[Existential Risk / Opportunity] Singularity Management

July 31, 2019

Contents:

- Introduction to Keith Henson's Open Letter p. 1
- Open Letter to Blue Origin/Jeff Bezos, by Keith Henson p. 2
-The Quest of Managing Existential Risk/Opportunity Singularities p. 2
by James Blodgett

Copyright © 2019 Global Risk SIG. Both authors of articles and Global Risk SIG may reprint. This publication is produced by the Global Risk Reduction Special Interest Group, a SIG within US and International Mensa. Content expressed here does not reflect the opinions of Mensa, which has no opinions. To join Mensa or just see what it is about, visit http://www.us.mensa.org . Past issues of this publication are available at: http://www.global-risk-sig.org/pub.htm .

Introduction to Keith Henson's Open Letter

By James Blodgett

Keith Henson is a practical engineer with EROSM-level ambitions. One of my strategies is to try a lot of things in the hope that something works. Keith is an example of that strategy. He is a serious engineer who has been trying hard to enable several types of singularities. In 1975 he co-founded the L5 society, which promoted Gerard O'Neill's ideas for rotating space habitats. I joined that society back in the day, although my only involvement was to read the newsletter. Keith has also been involved with cryonics (freezing people after death in the hope of technological resurrection.) Recently he has been interested not only in space-based solar power that could be beamed to Earth, but on a scale that would replace all fossil fuels, and with the design criteria that it should undercut their costs. He has not quite got there yet, but he and others working with him are making impressive progress. In 2015 he set up the Power Satellite Economics Discussion on Google Groups. The group has 213 members, including Keith Lofstrom and Paul Werbos. (Both have published in EROSM.) Many members know a lot about the underlying technology, and have explored many technical details and the economics of various scales of operation.

Keith's open letter to Blue Origin/Jeff Bezos contains some of their current ideas.

Open Letter to Blue Origin/Jeff Bezos

By Keith Henson

Hi Jeff

I was most impressed with your recent Blue Moon presentation. I too was strongly influenced by Dr. O'Neill's vision. https://en.wikipedia.org/wiki/L5_Society . You are right in that eventually materials to do everything, including building space colonies, will come from extra-terrestrial sources. However, I think there is a nearer-term step possible, power satellites built from earth materials. This opens a large market for reusable rockets and (eventually) a market for water from the Moon.

Power satellites make economic sense for reasonable specific power (6.5 kg/kW or better) if the cost to put parts in LEO can be brought down to about $100/kg. I don't know if this is possible for rockets, but Reaction Engines thinks their Skylon proposal will get there at high flight rates matching fast construction of power satellites (100,000 to a million flights per year). There is a market for about 3000 five GW power satellites (some $36 trillion) over a ten to 15 year time frame.

The effects of high flight rates on the ozone were a concern of mine for a long time. Some years ago, I asked NOAA to model the problem, which they did for the hydrogen-burning Skylon. https://agupubs.onlinelibrary.wiley.com/doi/full/10.1002/2016EF000399 . The last paragraph of the article mentions that they have not done this for any other rocket fuel. In recent months I have talked to them and the office that did that report is willing (given minimal input from Blue Origin propulsion engineers) to reproduce the study for an LNG booster. If you are interested in the consequence of high rocket traffic, I would be happy to connect you to the NOAA people.

Best wishes,
Keith Henson

The Quest of Managing Existential Risk/Opportunity Singularities

by James Blodgett

I am selling a quixotic quest, an attempt to improve human futures that will not be easy to make work. Most people are mundanely practical and don't buy it or at least don't put much into it because they don't see things they can do to help. However, there are thing that can be done, and some do buy. A few good people can do a lot. I say the quest I pitch is quixotic, but even Don Quixote's quest seems gallant to some. His song "The Quest" in Man of La Mancha is popular and has been covered by many singers including

Elvis. Don Quixote was crazy, but in this depiction his craziness has a gallant heroism to it. He knows he is attempting the impossible, but he does so anyway. "To dream the impossible dream, to fight the unbeatable foe..." You can hear the song in its several repetitions within the musical at: https://www.youtube.com/watch?v=RfHnzYEHAow&t=45s .

My quest is not crazy. It makes sense under strict decision theory. However, there is a quixotic impracticality to it, and the strength of its justification varies when evaluated by different philosophical systems and from different points of view. So at least there is intellectual fun in thinking about it, a topic for another time. If it does not work for you as a quest, it is still an interesting hobby.

