

The Ethics of Global Catastrophic Risk from Dual-Use Bioengineering

Seth D. Baum and Grant S. Wilson
Global Catastrophic Risk Institute
<http://sethbaum.com> * <http://gcri.org>

Ethics in Biology, Engineering and Medicine, 4(1):59-72 (2013). This version: 7 January 2014.

Abstract

Global catastrophic risks (GCRs) are risks of events that could significantly harm or even destroy civilization at the global scale. GCR raises a number of profound ethical issues, with a range of ethical theories suggesting that GCR reduction should be society's top priority. This paper discusses GCR ethics in the context of dual-use bioengineering: bioengineering that can cause either benefit or harm, including increases and decreases in GCR. Advances in bioengineering offer great promise, but also introduce new perils. Key ethical questions include what phenomena hold intrinsic value and how the phenomena are valued across space and time. Another key question is how decisions about bioengineering risks should be made. The global scope of bioengineering and GCR suggests a role for international law. Bioengineering does not fall neatly within existing international regimes such as the Convention on Biological Diversity, Cartagena Protocol, and Biological Weapons Convention. An international regime with comprehensive coverage of bioengineering would help address dual-use bioengineering as it relates to GCR.

Keywords: bioengineering; consequentialism; dual-use technology; ethics; global catastrophic risk; international law; risk analysis

I. INTRODUCTION

This paper looks at global catastrophic risks (GCR) from bioengineering from both an ethical and international law perspective, ultimately concluding that those involved in bioengineering research and regulation should make reducing GCR from bioengineering their top priority. GCRs are risks of events that could significantly harm or even destroy civilization at the global scale. In the extreme case, this can include human extinction. Common GCRs include climate change, nuclear war, pandemics, and disruptive emerging technologies, including certain forms of bioengineering. GCR is increasingly recognized as an important ethics topic, especially within the context of consequentialist ethics. A growing body of ethics research and ethics-inspired research explores the ethical issues surrounding GCR and the corresponding implications.¹⁻¹⁰ A basic insight is that GCR reduction is a key priority for humanity today from the standpoint of a range of basic consequentialist views.

This paper develops the ethics of GCR in the context of dual-use bioengineering, that is, bioengineering in which the same technological advance has both beneficial and harmful applications. Some bioengineering technologies have the potential to influence GCR, either by reducing the risk, increasing it, or both. Examples include bioengineered viruses, or bioengineered algae that absorb CO₂ and produce biofuels. Given the ethical significance of GCR, these technologies could and should be evaluated heavily in terms of their contribution to GCR. At first glance, a technology that causes a net decrease in GCR should be promoted, whereas a technology that causes a net increase in GCR should be discouraged. But subtleties of the ethics of GCR suggest that this basic conclusion does not always hold.

This paper fits in the broader context of more general calls for research to be evaluated in terms of its consequences for the world at large. This has been referred to as *extrinsic ethics*, in the sense of ethical issues that are extrinsic to the research itself.¹¹ Extrinsic ethics can be considered in terms of the National Science Foundation's Broader Impacts criterion, in which research is evaluated in terms of its broader impacts on society. In these terms, one possible broader impact of research is a change (increase or decrease) in GCR. We argue that those involved in bioengineering should make preventing a GCR their top priority. Given that advances in bioengineering are heavily driven by new research, this paper's analysis is largely oriented towards guiding and regulating bioengineering research.

The paper begins with an overview of the ethics of GCR (Section 2). The paper then discusses several ways in which bioengineering influences GCR, both by increasing and decreasing the risk (Section 3). Next, the paper discusses some options for decision-making about bioengineering, with emphasis on international law (Section 4). Finally, the paper concludes.

II. ETHICS OF GCR

The importance of GCR is fundamentally an ethics question. Different ethical views can reach different conclusions about GCR. Many views treat GCR reduction as an important goal or even a top priority for humanity today. However, some views are neutral to GCR, and some views would even seek to *increase* GCR. Here we review key ethical issues related to GCR and argue that GCR reduction should be the top priority for people alive today.

The views most closely associated with GCR are certain forms of consequentialism,¹² including common variants of utilitarianism.¹ Indeed, the concept of risk itself is commonly understood in consequentialist terms: Risk is commonly defined as the possibility of an event that has bad consequences. (But there are deontological conceptions of risk.¹³) The size of a risk is typically defined as the probability of that event multiplied by the magnitude of harm associated with the consequences if the event occurs. For example the risk of a disease outbreak is the probability of the outbreak multiplied by the harms of the deaths, illnesses, and other consequences of the outbreak if the outbreak occurs.

