Dissolving the Fermi Paradox

Anders Sandberg, Eric Drexler & Toby Ord
Future of Humanity Institute
University of Oxford

The Fermi ~~paradox~~ question: “where are they?”

- A paradox is a tension between experience and theory
- The “theory” is that the number of sites where intelligence can emerge is vast in time and space, that the prior probability is not tiny, and detection/settlement is relatively doable.
- Hence an empty sky (“the Fermi observation”) is odd.

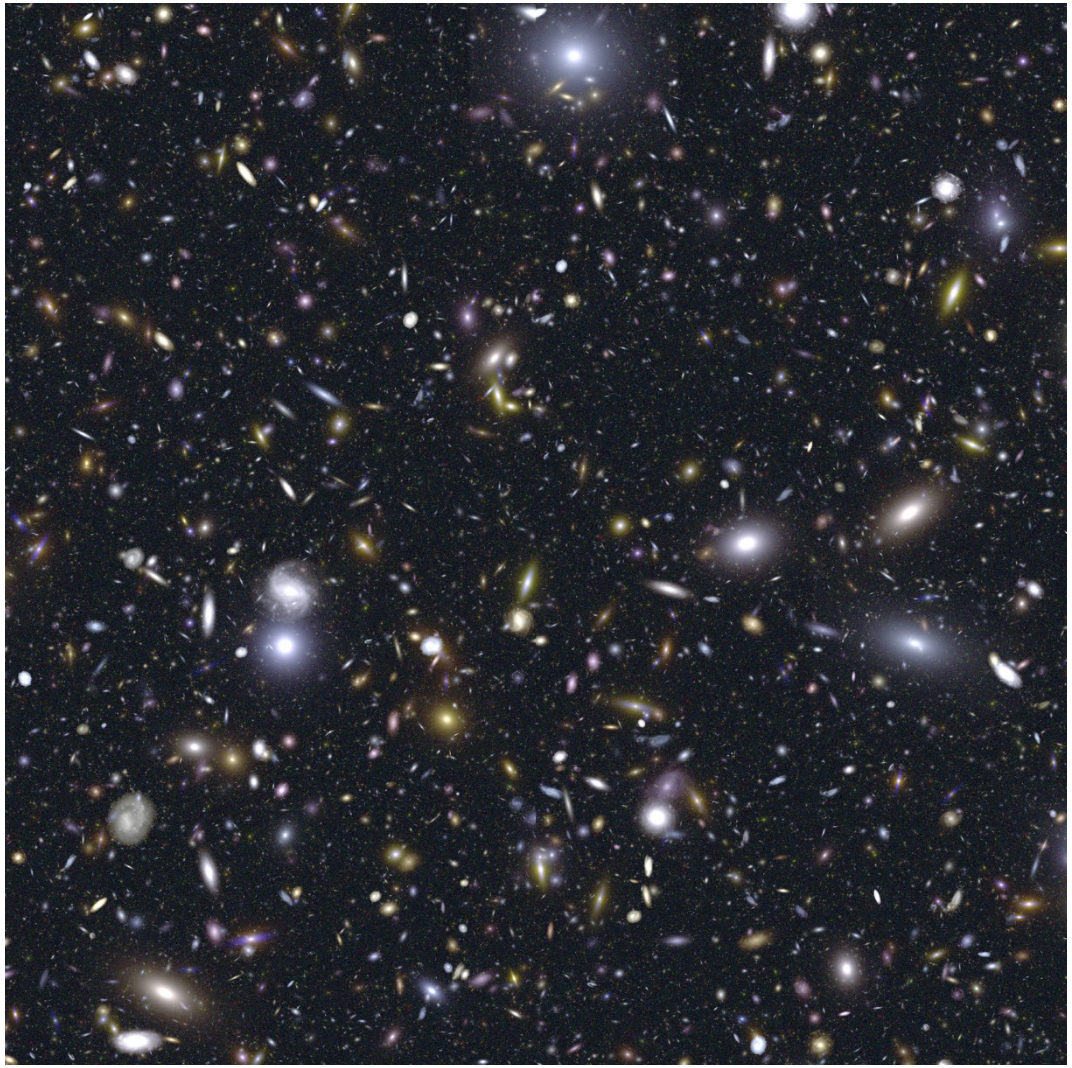

Doing Drake wrong

- Everybody makes up numbers
  - “Perhaps never in the history of science has an equation been devised yielding values differing by eight orders of magnitude. . . . each scientist seems to bring his own prejudices and assumptions to the problem.”
    - History of Astronomy: An Encyclopedia, ed. by John Lankford, s.v. “SETI,” by Steven J. Dick, p. 458.
- Suspiciously convenient conclusions
  - The N≈L and N≈1 schools
- Galactocentrism
  - Ignores intergalactic colonization.
  - N≈1 means 10^{10} civilizations in the visible universe.
    - And 42% chance of at least one more civilization in your galaxy

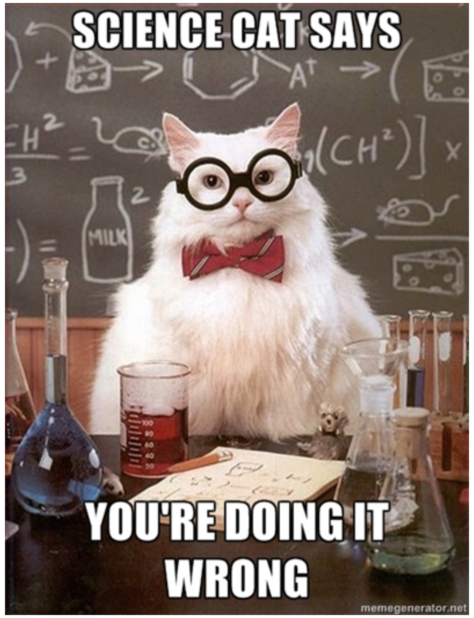

Point estimates considered harmful

- Use single value as best guess of unknown variable.
- Example:
- Assume nine factors x_i multiplied together to give $\text{Pr(life per star)} = \prod_i x_i$.
- Each x_i is a random real number drawn uniformly from [0, 0.2].
- The point estimate for each is 0.1:
  - The product of point estimates is 1 in a billion.
  - Given a 100 billion stars, it naively looks like it is *spectacularly* unlikely for life to have only happened once (3.7×10^{-42}) and the expected number of life-bearing stars would be ≈100.

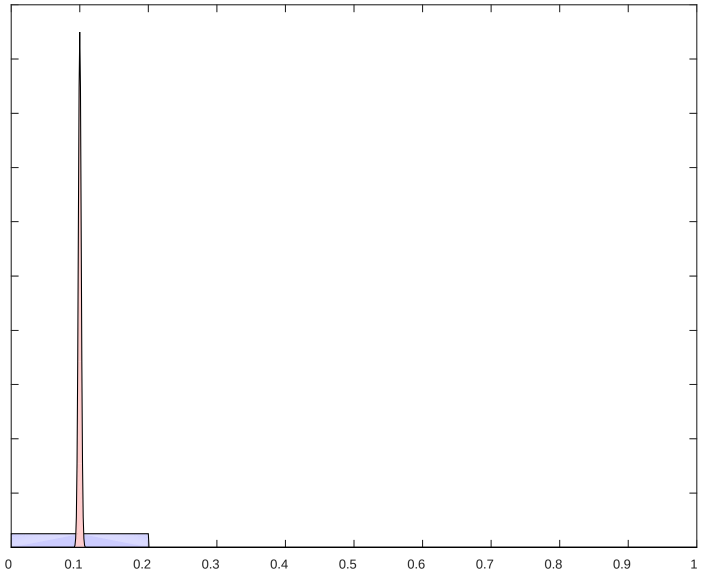

Point estimates considered harmful

- However, actually combining the probabilities as distributions:
    - The *median* number of life-bearing stars is just 8.7 (the mean is still 100).
    - “Life only once” actually occurs 8% of the time
- Multiplying point estimates can be incorrect and misleading, and we need to convolve probability distributions instead.
- Implicit certainty in use of Drake equation produces problematic conclusions.

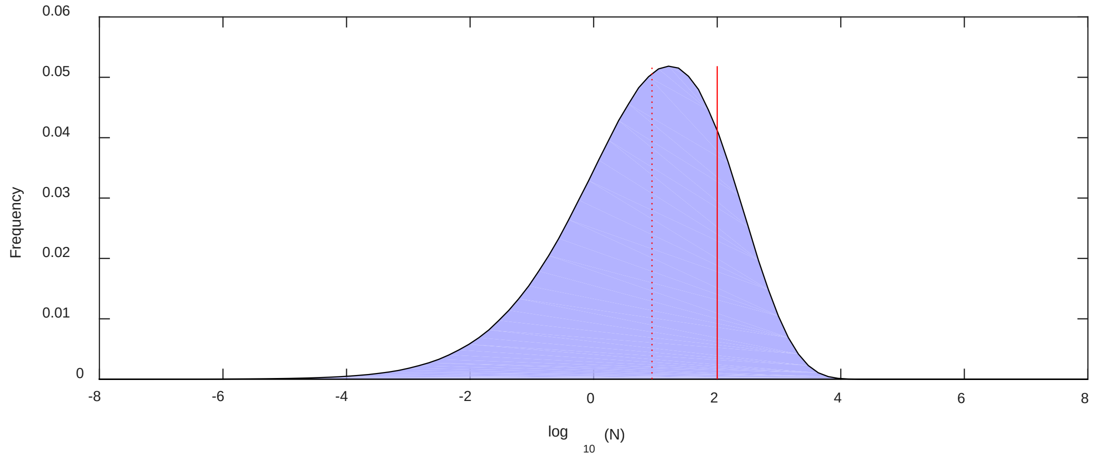

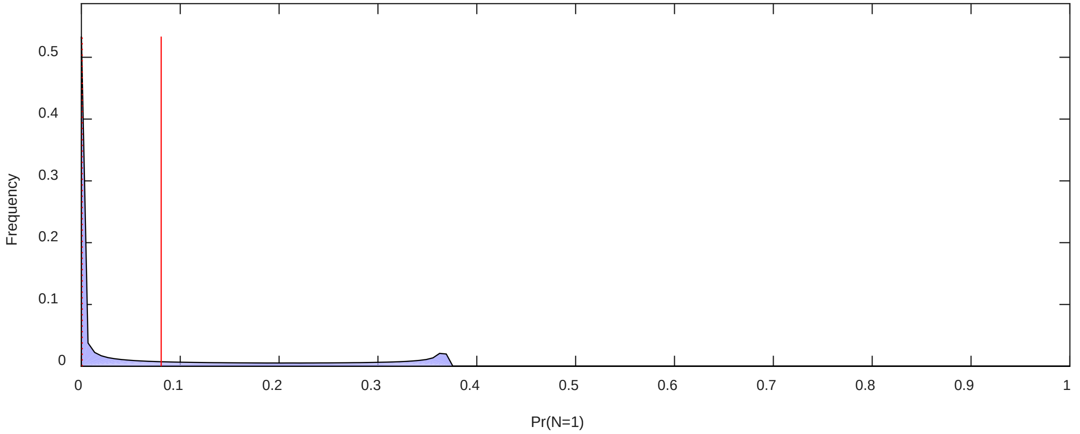

Distribution of claims: what is the view of researchers?

- Literature resampling: review parameter estimates from the SETI literature, produce a distribution by random resampling.
- Some issues of copying estimates bringing down variance.
- Bias in who makes estimates, of course.
  - We should expect strong optimism bias!

