing picking up pathogens that environmental sequencing misses, with each layer covering the blindspots of others.²⁹⁹

The Nucleic Acid Observatory Consortium has stated this succinctly:

“By searching for divergences from historical baseline frequencies at sites throughout the world, [a Nucleic Acid Observatory] could detect any virus or invasive organism undergoing exponential growth whose nucleic acids end up in the water, even those previously unknown to science. Continuously monitoring nucleic acid diversity would provide us with universal early warning, obviate subtle bioweapons, and generate a wealth of sequence data sufficient to transform ecology, microbiology, and conservation.”³⁰⁰

²⁹⁶ Thanks to Jake Swett for pointing to the importance of thinking beyond merely environmental monitoring.

²⁹⁷ “any system capable of detecting exponentially growing patterns of nucleic acid fragments should be capable of reliably detecting any and all catastrophic biothreats, including stealthy agents analogous to HIV that might otherwise infect most of humanity before exhibiting any visible clinical effects” Kevin M Esvelt, “Delay, Detect, Defend: Preparing for a Future in Which Thousands Can Release New Pandemics,” Geneva Papers (Geneva Centre for Security Policy, 2022), 30, https://dam.gcsp.ch/files/doc/gcsp-geneva-paper-29-22?_gl=1*xm44p1*_ga*ODOxMDI0NiY4LjE2ODM2NTg5MTg.*_ga_Z66DSTVXTJ*MTY4MzczNTIzNi4yLjAuMTY4MzczNTIzNi4wLjAuMA..#page=49&zoom=100.0.0.


²⁹⁸ See David Manheim, Naham Shapiro, and Rona Tobolsky, “How Agnostic Is Your ‘Threat Agnostic’ Biosurveillance System?,” SSRN Scholarly Paper (Rochester, NY, May 13, 2023), <https://doi.org/10.2139/ssrn.4447633>.

²⁹⁹ Thanks to Jake Swett for this point in a round of external reviews.

³⁰⁰ The Nucleic Acid Observatory Consortium, “A Global Nucleic Acid Observatory for Biodefense and Planetary Health” (arXiv, August 5, 2021), <https://doi.org/10.48550/arXiv.2108.02678>.



Detecting Threats via Deep Metagenomic Sequencing



Source: Nucleic Acid Observatory Consortium, “[A Global Nucleic Acid Observatory for Biodefense and Planetary Health](#).”

For a deeper discussion of the nuances and obstacles to widespread adoption of pathogen-agnostic biosurveillance systems, we recommend “[Managing the Transition to Widespread Metagenomic Monitoring: Policy Considerations for Future Biosurveillance](#),” “[Threat Net](#)” and “[Toward a Global Pathogen Early Warning System](#)”³⁰¹

Platform Technologies for Medical Countermeasures

As outlined by the Bipartisan Commission for Biodefense, it appears useful to develop platform technologies for medical countermeasures. Even if vaccines and medical countermeasures cannot be tested quickly enough to stop the spread of a catastrophic pandemic, the readiness to treat and find cures for a threat with a long incubation period remains important.

The ability to quickly pivot production facilities to the newest threat is an important part of layered defense. We remain highly uncertain about this recommendation, however, and expect that it is likely to be less neglected than other recommendations discussed here. Within the space of medical countermeasures there is also high potential for dual-use concerns. This applies, for example, to some vaccine technologies.³⁰² Some biosecurity experts have suggested that RNA vaccine

³⁰¹ Chelsea Liang, James Wagstaff, Noga Aharoni, Virginia Schmit, and David Manheim, “Managing the Transition to Widespread Metagenomic Monitoring: Policy Considerations for Future Biosurveillance,” *Health Security* 21, no. 1 (February 2023): 34–45, <https://doi.org/10.1089/hs.2022.0029>

³⁰² Jonas B. Sandbrink and Gregory D. Koblentz, “Biosecurity Risks Associated with Vaccine Platform Technologies,” *Vaccine, Pandemic Simulation, Pacific Eclipse*, 40, no. 17 (April 14, 2022): 2514–23, <https://doi.org/10.1016/j.vaccine.2021.02.023>.



technologies have lower dual-use concern than other technologies, and are therefore more suitable platform candidates.³⁰³

Open Questions for Further Research

- How tractable is each of these potential interventions in the current political environment?
- What are the remaining uncertainties on germicidal light and other transmission-blocking interventions? What safety and efficacy studies would be most useful?
- What, if anything, can philanthropists learn from efforts to protect isolated populations (e.g. astronauts) for pandemic preparedness?

