International Journal of Astrobiology 12 (2): 144–148 (2013)
doi:10.1017/S1473550412000511 © Cambridge University Press 2013

Bioterrorism and the Fermi Paradox

Joshua Cooper

Department of Mathematics, University of South Carolina, 1523 Greene St., Columbia, SC 29208, USA
e-mail: cooper@math.sc.edu

Abstract: We proffer a contemporary solution to the so-called Fermi Paradox, which is concerned with conflict between Copernicanism and the apparent paucity of evidence for intelligent alien civilizations. In particular, we argue that every community of organisms that reaches its space-faring age will (1) almost immediately use its rocket-building computers to reverse-engineer its genetic chemistry and (2) self-destruct when some individual uses said technology to design an omnicidal pathogen. We discuss some of the possible approaches to prevention with regard to *Homo sapiens*' vulnerability to bioterrorism, particularly on a short-term basis.

Received 7 October 2012, accepted 10 December 2012, first published online 16 January 2013

Key words: Fermi paradox, bioterrorism, genetic engineering, biomalware.

Introduction

'Where are they?' asked legendary physicist Enrico Fermi in 1950. Given even the most conservative estimates for the density of stars in our galaxy, the average number of satellites per star, the probability of 'Goldilocks zone' planetary conditions, the variety of possible chemistry for life processes, the chances of autocatalysis arising by random processes, and half a dozen other considerations, there should be a huge number of space-faring intelligent civilizations living right here in our own galaxy. Indeed, the famous 'Drake equation' attempts to quantify this number, and it is very hard to come up with plausible input values that lead to any conclusion other than the utter infestation[1] of space with alien societies (Drake & Sobel 1992). Although the Milky Way is about 10^5 light years wide, it is about 10^{10} years old – so any space explorers would have had eons and eons to reach us. Through the years, countless arguments have been put forth to explain the evident paucity of extraterrestrials in light of the powerful argument that the unfathomably enormous scale of the universe swamps any objection that life is unlikely to arise in any particular place, at any particular time[2]. Perhaps we are first! Maybe every alien civilization transcends corporeality (in ways anticipated – like a technological singularity, or full entanglement of tissue and machine – or in ways no one has thought of) so soon after reaching space-travel technology that we simply are not here at the right moment to meet anyone. Or, disturbingly, perhaps all the other Milky Way denizens who have reached the point of going through this very *Gedanken experiment* conclude that inter-civilization 'first contacts' must result in at least one of the parties destroying the other, so they stay quiet in order to avoid attracting the attention of bellicose passersby. Some naysayers invoke technological obstacles; others suggest that we have not been looking hard enough, that we have been looking in the wrong places, or that we have been subject to a massive conspiracy to cover up alien visitors. Numerous other solutions have been proposed, some more plausible than others.

One very natural proposal is the so-called 'doomsday hypothesis,' which posits that civilizations necessarily destroy themselves soon after reaching the technological sophistication necessary for interplanetary travel or communication. Visions of nuclear wars, 'gray-goo' nightmare scenarios[3], climate change disasters, Malthusian catastrophes, and similar phenomena give weight to this perspective. It is a depressing and unpalatable conclusion, because it means that the 'Great Filter' – whatever force inexorably turns the adaptivity of intelligence on its head – is in our near future. However, it is hard to see how even a collection of many different extinction mechanisms could have the *surefire* effect of the Great Filter, and, for the argument to have force, essentially all technologically advanced societies must be ground into rubble. If only 1 in 10^6 aliens survive these potential cataclysms, the surfeit of opportunities for a survivor to make it through the Filter guarantees that some will; that is, the argument from preponderance implies that the Great Filter must *almost always* cause *usually total* destruction. It is hard to imagine how even the worst calamities conceivable could ensure such utter devastation.

Presently, we flesh out one proposal for a contributor to the Great Filter: bioterrorism. Suppose a society of organisms has

[1] Relatively speaking: Drake's original estimate for the number 'N' of possible alien interlocutor societies in the Milky Way – 10^5 – implies an average separation on the order of 10^3 light years.

[2] We will henceforth refer to this theme – the Universe's sheer magnitude mooting arguments of improbability – as the argument from preponderance.

[3] Drexler's terminology for runaway nanotechnological self-replicating robots transforming all terrestrial matter into copies of themselves (Drexler 1986).

https://doi.org/10.1017/S1473550412000511 Published online by Cambridge University Press

reached space-faring capacity. We claim that, in the blink of geologic time, they will also:

1. Consist of a huge number of individuals.
2. Master the chemistry of (their version of) life.

Humanity certainly has the former, and we are in an age of shockingly rapid maturation toward the latter; biology is taking its rightful place at the hard sciences table. Why must an alien society also have these features at the moment they begin to explore their astronomical neighbourhood? We argue for the inevitability of both developments below.