GCRs are risks of the highest magnitude, regardless of probability. There exist several definitions of the magnitude of global catastrophe. Atkinson proposes a definition of global catastrophe as an event in which more than one quarter of the human population dies.^{14,15} Bostrom and Ćirković define global catastrophes as events causing at least X deaths or Y monetary damages, with $10^4 < X < 10^7$ and $10^{10} < Y < 10^{13}$.² Scholars have often emphasized the particular importance of global catastrophes that would cause human extinction or other permanent declines of human civilization.^{1,3,5-7}

Consequentialism comprises a broad set of ethical views. For GCR, two aspects of consequentialism merit attention: intrinsic value and space-time universalism.

A. Intrinsic Value

Intrinsic value is defined as that which is valuable for its own sake.¹⁶ In consequentialism, intrinsic value is the consequence to be promoted. For example, in utilitarianism, utility holds intrinsic value. Different variants of utilitarianism define utility differently, for example as happiness or as the satisfaction of preferences; they also differ on questions of whose utility to count, for example whether to count the utility of sentient non-human animals. Outside of utilitarianism, there are other consequentialist views such as biocentric ethics, which places intrinsic value on life itself, and ecocentric ethics, which places intrinsic value on ecosystems.

The relationship between intrinsic value and GCR depends on how GCR is defined. While GCR could be defined otherwise, existing definitions define GCR specifically in anthropocentric

terms, i.e. as a catastrophe to humans. With this definition, GCR would be considered very important if intrinsic value is placed on some property of humans – human utility, human lives, etc. – because then a global catastrophe would result in a large loss of this intrinsic value (or perhaps a large gain; on this point see below). However, a biocentric or ecocentric notion of intrinsic value could be indifferent to some GCRs. To see this, consider a pandemic that kills all humans and leaves other species intact. This event could cause minimal change in biocentric or ecocentric intrinsic value. Indeed, the event could even cause an *increase* in biocentric or ecocentric intrinsic value, in which case biocentric or ecocentric consequentialism would recommend trying to increase certain GCRs. Thus it would seem that much is at stake with what we place intrinsic value on. However, these different conceptions of intrinsic value – anthropocentric, biocentric, and ecocentric – may all point towards decreasing GCR, if they are universalist across space-time; see discussion below.

Another important consideration related to intrinsic value is whether existence is a net positive or a net negative. To clarify, let X be whatever we place intrinsic value on – human utility, ecosystems, etc. The question this is simply this: Is it better for X to exist? This might seem like an obvious question: Of course we want what we value to exist. But there are arguments to the contrary, arguments that it would be better to not exist.

Consider *negative utilitarianism*, the view that what matters is negative utility. Here the goal is along the lines of minimizing pain or suffering, perhaps inspired by a ‘do no harm’ principle. Benatar¹⁷ and Mayerfeld¹⁸ articulate views that include aspects of negative utilitarianism. But all people experience pain at least once in their lives. The only way to guarantee the absence of pain is to guarantee the absence of people, i.e., to kill everyone. The fact that this would also eliminate pleasure and happiness and all the other aspects of humanity we might consider good is irrelevant, because in negative utilitarianism, only the negative utility matters. Thus negative utilitarianism could be interpreted as concluding that we should *increase* GCR.

The same logic also applies for certain applications of the maximin principle.^{4,17} Suppose the aim is to maximize the minimum wellbeing of any member of society, which suggests that we should focus our efforts on helping the worst-off members. Suppose also that some members will have lives of net negative wellbeing, but we don’t know who this will be. In this case the minimum wellbeing would be maximized by killing everyone: This removes all lives of negative wellbeing, thereby increasing the minimum wellbeing up to zero. Alternatively, suppose the aim is to maximize the minimum level of wellbeing within each individual’s life. But everyone probably has at least some moments of negative value – moments of pain or suffering. Again, the minimum level of wellbeing could thus be increased by killing everyone. Thus maximin can also conclude in favor of increasing GCR.

These two anti-existence views – negative utilitarianism and maximin – both suffer from the problem that they recommend nonexistence no matter how much pleasure any given life might hold, as long as it also contains at least a tiny amount of pain. Would we prefer to die now, thereby denying ourselves a good life, in order to avoid a brief moment of suffering? We think not.^{19,20} Without denying the importance of pain and suffering, we would not give them total priority over everything else. Overall, we believe existence to be a net positive. The exception would be for events that cause a significant increase in total suffering, so as to outweigh what pleasures may still exist. We would consider these events, if they are at the global scale, to be global catastrophes as well.

B. Space-Time Universalism

In consequentialism, universalism can be defined as the view that any given phenomenon holds the same intrinsic value regardless of when or where it exists.²¹ For example, universalist

anthropocentric utilitarianism would value all human utility equally, regardless of where or when that utility exists. The utility of someone in India or Paraguay would hold the same intrinsic value as the utility of someone in Denmark or Senegal. Likewise the utility of someone alive in 2013 would hold the same intrinsic value as the utility of someone alive in 2014 or 2114 or even year 21114.