³⁰³ Jonas B. Sandbrink and Gregory D. Koblenz, “Biosecurity Risks Associated with Vaccine Platform Technologies,” *Vaccine*, Pandemic Simulation, Pacific Eclipse, 40, no. 17 (April 14, 2022): 2514–23, <https://doi.org/10.1016/j.vaccine.2021.02.023>.



Conclusion

Much of this report has been gloomy; pandemic risk is rising, malevolent actors are becoming more capable, and society is entirely unprepared for a global catastrophic biological event. As the previous section showed, however, **there is much we can do**. We are not clueless about risk-mitigation measures, but rather have a deep understanding of biological processes. Moreover, it is possible that advances in bioscience and biotechnology will enable a world with a far lower burden of infectious disease, and a truly pandemic-proof society.³⁰⁴ We can imagine the following future:

- Bioscience is regulated in a way that restricts malevolent actors' access to dangerous knowledge and capabilities without stifling scientific progress.
- Thanks to public-service fellowships, scientists and policymakers understand each other better, and a security mindset flourishes in the life sciences, helping to keep the rest of us safe.
- Monitoring systems around airports and other global hubs monitor for exponentially-growing nucleic acid fragments, and decision-makers take alerts seriously.
- Advanced low-wavelength germicidal light is shown to be safe and effective and installed in most public spaces at low cost, keeping airborne transmission low.
- Essential workers know that they are valued and would have access to truly pandemic-proof personal protective equipment in the event of a pandemic.
- States strengthen the Biological Weapons Convention, such that no state's leader feels that they could get away with violations of the treaty.
- A vast biomedical industrial base not only leads to rapid economic growth, but also contains the kinds of platform technologies that could quickly pivot to address emerging threats.
- Supply chains are actively designed and on-shored where necessary to ensure that pandemic response can continue even when global trade falters.
- Omnicidal and near-omnicidal groups know that there are impenetrable bio-bunkers spread across the world at undisclosed locations. Thanks to deterrence-by-denial, these groups

³⁰⁴ Tom Inglesby, the Director of the Johns Hopkins Center for Health Security, outlined one hopeful vision for the future in an interview for this report: "It is possible to imagine how indoor air in commercial buildings could be 100 times safer than it is now. It's possible to imagine PPE that is usable 1000 times as opposed to once and then thrown away. And it is possible to imagine scaling up diagnostics and having them in people's homes around the world. [...] It just requires the steady application of pressure and [...] a multi-step process and political advocacy and building coalitions and having smart people committed to it." (Interview with Tom Inglesby, 8 June 2023).



know that they cannot extinguish humanity with biological agents, try as they might; hence they don't try.

This is not an impossible world. It simply requires targeted investments by far-sighted philanthropists. Such a world would not only reduce the risk of a biological catastrophe, but it would likely also reduce the global burden of disease dramatically. Millions of people still die every year from infectious diseases. These diseases disproportionately affect the most vulnerable members of our society — poor and marginalized groups. Better preparedness for catastrophic biological risks entails a better world for everyone.



Appendix

A Grantmaker Dilemma Checklist

This checklist is a set of questions designed to stimulate careful thought about the possible risks associated with grantmaking surrounding global catastrophic biological events. We encourage philanthropists to edit and adapt this checklist to suit their process.

Note that this checklist is *not* designed to screen for traditional grant risks (e.g. legal risks), but for a specific subset of risks that are especially important when dealing with GCBRs. This checklist is also not designed for general science funders, who also need to consider the risks of funding dangerous scientific research (e.g. when working with enhanced potential-pandemic pathogens). The checklist is inspired in part by Jason Matheny’s “Questions for IARPA Program Managers,” as cited in [Technology Roulette](#).