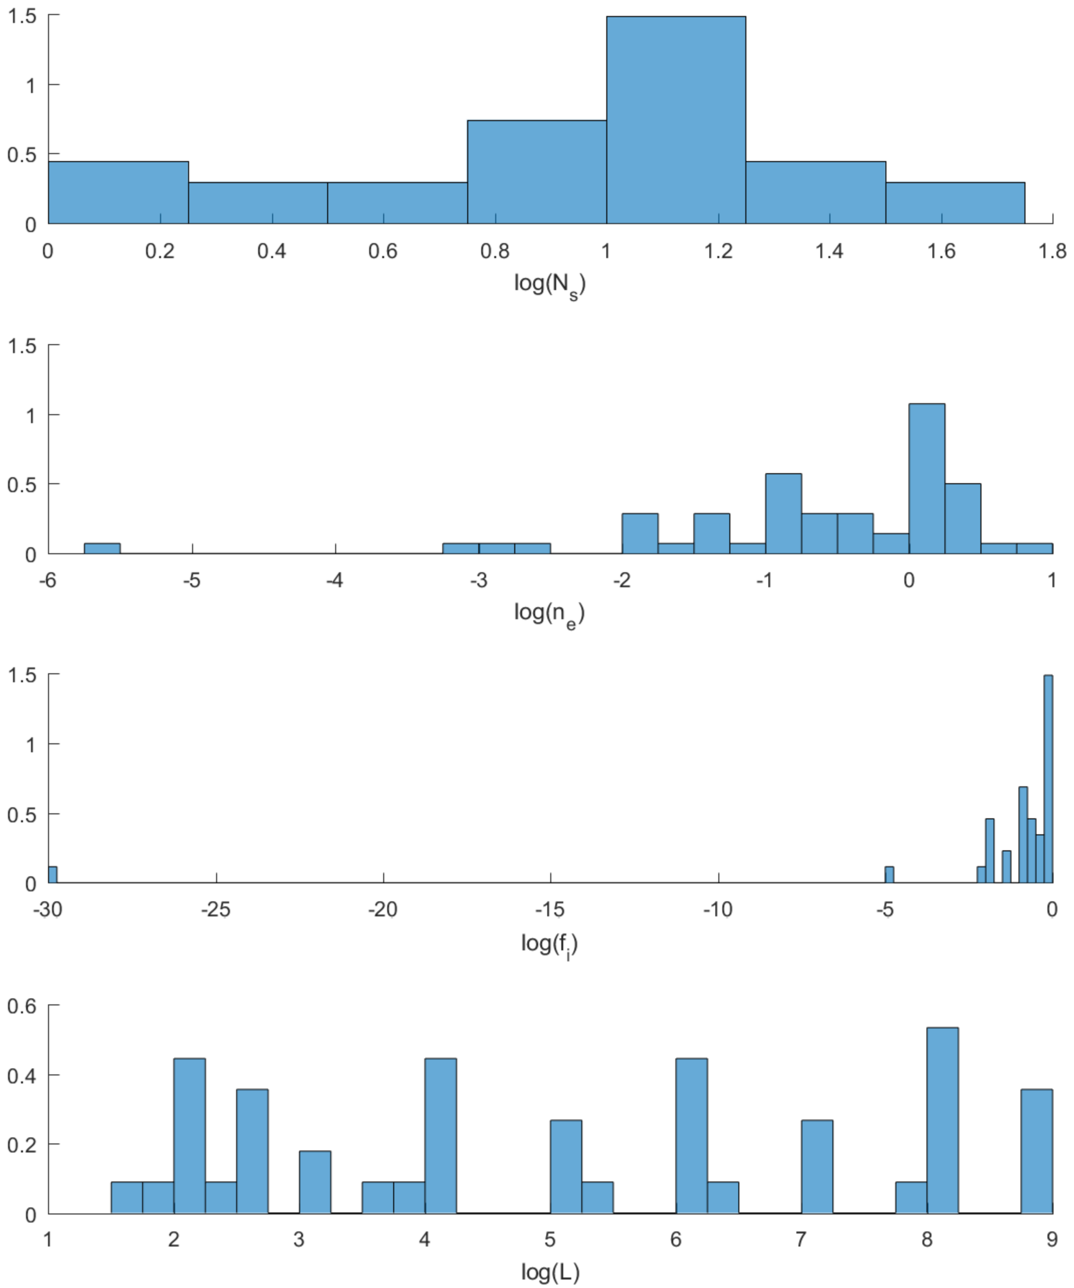

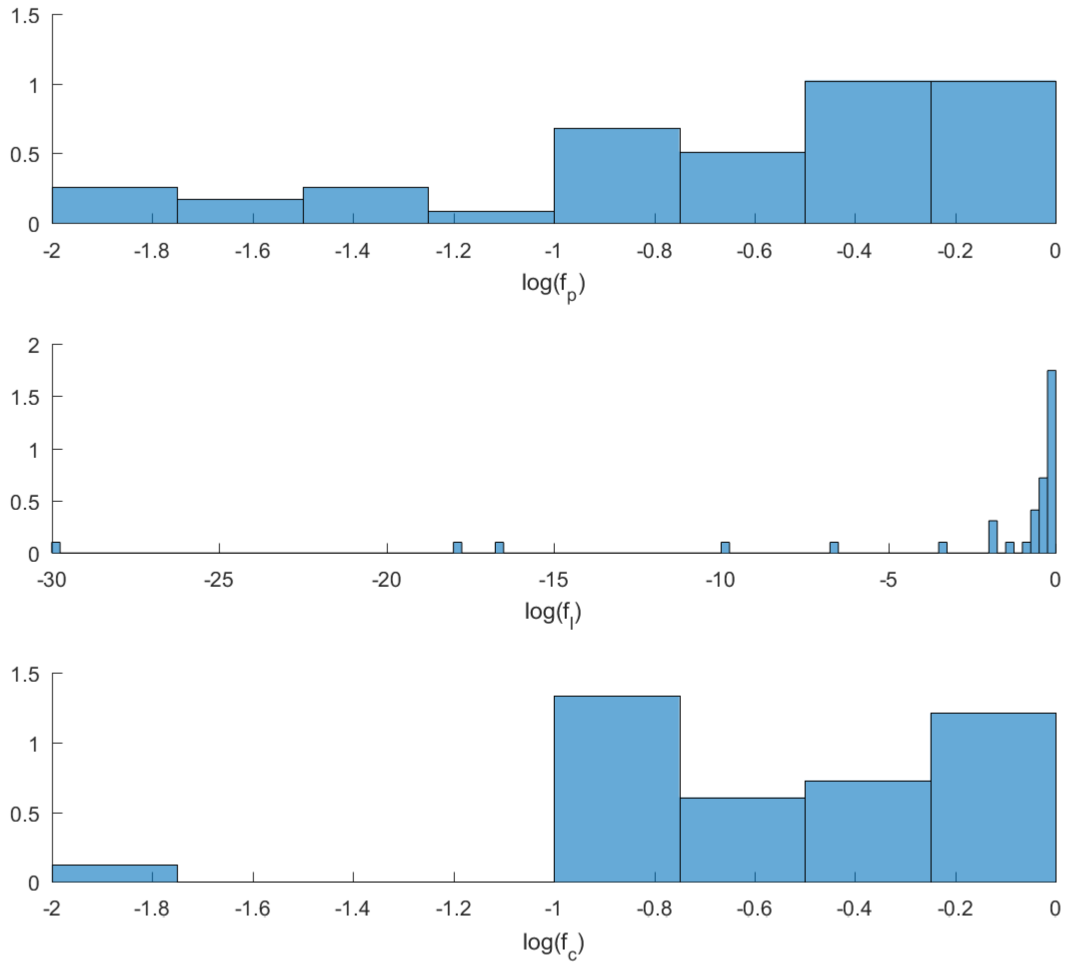

Histogram of parameter values, on log scales. Note the existence of extreme outlier for f_l due to (Behroozi & Peeples 2015).