Strong arguments can be made in favor of universalism. The simplest is to use a ‘God’s eye view’, looking at the world from ‘above’. From this standpoint there is no reason to value anyone any more or less than anyone else. Similar is the ‘veil of ignorance’, in which decisions are made as if the decision-maker doesn’t know which member of society she will end up being.^{22,23} In this case the decision-maker is forced to be equally fair to everyone since she has an equal chance of being any of these individuals. There is also a material logic to universalism: The same material phenomenon should have the same value in each of its instantiations, because each is materially identical. Finally, one might ask: Why is location in time or space a morally relevant criterion? Why should someone be intrinsically valued any differently just because of where or when she lives? We object to such bias,^{24,25} and indeed it reminds us of other objectionable biases such as racism and sexism. We therefore conclude in favor of universalism.

Universalism is important to GCR because global catastrophes affect so many people – or ecosystems, or whatever global catastrophes are defined to affect. Suppose global catastrophes are defined in the usual anthropocentric fashion. Then global catastrophes affect millions or even billions of people throughout the world. Furthermore, global catastrophes (especially the largest ones) can affect future generations, denying many people the chance to exist. The extreme case of human extinction would kill everyone around the world and prevent everyone in all future generations from existing.

A selfish person would care about none of this loss of life, except her own. A global catastrophe would be no more important than her own death. A car accident that kills her would be equally tragic as a nuclear war from which humanity never recovers. On the other hand, a universalist would care deeply about global catastrophes. Indeed, the universalist would not care much about her own death relative to all the others dying in a global catastrophe. As long as other people continue to exist, her values would be realized. A global catastrophe would prevent many other people from continuing to exist, and so she would care deeply about this.

An important result is that GCR is considered equivalently for universalist variants of anthropocentric utilitarianism and several non-anthropocentric consequentialisms, including non-anthropocentric utilitarianism, biocentrism, and ecocentrism. As long as each ethical view places positive intrinsic value on the continued existence of Earth-life (in contrast to e.g. negative utilitarianism), then these views all reach the same conclusion about GCR, namely that preventing the permanent destruction of human civilization is the key priority for this era.

Reaching this conclusion requires looking to the very far future. First, note that life on Earth can continue to exist for about five billion more years, until the Sun becomes too hot and large for Earth life. Second, note that the rest of the galaxy could support life for many orders of magnitude longer,²⁶ if Earth-originating life is able to colonize space. A similar orders of magnitude difference exists across space between Earth and the galaxy. If we are truly universalist – if we truly do assign intrinsic value equally for all locations in space and time – then this includes spaces and times beyond that of this planet. Most of the opportunity for intrinsic value exists beyond Earth, and so space colonization becomes a crucial objective, to spread intrinsic value throughout the galaxy.

The importance of space colonization holds regardless of whether intrinsic value is placed on humans or non-humans. But presumably only humans can colonize space, as our species is the only one on the planet with the requisite advanced technology. And so, universalist consequentialism will recommend that humanity *eventually* colonize space to spread intrinsic value

throughout the galaxy. However, space colonization need not be an immediate goal. Humanity still has about five billion years left on Earth to figure space colonization out. Instead, the immediate goal should be to prevent global catastrophe, so that future generations can go on to colonize space.^{5,8} The importance of eventual space colonization provides some focus on which global catastrophes are most important, namely the global catastrophes that would prevent Earth-originating civilization from eventually colonizing space. This focus is worth bearing in mind when evaluating specific GCRs.

Finally, it should be noted that concern about GCR does not require universalism. As a simple case, imagine weighting everyone equally except yourself, and weighting yourself twice as much as everyone else. In this case GCR is only negligibly less important than under universalism. Another case involves applying a temporal discount rate (e.g. an exponential decay function) such that future times are valued less. In this case GCR can still be quite important, especially if the global catastrophe were to occur relatively soon. Similar logic holds for applying a spatial discount rate such that distant places are valued less. And so even of the case for universalism is not accepted, there still can be much concern for GCR.

III: BIOENGINEERING AND GCR

Bioengineering can be a factor in GCR, both to increase and decrease the risk. In this and other regards it is a dual-use technology, meaning that it can be used both for beneficial and harmful purposes. Bioengineering has already led to breakthroughs in fields like agriculture and medicine. However, bioengineering has also been used to produce weapons, such as the Soviet Union's efforts to engineer a highly virulent, antibiotic resistant form of anthrax.²⁷ And as biotechnology improves, its capacity to cause benefit or harm improves, too. Ongoing advances are making the technology more accessible, more affordable, and more powerful than ever. All this raises the stakes for bioengineering, including for bioengineering as a contributor to GCR.