- ☐ Are you following [best practices](#) with regard to [information hazards](#)? Could this grant disclose or amplify information that could be misused by malicious actors or point to specific vulnerabilities that malicious actors could exploit?
- ☐ What are the possible dual-use implications of your grantmaking? Could this grant create novel capabilities that malicious actors could misuse?
- ☐ How might different countries and other stakeholders misinterpret (intentionally or unintentionally) this grant? Could this grant be misconstrued as helping to support an offensive biological weapons activity?
- ☐ Is there a clear process in place for screening for hazards for any parts of the project that will be made public?
- ☐ Can you ensure that researchers will share findings in a limited and responsible manner (e.g. only in closed-door briefings) if these findings involve potentially dangerous information? What would happen if the information is made public anyway?
- ☐ Could this grant politicize the discourse around extreme biological risks in unhelpful ways?
- ☐ If you are uncertain about the risks associated with your grantmaking, are there other actors who are not potential grantees (e.g. program officers from peer institutions) whom you can ask for a second opinion?



Evaluating Interventions

The “impact multiplier” framework above helps to evaluate different kinds of interventions. For example, we can use it to compare two fictional organizations:

1. **Organization A (GUV)** — a 501(c)(3) policy advocacy organization focusing on germicidal UV light, helping to stimulate government funding for better funding on real-world transmission and safety studies, which the organization has determined are the crucial considerations for whether GUV could be useful for an extinction-level event. In policy advocacy, Organization A also emphasizes the economic benefits of reduced pathogen spread in offices and public spaces. Organization A has several distinguished but low-profile national security experts on its board and specifically focused on GUV because of its apparent low dual-use potential.
2. **Organization B (Virus Hunting)** — a university-affiliated research group that tries to find and catalog existing strains of coronavirus to “prevent future natural spillover events and help prevent any future COVID-like pandemic from ever occurring again. Organization B is staffed by world-class scientists who prioritize the importance of free and open exchange of information and are suspicious of the securitization of their work. They prefer to work with international organizations like the WHO and large traditional global health organizations. They believe that the scientific community, working collaboratively across borders, can help eliminate the threat of natural pandemics. Organization B has endorsements from Nobel Prize-winning scientists and has former Presidents and Prime Ministers on its board.

Impact Multiplier	Organization A (GUV)	Organization B (Virus Hunting)
Focus on worst-possible scenarios.	1/1 The organization specifically worries most about existential risks from engineered pathogens.	0/1 The organization focuses on natural spillover from coronaviruses, apparently not for impact-related reasons.
Fund interventions that are robust to the entire spectrum of risk, up to and including extinction-level pandemics.	1/1 GUV can be a “passive defense” that can be useful for preventing and responding to extinction-level threats.	0/1 The organization focuses only on natural spillover events for vaccine candidates.
Within anthropogenic events, generally pursue pathogen- and threat-agnostic approaches , and beware mere	1/1 GUV is a pathogen-agnostic intervention that inactivates many different kinds of threats,	0/1 Again, the organization focuses only on coronaviruses.



risk-shifting and “pandemic Maginot Lines.”	suitable to a risk-general approach.	
Leverage existing societal resources using advocacy-based interventions.	1/1 Rather than conduct studies themselves, Organization A seeks to mobilize much larger government resources for major work on transmission and safety.	?/1 The organization does direct work, but collaborates with the WHO. We do not have enough information to assess whether this collaboration also seeks to affect the priorities of the WHO and other organizations.
<i>All else equal, prioritize interventions with near-term positive externalities</i> that can garner sustained and broad political support.	1/1 Organization A emphasizes near-term economic benefits of GUV installment to incentivize political action. Their brand focus on “existential risk” may undermine this effort somewhat.	1/1 Given the memories of the COVID-19 pandemic, Organization B work does have the benefits of political tractability.
Prioritize offense-defense distinguishability and avoid information hazards when selecting funding opportunities to avoid fuelling dangerous security dilemmas and doing more harm than good.	1/1 Organization A is clearly concerned about these hazards and has an advisory board of national security experts.	0/1 Organization B emphasis on openness and suspicion of security apparatus, as well as their pursuit of “virus hunting,” suggest that they do not prioritize these hazards.

