Lecture 12: Science fiction meets science..

David Aldous

October 12, 2017

"Probability and X" where today X = science fiction.

The Fermi paradox is:

> *The Universe is very big and very old; given there is a human technological civilization on Earth, why don't we see evidence of technologically advanced extraterrestrial civilizations?*

[what data do we have??? – later]
Devising possible explanations is an interesting exercise, as organized logic. A top-down organization might start with the alternatives

- they have (almost) never arisen anywhere.
- they don't last long in a form we would recognize, so none are currently close enough to detect.
- they do currently exist but we can't detect them for some reason.

75 more detailed possible explanations are given in the non-technical book *If the Universe Is Teeming with Aliens . . . Where Is Everybody?* by Stephen Webb (and mostly copied to Wikipedia *Fermi Paradox*).

[show contents and Wikipedia]

This book is ab
apparent absen
see evidence of
met it some 17
authors (too ma
erence list at th
about the parac
discussed the p
are too numero

Several peo
and I would li
Praxis Publish
Copernicus Bo
would not hav
encouragemen
resolution of th
provided man
Bob Marriott a
later draft (Bob
of which I agr
putting me rig
for those erro
observant and
Assistant Edit
authors and or
I am particula

This topic illustrates a general point: before you consider probabilities, you should consider the possibilities. I won't go throughout the 75 possible explanations in detail – may be interesting to do yourself and think which are plausible? other explanations?

The first alternative – "they (almost) never arise" – splits into two cases

- physical conditions favorable to Life getting started are rare
- evolution from simple to complex Life is rare.

Amongst the few bits of relevant data

- We now know that many stars have planets, including planets in the presumed "habitable zone".
- Life on Earth originated fairly soon after the formation of Earth 4.5 billion years ago.
- But no discovered extra-solar planets have detectable free oxygen in atmosphere [free oxygen in Earth atmosphere created by photosynthesis].

Although we know a lot about the actual history of evolution of life on Earth (except the beginning), there is no way to go back and assess, given the physical state of newly-formed Earth 4.5 billion years ago, the chances of

- simple Life getting started
- evolution of Life from simple to complex to human-type intelligent.

People often try to do this – the **Drake equation** – but the numbers are totally guesswork. However one topic in this lecture – *the Great Filter* – attempts to say something relevant without assessing chances, based on a simple STAT134 exercise.

Separate from biology, the book Ward - Brownlee *Rare Earth: why complex life is uncommon in the universe* argue that the entire geophysical history of Earth has been fortuitously favorable to Life.

The possibilities within the third alternative “they do currently exist but we can’t detect them” seem pure speculation – we can’t imagine capabilities or motivations of hypothetical technologically much more advanced extraterrestrials.

Any sufficiently advanced technology is indistinguishable from magic.
Arthur C. Clarke

The puzzle is that there are many ways we **can** imagine an advanced technological civilization would be detectable to us, but we don’t see any of them.

[board]

How can we think about the second alternative

> *technological civilizations don't last long (in a form we would recognize).*

We can think about how our own civilization might end. A book *Global Catastrophic Risks* contains 15 chapters discussing different such risks. As the second topic of this lecture, I will describe some of these risks (later).

The Great Filter

This is a very speculative line of thought, due to Robin Hanson. Consider the product

$$Npq \tag{1}$$

where

- N is the number of Earth-like (loosely, and at formation) planets in the galaxy
- p is the chance that, on such a planet, an intelligent species at a technological level comparable to ours will arise at some time
- q is the chance that such a species would survive in such a way as to be observable (via communication or exploration) to other galactic species for an appreciable length of time.

The point is that Npq indicates (after some more time scaling) the number of *other* intelligent species we expect to observe in the galaxy. Because we don't observe any, we conclude *prima facie* (treating absence of evidence as evidence of absence) that it cannot be true that $Npq \gg 1$. Since it would be a bizarre coincidence if $Npq \approx 1$, we should conclude (*according to this argument*) that $Npq \ll 1$ and so humans are most likely to be the only technological species in the galaxy.

A paradoxical argument.

Human beings did not create the Universe or direct the course of evolution, so N and p are not our responsibility. But q, as applied to us (i.e. will our species leave its mark on the galaxy?) is presumably under our control. Viewing q very roughly as the chance that a hypothetical technological species arising across the galaxy 25 million years in the future would then be able to observe the then-current or previous existence of humans, being told that $q = 10^{-6}$ would be rather depressing. Depressing, because of the ways this might come about, for instance if humans soon become extinct, or change and cease to interact with the macroscopic physical universe. Knowing $q = 10^{-6}$ would be knowing that something like this is *almost certain* to happen. Now having decided that Npq is small, implying pq is very small, the only way to avoid the depressing possibility of q being very small is to for p to be very small.