1. Plenitude: First of all, the civilization under consideration must have gone through the process of (Darwinian) evolutionary selection in order to reach the point of space-travel. Evolution is the only known means by which the second Law of Thermodynamics can be flouted so brazenly. One might even define 'life' in terms of this (local) reversal of the inexorable march of work-to-entropy conversion. On the other hand, evolution requires a long-term competition for resources alongside copious genetic recombination and mutation. In particular, in order for a civilization to reach the sort of technological sophistication Fermi's Paradox posits, it must have consisted of many independent organisms.

One might make the following objection: strong evidence exists that even the superpredator *Homo sapiens* went through a Pleistocene population bottleneck of at most a few tens of thousands of individuals (Hawkes *et al*. 2000). However, our ancestors of this epoch were still rather far from space flight! It was not until huge numbers of individuals were available (just over 3 billion by the time Yuri Gagarin made his historical orbital flight in 1961) to generate enough capital that sufficient resources were in play for a government to pool it into space-technology-grade research funding (US Census Bureau http://www.census.gov/population/international/). The frustrating reality of space flight is that it is prohibitively expensive; just to escape the gravitational field of an astronomical body large enough to hold onto an atmosphere requires a huge amount of energy, and collecting that much energy in one place – another audacious upstream battle against entropy – requires the collective efforts of thousands of individuals, funded by billions of indirect participants[4].

Therefore, it would be very surprising if a collection of living organisms were capable of reaching the stars without being positively legion. Perhaps technology would allow for the consolidation of these organisms into fewer, more powerful ones once such scientific advances were made – but there must have been an extended period when many organisms existed during the dawn of their space-faring epoch. The seeming unlikelihood of every advanced civilization falling under the complete control of just a few independent minds makes such centralization eminently vulnerable to the argument from preponderance.

2. Mastery of Biochemistry: The second point we argue is that, once a civilization is in the chronological vicinity of achieving space travel, it will also develop a good deal of technology to tinker with its own chemistry. The advances that have made space accessible to humans – primarily, the harnessing of energy and quantum physics for ever-faster computation – are precisely those that are fuelling the current avalanche of discovery in biology. *Homo sapiens* is starting to understand the inner workings of its brain, its cells, and, most importantly, its genome, because it is learning how to use massive computing power to decipher the functionality of its informational macromolecules. The incredible snowballing business of modern computing – borne of a desire to aim missiles more accurately and crack military encryption during World War II – is being used for all sorts of applications, including biochemistry. In order to produce computers small and efficient enough to facilitate space exploration, science has learned how to study the tiniest units of matter; when those microscopes, X-ray crystallographic tools, and sophisticated understandings of electron behaviour were turned inward on the cell nucleus, the structure of DNA revealed itself almost immediately. It did not take long before the chemistry of nuclear polymerization was worked out, mastered and imitated *in vitro* on an industrial scale (fewer than 30 years after the advent of the modern computer), and then it was only another 20 years before the human genome was deciphered. Generating specified DNA sequences to order has become astonishingly easy, as evidenced by Church, Gao, and Kosuri of the Harvard Wyss Institute's recent announcement of their having stored 700 terabytes of data – the equivalent of 14,000 Blu-ray disks – in a single gram of DNA (Church *et al*. 2012). A civilization would have to wilfully ignore or suppress the subject matter of its own chemical constitution to avoid spilling some of the runaway torrent of science and technology prerequisite to space travel onto the plate of biology.

Multitude, variation, and genetic technology

Approximately at the moment a civilization reaches its space age, then, it will consist of an enormous collection of individuals and have sophisticated technology for the analysis and manipulation of their own genomes. This can have tremendously beneficial consequences: for example, a multitude of researchers using increasingly powerful tools to cure disease, extend lifespan, improve quality-of-life, etc. However, it also enables the possibility of great catastrophe in the form of bioterrorism. When there are so many independent organisms, at least a few are sure to be some version of hateful (The American Front), vengeful (Timothy McVeigh), insane (Aum Shinrikyo), radical Luddite (Theodore Kaczynski), power

[4] Although the USSR and USA of the early 1960s consisted of only about 300M living persons, their treasuries were filled through a long history of colonial and neocolonial resource-extraction, outright conquest, and the legacies of residents past in the form of intellectual, cultural, technological and physical assets they or their ancestors created or acquired. Billions of people worked and established control over the physical materiel needed to create stable spacecraft and gather sufficient chemical energy in one spot to propel them out of Earth's gravity, even if the actual tax structures providing pay checks for all the NASA and OKB employees showed 'only' 300M accounts receivable.

https://doi.org/10.1017/S1473550412000511 Published online by Cambridge University Press

hungry (Osama bin Laden) or some combination of the above[5].

