

The bioethicist who cried 'Synthetic biology' – An analysis of the function of bioterrorism predictions in bioethics

Søren Holm

Centre for Social Ethics and Policy, School of Law, University of Manchester, UK

&

Centre for Medical Ethics, HELSAM, Oslo University, Norway

&

Institute for Health Science and Technology, Aalborg University, Denmark

soren.holm@manchester.ac.uk

Abstract

This paper analyses a spectre that has haunted bioethics almost since its inception, namely the spectre of the misuse of biotechnology by maleficent agents bent on mass destruction or the complete eradication of human kind and life as we know it.

It provides a general account of why bioethicists cry 'catastrophic bioterrorism potential' when new biotechnologies emerge, and an analysis of the arguments that flow from the prediction, especially in relation to synthetic biology.

Key words

Bioethics, Bioterrorism, Bioterrorism argument, Catastrophic risk, Existential risk, Recycling of arguments, Slippery slope, Synthetic biology

Introduction

Many, perhaps most biotechnologies have potential for 'dual use'. They can be used to great benefit but they can also be used deliberately to cause harm. Some biotechnologies do not only have dual use possibilities in the hands of state actors, but also have the theoretical potential to be deliberately used by non-state actors to develop novel infectious organisms, or to 'resurrect' old ones like smallpox or the 1918 Spanish flu. Let us call the persons who attempt to develop such organisms 'bioterrorists', and their final goal 'bioterrorism'. Bioterrorists may target other living organisms than humans, but we will here focus on humans as targets. Depending on the characteristics of the pathogen produced and the method of dissemination the potential/predicted effects of releasing it can range from very limited to the complete extinction of human kind. In the literature on catastrophic risk an existential risk is defined as the risk of extinction of human kind, and a catastrophic risk as the risk of an event leading to the death of 10% or more of the global population (1).

The prediction of significant bioterrorist use of a particular biotechnology has been a recurrent theme in bioethics over the years, and it has been raised again in relation to synthetic biology (2-6). The general structure of the argument and the content of the premises supporting the predictions of imminent bioterrorism and the need for bioethicists to take this risk seriously (let us call them the 'bioterrorist argument' and the 'bioterrorism problem' respectively) have been remarkably stable over the years and can be reconstructed in outline as follows:

P1 There are evil people in the world bent on mass destruction

P2 One of the ways in which to achieve mass destruction is through bioterrorism

P3 Technology X will make it easier/easy for bioterrorists to make a dangerous pathogen

P4 Technology X will make it cheaper for bioterrorists to make a dangerous pathogen

P5 Technology X will make it easier for bioterrorists to hide from the authorities while developing a dangerous pathogen

Therefore

C1 It is likely that Technology X will be used by bioterrorists to achieve their nefarious aims

P6 A significant bioterrorist attack is a very bad thing

Therefore by C1 & P6

C2 Bioethicists ought to take this risk very seriously

I take it that P1, P2 and P6 are uncontroversial and true, and that the soundness of the argument therefore hinges on the truth of P3-P5 and the validity on the inference from P3-P5 to the likelihood of bioterrorism actually happening.

The argument from P1 - P5 to C1 is generic to debates about bioterrorism, and the only bioethical contribution is the introduction of P6 as the justification for C2. The bioterrorist argument is closely related to an argument with premises predicting a possible accidental release of novel pathogenic microorganisms created during scientific experiments, and to similar arguments focusing not on biotechnology, but on nanotechnology. The predictions of a likely major catastrophe through the use of biotechnology and the similar nanotechnology based arguments have already formed the basis for Science Fiction novels in the biopunk and nanopunk sub-genres (most famously perhaps in Michael Crichton's novel 'Prey' (7)).

The generic bioterrorism argument is quite old, and has been applied to technologies as simple as bacterial culture of unmodified pathogenic bacteria (8). Simple bacterial culture was also the technique used in the only verified bioterrorist attacks aimed at causing mass casualties. These, completely unsuccessful attacks were carried out using aerosolised anthrax by the Aum Shinrikyo cult in Japan in 1993 (9). They were so unsuccessful that they were only identified later when the cult was investigated by the police after its successful Sarin gas attack in the Tokyo Metro in 1995.