Our human future could be very good or very bad. Existential risks, risks to our existence, risks like nuclear war, could destroy civilization, and in the worst case could make us all extinct. There are many existential risks. On the other hand, exponentially increasing technology might allow us to do amazing things in the future, more beyond what we can do today than we are beyond cave men, and in a shorter time. It might also enable trillions of human lives within our solar system, and trillions upon trillions if we get to the point of being able to settle galaxies. There are methods of settling galaxies that require things not invented yet, but some methods at least don't require speculative physics, and there are existence proofs that suggest that such things are possible.

My quest is to try to tweak a few of the odds away from the bad futures and towards the good. It is not easy to make things happen that affect our entire species. However, heroes of history have done so. One reason this is not crazy is the mathematics of expected value, a standard metric of decision theory. The expected value of a choice is the value of an outcome multiplied by the probability of its occurrence, summed over all possible outcomes. If we win a dollar if a coin flip results in heads, and nothing otherwise, the expected value of being able to flip that coin is the fifty percent probability of winning multiplied by the dollar won, or fifth cents. If we do many flips with these odds, it is highly probable that the amount we win will be close to fifty cents times the number of flips. If we have a one-in-a-million chance of preventing a nuclear war that would exterminate our species, the expected value of trying is at least the population of Earth, as I write about 7.7 billion people, times the odds of succeeding, or 7,700 expected value lives. Few heroes of legend have saved more. We should also add future lives, in some scenarios a very large number. And our odds of succeeding are sometimes better than one in a million. And even if we fail it is better to try than to lie down and die.

I can claim to have already tweaked the odds, and so did others. A theme of the following history is that I and others did things that made a difference, so note when that happens.

My first involvement was the collider controversy. Particle colliders pump tremendous energy into subatomic particles, accelerating them to very near the speed of light. Two streams of particles are accelerated in opposite directions around a ring of evacuated pipe by superconducting magnets. At points in the ring, the particle beams intersect, and particles going in opposite directions hit each other. When they hit, they stop each other, and the energy pumped into them turns into other particles that explode from the collision point. Radiation-hardened sensors like a giant digital camera surrounding the collision point image the paths of the particles, and physicists extend science and win Noble prizes by discovering new particles. Before CERN's Large Hadron Collider, over ten times more powerful that preceding colliders, was first used in 2008, scientists were hoping to see physics beyond the standard model. They were exploring unknown territory, so they didn't know what they would see. However, there were speculations, and some sounded dangerous. Some postulated production of vacuum transitions, strangelets, or mini black holes. Any of these might have the potential to destroy Earth. Because of this, scientists developed reasons to prove that these things would not be produced. For example, they did the math, based on Newton's theory of gravity, and showed that the energy of a particle collision was inadequate to produce a mini black hole.

Shortly afterward, string theorists, who postulate extra dimensions on a subatomic scale, published several papers in physics journals in which they redid this math. Extra dimensions would change Newton's inverse square law for gravity into an inverse hypercube law on the scale in which those dimensions are active, making gravity much much stronger at its point of origin. This would help with Einstein's quest to unify gravity with other forces, which were hard to unify with gravity because they were thought to be much stronger than gravity. It would also enable production of mini black holes. But don't worry about the collider, because mini black holes would dissipate instantly in a burst of Hawking radiation, which would be seen by the sensors and provide the first ever empirical proof of string theory!

The problem with this is that Hawking radiation is theoretical. Several papers published in physics journals questioned whether it would work as Hawking expected. If black holes are produced and don't dissipate we could be in trouble.

Another safety consideration was an analogy between colliders and cosmic rays. Ground measurements of secondary radiation from cosmic ray collisions with Earth particles in the upper atmosphere, and a few measurements of primary cosmic rays, seemed to show that some cosmic rays have as much energy as CERN particles. Cosmic rays have been hitting Earth since its formation four billion years ago. If collider collisions could make mini black holes, it would seem that cosmic ray collisions with Earth particles could make them too. Since Earth is still here, this seems to show that CERN particles would not cause this kind of trouble.