Probability density from resampling

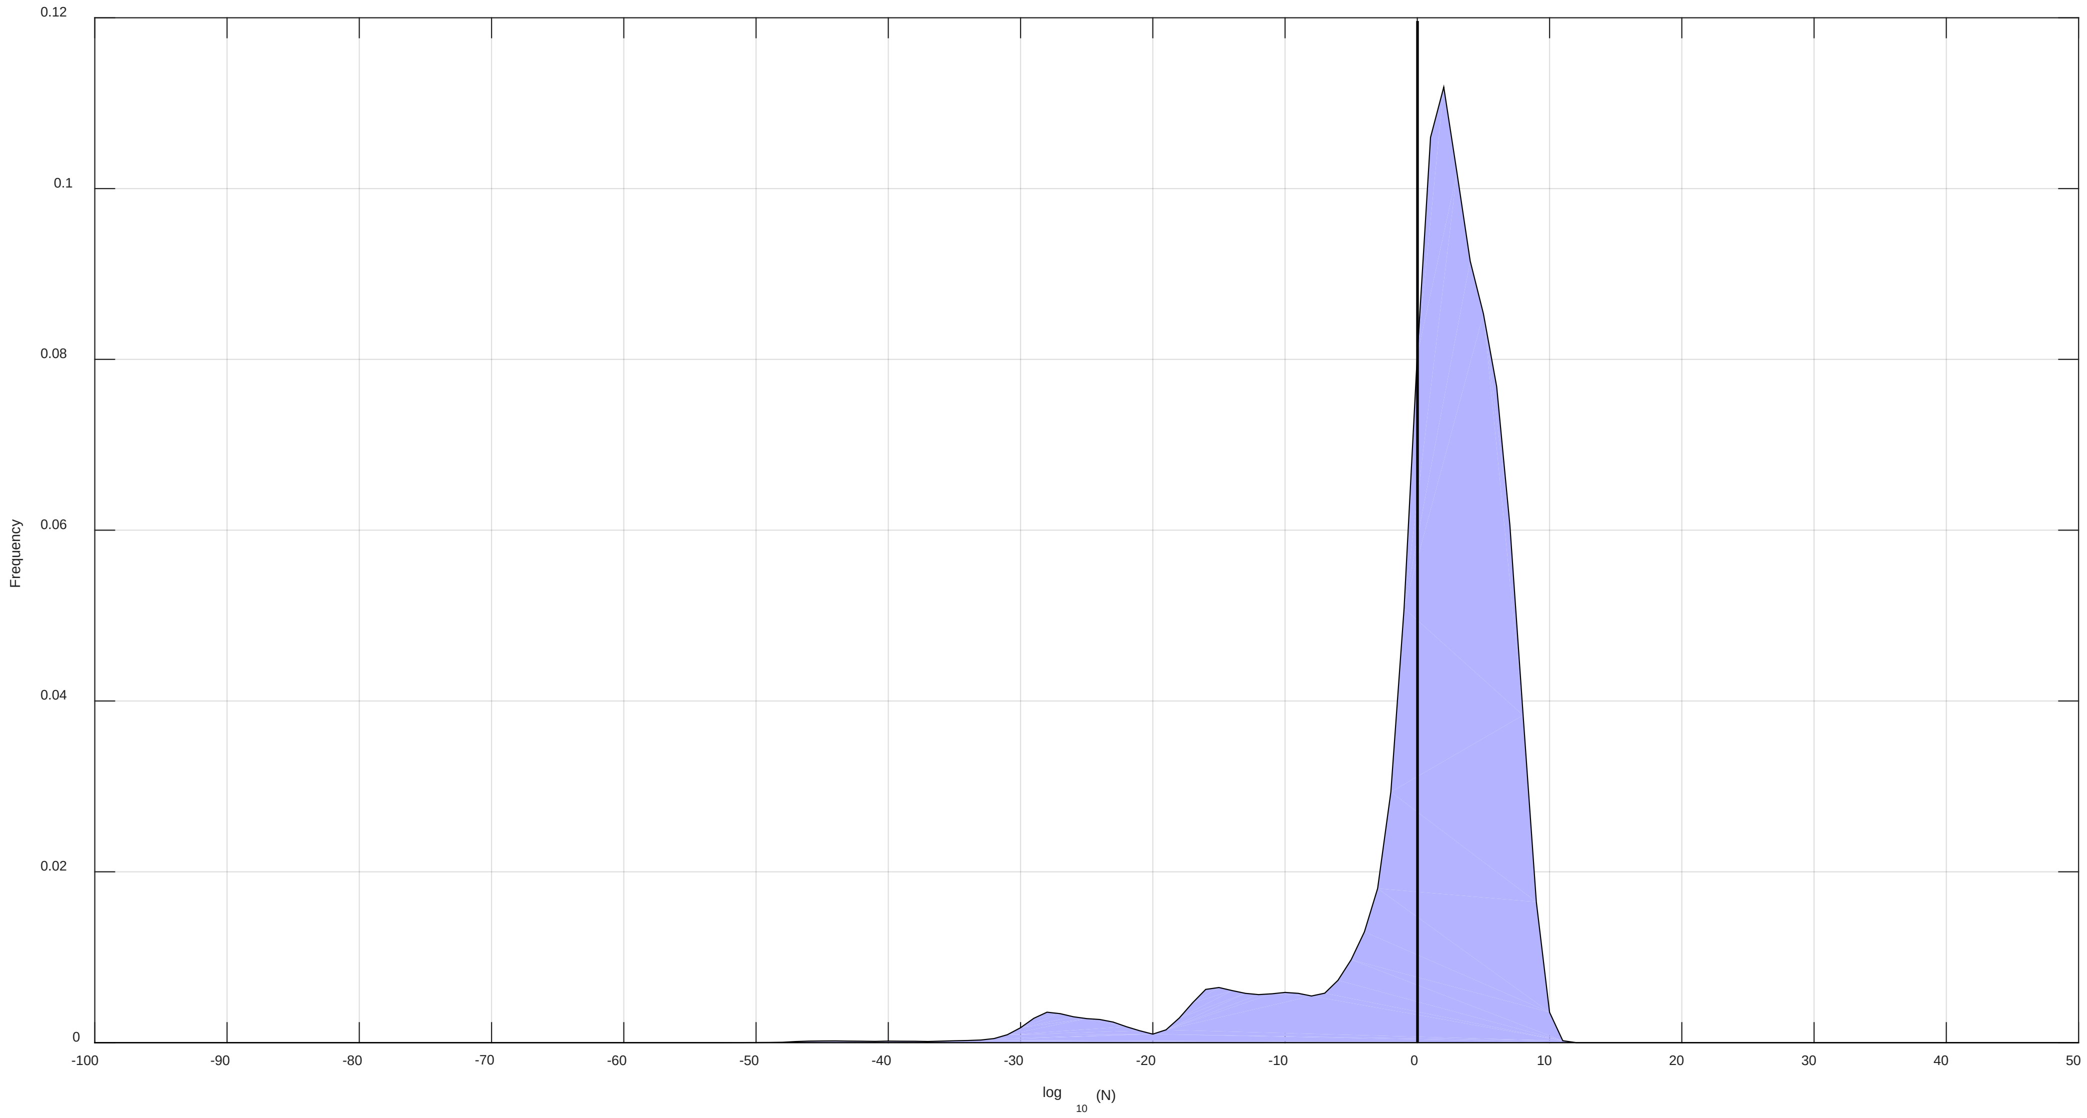

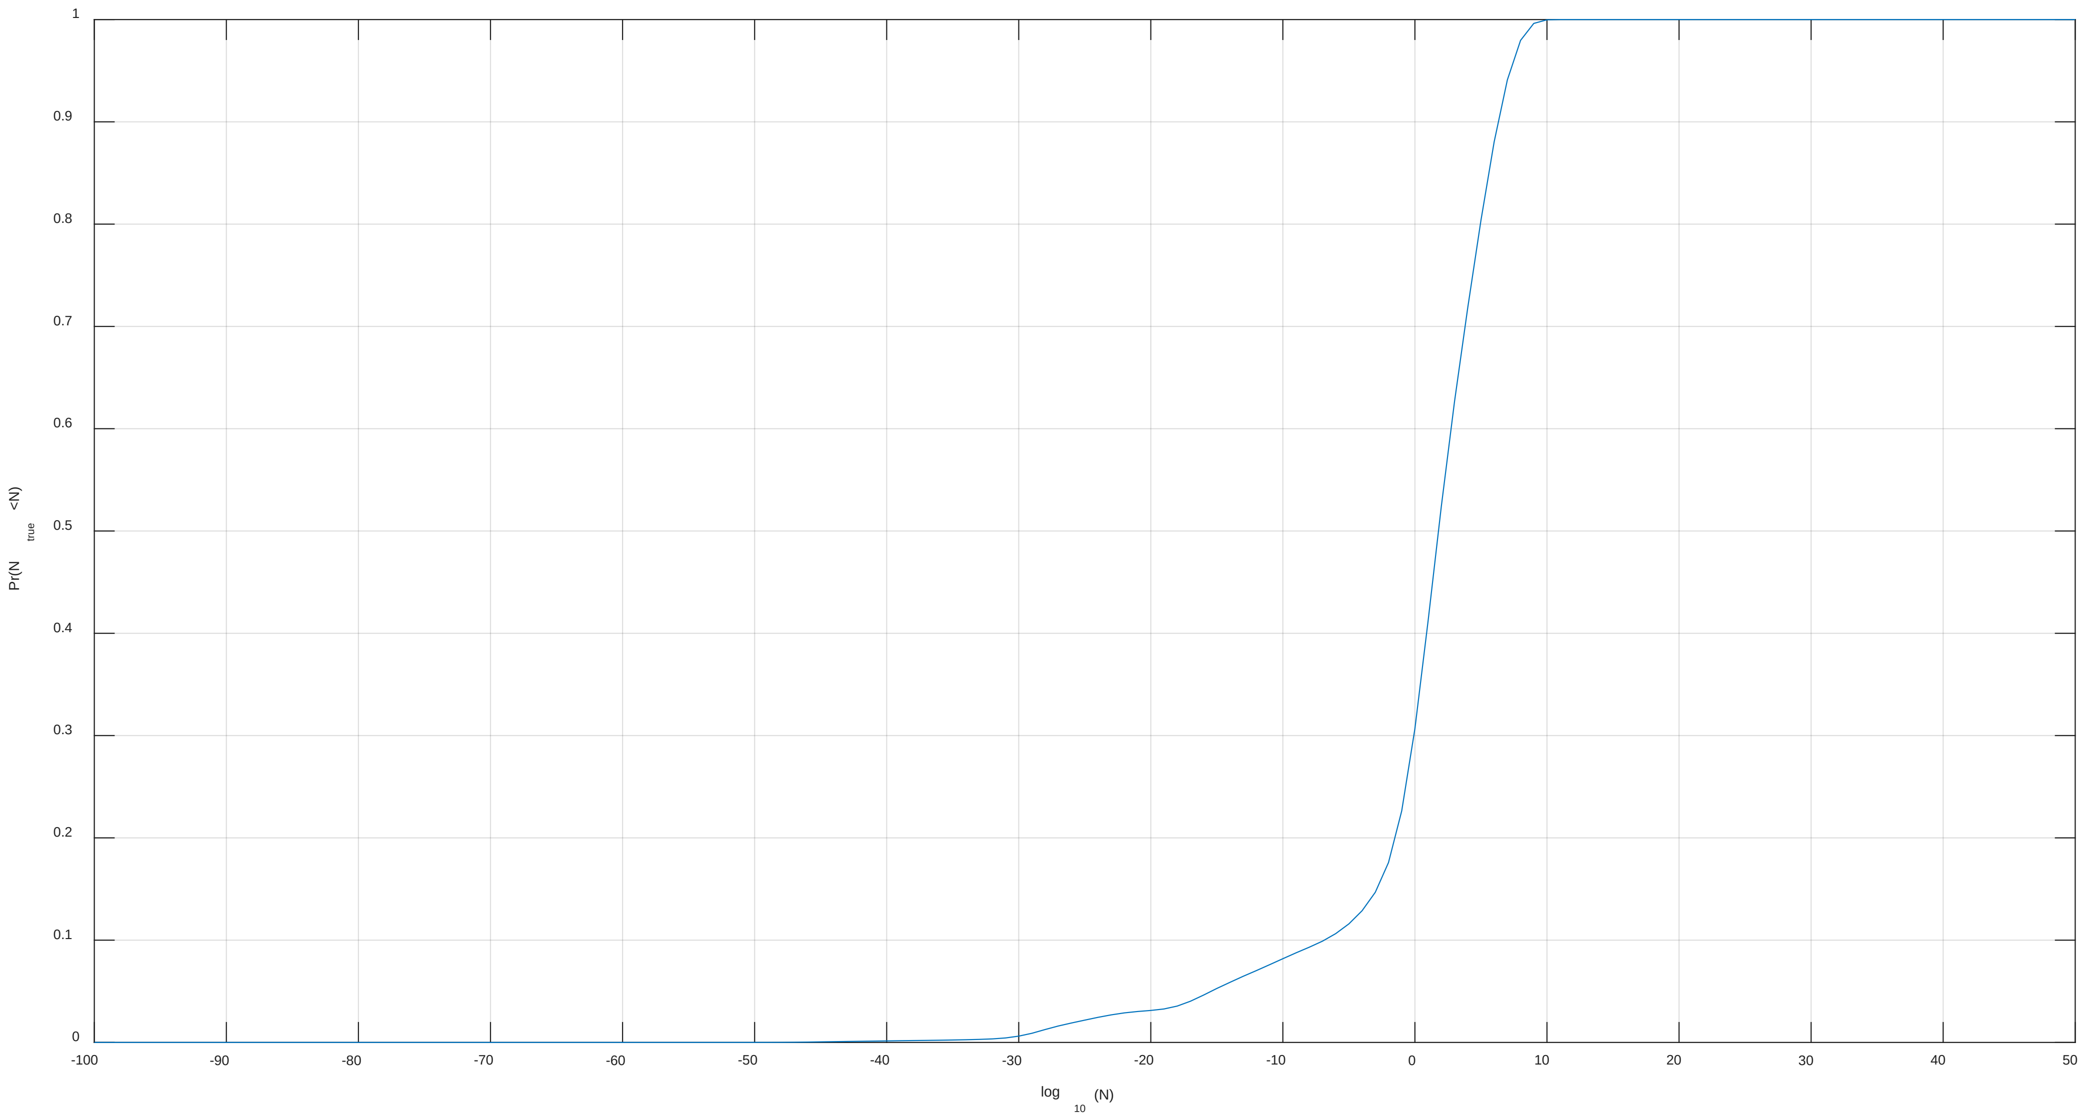

1
0.9
0.8
0.7
0.6
0.5
0.4
0.3
0.2
0.1
0
Pr(N true <N)
-100
-90
-80
-70
-60
-50
-40
-30
-20
-10
0
10
20
30
40
50
log 10 (N)