This argument leads to the counter-intuitive conclusion that

$$\text{we should } \textbf{want } p \text{ to be very small} \tag{2}$$

where the sense of **want** is, as above, “to be consistent with humanity surviving long enough to have at least a tiny chance of leaving its mark on the galaxy”.

This is a strange argument at first sight but (to me) is actually reasonable. It suggests the following **Program;**

> try to ascertain whether what we know about the evolution of life on Earth is <u>consistent with</u> the possibility that p is very small. (3)

Can we actually carry out anything like this program?

Let me set up the key mathematical idea via a story with more familiar ingredients. Consider two unlikely events for an individual – say, being struck by lightning and winning the lottery. Suppose I point to a 50-year-old man and tell you that sometime since age 20 he won the lottery and sometime later he was struck by lightning. And suppose the relative probabilities of these two events are not known, except that each is very unlikely. What can you say about the probability distribution of the ages at which these two events occurred?

In thinking about this question with (non-mathematical) intuition, you might think as follows. The two events split the interval $[20, 50]$ into three sub-intervals. In general you have to wait longer for a more-unlikely event than for a less-unlikely event, so the interval ending with the less-likely of the two events will probably be longer.

But, within the simplest probability model one might devise for this story, that conclusion is wrong. Counter-intuitively, *regardless of the relative probabilities*, the three intervals have the same mean length.

[board]

Robin Hanson pointed out that one can try to apply the same argument to the evolution of intelligent technological species on Earth. Suppose we identify some key steps between the formation of an Earth-like planet and technological civilization. He suggested six intermediate steps (I quote five of his and have modified the sixth).

- Reproductive something (e.g. RNA)
- Simple (prokaryotic) single-cell life
- Complex (archaeatic & eukaryotic) single-cell life
- Sexual reproduction
- Multi-cell life
- Animal-level intelligence
- Technological civilization

The relevant time interval is the interval when life on Earth is possible, roughly 4 billion years ago to 1 billion years in the future. If we suppose each of these 7 steps was an **unlikely** random event then, under the simplest probability model, the math argument shows that **regardless of the seven actual probabilities**, the eight sub-intervals would have random lengths with equal means. And the data we have on the dates of these events is quite consistent with this model prediction, and so in particular is quite consistent with the possibility that p is very small.

The metaphor of **filter** is as stages to pass through before becoming visible in the galaxy; if we decide that there are no other technological species now, then either the filter lies in humanity's past (p very small) or in the future. As said earlier

This argument leads to the counter-intuitive conclusion that

$$\text{we should } \textit{want } p \text{ to be very small} \tag{4}$$

where the sense of *want* is, as above, "to be consistent with humanity surviving long enough to have at least a tiny chance of leaving its mark on the galaxy".

[board – math variant – branching history - answers one of many objections]
– see my paper *The Great Filter, Branching Histories and Unlikely Events*.

One suggested answer to the Fermi paradox is that technological civilizations do arise but don't survive very long. This brings us to the second topic

Global Catastrophic Risks (GCR)

Thinking about this involves an issue of time-scale. We know the world (human society) has changed in the last 100 years, and a default is to assume some comparable amount of change in the next 100 years. We can't imagine 1 million years ahead (which was relevant to Fermi paradox). So let's fix on 500 years.

Question: How might it happen that in 500 years there might be no recognizable "human technological civilization"?

The book GCR contains 15 chapters analyzing particular risks. But what possibilities can you suggest?

[board]

(A): Asteroid or comet impact. Roughly, *asteroid* = solid rock, *comet* = gravel/ice slush, but at 55km/second there is little difference.

This possibility is familiar from both science fiction (SF) and the *K-Pg extinction event* 65 million years ago. But how often do such major events happen?

This is scientifically interesting because we have three separate sources of data.

(i) Impact craters. 170 are known.

[next slide]

So we can estimate the crater size from a "once in 1 million years" impact.