However, this analogy is incomplete. If a cosmic ray hits an Earth particle and the result is a mini black hole, the black hole will capture the momentum of the cosmic ray. It will travel more slowly than the cosmic ray because it will have the mass of two particles, but it will still be traveling much faster than escape velocity from Earth. If it is electrically neutral and if it behaves like a neutrino, almost all will zip right through Earth and only one in millions will hit anything. Even if a mini black hole behaving like a neutrino does hit something, it will only swallow the particle it hit, slow a bit, and keep going. It would have to swallow many Earth particles in order to slow below escape velocity from Earth and the probability of this ever happening in the entire history of Earth seems low. On the other hand, a mini black hole made by a collider would be made by two particles moving in opposite directions. Most of their momentum would cancel. There would be enough difference in the energy of the two particles so that most black holes made this way would still be traveling faster than escape velocity, but a few would not. These would be captured by Earth, would orbit within Earth, and would have forever to accrete particles and grow.

I developed this idea about the problem with the analogy between colliders and cosmic rays on my own. I discussed it during an email exchange between myself and a physicist, Greg Landsberg. I never published it in a physics journal, and I am not sure that someone else did not develop the same idea, but I did make the point rhetorically several times both with collider advocates and with collider critics. The idea eventually motivated Giddings and Mangano's search for an astronomical object that would reliably capture a mini black hole and save the collider/cosmic ray analogy.

Several things happened to motivate that search. In the year 1999, Walter Wagner published a letter to the editor of Scientific American questioning the safety of the Relativistic Heavy Ion Collider at Brookhaven, and in 2000 Francesco Calogero, an Italian physicist, published "Might a laboratory experiment destroy planet Earth?" in Interdisciplinary Science Reviews. Calogero expressed some concerns, but thought the danger was improbable. I was not aware of these publications at the time. My first awareness of the issue was when I read an article in the August 2000 issue of Scientific American by string theorists predicting black hole formation at colliders. They also predicted that tiny black holes would not be dangerous, because they would evaporate instantly with a burst of Hawking radiation. I knew a bit about Hawking. One thing I knew was that his radiation was theoretical, and had never been seen. It seemed inappropriate to rely on speculative physics as a safety factor for something as important as Earth. So I got involved in the collider controversy.

Several collider critics corresponded over a period of years. Around 2007 several communicated with Calogero, including astrophysicist Rainer Plaga who talked to him in person, and risk analyst Mark Leggett and myself who communicated via email. We convinced him that existing safety considerations were not adequate. Calogero had been

Secretary General of the Pugwash Conferences on Science and World Affairs, which was concerned about nuclear war. The Pugwash Group had recently won a Nobel Peace Prize, so Calogero had some clout. CERN had already done a safety study, which relied on Hawking radiation and the cosmic ray analogy as safety considerations. When Calogero became convinced that better safety considerations were needed, he went to the director of CERN and convinced him that they needed another safety study. So CERN established the Large Hadron Collider Safety Assessment Group (LSAG).

I was disappointed because most of the members of the LSAG disparaged the problem, but Michelangelo Mangano did not. He gave lectures detailing the issues, and he partnered with Steven Giddings to develop better safety considerations. Giddings and Mangano's main concern was answering the idea that I had promoted about Earth not being able to capture mini black holes made by cosmic rays. They postulated that white dwarf stars and neutron stars, astronomical objects with extreme density and extreme gravity, would reliably capture mini black holes. They found several white dwarf stars that were thought by astronomers to have had relatively long lifetimes. Therefore they proposed that the analogy between colliders and cosmic rays can be made to work by considering those stars.

After this, Otto Rössler, Rainer Plaga and others published ways that Giddings and Mangano's findings might be circumvented. (Rössler, known for the Rössler attractor in chaos theory, was a collider critic for years, and recruited many critics in Europe. He recently told me that he had learned about the issue from me, so this is another demonstration that it is possible for a single person, me in this case, to make a difference, because Rössler also made a difference.)

My take on the implications of these safety considerations are that we still do not have the absolute assurance required by the precautionary principle (the precautionary principle says not to do something until it has been proved to be safe) but we are safer. We were fairly safe from the beginning of the debate because black hole production by colliders required speculative physics that was most likely not true. However, that level of safety was not adequate to protect something as valuable as Earth. An airplane with those odds of crashing would not be allowed to fly civilians, and airplanes carry at most a few hundred people, whereas Earth carries billions. Giddings and Mangano's safety considerations reduced the odds of disaster, not because they were perfect, but because it required another level of speculative and therefore somewhat improbable physics to circumvent those safety considerations.