In relation to synthetic biology the focus of the argument has been the putative malevolent biohacker doing DIY synthetic biology in his or her garage. We are asked to imagine "... a world where practically anybody with an average IQ would have the ability to create novel organism in their home garage" (10, p 2). The biohacker has been endowed with very considerable powers in the literature. Not only is s/he able to design and synthesise (or edit) the DNA or RNA of the novel pathogen, s/he is also able to culture it in the necessary quantities, weaponise it as an aerosol or a powder, package it in a suitable delivery vehicle and spread it, all within the enclosure of the suburban garage and without becoming infected and without anyone becoming suspicious. All of these powers of the biohacker are to some extent mythical in that they are not based in the current reality of lab work needed to perform any meaningful synthetic biology (11).

More recently a version of the argument has appeared where the malevolent agent is no longer a human being, but a superintelligent artificial intelligence (see the "AI takeover scenario" in (12)).

In the bioethics literature it is not only the structure of the bioterrorist argument that is recycled, it is also often the specific examples given to give credence to the premises. The following quote from a paper by Douglas and Savulescu is typical in this regard (original references removed):

"A third concern about synthetic biology has been that the knowledge it produces might be intentionally misused. Two ground breaking studies last decade reported the de novo synthesis of human pathogens—the poliovirus, and the 1918 Spanish influenza virus. These studies raise the prospect that states or terrorist cells without access to wild-type viruses might acquire the ability to create human pathogens for use in biological weapons. Further

into the future, synthetic biology may enable the creation of novel pathogens more deadly than any known before. Scientists have already accidentally created a vaccine resistant strain of the mousepox virus through more traditional genetic engineering techniques. This killed 100% of mice infected. The mousepox virus is not a human pathogen, but the same technique might also render viruses that do affect humans, such as smallpox, vaccine resistant. Prior to eradication by vaccination, smallpox was the scourge of humanity killing about a third of all affected people in a horrible manner. If the results of the mousepox experiment applied to smallpox, it would be possible to significantly increase its lethality and render it vaccine resistant. If the incubation period, almost ideal already for terminal pandemic, could be further increased, it may become possible to produce a bioweapon capable of devastating humanity. The virus could be released at airport terminals, stadiums, train stations and public events. Thousands could potentially be infected at such loci and the virus might be spread during the incubation period all around the world, before any pandemic could be detected. " (2, p 689)

The polio virus, Spanish flu and mousepox examples can also be found in many other papers, as can the expansion of the actual mousepox research to create a full human smallpox bioterrorism scenario (see, for instance 4-6). Both the argument and the examples used to lend credence to it has thus become a stock, routinely applied item in the bioethics literature. Why has this happened? Are there functions of the bioterrorist argument that can help to explain its attractiveness?

The functions of the argument

What role does the bioterrorist argument play in bioethics? A pragmatic role played by the argument is that the continuing development of new technologies to which the argument can be applied allows the middle aged or older bioethicist to recycle arguments in an eternal recurrence (13, 14). The arguments used in discussing the bioterrorist potential of recombinant DNA in the 1970s can with very limited modification be recycled to write a new article on synthetic biology or gene editing (15, 16). And, every new emerging terrorist group can become the occasion for reconsidering whether that specific group will engage in bioterrorism (17).

But the bioterrorist argument can be used for three more substantial purposes as well. It can be used:

1. as the basis for a precautionary argument in relation to the development of a particular technology,
2. as a marker for how important the argument that follows is (almost irrespective of what that argument is); and

3. as a bogeyman frightening people into agreeing with your arguments and solutions (almost irrespective of what your arguments and solutions are)

The first of these uses is the most straightforward and legitimate. If it is really true that synthetic biology makes catastrophic or for that matter non-catastrophic bioterrorism likely, then we have a good reason to proceed with caution in the development of synthetic biology. However, since synthetic biology is an umbrella term covering a wide range of different activities we need to be careful when we target our (pre)caution. There are activities and developments within the field of synthetic biology that are very unlikely to make bioterrorism easier/cheaper/more likely and our regulatory efforts have to be precisely targeted to avoid hampering developments in the 'safe' areas. I have previously written on the problems inherent in the precautionary principle, so will not expand on that here (18-20).