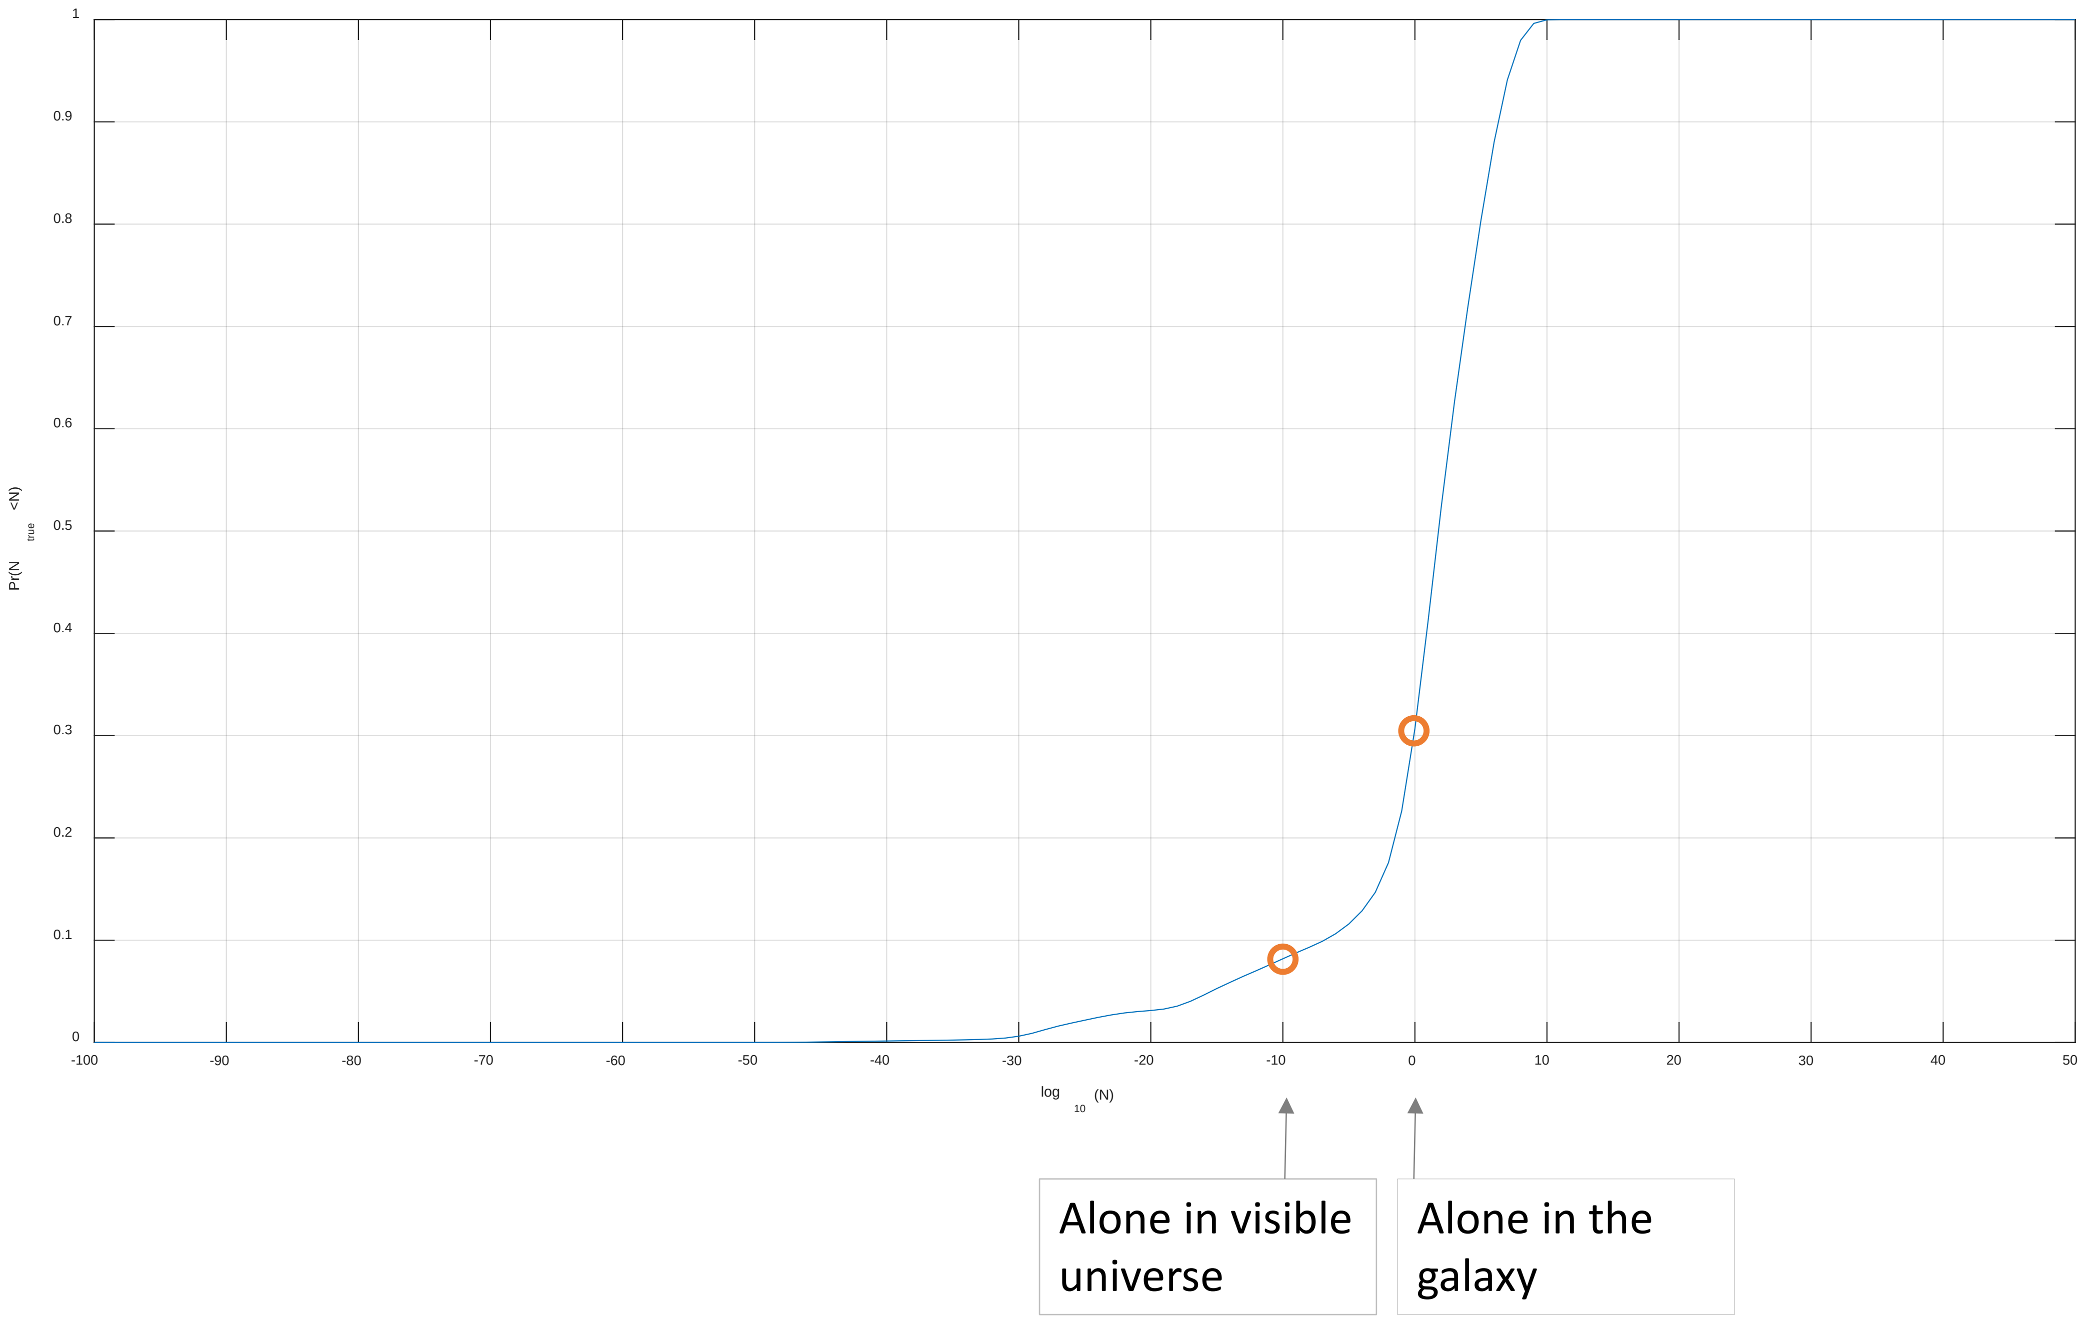

1
0.9
0.8
0.7
0.6
0.5
0.4
0.3
0.2
0.1
0
Pr(N true <N)
-100
-90
-80
-70
-60
-50
-40
-30
-20
-10
0
10
20
30
40
50
log 10 (N)
Alone in visible universe
Alone in the galaxy

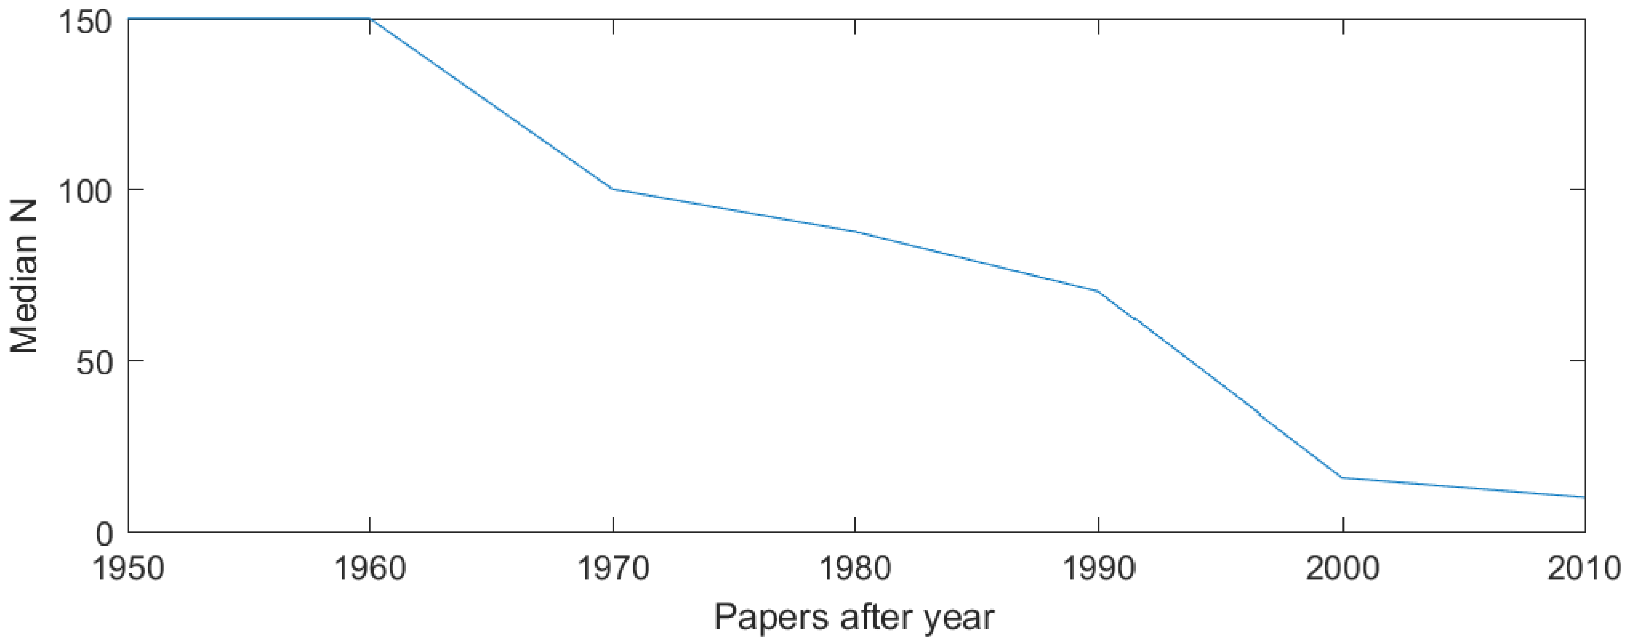
150
100
50
0
Median N
1950
1960
1970
1980
1990
2000
2010
Papers after year

Are there good prior distributions for the Drake equation?