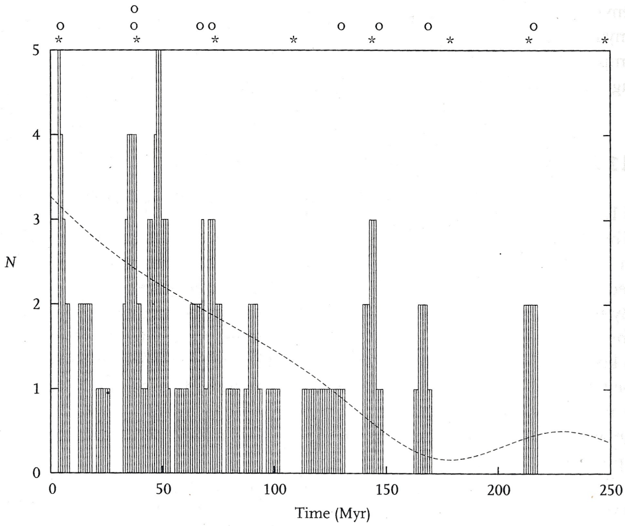

Fig. 11.1 The age distribution of 40 impact craters 3 km or more in diameter, with ages less than 250 Myr, known to precision better than 10 Myr. A rectangular window of width 8 Myr corresponding roughly to the mean uncertainty in ages has been passed over the data, and the smooth curve is a fit to the overall trend. The impression that impacts occur in discrete episodes of bombardment and is confirmed by detailed statistical analysis (Napier 2005).

(ii) Sizes and orbits of asteroids/comets (orbits very different).

Telescopes see only comparatively large asteroids but we can extrapolate to estimate numbers of smaller ones (power law distribution). Intensive recent study of **near-Earth objects (NEO)** gives about 1,000 with diameter over 1km.
[show page]

One estimate is that objects with a diameter of one kilometer hit the Earth an average rate of 2 per 1 million years. This is roughly consistent with observed sizes of impacts craters.

[board – discuss "dark comets"]

(iii) Extinction events in evolutionary history.

This is a puzzle – aside from the famous K-Pg extinction event 65 million years ago, we don't see any smaller extinction events, which one would have expected.

Effect on human civilization?
This becomes speculative . . . . . .

[next slide: table from GCR]

Table 11.2 Possible Impact Effects

Megatons	10,000	1 million	100 million
Impactor	500 m	1–2 km	10 km
Scope	Regional	Civilization-destroying	Species-destroying
Land	Fires, blast and earthquake over 250–1000 km	Destructive blast, quake and possibly fire over continental dimensions	Global conflagration and destructive earthquake
Sea	Uncertain	(Mega)tsunamis around ocean rims	Ocean rim devastation; cities replaced by mudflats; oceans acidified
Air	Sun obscured; possible 536 AD event	Agriculture collapses; ozone depletion; acid rain; sky darkened for years	Land and sea ecologies collapse; ozone depletion; acid rain; sky black for years
Climate	Possible brief cooling	Global warming followed by sharp cooling	Global warming followed by cosmic winter

Note: Uncertainties attend all these thumbnail descriptions, and are discussed in the text. Timescales are also open to debate, but a recurrence time of one megaton per annum is a reasonable rule of thumb to within a factor of a few.

Main effects are from physical destruction of economic infrastructure and then from "nuclear winter" – dust in atmosphere cuts agricultural food production for several years. The "food" issue is relevant to other possible catastrophes, so let's consider it.

Some facts:
World population today around 7 billion.
Almost all food comes, directly or indirectly, from annual crops (rice, corn,. . .).
60 days supply "available" in production and storage.

World population in 1800 around 1 billion – all that could be fed with that level of technology. Today's ability to feed 7 billion depends on

industrial technology
functioning economic system
international trade.

Usually we perceive economic crises in terms of finance (stock prices, budget deficits, house prices) or inflation/unemployment (living standards). But the bottom line is maintaining world food production. No-one has any way to understand how much the world economic system is fragile/robust to external shocks.

SF movie scenario: You wake up tomorrow; some disaster has happened (or will happen, unstoppable) which will cut world food production in half for a few years. What do you think will happen next?

More plausible scenario: Food shortages in one year due to extreme weather events. This scenario is considered in detail in a 2015 report *Extreme weather and resilience of the global food system.*

(B) Gamma-ray bursts (GRBs) are the most powerful known physical events. Satellites detect around one per day, because they can be detected from (literally) halfway across the Universe. A typical GRB generates as much energy in a few seconds as a star will in its entire multi-billion-year lifetime. So bad news if one occurs nearby. Here "nearby" is relative to Universe – perhaps 10,000 light years, a proportion of our own galaxy.

A recent paper (see posted *Bolts from the blue* Economist article) estimates that a GRB near enough to Earth to destroy life on land will occur once per **billion** years. From evolutionary record, this has not happened over the last 500 million years.

Relevant to Fermi paradox – maybe rate higher in other galaxies/nearer galactic center/earlier in history of Universe – fits the "Earth fortuitously hospitable to Life" possibility.

[show Technological singularity]

There is a separate category of *Slowly Developing* catastrophic risks, exemplified by **climate change**. Here the issue is whether the human political system will react sensibly. A posted paper *Preparing for future catastrophes* discusses various risks, with brief mention of some recent mathematical theories.

Project: Read and report on some scientific literature cited in that paper.