The second use of the bioterrorism argument is a version of the much more general rhetorical figure of 'my X is bigger than your X', in this case completed as:

'I am thinking about the really big problem of bioterrorism - we could all go extinct! - while you are just thinking about small problems like justice between the first and the third world or the euthanasia of people with dementia, so I am (or perhaps slightly more modestly my arguments are) much more important than you(rs)'

This claim of comparative importance is then often linked to the claim 'that apart from my work very little work is being done on the bioterrorism problem' and/or 'more resources should be allocated to the bioterrorism problem', i.e. someone should make grants available to me (Douglas, Daniel). There is a significant amount of work on the general dual use problem raised by biotechnologies (see for instance the substantial body of work by Seumas Miller and Michael J Selgelid individually and jointly), so both claims about too little work being done are empirically dubious.

The philosophical problems with this use of the bioterrorism argument are that 1) it can be very difficult to correctly assess the comparative importance of the analysis of different philosophical questions, and 2) that even if we could establish comparative importance there is a fallacy involved in eliding the importance of a particular question and the importance of my thinking and writing about that question (the question could be genuinely important, but what I write about it could be crap!).

What is the best analysis of the comparative importance of philosophical questions, and what does the best analysis tell us about the importance of the bioterrorist argument?

Let us first consider what we mean by saying that a philosophical question is important *simpliciter* before tackling the issue of comparative importance. One initial relevant distinction is whether:

1) a question is important because simply trying to analyse and answer it is likely to lead to major philosophical advances (Philosophical importance),

2) a question is important because the answer to the question is important (Practical importance), or

3) a question is important because we urgently need an answer (Urgency).

A question may fall into the first category either because analysing is likely to lead to a major advance within this particular field of philosophical inquiry, or because analysing it is likely to lead to advances in a number of fields because it is a very fertile question. And, a question may fall within the second category either because an answer will have significant practical implications, or because it will require all of us, including non-philosophers to reconsider commonly held assumptions about ourselves or our place in the world. Many philosophical questions will be important in all three senses, but others will only be important in one of them.

The claim that a philosophical question can be important because of its fertility is very similar to Ryle's claim that questions of principle are important:

"A philosophical question is a question of principle when it is philosophically much more important than most other questions. And the relative importance of philosophical questions could be explained on these lines that when, given the answer to one question, it is at once clear what are the answers, or of what sort are the answers, to an expanding range of other questions, while the answers to any of the latter do not in the same way throw light on the former, then the former is a question of principle relative to the latter."

(21, p 327)

How should we categorise the philosophical questions that arise if the bioterrorism argument is accepted? The putative bioterrorism problem created by synthetic biology does not seem to be important in the first sense, analysing it or even solving it is unlikely to lead to any major philosophical advance which could not be achieved simply by considering the undramatic, general 'dual use' problem, and the question does not seem particularly fertile either.

If the bioterrorism argument is valid and sound the bioterrorism problem does, however seem to be a *prima facie* important problem in the two other ways identified. The answer is practically important, we want to prevent catastrophe or extinction, and given that bioterrorism is claimed to be imminent the answer also seems to be urgent. But how important it is comparatively in relation to other problems crucially depends on the real probability of bioterrorism and the likely magnitude of bioterrorist events. If Aum Shinrikyo had been successful in 1993 they would have infected hundreds, perhaps even thousands of people with anthrax and would have killed a large proportion

of them, but even though that would have been very successful bioterrorism it is far from being a catastrophic or existentially threatening event. Even if we postulated as likely a future scenario where bioterrorism killed an average of 100,000 people every year, the bioterrorism problem would be fairly low down the list of important problems that bioethicists could spend their time analysing. 842,000 people, of which 361,000 are children die from lack of access to clean drinking water every year and 303,000 women die from preventable complications in childbirth (22, 23). It is thus only if bioterrorism is a catastrophic or existential risk that it is obviously a more important question than many other questions in bioethics.