- N_s: fairly well constrained.
  - Max 5 orders of magnitude given other galaxies; actual current uncertainty likely <1 order magnitude .
  - Time variation issues!
- f_p: increasingly clear ≈1
- n_e: from rare earth arguments ($<10^{-12}$) to >1
- f_l: Very uncertain; [Next slides]
  - Absolute lower limit due to ergodic repetition: $10^{-10^{115}}$
- f_i: Very uncertain; [Next slides]
  - $5 \cdot 10^9$ species so far, 1 intelligent: $2 \cdot 10^{-10}$
  - But also around 10^7 species at a time, 1/500 per assemblage
- f_c: Very uncertain; (human case 0.000615 so far)
- L: Uncertain; $50? < L < 10^9 - 10^{10}$ years (upper limit because of Drake applicability)

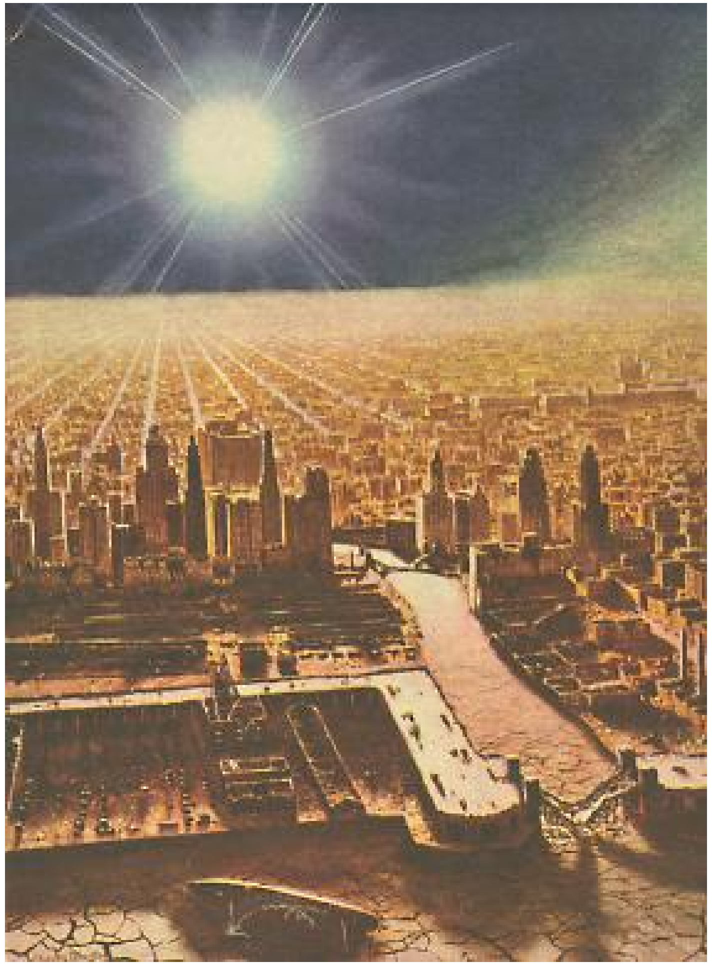

Abiogenesis as a physical process

- Instead of thinking in terms of fraction of planets having life, consider a rate of life formation in suitable environments: what is the induced probability distribution?
- The emergence of a genetic system is a physical/chemical transition
- Transition events occur in some medium at some rate per unit volume: $f_L \approx \lambda V t$
  - High rates would imply that almost all suitable planets originate life
  - Low rates would imply that almost no suitable planets originate life
- The nature of transitions and suitable media are largely unknown.

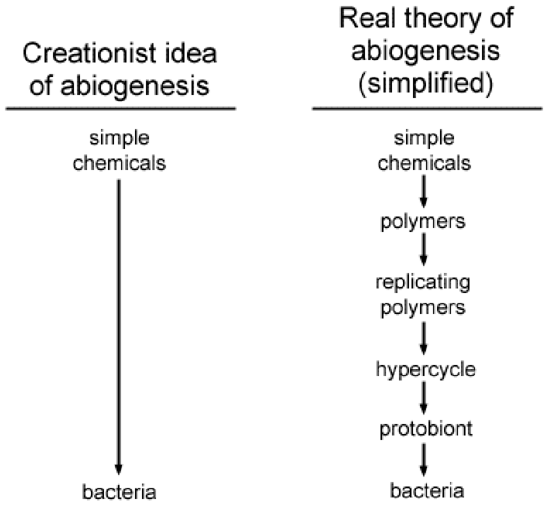

Abiogenesis

- What range of rates is possible given current knowledge?
  - Uncertainty regarding time when possible at least 3 orders of magnitude (10^7-10^{10} years)
  - Uncertainty regarding volumes spans 20+ orders of magnitude
  - Uncertainty regarding rates can span 100+ orders of magnitude
    - Combinatorial flukes? Protein folding? Reaction rates vary a lot.
- Spontaneous generation *could* conceivably be common and fast!

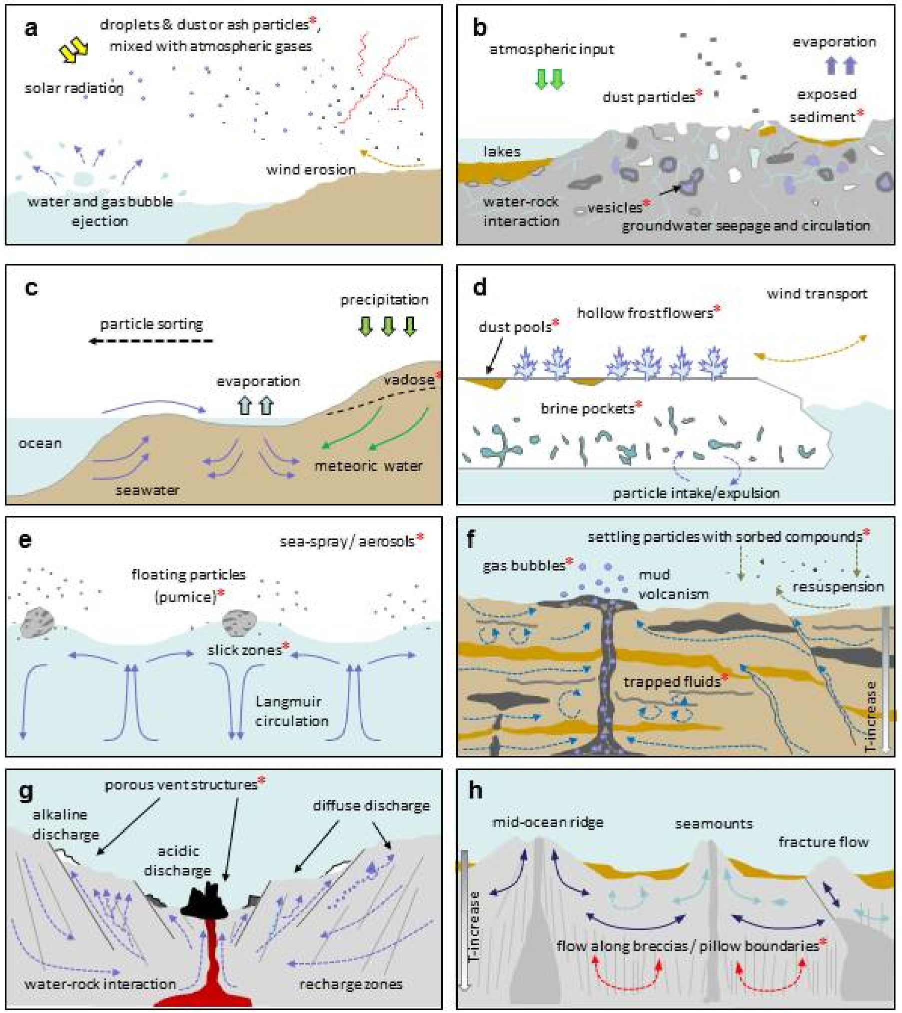

Genetic transitions: potential alternative forms of life

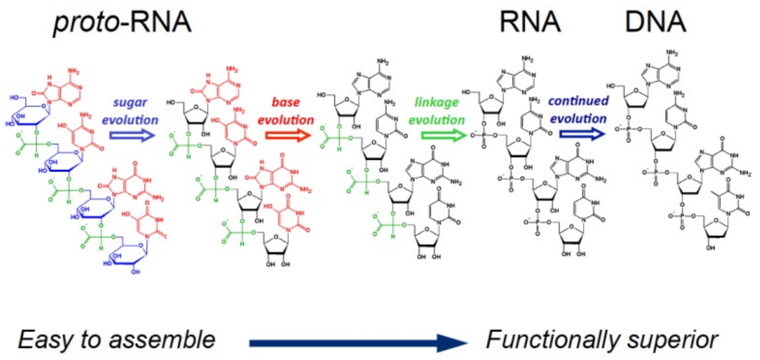

- All life on Earth shares almost exactly the same genetic systems
  - Only rare and minor changes have occurred in $\approx 10^{40}$ cell divisions
  - Nonetheless, other genetic systems preceded the modern form
- The transition to the modern form required major changes
  - It would be unsurprising if the rate were < 1 per 10^{100} cell divisions
  - Modern genetics required >1/5 the age of the universe to evolve intelligence
- A genetic system like the one that preceded ours might
  - (1) Be stable across $>10^{100}$ cell divisions
  - (2) Evolve more slowly by a factor of 10, and run out the clock
- If the rate of discovering it is λ_B and the rate of discovering "our" kind of capable life is λ_A, then the fraction of A-life is λ_A/λ_B.
  - Rates can differ many orders of magnitude, producing a life-rich but evolution/intelligence-poor universe.
  - Multiple step models add integer exponents to rates: *multiply* order of magnitude differences.

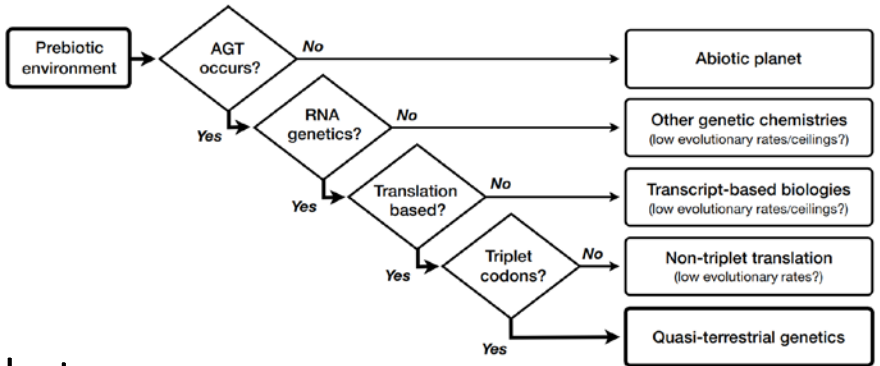

Drake equation in a Bayesian framework

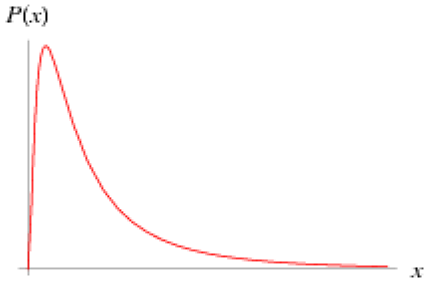

- Work on log-space: sum of bunch of log-distributions
  - Suitable because order of magnitude uncertainties
  - Log-uniform and lognormal are rather natural and simple
    - Log-uniform is scale free; lognormal is maximum entropy and stable
  - "Log-Drake": $\log N = \sum lx_i$
- Monte Carlo sample resulting distribution
- Our priors (mostly for illustration): $lN_s \sim U(0,2), lf_p \sim U(-1,0), ln_e \sim U(-1,0), \lambda_l \sim N(0,50), lf_i \sim U(-3,0), lf_c \sim U(-2,0), lL \sim U(2,9)$