In the end CERN accepted the findings of the LSAG and Giddings and Mangano, and declared the probability of disaster to have the silly value of zero. A zero probability of disaster translates precisely into that famous assertion of scientific hubris, "Nothing can possibly go wrong." The Daily Show made fun of this, comparing CERN's zero probability estimate with the estimate of a collider critic who said that disaster had a probability of 50%, because that is a proper value for something we don't know. I disagree with both estimates. At least we know that disaster requires speculative physics, which is somewhat unlikely because there are many speculations, so I would estimate its probability as less than 50%. I also think that running the collider was the wrong decision, as per the precautionary principle, but the precautionary principle has the flaw of requiring an unavailable perfection. The precautionary principle has been endorsed by many countries, a motivating exemplar for us because it was promoted effectively by environmentalists whose cause is in some ways analogous to the quest of advocating for the future that I am promoting here. However, endorsing countries have not always set up effective enforcement, and the precautionary principle is often ignored in practice when it doesn't work for the group that wants to do something it would prohibit. I can see the point of ignoring the precautionary principle when the benefit is important and the risk minor, when safety considerations are good if not perfect, and when absolute proof is unavailable. Absolute proof is never available in a universe that may contain as yet undiscovered physics. I would prefer requiring a result closer to perfection when Earth is in the balance, but then others have votes in these matters.

I tell this story to motivate work in this field. I, and others, did help to reduce risk.

Now they are thinking of building a more powerful collider. It would give physicists more data to analyze, and might enable breakthroughs. Depending on how one computes the risk, a good chance of a breakthrough might have enough expected value to balance a small chance of destroying Earth. Physicists have somewhat of a conflict of interest in evaluating this, because a more powerful collider would also help continue both their scientific quest and their jobs. A few physicists think that a more powerful collider is not a good idea because the LHC has not yet discovered physics beyond the standard model as hoped, and a more powerful version might not either. Tom Kerwick, who has published in EROSM twice, thinks a new collider is not a good idea because some safety considerations that were relevant for the LHC do not scale to higher energy levels. On the other side of the balance, science is sometimes a good thing and might save us from other disasters, so it is not optimal to put too many restrictions on science. What is needed is a balance, and both critics and advocates of science experiments play an important role in finding that balance.

I wrote this essay to motivate readers, so it is time to pitch things that readers can do to help human prospects. Smart readers will learn relevant fields and find things to do by themselves. A good way to learn relevant fields is to start working on someone else's related project. There are good ideas flying around. Find one and help to make it work.

I suggested a project a few years ago in an essay in Lifeboat Foundation's anthology, Visions of the Future. I have been too busy to do much with it. I have family responsibilities and I am trying to write a book, so I am unlikely to be able to work on it soon. I would applaud someone who could take it over and do a good job with it.

My idea was to help with the ideas of Philip Metzger et al in their paper, "Affordable, Rapid Bootstrapping of the Space Industry and Solar System Civilization," Journal of Aerospace Engineering, January 2012 pgs18-29. A preprint is available at https://arxiv.org/abs/1612.03238 .

Their idea is to start by sending a few tons of equipment to the moon, machines that can dig up and process lunar material, automatic machine tools, 3D printers, teleoperated robots, and solar power, machines that can make other machines. These machines would be built by the best machine shops on Earth, but at first the machines they could make would be crude. Sophisticated parts like computer chips would have to come from Earth. However, each generation of machines would be designed with more capacity, and they would eventually achieve parity with Earth's machine shops and even with Earth's billion-dollar chip fabrication facilities. Design engineers and machine operators could be employed by NASA or a big company, but when there are enough machines on the moon it might be a good idea to involve others by allowing them to rent time on machines and allowing smaller companies to produce products that could be sold to other groups doing similar work. The number and capacity of machines could grow exponentially as they did during the industrial revolution and industrialize the moon and the solar system in a fairly short time. This would also help protect us from existential risks because an industrialized solar system with robust human settlement would provide a backup for Earth.

Challenge prizes often encourage work that would cost many times the prize value. Lindbergh flew from New York to Paris in response to a challenge prize, after several other groups had tried and some died. I think we could encourage the ideas of Metzger et al by offering challenge prizes for groups that build and demonstrate machines that could build other machines using relevant materials. One group that might try for a prize is the people who belong to makerspaces. Industry and university groups might also be interested. Challenge prizes would also publicize Metzger's ideas, and perhaps persuade governments to get involved and to provide launch capacity.

If someone wants to work on this I don't even need to know, but let me know anyway so I can add another effectuality feather to my cap when you make it work, and perhaps I can help a little bit. If more than one person is interested, there is room for more than one person and even for more than one project. Cooperation is a good thing, and so is competition.

Finally, let's keep working on the myriad of other futurist issues.