The third use of the bioterrorism argument is perhaps the most problematic, because it can have real life consequences. The spectre of the bioterrorist ready to unleash microbial havoc on the unsuspecting world can be and is used to argue for the necessity for wide ranging and privacy infringing surveillance and draconian legal powers (24), or for the necessity for (compulsory?) moral enhancement (25). The claimed likelihood of an imminent attack and the predicted enormous magnitude of the harm caused combine to make what looks like an unassailable pragmatist or consequentialist case for strong and effective intervention. Almost any loss of welfare / utility caused by the intervention can be outweighed by the continued benefit of human existence.

But is all what it seems? In order for a bioterrorist attack to be catastrophic in the sense defined above it has to be at least as successful as the most 'successful' historical, natural epidemics (1), and about twice as deadly as the 1918 Spanish Flu (26). So we are setting the bar quite high for true catastrophic success, and this of course entails that the likelihood of achieving success becomes correspondingly lower. Even a single attack causing 1 million deaths, thus making it the most successful single terrorist attack ever by some considerable margin, is three orders of magnitude away from being catastrophic. And if there was such an attack, it would immediately lead to very, very significant regulation, surveillance and enforcement that would make the 'War on Terror' pale into insignificance, thereby significantly decreasing the likelihood of any further bioterrorism. What we therefore have to postulate is thus not only the possibility of bioterrorism developing towards being a catastrophic or existential risk, but that that development happens before any major, but non-catastrophic bioterrorism attack has taken place. This is possible, but not very likely.

These considerations also point to an inherent issue with the bioterrorism argument, i.e. that it does not account for the government actions that would follow from a significant, but not catastrophic bioterrorism attack. If there was such an attack governments would immediately act to ensure that P3-5 were no longer true. The bioterrorism argument is thus only sound as long as no successful, significant bioterrorism has occurred.

Now, a bioterrorist attack leading to 1 million deaths is, of course still 'catastrophic' in the normal sense of this term and something that governments should try to prevent. But given that it would likely be a one off event a serious cost-benefit analysis would have to be made to estimate whether the ongoing costs imposed on citizens by a particular regulatory and surveillance regime were outweighed by the benefits.

But, what about the non-zero risk of a truly catastrophic or existentially threatening bioterrorist attack? Should we not take that into account?

Here it is important to distinguish between the catastrophic and the existential risk scenarios. Let us assume that a bioterrorist group managed to mount a catastrophic attack killing 10% of world population. This is, obviously a very bad outcome involving a large immediate loss of welfare, but the underlying causation is time limited and does not leave any longer term environmental problems unlike some other postulated catastrophic risk scenarios (e.g. asteroid strikes, super-volcano eruptions and global warming (1)). This means that the longer term loss of welfare / utility caused by a catastrophic bioterrorist attack will be limited. We have good evidence that near catastrophic epidemics like the Spanish Flu have limited long term economic effects (27,28). A catastrophic bioterrorist attack will thus lead to a very large, but finite loss of welfare / utility. Because the loss is finite, this in turn entails that it does not generate a blank cheque in so far as prevention is concerned. The long term costs of prevention have to be assessed and balanced against the benefits of preventing a catastrophic bioterrorist attack, taking into account also that prevention might not be 100% effective.