A. Non-GCR Benefits of Bioengineering

Bioengineering has already led to an array of peaceful technologies that are pervasive in many sectors of society. In agriculture, bioengineering has led to pest resistant crops that need fewer pesticides, drought resistant crops have increased the amount of viable farming land, herbicide resistant crops that helps eradicate weeds without killing crops, and virus tolerant crops. Bioengineering has also resulted in increased crop yields, with a 2010 survey showing that genetically engineered crops increased crop yield by 6 to 30 percent.²⁸ Some crops can have increased nutritional value, such as "golden rice" that can help remedy Vitamin A deficiency for millions of children and pregnant women, although this particular food has been met with opposition and is not yet widespread on the market. Finally, the net economic benefit is massive: fueled by increased crop production, biotechnology boosted net income for farmers by \$65 billion from 1996 through 2009 and \$10.8 billion in 2009 alone.²⁹

Bioengineering has also revolutionized many aspects of the medical field. First, bioengineering has been used to create an array of medicines that were previously unavailable. Insulin (formerly available only from sheep and cows), a Hepatitis B vaccine, and a HPV vaccine are just a few of the examples. Bioengineering has also opened the door to medical treatments for previously incurable diseases, such as through gene therapy techniques, which uses genes to treat diseases. For example, one recent clinical trial successfully used gene therapy to attack the tumor cells of patients with recurrent B-cell acute lymphoblastic leukemia, and three out of five patients went into remission, with the two other patients being subject to unrelated complications.³⁰

Finally, bioengineering shows great potential to promote environmental health. For example, scientists are developing bioengineered algae that convert CO₂ into an efficient fuel source for

motor vehicles. Scientists have also developed bioengineered microbes that are customized to gobble up pollution, such as oil slicks or toxic waste, although these have not successfully entered the marketplace. Bioengineered crops can also have a positive impact on the environment because farmers can use less harmful tilling practices, fewer pesticides and herbicides, and can produce more crops on the same amount of land.

On the other hand, these benefits of bioengineering may have unknown effects on human, animal, and environmental health. For example, pesticide and drought resistant crops may spread into unintended areas and be difficult to eradicate. Furthermore, the overall environmental effect of GE crops is still unknown, and may not be known for decades, until the impact is irreversible. Bioengineered microbes in the environment could also have vast unknown environmental implications, potentially outcompeting their natural counterparts and transferring novel DNA to other microbes,³¹ although these risks likely fall short of a GCR.

B. Bioengineering as a GCR

While bioengineering has many benefits, bioengineering can also increase GCR. There are two general risks that arise from bioengineering: first, the risk of an accident involving a dangerous bioengineered organism (a biosafety issue), and second, the risk of the purposeful release of a dangerous bioengineered organism (a bioweapons issue). Biosafety is a major concern across the globe. Biosafety lapses resulted in the apparent accidental release of Foot and Mouth Disease from a leaky pipe in the UK in 2007 and several instances of SARS infections from laboratories in 2003.³² Meanwhile, scientists are now able to use bioengineering techniques to create incredibly deadly organisms, such as a genetically modified H5N1 virus (bird flu), which scientists engineered to be airborne. The natural H5N1 virus killed an alarming 60 percent of reported infected individuals,³³ which is a higher fatality rate than the 1918 influenza pandemic (Spanish flu).³⁴ Although the airborne H5N1 virus never escaped from the lab, its accidental release, or the accidental release of a similarly dangerous virus, could cause a loss of life significant enough to qualify as a global catastrophe by its normal definition.

Bioweapons are also a growing threat. The technology to create dangerous bioengineered organisms is increasingly cheap and accessible. Custom DNA strands are available for shipment online for several thousand dollars, and consumers can even purchase DNA synthesis machines online, which can “print” customized strands of DNA.³⁵ Individuals have also shown a willingness to engage in attacks using biological agents, such as the postal anthrax attacks of 2001, which killed 5 people in the United States. Furthermore, publicly available information on how to bioengineer extremely deadly viruses can be used as an instruction manual for individuals, organizations, or governments to create extremely deadly bioweapons.³⁶ So with bioengineering technologies advancing, bioterrorists or other actors could create a biological weapon that is more deadly than anything existing in the natural world. Therefore, bioweapons that utilize bioengineered organisms pose a GCR.

C. Bioengineering to Reduce other GCRs

In addition to itself being a GCR, bioengineering can also reduce the chances that other GCRs will occur. One such GCR is climate change. Catastrophic climate change scenarios could involve sea level rise of up to 10 meters, droughts, increased extreme weather events, loss of most threatened and endangered species, and temperature increases of 6 degrees Celsius.³⁷ Still worse than that would be outcomes in which large portions of the land surface on Earth become too warm for mammals (including humans) to survive.³⁸ And the worst scenario could involve climate engineering backfiring to result in extremely rapid temperature increase.³⁹

Despite the risks of climate change, the international community has struggled to satisfactorily address the issue, for a variety of political, technological, and economical reasons. Bioengineering may be able to help. An army of bioengineered algae that is specifically designed to convert carbon dioxide into a “biocrude” fuel ready to be made into fuel for any vehicle type – a technology that Craig Venter’s Synthetic Genomics, Inc. is developing with a \$600 million investment from ExxonMobil – could remove greenhouse gases from the atmosphere and provide a plentiful, carbon-neutral fuel source that does not pose many of the downsides of today’s biofuel options (although this technology has its own risks).⁴⁰ Or, despite being a bizarre proposition, humans could be genetically engineered to reduce our CO₂ output, such as by engineering humans to be intolerant to meat or to be smaller in size.⁴¹