In short, although Organization B is an organization staffed with admirable scientists doing interesting work, they fail to meet the impact multipliers specified above, and may be pursuing some actively harmful work.



Plain Text R Code

Carlson Curve on DNA Synthesis ([data](#))

```
# Description: This is an R script to analyze the falling costs of DNA
synthesis

# Load necessary packages

library(tidyverse)  # includes ggplot2, tidyr, etc.

library(scales)     # For the comma format

# Data input

data <- read.csv(text="

Year,Gene.synthesis.cost.Carlson,Gene.synthesis.cost.Potomac,Today.cheapest
.cost

1999,25,NA,NA

2000,NA,NA,NA

2001,12,11,NA

2002,8,10,NA

2003,4,6.3,NA

2004,NA,NA,NA

2005,NA,NA,NA

2006,1,2,NA

2007,0.5,0.63,NA

2008,NA,NA,NA

2009,0.39,0.63,NA


```



```
2010,0.35,0.5,NA
2011,0.29,0.4,NA
2012,0.2,0.3,NA
2013,0.18,0.25,NA
2014,0.15,0.2,NA
2015,NA,NA,NA
2016,0.03,0.1,NA
2017,0.02,0.08,NA
2018,NA,0.06,NA
2019,NA,NA,NA
2020,NA,NA,NA
2021,NA,NA,NA
2022,NA,NA,NA
2023,NA,NA,0.07
")
```

#NB: I have left the 7 cent from 2023 estimate in the data here, but remove it for the visualization.

```
# Transform data, filter out NA values, and remove 2023 data
```

```
data_long <- data %>%
  select(-Today.cheapest.cost) %>% # Removes the "Today cheapest" column
  tidyr::gather("Source", "Cost", -Year) %>%
  filter(!is.na(Cost) & Year != 2023)
```




```
# Create the scatter plot with a logarithmic y-axis and specified
formatting

p <- ggplot(data_long, aes(x = Year, y = Cost)) +

  geom_point(aes(color = Source), size=3) +

  geom_smooth(method = "lm", color = "#4fc27a", aes(group = 1)) +

  scale_color_manual(values = c("Gene.synthesis.cost.Carlson" = "#edd2b7",
"Gene.synthesis.cost.Potomac" = "#a3d6b0")) +

  scale_y_log10(labels = comma) +

  labs(title = "DNA Synthesis Cost 'Carlson Curve'",

        subtitle = "USD per base, 1999-2018. Data from Carlson (2023) and
Potomac (2018)",

        x = "Year",

        y = "Cost per Base (log scale)",

        caption = "Plot generated in R using assistance from GPT4. Data from
Carlson (2023) and Potomac (2018).

        With thanks to Max Langenkamp (see in-text links).") +

  theme_minimal() +

  theme(plot.title = element_text(size = 18, face = "bold", family =
"Georgia"),

        plot.subtitle = element_text(size = 14),

        axis.title.x = element_text(size = 14),

        axis.title.y = element_text(size = 14),

        plot.caption = element_text(size = 10),

        legend.position="none")


# Show the plot

print(p)
```



DNA Synthesis Cost 'Carlson Curve'

USD per base, 1999-2018. Data from Rob Carlson (2023) and Potomac (2018)



Plot generated in R using assistance from GPT4. Data from Carlson (2023) and Potomac (2018).
With thanks to Max Langenkamp (see in-text links).