At the same time that the argument from preponderance is kicking in demographically with regard to the emergence of individuals bent on destruction, computing resources will be engaged in a self-propelling race towards unimaginable power, and biochemistry will be making base-pair I/O per unit energy plummet to within reach of even impoverished amateurs. The Human Genome Project cost $3 billion in US taxpayer money (among other sponsors) and took thousands of scientists 13 years (Department of Energy, http://www.ornl.gov/sci/techresources/Human_Genome/project/about.shtml). This year, Oxford Nanopore Technologies Ltd is expected to make commercially available a USB-powered DNA sequencer for less than $1000, one of which can purportedly sequence the entire human genome in 5 hours (Hayden 2012).

Watson, Crick and Franklin were at the cutting edge of science when they deduced the structure of DNA in 1953. By 1988, the Cetus/Hoffman-La Roche PCR-1000 enabled any well-funded biochemistry laboratory to sequence thousands of DNA base pairs. By the mid-90s, undergraduate biology students at top universities in the US and Europe were able to sequence DNA themselves. Today, a middle class American high school student could likely self-fund a sequencing-based science project. Currently, biology is making rapid progress towards creating an organism from scratch, wrapping functioning organelles, a genome and cytoplasm inside a phospholipid bilayer bubble or packing nuclear material into protein-tile capsids. In 20 years, this also will be part of standard undergraduate curricula. See Figure 1 for a sobering depiction of the plunging cost of DNA sequencing.

One could celebrate the virtues of continued scientific progress unreservedly, were it not for the fact that the information contained in the genomes of some very dangerous organisms is now available for examination and experimentation to just about anyone and will be increasingly so in coming years. With information networks spanning the globe, this extremely perilous biosoftware is essentially cost-free to transmit. 'Information wants to be free,' in the iconic words of Stewart Brand (Brand 1985). Indeed, the genome of *poliovirus* became publicly available in 1981 (Kitamura *et al.* 1981; Racaniello & Baltimore 1981) and, in the first half of 2012, *Nature* (Herfst *et al.* 2012) and *Science* (Imai *et al.* 2012) published recipes to breed an artificial variant of H5N1 avian flu with a very high expected human mortality rate and transmissibility. At almost the same moment, Stanford University researchers and the J. Craig Venter Institute announced the first digital simulation of the entire lifespan of a complete organism (Karr *et al.* 2012).[6] Although their software ran on a

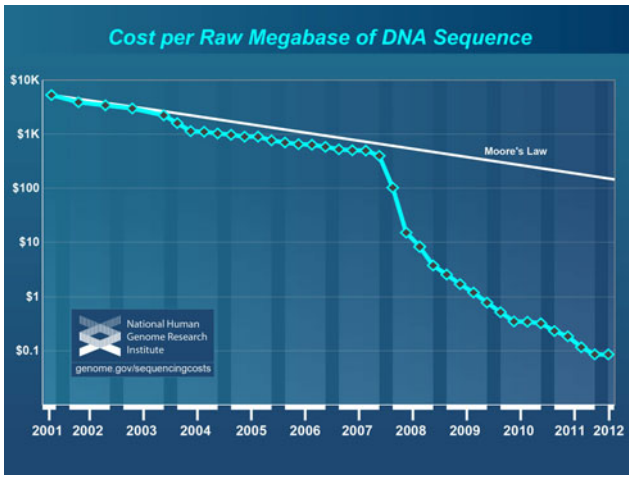


Fig. 1. The past decade of progress in reducing the cost of genomic sequencing, in dollars per megabase; note the logarithmic scale. (National Human Genome Research Institute, http://www.genome.gov/sequencingcosts/).

large computing cluster, the brute force and moral haphazardness of human ingenuity virtually ensures that it will not be long before such tools are widely available. Any devious bioprogrammer will be empowered to accurately predict the deadliness and virulence of human infectious agents modified by a billion randomly chosen mutations; they need not understand a mote of biochemistry in order to invent genomes of extinction-grade pathogenicity. Even with the most draconian attempts at censorship, in 2050, every angry peasant in the developing world will be able to download and print an engineered-to-be-airborne *ebolavirus* variant from their mobile phone, and perhaps even tweak it themselves to kill carriers of their most reviled genotype.[7] The time-scales under consideration here are miniscule on a cosmic scale; assuming we do not destroy ourselves by another mechanism before then, the tools to carry out serious bioprogramming experiments will be available to almost everyone at some point geologically soon.