Likewise, while a deadly bioengineered virus has the potential to escape from a laboratory and cause a global catastrophe, such research may be necessary to create vaccines for viruses that could cause worldwide pandemics. For example, the Influenza Pandemic of 1918-1919 (the Spanish flu) killed about 50 million people worldwide.⁴² Would modern bioengineering technology have been able to avoid this global catastrophe? In fact, researchers justified the airborne H5N1 virus, discussed above, as helping to prevent the spread of a similar strain that could mutate naturally. Overall, there is a dynamic relationship between bioengineering and other GCRs that should be assessed when considering how to respond to these risks.

IV. INTERNATIONAL REGULATION

International regulation of bioengineering is important for at least two reasons. First, bioengineering can influence (increase or decrease) GCR, which is an inherently global issue. Thus the entire world has a stake in ensuring that this technology is used safely. Second, bioengineering research and development can be conducted anywhere in the world. With only scattered national or regional regulation, harmful bioengineering could be conducted in unregulated areas. A deadly virus could be bioengineered by terrorists in an unregulated area, but it would have an impact on the entire planet. Thus international regulation has an important role to play in the overall management of bioengineering, as with GCR in general.

Several international treaties already attempt to regulate various GCRs. The U.N. Framework Convention on Climate Change and the Kyoto Protocol regulate climate change; the U.N. Convention on Biological Diversity (CBD) regulates biodiversity loss; the Treaty on the Non-Proliferation of Nuclear Weapons (NPT) regulates nuclear weapons; and the Biological Weapons Convention (BWC) regulates biological weapons. However, there is no international treaty specifically for bioengineering.

Nonetheless, several other treaties regulate certain aspects of bioengineering, including the CBD, the Cartagena Protocol, and the BWC. Still, the drafters of these treaties did not design them with all the risks of modern dual-use bioengineering in mind. And so these treaties provide imperfect regulations that fall well short of the comprehensive international scheme that would be necessary to minimize the GCR arising from bioengineering.⁴³

The existing treaties fall short in several ways. First, while the Convention on Biological Diversity (CBD) requires parties to “regulate, manage or control the risks associated with the use and release of living modified organisms,” the treaty does not establish specific steps that countries must take, such as requiring high-security laboratories or safety training for scientists handling dangerous bioengineered pathogens (CBD, Article 8(g)). Second, the Cartagena Protocol on Biosafety expands on the CBD, requiring parties to conduct risk assessment and risk management for living modified organisms that may adversely affect biological diversity. However, these requirements only extend as far as a country’s self-determined protective goals,⁴⁴ and the treaty is overall too trade-focused to provide sufficient regulation of bioengineering. Third, the Biological

Weapons Convention takes measures to ensure that countries do not stockpile, acquire, or retain dangerous microbial or biological agents, including dangerous bioengineered organisms, but there is an exception for “prophylactic, protective or other peaceful purposes” (BWC, Article 1), so while the BWC may be useful in preventing acts of bioterrorism, the treaty does not sufficiently cover biosafety issues. Overall, while these treaties regulate certain risks arising out of bioengineering, they fail to provide a comprehensive regulatory regime that addresses all of the risks that arise from dual-use bioengineering technologies.

A. International Legal Regime

There are a variety of options of how to create an international legal regime for bioengineering. The first is to create a framework treaty that covers an assortment of emerging technologies such as bioengineering, nanotechnology, artificial intelligence, and geoengineering. Having one overarching governance scheme is perhaps ideal because of the fundamental similarities amongst emerging technologies: they are all dual-use technologies; they all have risks that increase with research and development; they call for similar forecasting techniques; and laboratory transparency is central to their governance—just to name a few.⁴⁵

Another approach to create a treaty is to begin by first creating “soft law.” Soft laws are nonbinding laws, meaning that countries have discretion over whether to follow them.⁴⁶ Hard laws, on the other hand, are legally binding, such as treaties. This approach would involve creating soft laws, either for bioengineering or emerging technologies in general, until there is the necessary political will for a legally binding treaty, at which point the soft law could be integrated into the treaty.⁴⁷ Beginning with soft laws is a prudent approach when countries are unlikely to otherwise agree to a legally binding treaty. However, soft laws do not provide the sort of robust protections that may be necessary to prevent a GCR from bioengineering.