Cost of Sequencing a Human Genome over Time

Description: This is an R script to analyze the falling costs of sequencing a human genome over time.

```
# Import libraries
```

```
library(ggplot2)
```

```
library(dplyr)
```

```
library(scales)
```

```
# Data
```

```
dates <- c('2007-01-01', '2008-01-01', '2009-01-01', '2010-01-01',  
           '2011-01-01', '2012-01-01', '2013-01-01', '2014-01-01',  
           '2015-01-01', '2016-01-01', '2017-01-01', '2018-01-01',  
           '2019-01-01', '2020-01-01', '2021-01-01', '2022-01-01',
```



```
'2026-01-01', '2031-01-01')
costs <- c(9408739, 3063820, 232735, 46774, 20963, 7666, 5570, 5004, 1363,
          1356, 1015, 1232, 993, 645, 851, 525, 114, 56) # median costs
lower <- c(NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA,
          525, 66.6, 18.8) # lower 25% costs
upper <- c(NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA, NA,
          525, 191, 122) # upper 75% costs

df <- data.frame(dates = as.Date(dates), costs = costs, lower = lower,
                 upper = upper)

# Split data into before and after 2023
df_before <- df[df$dates <= as.Date('2022-12-31'),]
df_after <- df[df$dates >= as.Date('2022-01-01'),]

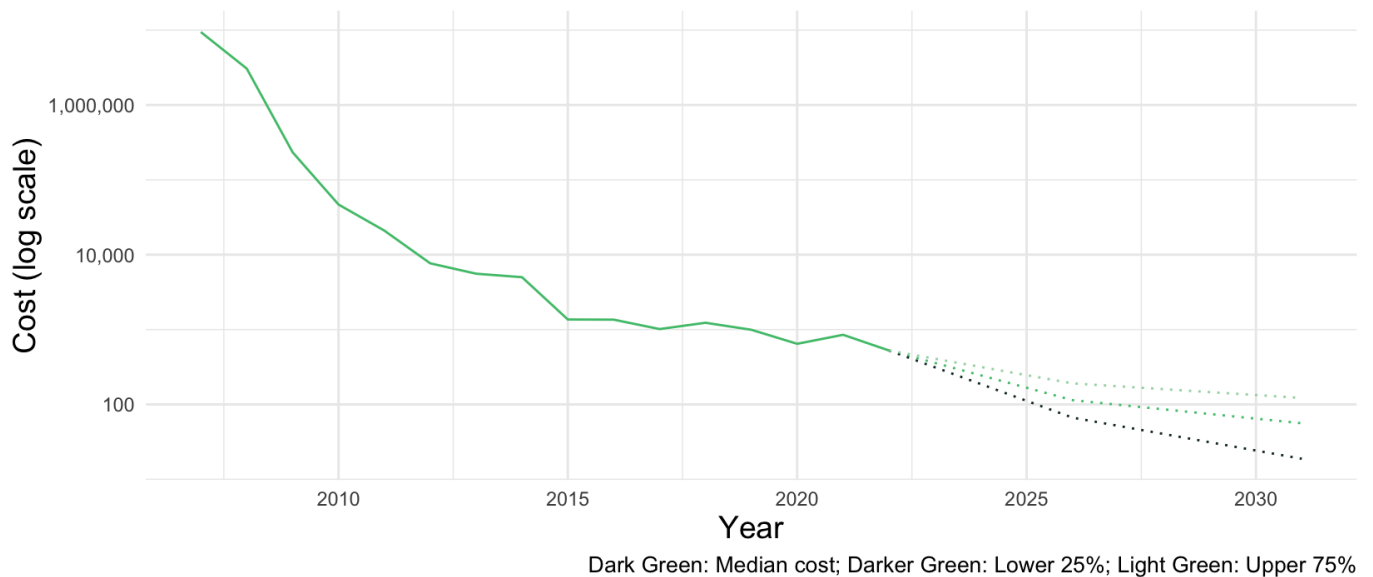
# Visualization
# Visualization
ggplot() +
  geom_line(data = df_before, aes(x = dates, y = costs), color = "#4fc27a") +
  geom_line(data = df_after, aes(x = dates, y = costs), color = "#4fc27a",
            linetype = "dotted") +
  geom_line(data = df_after, aes(x = dates, y = lower), color = "#12312d",
            linetype = "dotted") +
  geom_line(data = df_after, aes(x = dates, y = upper), color = "#a3d6b0",
            linetype = "dotted") +
  scale_y_log10(labels = scales::comma) + # logarithmic scale
  labs(title = "Cost of Sequencing a Human Genome Over Time",
       subtitle = "Combining National Human Genome Research Institute Data
with Metaculus Forecasts",
       x = "Year",
       y = "Cost (log scale)",
       caption = "Dark Green: Median cost; Darker Green: Lower 25%; Light
Green: Upper 75%") +
  theme_minimal() +
```



```
theme(plot.title = element_text(size = 18, face = "bold", family =  
"Georgia"),  
      plot.subtitle = element_text(size = 14),  
      axis.title.x = element_text(size = 14),  
      axis.title.y = element_text(size = 14),  
      plot.caption = element_text(size = 10))
```

