SSC Journal Club: Dissolving The Fermi Paradox

Posted on July 3, 2018 by Scott Alexander

I'm late to posting this, but it's important enough to be worth sharing anyway: Sandberg, Drexler, and Ord on Dissolving the Fermi Paradox.

(You may recognize these names: Toby Ord founded the effective altruism movement; Eric Drexler kindled interest in nanotechnology; Anders Sandberg helped pioneer the academic study of x-risk, and wrote what might be my favorite *Unsong* fanfic)

The Fermi Paradox asks: given the immense number of stars in our galaxy, for even a very tiny chance of aliens per star shouldn't there should be thousands of nearby alien civilizations? But any alien civilization that arose millions of years ago would have had ample time to colonize the galaxy or do something equally dramatic that would leave no doubt as to its existence. So where are they?

This is sometimes formalized as the Drake Equation: think up all the parameters you would need for an alien civilization to contact us, multiply our best estimates for all of them together, and see how many alien civilizations we predict. So for example if we think

there's a 10% chance of each star having planets, a 10% chance of each planet being habitable to life, and a 10% chance of a life-habitable planet spawning an alien civilization by now, one in a thousand stars should have civilization. The actual Drake Equation is much more complicated, but most people agree that our best-guess values for most parameters suggest a vanishingly small chance of the empty galaxy we observe.

SDO's contribution is to point out this is the wrong way to think about it. Sniffnoy's comment on the subreddit helped me under-stand exactly what was going on, which I think is something like this:

Imagine we knew God flipped a coin. If it came up heads, He made 10 billion alien civilization. If it came up tails, He made none be-sides Earth. Using our one parameter Drake Equation, we deter-mine that *on average* there should be 5 billion alien civilizations. Since we see zero, that's quite the paradox, isn't it?

No. In this case the mean is meaningless. It's not at all surprising that we see zero alien civilizations, it just means the coin must have landed tails.

SDO say that relying on the Drake Equation is the same kind of er-ror. We're not interested in the *average* number of alien civiliza-tions, we're interested in the distribution of probability over number of alien civilizations. In particular, what is the probability of few-to-none?

SDO solve this with a "synthetic point estimate" model, where they choose random points from the distribution of possible estimates suggested by the research community, run the simulation a bunch of times, and see how often it returns different values.

According to their calculations, a standard Drake Equation multiplying our best estimates for every parameter together yields a probability of less than one in a million billion billion billion that we're alone in our galaxy – making such an observation pretty paradoxical. SDO's own method, taking account parameter uncertainty into account, yields a probability of one in three.

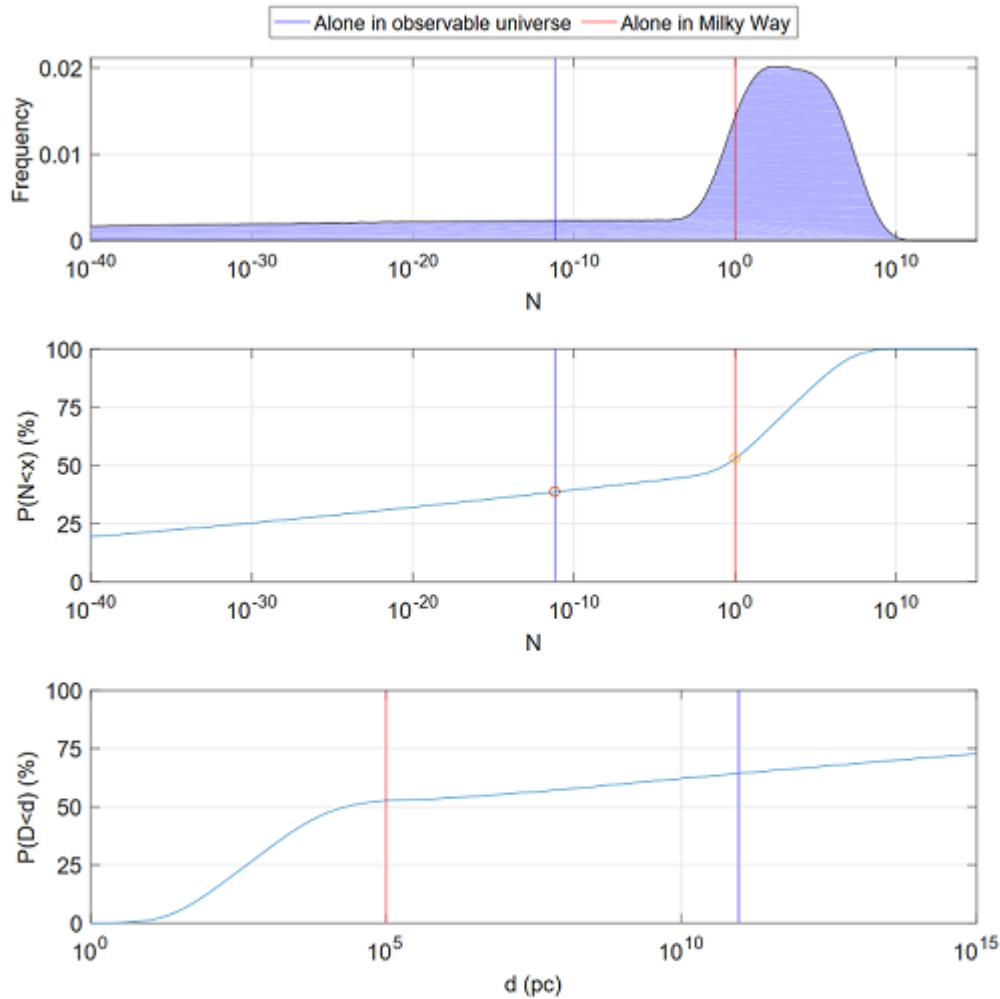They try their hand at doing a Drake calculation of their own, using their preferred values, and find:

Figure 2: (A) A probability density function for N generated by Monte Carlo simulation based on the authors' best estimates of our current uncertainty for each parameter. (B) The corresponding cumulative density function. (C) A cumulative density function for the distance to the nearest detectable civilization.

N is the average number of civilizations per galaxy

If this is right – and we can debate exact parameter values forever, but it's hard to argue with their point-estimate-vs-distribution-logic – then there's no Fermi Paradox. It's done, solved, kaput. Their title, "Dissolving The Fermi Paradox", is a strong claim, but as far as I can tell they totally deserve it.

"Why didn't anyone think of this before?" is the question I am only slightly embarrassed to ask given that *I* didn't think of it before. I don't know. Maybe people thought of it before, but didn't publish it, or published it somewhere I don't know about? Maybe people intuitively figured out what was up (one of the parameters of the Drake Equation must be much lower than our estimate) but stopped there and didn't bother explaining the formal probability argument. Maybe nobody took the Drake Equation seriously anyway, and it's just used as a starting point to discuss the probability of life forming?

But any explanation of the "oh, everyone knew this in some sense already" sort has to deal with that a lot of very smart and well-credentialled experts treated the Fermi Paradox very seriously and came up with all sorts of weird explanations. There's no need for sci-fi theories any more (though you should still read the Dark Forest trilogy). It's just that there aren't very many aliens. I think my past speculations on this, though very incomplete and much inferior to the recent paper, come out pretty well here.

(some more discussion here on Less Wrong)

One other highlight hidden in the supplement: in the midst of a long discussion on the various ways intelligent life can fail to form, starting on page 6 the authors speculate on "alternative genetic systems". If a planet gets life with a slightly different way of encoding genes than our own, it might be too unstable to allow complex life, or too stable to allow a reasonable rate of mutation by natural selection. It may be that abiogenesis can only create very weak ge-

netic codes, and life needs to go through several "genetic-genetic transitions" before it can reach anything capable of complex evolution. If this is path-dependent – ie there are branches that are local improvements but close off access to other better genetic systems – this could permanently arrest the development of life, or freeze it at an evolutionary rate so low that the history of the universe so far is too short a time to see complex organisms.
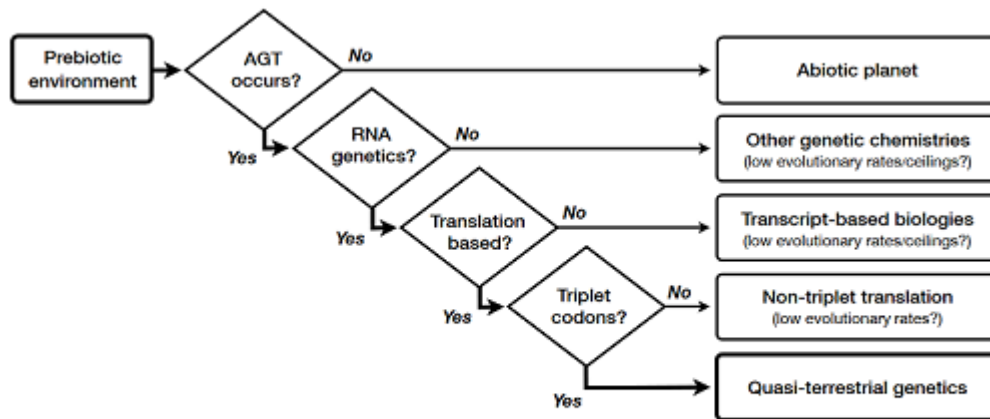
Figure 1. A schematic diagram of potential AGT and GGT branches with aggregated states and transitions. At each branch point, outcome ratios are weakly constrained, with LU values arguably > 100.

I don't claim to understand all of this, but the parts I do understand are fascinating and could easily be their own paper.