We may conclude that, when a civilization reaches its space-faring age, it will more or less at the same moment (1) contain many individuals who seek to cause large-scale destruction and (2) acquire the capacity to tinker with its own genetic chemistry. This is a perfect recipe for bioterrorism and, given the many very natural pathways for its development and the overwhelming evidence that precisely this course has been taken by humanity, it is hard to see how bioterrorism does not provide a neat, if profoundly unsettling, solution to Fermi's paradox. One might object that, if omnicidal individuals are successful in releasing highly virulent and deadly genetic malware into the wild, they are still unlikely to succeed in

[5] Obviously, there is a certain danger in the attribution of human emotions to extraterrestrial organisms. However, the point that there are plenty of conceivable motivations for an alien to want to kill its brethren is not anthropocentric, and can be seen as a very general statement about the nature of evolution.

[6] Their species of interest was the bacterium *Mycoplasma genitalium*, chosen because it possesses the fewest number of genes of any known organism. It is difficult to imagine that no one involved with the project noticed the absurdity of the fact that *M. genitalium* is a (sexually transmitted) human pathogen, the elucidation of whose genome may be something of a Pandora's box.

[7] High skin melanin? Double X-chromosome? Large nasal bridge? It requires little strain to imagine what the hideous creative depths of human violence could make possible.

https://doi.org/10.1017/S1473550412000511 Published online by Cambridge University Press

killing everyone. However, even if every such mass death event results only in a high (i.e., not total) kill rate and there is a large gap between each such event (so that individuals can build up the requisite scientific infrastructure again), extinction would be inevitable regardless. Some of the engineered bioweapons will be more successful than others; the inter-apocalyptic eras will vary in length; and post-apocalyptic environments may be so war-torn, disease-stricken, and impoverished of genetic variation that they may culminate in true extinction events even if the initial cataclysm 'only' results in 90% death rates, since they may cause the effective population size to dip below the so-called 'minimum viable population.' This author ran a Monte Carlo simulation using as (admittedly very crude and poorly informed, though arguably conservative) estimates the following Earth-like parameters: bioterrorism event mean death rate 50% and standard deviation 25% (beta distribution), initial population 10^{10}, minimum viable population 4000, individual omnicidal act probability 10^{-7} per annum, and population growth rate 2% per annum. One thousand trials yielded an average post-space-age time until extinction of less than 8000 years. This is essentially instantaneous on a cosmological scale, and varying the parameters by quite a bit does nothing to make the survival period comparable with the age of the universe.

Going forward

Although the survival period is most relevant to considerations bearing on Fermi's paradox, the time until the first such event is of greatest immediate concern. Even with an omnicidal act probability of 10^{-10} per annum, the average time until the first bioterrorism event is less than 2 years.

Suppose[8] that tomorrow, a bioengineered weapon is released by a disgruntled (say) South-Ossetian. Most likely, the first such event that really gets out of control will not be as sophisticated as it could be, killing say, only a tenth of the world's population. Most governments of the world would survive and probably take the opportunity to declare an emergency and enforce martial law, with the cheering help of their frightened citizens. The sudden uptick in unwatched loot and the subsequent clashes between civilians and police/military would result in an unstable, atavistic, hot-iron political situation that could easily lead to widespread war and further destruction.

One might reasonably try to think of ways to prevent this or similar situations. Censorship – attempts to ensure that only people with 'good' intentions are able to obtain the requisite information to perform dangerous biochemistry experiments – could never work, since someone will always leak and/or rediscover information eventually, and humanity would then live out its last days writhing in agony under the fist of authoritarian governments for no reason. Preventing the spread of the equipment needed to perform genetic engineering and to 'boot' pathogens is also an impossible task, as the raw machinery needed is actually quite simple: small, cheap, easy to assemble, and already ubiquitous, the most sophisticated chemistry it uses is essentially available in every living cell on the planet. We could try to avert catastrophe by somehow anticipating every possible killer pathogen and treating preemptively, but Cohen's Theorem[9] and its relatives appear to make vaccines, cures, and treatments a losing strategy in the long run; indeed, one need only glance at the miserable failure of modern computer immunoprogramming to control today's primitive malware or consider the helplessness of our most sophisticated medical technology against seasonal rhinoviruses to get an idea of how little such efforts could do. Furthermore, given that computer malware programmers began from scratch, whereas biomalware programmers have a three billion year-old evolutionary mine of ingenious tricks and tools to pick from, the situation is unfathomably more difficult for a medical approach to preventing bioterrorism. High-speed, heavy-handed global alert and quarantine mechanisms might help to squelch the first few attempts at omnicide, but the possibility of atmospheric bioweapon deployment and the famously unstoppable penetrative reach of life render such measures temporary at best.