Finally, another option is to create a new international organization dedicated to mitigating GCR and other risks from emerging technologies, including bioengineering. This international organization could be established by an existing international body like the U.N. General Assembly, the World Health Organization (WHO), the Economic and Social Council (ECOSOC), the United Nations Environment Programme (UNEP), the U.N. Educational, Scientific and Cultural Organization (UNESCO), the Organisation for Economic Co-operation and Development (OECD), or the Strategic Approach for International Chemical Management (SAICM), for example. Perhaps the best option is for several existing international organizations to jointly create and/or chair a new body; for example, UNEP and the WHO co-chair SAICM. A new organization that covers emerging technologies could attempt to quickly begin treaty negotiations or else create soft law that is later embodied in an international treaty, as described above.

B. Universal Treaty

Finally, a treaty that regulates bioengineering (or several emerging technologies) should have universal participation, meaning that every country is a party to the treaty. As Charney⁴⁸ argues, global threats like environmental degradation, climate change, terrorism, and nuclear weapons require treaties that are binding on all countries, or else the unregulated countries could impede with the purpose of the treaty. For bioengineering, universal participation is important because otherwise dangerous bioengineering technologies could be freely developed in the non-signatory countries, putting every other country in the world at risk. This type of situation threatened to play out when physicist Richard Seed threatened to clone a human in Japan or Mexico if the technique was made illegal in the United States. If near-universal participation is achieved, another option is to impose sanctions on countries that do not comply with the terms of the treaty,

much like the Security Council imposes sanctions on countries that violate certain provisions of the NPT.

Another reason that a treaty on bioengineering should have universal participation is because this would help prevent an arms race for dangerous biotechnologies. According to Pinson,⁴⁹ an arms race could be triggered by advanced nanotechnologies that assemble novel biological or chemical compounds to be used as weapons. And Metzl⁵⁰ argues that the genetic manipulation of humans could spark an arms race to create populations that are smarter, physically superior, and more advanced in other ways. A global treaty with universal participation may help prevent similar arms races involving extremely dangerous and powerful bioengineering technologies, because all countries would be subject to the same obligations and therefore would not be threatened with comparative disadvantages.

Of course, achieving universal participation in a treaty is extremely difficult: Most of the international community suffers from treaty fatigue, and developing countries have reservations about diverting even more resources to implement treaties. Furthermore, an outright ban on dangerous technologies does not always work, because this could hand over a monopoly to rogue actors⁴⁹ and because countries want make beneficial use of dual-use technologies, including as defensive mechanisms against noncompliant countries.⁵¹ Despite these hurdles, forces such as the threat of an imminent global catastrophe from bioengineering or massive public pressure could create the requisite political will for countries to conclude a global treaty on bioengineering or all emerging technologies.

C. Precautionary Principle

If the international community does decide to create an international treaty on bioengineering or all emerging technologies, they should consider the precautionary principle (or the “precautionary approach”) as a tool to reduce GCR.⁵² There are many interpretations of the precautionary principle, but the crux of the meaning is that precautionary measures should be taken in the face of scientific uncertainty, or, inversely, that scientific uncertainty should not be an excuse for forgoing precautionary measures. Applied to bioengineering, this could require, for example, that scientists develop dangerous bioengineered pathogens only once they can affirmatively demonstrate that there is a satisfactorily small possibility that the pathogens will escape from the laboratory. Similarly, a treaty could impose moratoriums on certain bioengineering applications, such as the release of bioengineered microbes into the environment, until proponents of these technologies can satisfactorily prove their safety or, at minimum, that they will not cause a global catastrophe.

V. CONCLUSION

Advances in bioengineering significantly increase humanity’s capacity to affect itself and the world it lives in – for better or for worse. Bioengineering could revolutionize medicine and help slash global greenhouse gas emissions, which would allow humans to live long, healthy lives on a clean planet. Or, it could cause an unprecedented pandemic from a biosafety lapse or bioterrorist attack, which could devastate the global population and have other consequences like societal collapse.

This paper develops GCR as a key consideration for dual-use bioengineering. According to a range of consequentialist ethics views, reducing GCR should be the top priority for society today. However, other ethics views do not consider GCR to be a priority. Indeed, some views even suggest increasing GCR. The question of how bioengineering should proceed is fundamentally an ethics question, with the issue of GCR being a major factor.

Given the global nature of bioengineering and of GCR, regulating bioengineering should be a global endeavor. Several existing international legal structures can help, including the CBD, the Cartagena Protocol, and the BWC. However, these structures are insufficient to cover the entirety of bioengineering GCR. New international laws could help with this, whether applied solely to bioengineering or all emerging technologies, but creating them could be difficult. In the meantime, other options include soft law and new international organizations, either of which could eventually transition into a new international treaty.

Perhaps no formal governance structure could ever offer complete protection from bioengineering-related GCRs. Only the most intrusive surveillance regime could hope to watch over every decision made by every researcher; such a regime would itself pose many ethics issues and could even lead to global catastrophe.⁵³ In the absence of such a regime, it is incumbent that everyone in the bioengineering community – researchers, institutional homes, funders, and others – keep all the implications of bioengineering in mind, making preventing a GCR from bioengineering their top priority. The survival of human civilization could well be at stake.