This leaves bioterrorism leading to existential risk as the final issue for discussion. The extinction of human kind might initially seem to lead to an infinite loss of welfare / utility, since the production and experience of welfare would stop. But this is not actually the case. As philosophers of animal rights have shown, humans are not the only beings producing and experiencing welfare, so a world without humans would still be a world full of welfare / utility. Exactly how much would be lost, and how much would be regained when animals gradually filled the void left by humans is difficult to say, because we have no real handle on how to compare welfare between species in general or more specifically between humans and animals. Let us, as an illustration of this point consider the inter-species comparison of welfare between lions and wildebeest. We are able to make intra-species comparisons, i.e. estimating with some degree of accuracy whether a particular lion or wildebeest is doing better or worse than some other lion or wildebeest. But there seems to be no obvious way of making inter-species comparisons between lions and wildebeest when both have normal levels of welfare for their species. Even if we can make plausible comparative welfare claims when we compare a happy and well fed lion to a miserable and starving wildebeest, this does not entail that

we can make comparative claims between happy lions and happy wildebeest. It might be objected that we can make other inter-species comparisons of welfare, e.g. between lions and mosquitos, where we would be fairly certain that a non-suffering lion had a higher level of welfare than even the best fed mosquito. This is true, but the reason it is true is that most of us do not tend to think that mosquitos and other insects have a welfare in a subjective sense because their nervous systems are not sufficiently developed to sustain an inner life. So, would a world where humans had been extinct for 1000 years be better or worse and by how much? We can estimate the loss in human welfare, but can we estimate the countervailing gains and would we also have to take into account non-welfare gains as the value of the global re-establishment of wilderness (29)?

In relation to existential risks we also have to remember that they are in an important sense mutually exclusive. Humankind can only go completely extinct once, any particular extinction is subject to causal pre-emption, so the probability we have to estimate to assess the size of the bioterrorism problem is not the probability of an existentially threatening bioterrorist attack, but the probability of an existentially threatening bioterrorist attack being the first extinction event out of all the possible extinction events.

It may finally be interesting to note that the bioterrorism argument is very similar in structure and function to empirical slippery slope arguments (30, 31). In both cases we have an innocuous initial development predicted to lead to catastrophic consequences through a development that will automatically proceed if we take the first fateful step. Because of this similarity it is quite surprising to see the bioterrorism argument used by bioethicists who reject the validity of slippery slope arguments and decry anyone who uses them.

Conclusion

What are we to make of the use of the bioterrorism argument in bioethical analysis of synthetic biology?

As shown above the argument fulfils a range of functions that do not depend on it being valid and sound, but simply on it being (superficially?) rhetorically convincing. The use in bioethics is thus very similar to the more general use of the argument, i.e. to establish synthetic biology as an important locus of activity. Proponents of synthetic biology, paradoxically use it to illustrate the power and usability of this new technology in which they are invested, and politicians use it to argue for non-liberal security and surveillance policies. We should thus expect the bioterrorism argument to recur again and again as new biotechnologies are developed.

As argued above the possibility of bioterrorism is an important problem that needs attention, as does the prevention of ordinary terrorism which kills hundreds every month (although mostly in countries outside of the media gaze of the Western world). But, the bioterrorism problem is probably not as important as it is often made out to be in the more excitable parts of the literature. Bioethicists should not cry 'Bioterrorism!' any louder than they cry 'Systemic injustice!' or 'Abuse of power!'

References

1. Global Challenges Foundation. *Global Catastrophic Risks 2016*. Stockholm: Global Challenges Foundation.
2. Douglas T, Savulescu J. Synthetic biology and the ethics of knowledge. *Journal of medical ethics*. 2010; 36(11):687-93.
3. Daniel YG. Ethical Theory for "Dual-Use" Dilemmas in Synthetic Biology. *Asian Bioethics Review*. 2012; 4(2):150-9.
4. Kwik G, Fitzgerald J, Inglesby TV, O'Toole T. Biosecurity: responsible stewardship of bioscience in an age of catastrophic terrorism. *Biosecurity and Bioterrorism: Biodefense Strategy, Practice, and Science*. 2003; 1(1):27-35.
5. Miller S, Selgelid MJ. Ethical and philosophical consideration of the dual-use dilemma in the biological sciences. *Science and engineering ethics*. 2007; 13(4):523-80.
6. Resnik DB. H5N1 avian flu research and the ethics of knowledge. *Hastings Center Report*. 2013; 43(2):22-33.
7. Crichton, M. *Prey*, New York: HarperCollins; 2002.
8. Wohlstetter R. Terror on a grand scale. *Survival*. 1976; 18(3):98-104.