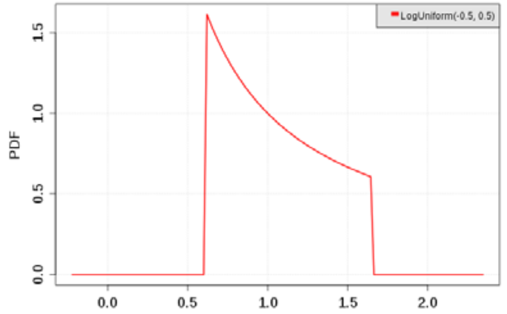

Monte Carlo results

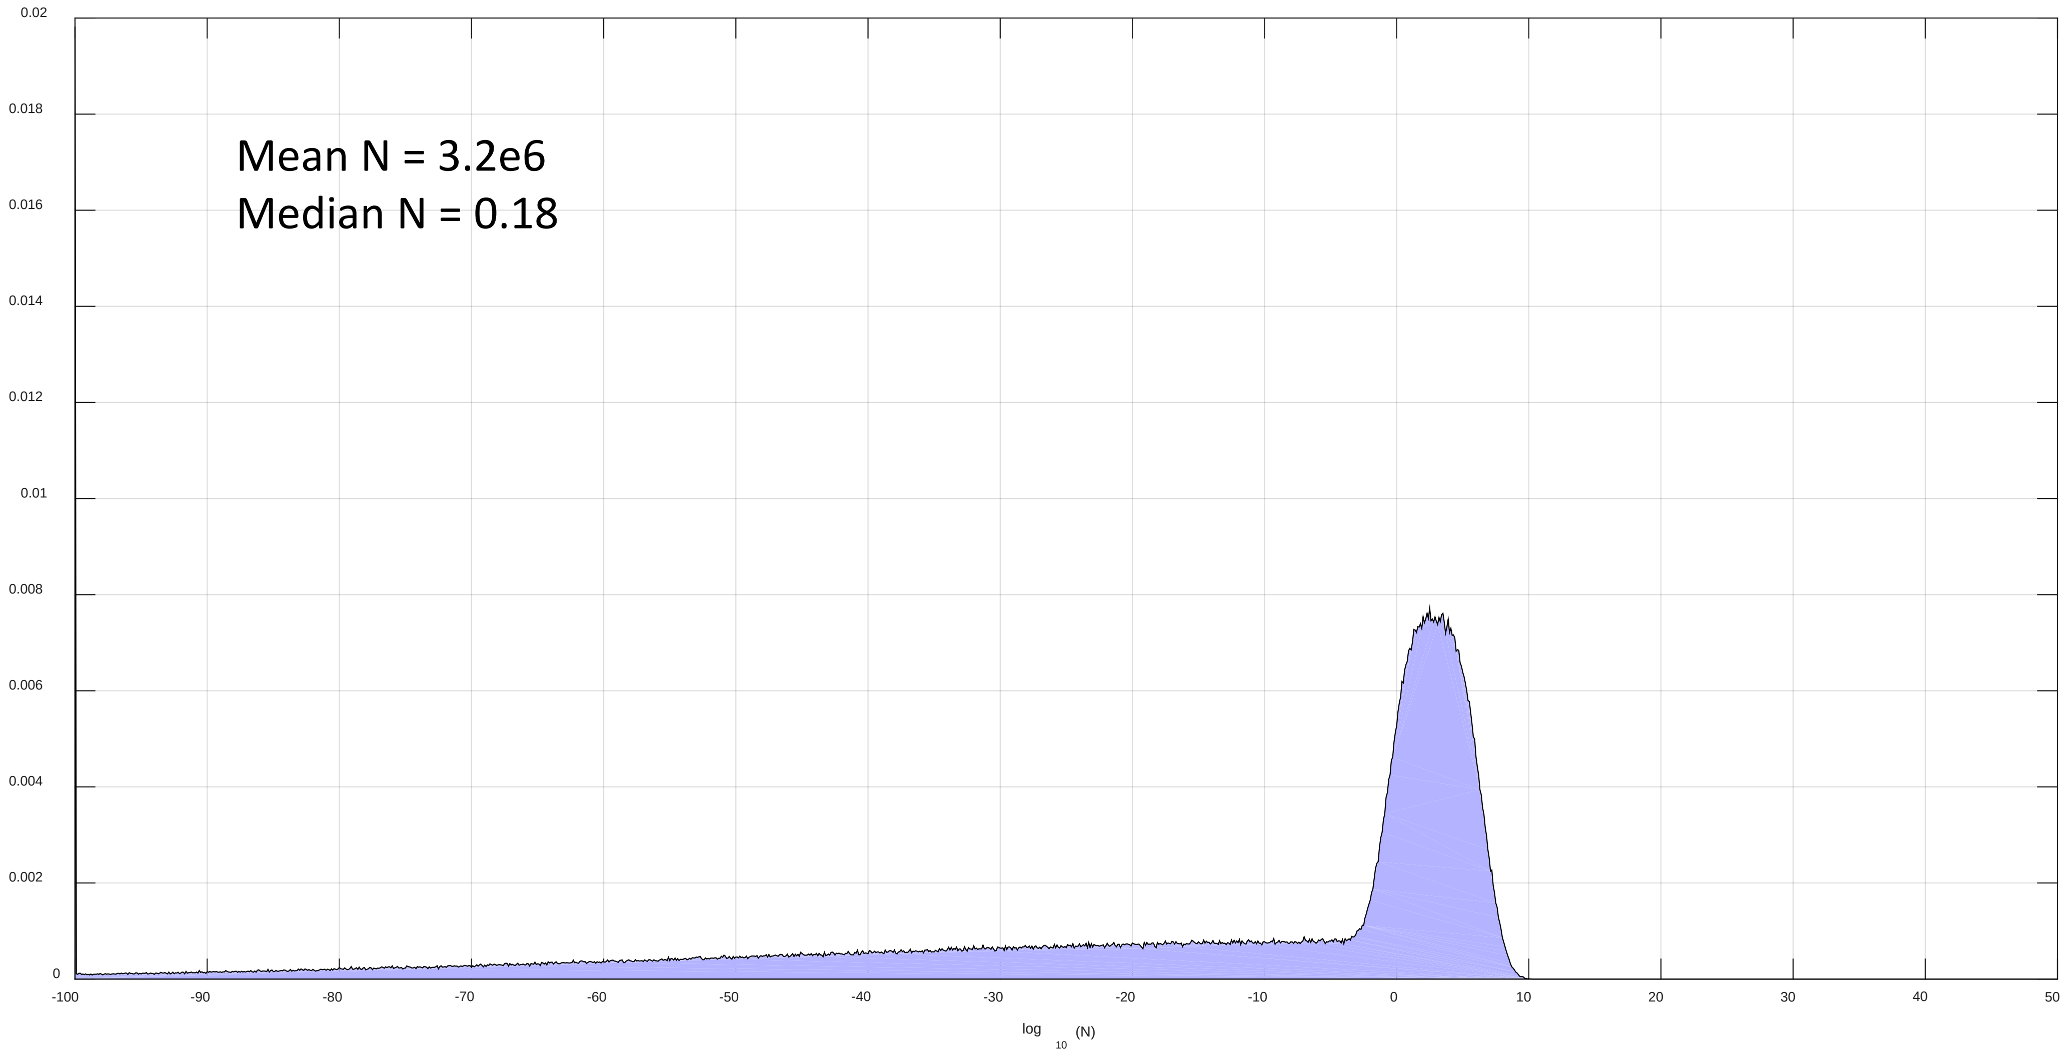

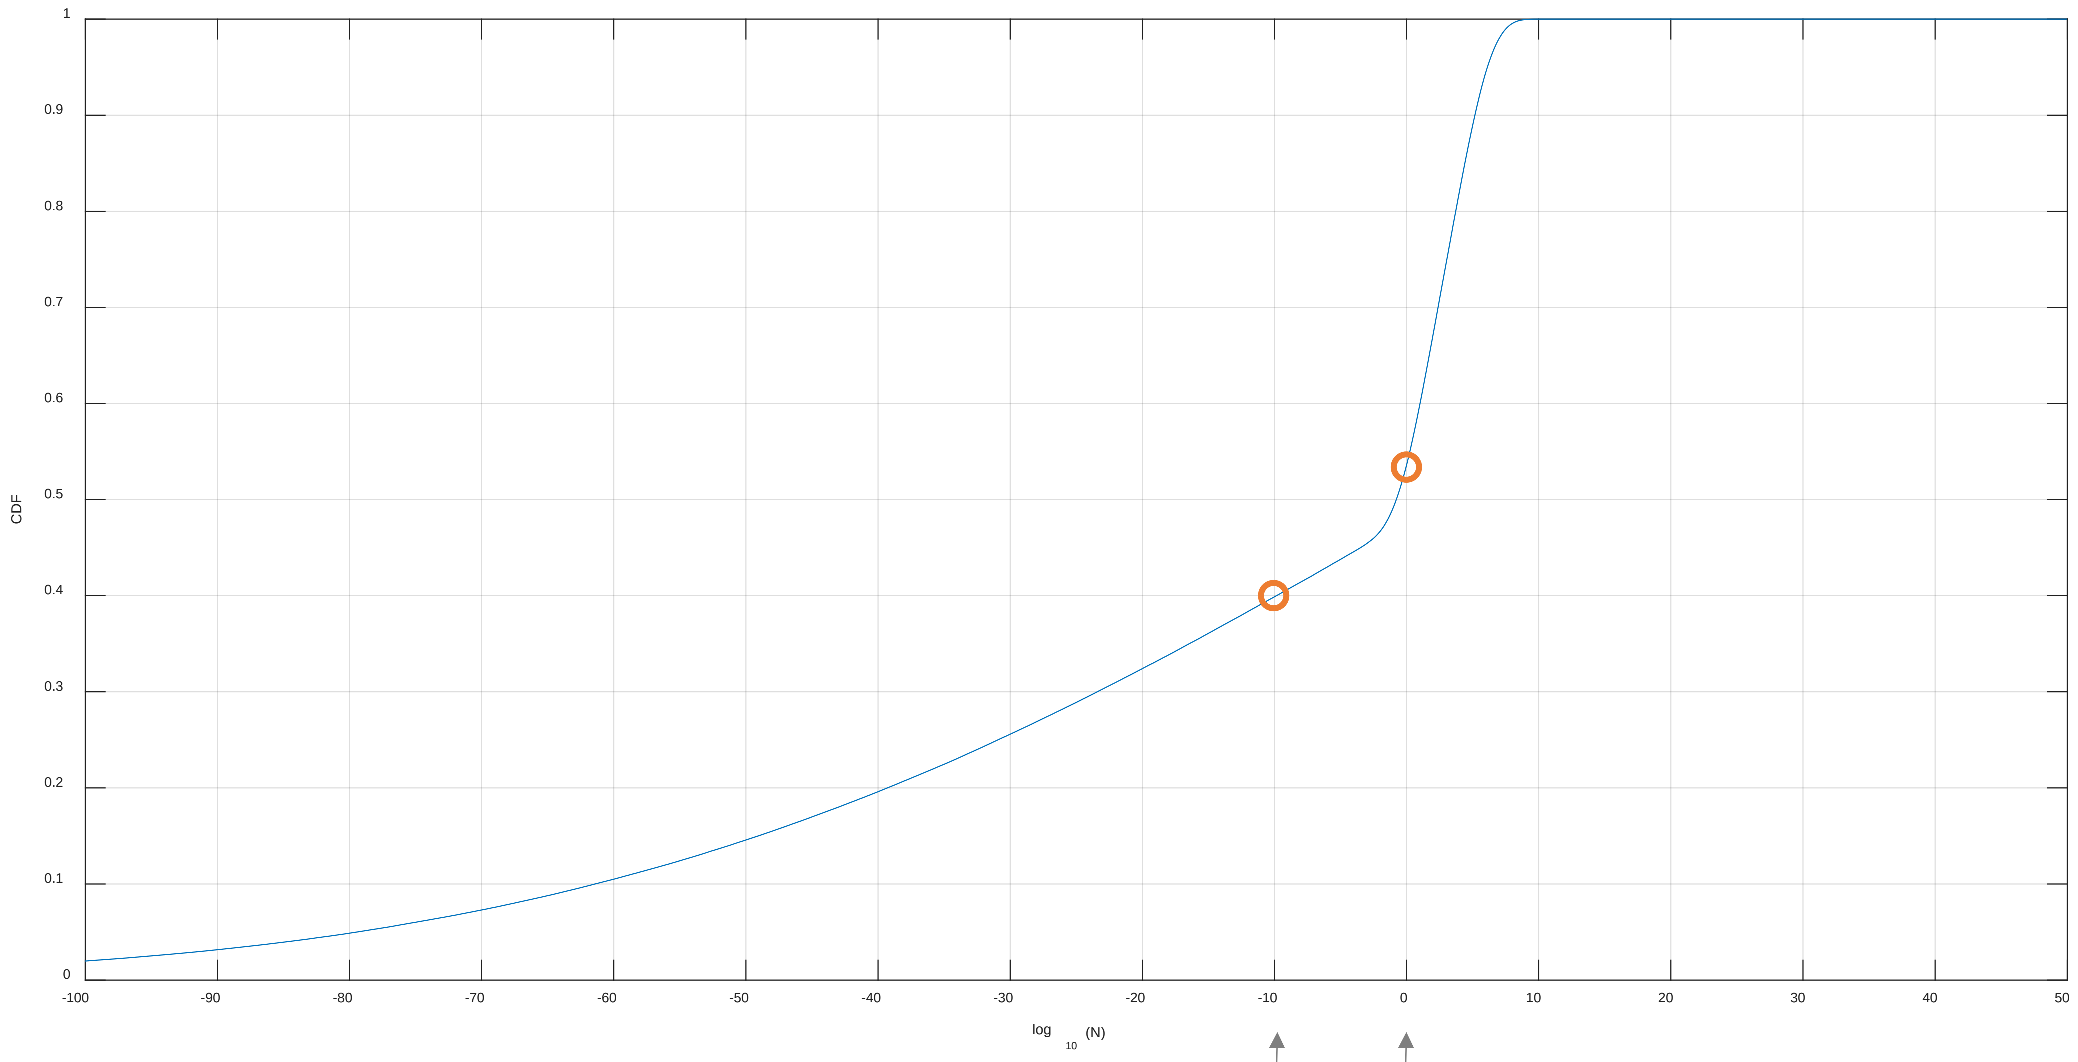

40% chance we are alone, despite *very* optimistic mean!