Cost of Sequencing a Human Genome Over Time

Combining National Human Genome Research Institute Data with Metaculus Forecasts



Philanthropic Funding for Biosecurity

Description: This is an R script to analyze a dataset of biosecurity funders using Philanthropic Biosecurity Funding data, exported as a .csv file.

```
# Load the required packages  
library(tidyverse)  
library(scales)
```



```
# Set the working directory
setwd("~/Desktop/r_data_ruhl")

# Read the CSV file
biosecurity_data <- read.csv("biosecurity.csv", stringsAsFactors = FALSE)

# Convert 'Amount' column to numeric (remove '$' and ',')
biosecurity_data$Amount <- as.numeric(gsub("\\$|,", "",
biosecurity_data$Amount))

# Sum amounts by year for all funders
total_funding_by_year <- biosecurity_data %>%
  group_by(Year) %>%
  summarise(Total = sum(Amount))

# Filter data for years from 2015 up to 2022
total_funding_by_year <- total_funding_by_year %>%
  filter(Year >= 2015 & Year <= 2022)

# Calculate the error for each year (100% of the Total amount)
total_funding_by_year$error <- total_funding_by_year$Total

# Create bar chart
ggplot(total_funding_by_year, aes(x = Year, y = Total)) +
  geom_bar(stat = "identity", fill = "#a3d6b0") +

  # Label inside the green bars
  geom_text(aes(label = dollar(Total)), vjust = 1.5, color = "white") +

  # Add gray error bars
  geom_errorbar(aes(ymin = Total, ymax = Total + error), width = 0.25,
color = "gray") +

  # Label above the error bars
```




```
geom_text(aes(y = Total + error, label = dollar(Total + error)), vjust =
-1, color = "gray") +

scale_y_continuous(labels = dollar) +
labs(x = "", # Remove "Year" label
      y = "Total Funding",
      title = "Philanthropic Biosecurity and Biodefense Funding",
      subtitle = "Error bars indicate under-counting by up to 50%.
Excluding direct response (e.g. COVID relief), and excluding HIV-related
funding",
      caption =
        "Note delayed funding spikes after Ebola, Zika, and after
COVID-19.
        Data compiled in linked spreadsheet from Peace and Security
Funding Index and other grants databases.") +
theme_minimal() +
theme(
  plot.title = element_text(size = 22, face = "bold", family = "Georgia")
)
```



Philanthropic Biosecurity and Biodefense Funding

Error bars indicate under-counting by up to 50%. Excluding direct response (e.g. COVID relief), and excluding HIV-related funding



Note delayed funding spikes after Ebola, Zika, and after COVID-19.
Data compiled in linked spreadsheet from Peace and Security Funding Index and other grants databases.



About Founders Pledge

Founders Pledge is a global nonprofit empowering entrepreneurs to do the most good possible with their charitable giving. We equip members with everything needed to maximize their impact, from evidence-led research and advice on the world's most pressing problems, to comprehensive infrastructure for global grant-making, alongside opportunities to learn and connect. To date, they have pledged over \$10 billion to charity and donated more than \$950 million. We're grateful to be funded by our members and other generous donors. founderspledge.com