Perhaps the solution to Fermi's Paradox is indeed that space-age civilizations self-destruct once they use their rocket-building computers to reverse-engineer their own genetic chemistry. Or perhaps even bioterrorism is not sure enough a cause of extinction that the Great Filter that silences *Homo* will be something else. In the near-term, we might attempt to avert catastrophe for as long as possible without making life substantially more miserable. A greater number of broad spectrum antibiotics and much greater discretion in their use might help. Serious attempts to block all virus activities by preventing the transcription of double-stranded RNA have actually shown promise[10] (Rider *et al*. 2011). Much work has gone into broad-spectrum mycocidal pharmacology, because it is difficult to identify biochemical processes to disrupt that are not common also to eukaryotes – i.e., us – and it is therefore almost impossible to identify a drug target that is specific to one fungal species. Hopefully, 'Homeland Security' and its sometimes dubious analogues elsewhere in the world have used a few dollars to study how to scale up drug, vaccine and quarantine equipment production to universal-availability levels on an emergency timeline.

Humanity might consider committing its resources to a serious universal education effort. As evidenced by, for example, the historically unusual calm that has prevailed (for the most part) in Europe since World War II, education and the cultivation of empathy it entails is incredibly effective at sublimating violent inclinations into more productive energies.

[8] Doing one's best to expel temporarily the ghosts of Boyle, Gilliam, Lawrence, Petersen, Soderbergh, etc.

[9] Fred Cohen proved that, in any reasonable model of computing, no algorithm can stop every possible piece of self-replicating software (Cohen 1987). In this case, one can consider the biosoftware of informational macromolecules and their transcriptional, translational, and regulatory accoutrements to be an extremely complex – and certainly Turing-complete – computer; 'algorithm' can be read as 'medical treatment'.

[10] Although, in a cruel irony, it would not work on Herpes Simplex Virus (HSV) or any of the poxes.

https://doi.org/10.1017/S1473550412000511 Published online by Cambridge University Press

Education also lowers birth rates, raises standards of living, and has many other effects which confer on people enough to lose that they are generally reluctant to engage in criminal behaviour. The funds for such a massive enterprise could conceivably be recovered by nation-states scaling down their more destructive or coercive capital expenditures. In particular, a moratorium on raining death on people's houses from flying *ha-mashchit* robots or expropriating others' natural resources via transactions in the Minuit-Lenape tradition would have the real potential to reduce the number of individuals in the world angered to the point of omnicidal misanthropy.

Acknowledgements

The author wishes to thank J. Hemphill, G. David Poznik, A. W. N. Riasanovsky and the anonymous reviewer for the many valuable contributions to this article they made.

References

Brand, S. (1985), *Whole Earth Review*, May, p. 49.
Church, G.M., Gao, Y. & Kosuri, S. (2012). *Science* **337**(6102), 1628.
Cohen, F. (1987). *Comput. Secur*. **6**, 22–35.
Drake, F. & Sobel, D. (1992). *Is Anyone Out There? The Scientific Search for Extraterrestrial Intelligence*. Delacorte Press, New York. pp. 55–62.
Drexler, K.E. (1986). *Engines of Creation: The Coming Era of Nanotechnology*. Anchor Books, New York.
Hawks, J., Hunley, K., Lee, S.-H. & Wolpoff, M. (2000). *Mol. Biol. Evol*. **17**(1), 2–22.
Hayden, E.C. (2012). Nature News, http://www.nature.com/news/nanopore-genome-sequencer-makes-its-debut-1.10051.
Herfst, S. *et al.* (2012). *Science* **336**, 1534–1541.
Imai, M. *et al.* (2012). *Nature* **486**, 420–428.
Karr, J.R. *et al.* (2012). *Cell*, **150**(2), 389–401.
Kitamura, N. *et al.* (1981). *Nature* **291**(5816), 547–553.
Racaniello, V.R. & Baltimore, D. (1981). *Proc. Natl. Acad. Sci. USA* **78**(8), 4887–4891.
Rider, T.H. *et al.* (2011). *PLoS ONE* **6**(7), e22572.

https://doi.org/10.1017/S1473550412000511 Published online by Cambridge University Press