Alone in visible universe

Alone in the galaxy

Conclusion 1: the Fermi paradox isn't very paradoxical

- Overconfident guesses makes it seem hard to get empty universe
- When our uncertainty is properly accounted for in the model, we find a substantial *a priori* chance that there is no other intelligent life in our observable universe, and thus that there should be little or no surprise when this is what we see.
- Reasonable priors (or even the literature!) give enough uncertainty to make empty universe fairly likely
  - In order to produce a non-empty universe but not an overabundant one parameters need to lie in a small interval (Carter)
  - Also similar to Tegmark's argument, but with more process
- Note that this conclusion does not mean we *are* alone! Just that we should not be surprised if this is the case.
  - This is a statement about knowledge and priors, not a measurement: armchair astrobiology

The Fermi observation and oblong distributions

- No visible aliens: what is the effect on parameters?
  - $f(x_i | \sum_i x_i < \theta)$
- Oblong joint distributions
  - $f_X(x_1, x_2, \dots, x_n) = \prod_i f_i(x_i)$ (conveniently independent)
  - $Var(x_1) \gg Var(x_j)$
  - Oblong distributions react most with their most uncertain component!

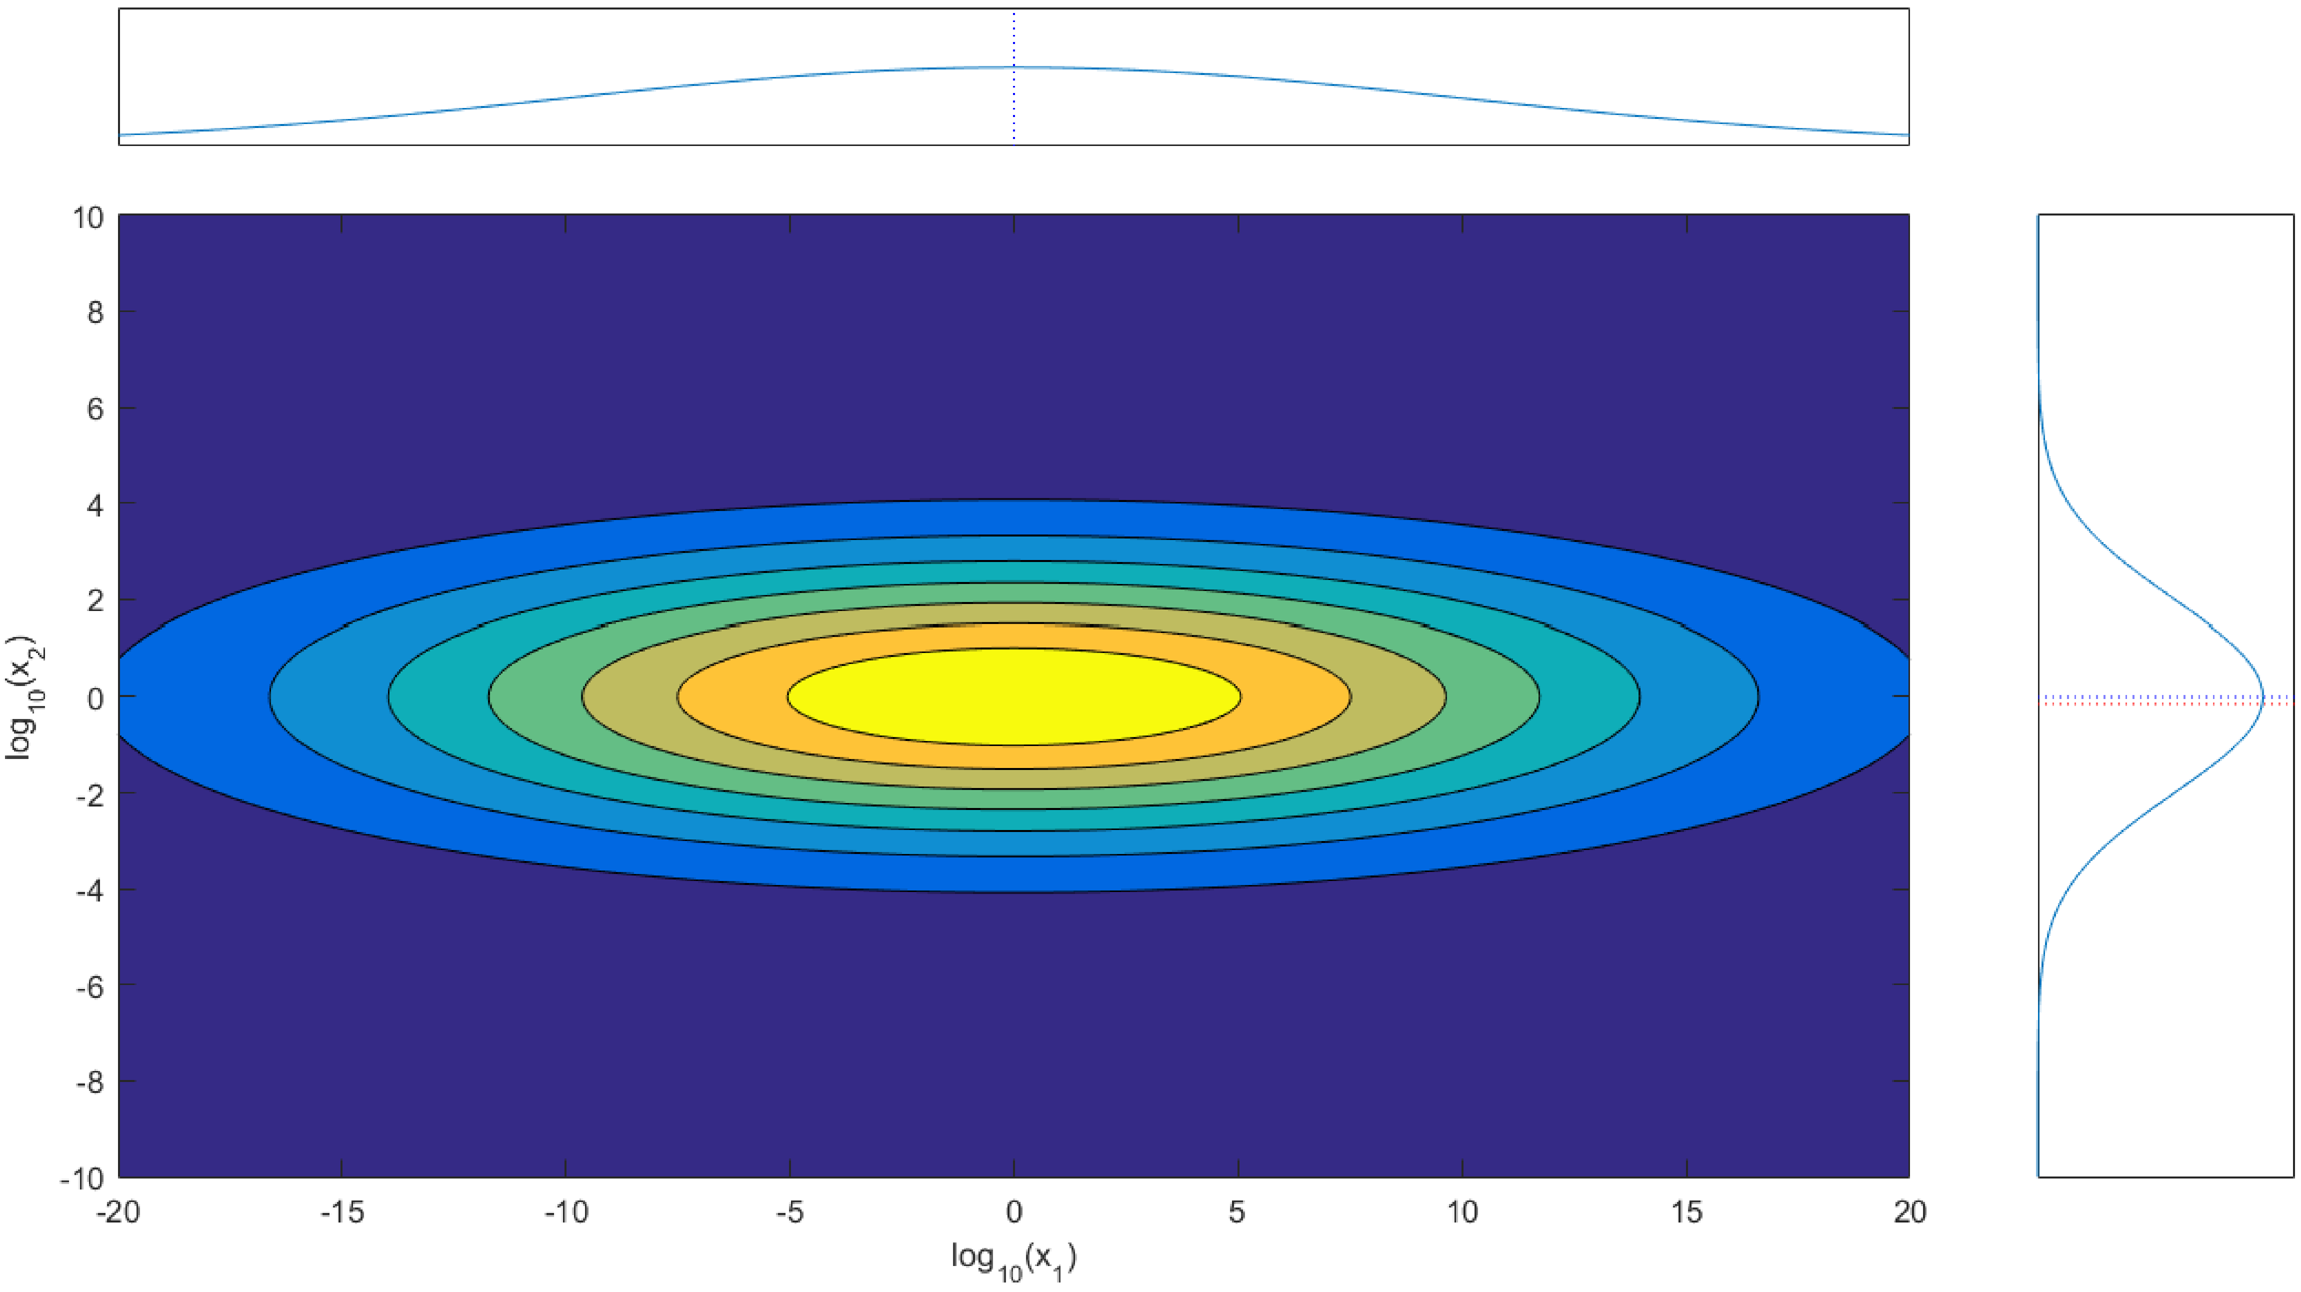

10
8
6
4
2
0
-2
-4
-6
-8
-10
log10(x2)
-20
-15
-10
-5
0
5
10
15
20
log10(x1)