9. Rosenau W. Aum Shinrikyo's Biological Weapons Program: Why Did it Fail?. *Studies in Conflict and Terrorism*. 2001; 24(4):289-301.
10. Schmidt M. Diffusion of Synthetic Biology: a Challenge to Biosafety. *Systems and Synthetic Biology*. 2008; 2(1-2):1-6.
11. Jefferson C, Lentzos F, Marris C. Synthetic biology and biosecurity: challenging the “myths”. *Frontiers in Public Health*. 2014; 2(115): 1-15.
12. Bostrom N. *Superintelligence: Paths, dangers, strategies*. Oxford: Oxford University Press; 2014.
13. Caplan AL, Magnus D. New life forms: new threats, new possibilities. *The Hastings Center Report*. 2003; 33(6):7.
14. Caplan AL, Parent B, Shen M, Plunkett C. No time to waste—the ethical challenges created by CRISPR. *EMBO reports*. 2015; 16(11):1421-6.
15. Cohen C. Restriction of research with recombinant DNA: the dangers of inquiry and the burden of proof. *S. Cal. L. Rev.*. 1977; 51:1081-1113.
16. McCaull J. Research in a Box. *Environment: Science and Policy for Sustainable Development*. 1977; 19(3):31-7.
17. Cenciarelli O, Gabbarini V, Pietropaoli S, Malizia A, Tamburrini A, Ludovici GM, Carestia M, Di Giovanni D, Sassolini A, Palombi L, Bellecci C. Viral bioterrorism: learning the lesson of Ebola virus in West Africa 2013–2015. *Virus Research*. 2015; 210:318-26.
18. Holm S, Harris J. Precautionary principle stifles discovery. *Nature*. 1999; 400(6743):398.

19. Harris J, Holm S. Extending human lifespan and the precautionary paradox. *Journal of Medicine and Philosophy*. 2002; 27(3):355-68.
20. Holm S. Reply to Sandin: The paradox of precaution is not dispelled by attention to context. *Cambridge Quarterly of Healthcare Ethics*. 2006; 15(02):184-7.
21. Ryle G. Taking sides in philosophy. *Philosophy*. 1937; 12(47):317-32.
22. World Health Organisation. *Preventing Diarrhoea through better Water, Sanitation and Hygiene*. Geneva: World Health Organisation; 2014.
23. World Health Organisation. *Maternal Mortality - Fact Sheet 348*. Geneva: World Health Organisation; 2015.
24. Posner RA. *Catastrophe: risk and response*. Oxford: Oxford University Press; 2004.
25. Persson I, Savulescu J. *Unfit for the Future – The Need for Moral Enhancement*. Oxford: Oxford University Press; 2012.
26. Taubenberger JK, Morens DM. 1918 Influenza: the mother of all pandemics. *Rev Biomed*. 2006; 17:69-79.
27. Brainerd, Elizabeth and Siegler, Mark V., The Economic Effects of the 1918 Influenza Epidemic (February 2003). CEPR Discussion Paper No. 3791. Available at SSRN: <http://ssrn.com/abstract=394606>
28. Garrett TA. *Economic Effects of the 1918 Influenza Pandemic - Implications for a Modern-day Pandemic*. St Louis: Federal Reserve Bank of St. Louis; 2007.
29. Godfrey-Smith W. The value of wilderness. *Environmental Ethics*. 1979; 1(4):309-19.

30. Lamb D. Down the slippery slope: arguing in applied ethics. Beckenham: Croome Helm; 1988.

31. Holm S, Takala T. High hopes and automatic escalators: a critique of some new arguments in bioethics. *Journal of medical ethics*. 2007; 33(1):1-4.