Acknowledgments

We thank Tony Barrett, Jacob Haqq-Misra, Karim Jebari, and Marc Saner for helpful comments on an earlier draft of this paper. Any remaining errors or shortcomings are our own.

References

1. Leslie J. *The end of the world: the science and ethics of human extinction*. London: Routledge; 1996.
2. Bostrom N, Ćirković M. *Global catastrophic risks*. Oxford: Oxford University Press; 2008.
3. Bostrom N. Existential risk prevention as a global priority. *Global Policy* 2013 Feb; 4(1): 15–31.
4. Baum SD. Better to exist: a reply to Benatar. *J Med Ethics* 2008 Dec; 34(12): 875-876.
5. Baum SD. Is humanity doomed? Insights from astrobiology. *Sustainability* 2010 Feb; 2(2): 591-603.
6. Maher TM Jr., Baum SD. Adaptation to and recovery from global catastrophe. *Sustainability* 2013 April; 5(4): 1461-1479.
7. Matheny JG. Reducing the risk of human extinction. *Risk Anal* 2007 Oct; 27(5): 1335-1344.
8. Tonn BE. Distant futures and the environment. *Futures* 2002 March; 34(2): 117-132.
9. Tonn BE. Obligations to future generations and acceptable risks of human extinction. *Futures* 2009 Sep; 41(7): 427–435.
10. Posner R. *Catastrophe: risk and response*. Oxford: Oxford University Press; 2004.
11. Schienke EW, Tuana N, Brown DA, Davis KJ, Keller K, Shortle JS, Stickler M, Baum SD. The role of the NSF Broader Impacts Criterion in enhancing research ethics pedagogy. *Soc Epistemology* 2009 July-Dec; 23(3-4): 317-336.
12. Sinnott-Armstrong W. Consequentialism. *Stanford Encyclopedia of Philos* 2011. Available from: <http://plato.stanford.edu/entries/consequentialism>.
13. Hansson SO. Ethical criteria of risk acceptance. *Erkenntnis* 2003 Nov; 59(3): 291-309.
14. Atkinson A. *Impact Earth: asteroids, comets and meteors—the growing threat*. London: Virgin; 1999.
15. Hempsell CM. The potential for space intervention in global catastrophes. *J Br Interplanetary Soc* 2004; 57: 14-21.
16. Baum SD. Value typology in cost-benefit analysis. *Environ Values* 2012 Nov; 21(4): 499-524.
17. Benatar D. *Better never to have been: the harm of coming into existence*. Oxford: Clarendon Press; 2006.

18. Mayerfeld J. *Suffering and moral responsibility*. Oxford: Oxford University Press; 2005.
19. Smart R. Negative utilitarianism. *Mind* 1958 Oct; 67(268): 542–543.
20. Ord T. Why I'm not a negative utilitarian. 2013. Available from: <http://www.amirrorclear.net/academic/ideas/negative-utilitarianism>.
21. Baum SD. Universalist ethics in extraterrestrial encounter. *Acta Astronautica* 2010 Feb-Mar; 66(3-4): 617-623.
22. Harsanyi JC. Cardinal utility in welfare economics and in the theory of risk-taking. *J Polit Econ* 1953 Oct; 61(5): 434–435.
23. Rawls J. *A theory of justice*. Cambridge, MA: Belknap Press; 1971.
24. Cowen T, Parfit D. Against the social discount rate. In: Laslett P, Fishkin J, editors. *Justice between age groups and generations*. New Haven: Yale University Press; 1992. p. 144-161.
25. Stern N. The economics of climate change. *Amer Econ Rev* 2008 May; 98(2): 1-37.
26. Lineweaver CH, Fenner Y, Gibson BK. The galactic habitable zone and the age distribution of complex life in the Milky Way. *Science* 2004 Jan 2; 303(5654): 59-62.
27. Tucker J. Biological weapons in the former Soviet Union: an interview with Dr. Kenneth Alibek. *Nonproliferation Rev* 1999 Spring-Summer: 1-10.
28. European Association for Bioindustries. GM crops: reaping the benefits, but not in Europe. 2011 Available from: www.europabio.org/sites/default/files/europabio_socioeconomics_may_2011.pdf.
29. Brookes G, Barfoot P. The income and production effects of biotech crops globally 1996–2009. *Int J Biotechnol* 2012 Oct-Dec; 3(4): 265-272.
30. Brentjens RJ, Davila ML, Riviere I, Park J, Wang X, Cowell LG, Bartido S, Stefanski J, Taylor C, Olszewska M, Borquez-Ojeda O, Qu J, Wasielewska T, He Q, Bernal Y, Rijo IV, Hedvat C, Kobos R, Curran K, Steinherz P, Jurcic J, Rosenblat T, Maslak P, Frattini M, Sadelain M, 2013. CD19-targeted T cells rapidly induce molecular remissions in adults with chemotherapy-refractory acute lymphoblastic leukemia. *Sci Translational Med* 2013 March; 5(177): 177ra38.
31. Dana GV, Kuiken T, Rejeski D, Snow AA. Four steps to avoid a synthetic-biology disaster. *Nature* 2012 March 1; 483(7387): 29.
32. Gaudioso J, Gribble LA, Salerno RM. Biosecurity: progress and challenges. *J Lab Automation* 2009 June; 14(3): 141-147.
33. Centers for Disease Control and Prevention. Highly pathogenic avian influenza A (H5N1) in people. 2012 June 21. Available from: <http://www.cdc.gov/flu/avianflu/h5n1-people.htm>.
34. Gottron F, Shea DA. Publishing scientific papers with potential security risks: issues for Congress. *Congressional Research Service* 2013 March 18; p. 5. Available from: www.fas.org/sgp/crs/secretary/R42606.pdf.
35. Garfinkel S, Friedman RM. Synthetic biology and synthetic genomics. In: Leary D, Pisupati B, editors. *The future of international environmental law*. Tokyo: United Nations University Press; 2010. p. 269-291.
36. Berns KI, Casadevall A, Cohen ML, Ehrlich SA, Enquist LW, Fitch JP, Franz DR, Fraser-Liggett CM, Grant CM, Imperiale MJ, Kanabrocki J, Keim PS, Lemon SM, Levy SB, Lumpkin JR, Miller JF, Murch R, Nance ME, Osterholm MT, Relman DA, Roth JA, Vidaver AK, 2012. Adaptations of avian flu virus are a cause for concern. *Nature* Feb 9; 482(7384): 153-154.
37. Schneider SH. The worst-case scenario. *Nature* 2009 April 30; 458: 1104-1105.
38. Sherwood SC, Huber M. An adaptability limit to climate change due to heat stress. *Proc Natl Acad Sci USA* 2010 May 25; 107(21): 9552–9555.
39. Baum SD, Timothy M. Maher, Jr., and Jacob Haqq-Misra. Double catastrophe: Intermittent stratospheric geoengineering induced by societal collapse. *Environ Systems Decisions* 2013 Mar; 33(1): 168-180.

40. Schmidt CM. Synthetic biology environmental health perspectives of a new field. *Environ Health Perspect* 2010 March; 118(3): A118-A123.
41. Liao SM, Sandberg A, Roache R. Human engineering and climate change. *Ethics, Policy & Environ* 2012 June; 15(2): 206:221.
42. Taubenberger JK, Morens DM. 1918 Influenza: the mother of all pandemics. *Emerging Infect Dis* 2006 Jan; 12(1): 15-22. Available from: <http://wwwnc.cdc.gov/eid/article/12/1/pdfs/05-0979.pdf>.
43. Wilson G. Minimizing global catastrophic and existential risks from emerging technologies through international law. *Va Environ Law J*. In press 2013.
44. Report of the Third Meeting of the Ad Hoc Technical Expert Group On Risk Assessment and Risk Management Under the Cartagena Protocol on Biosafety. U.N. Doc. UNEP/CBD/BS/AHTEG-RA&RM/3/4; 2011. Available from: <http://www.cbd.int/doc/meetings/bs/bsrarm-03/official/bsrarm-03-04-en.pdf>.
45. Baum S. Seven reasons for integrated emerging technologies governance. Institute for Ethics & Emerging Technologies 2013 Jan 23. <http://ieet.org/index.php/IEET/more/baum20130123>.
46. Grnchalla-Wesierski T. A framework for understanding “soft law”. *McGill Law J*. 1984; 37-88.
47. Abbott KW, Marchant GE, Sylvester DJ. A framework convention for nanotechnology? *Environ Law Reporter* 2006 Aug; 38(8):10507.
48. Charney JI. Universal international law. *Am J Int Law* 1993 Oct; 87(4): 529-551.
49. Pinson RD. Is nanotechnology prohibited by the biological and chemical weapons conventions? *Berkeley J Int Law* 2012; 22(2): 279-309.
50. Metzl J. Brave new world war. *Democracy* 2008 Spring; 8. Available from: <http://www.democracyjournal.org/8/6586.php?page=all>.
51. Hughes JJ. Global technology regulation and potentially apocalyptic technological threats. In: Allhoff F, Lin P, Moor J, Weckert J, editors. *Nanoethics: the ethical and social implications of nanotechnology*. Hoboken, NJ: John Wiley; 2007. p. 201-214.
52. Saner MA. An ethical analysis of the precautionary principle. *Int J Biotechnology* 2002; 4(1): 81-95.
53. Caplan B. The totalitarian threat. In: Bostrom N, Ćirković M, editors. *Global catastrophic risks*. Oxford: Oxford University Press; 2008. p. 504-519.