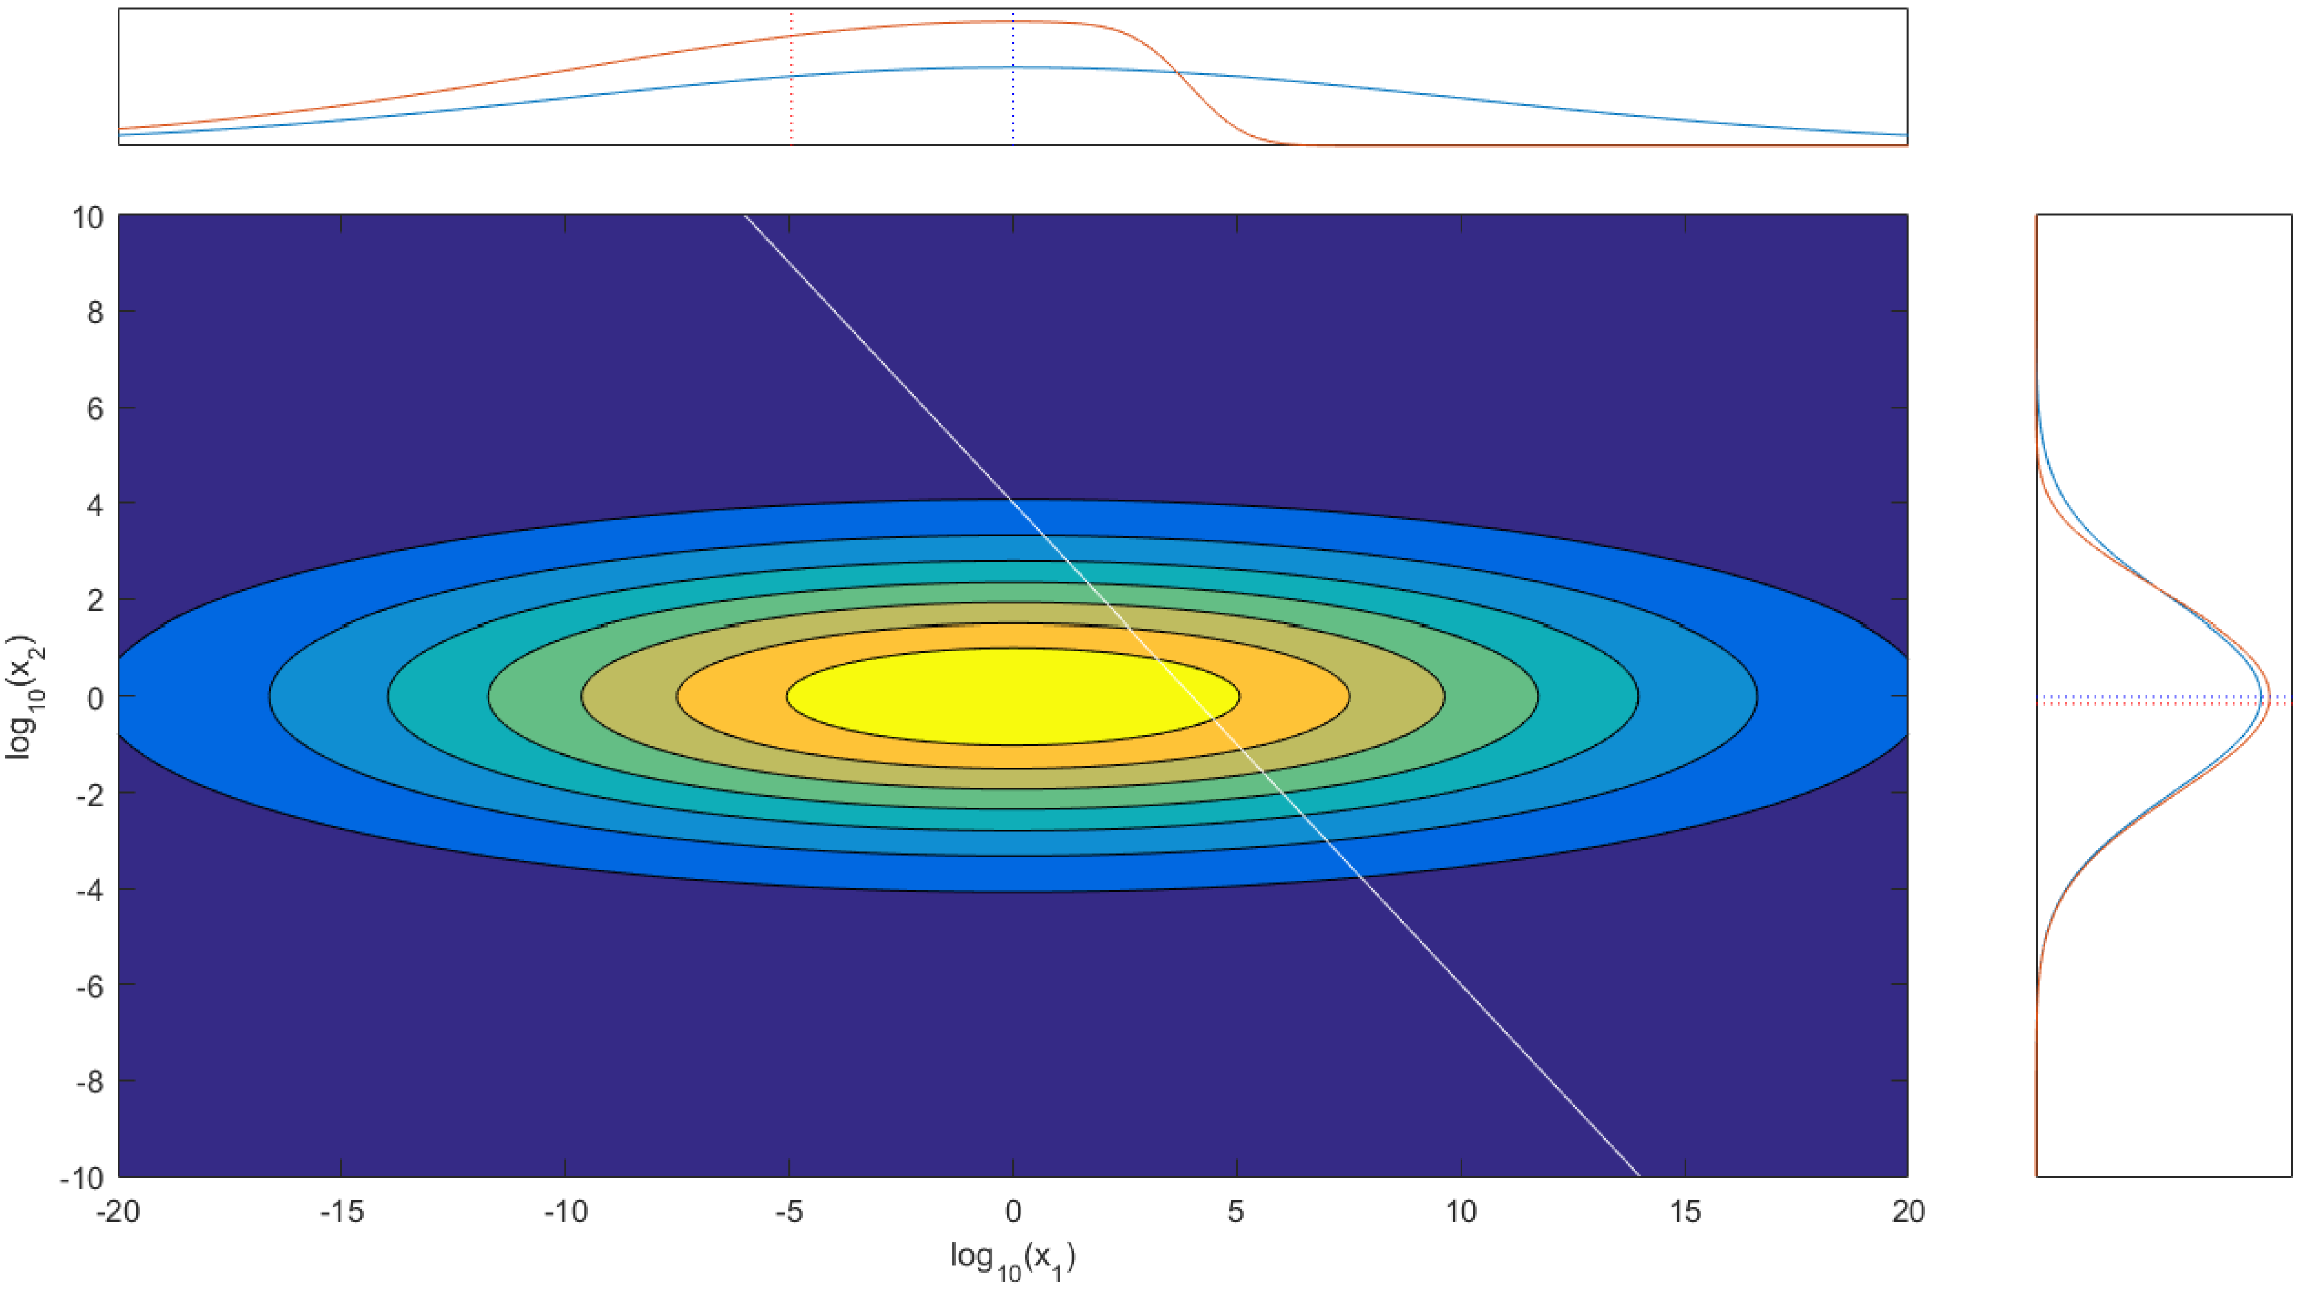

10
8
6
4
2
0
-2
-4
-6
-8
-10
$\log_{10}(x_2)$
-20
-15
-10
-5
0
5
10
15
20
$\log_{10}(x_1)$

Conditioning on a small joint tail moves the most uncertain component the most

- Narrow distributions experience conditioning as rescaling (removed by normalization), while broad experience it as cutting off tail.
- "Easy" to prove for rectangular distributions
- Gaussians messy but analytically doable
- Some counterexamples for special distributions

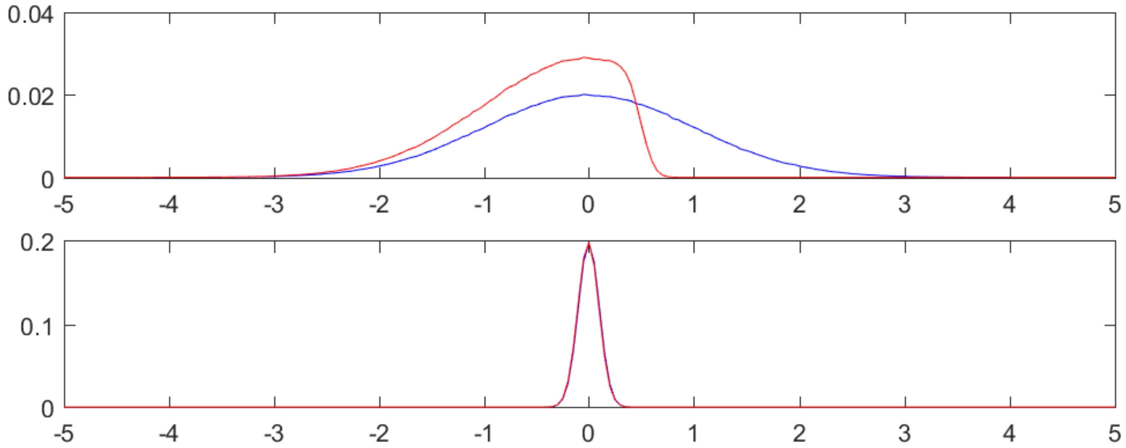

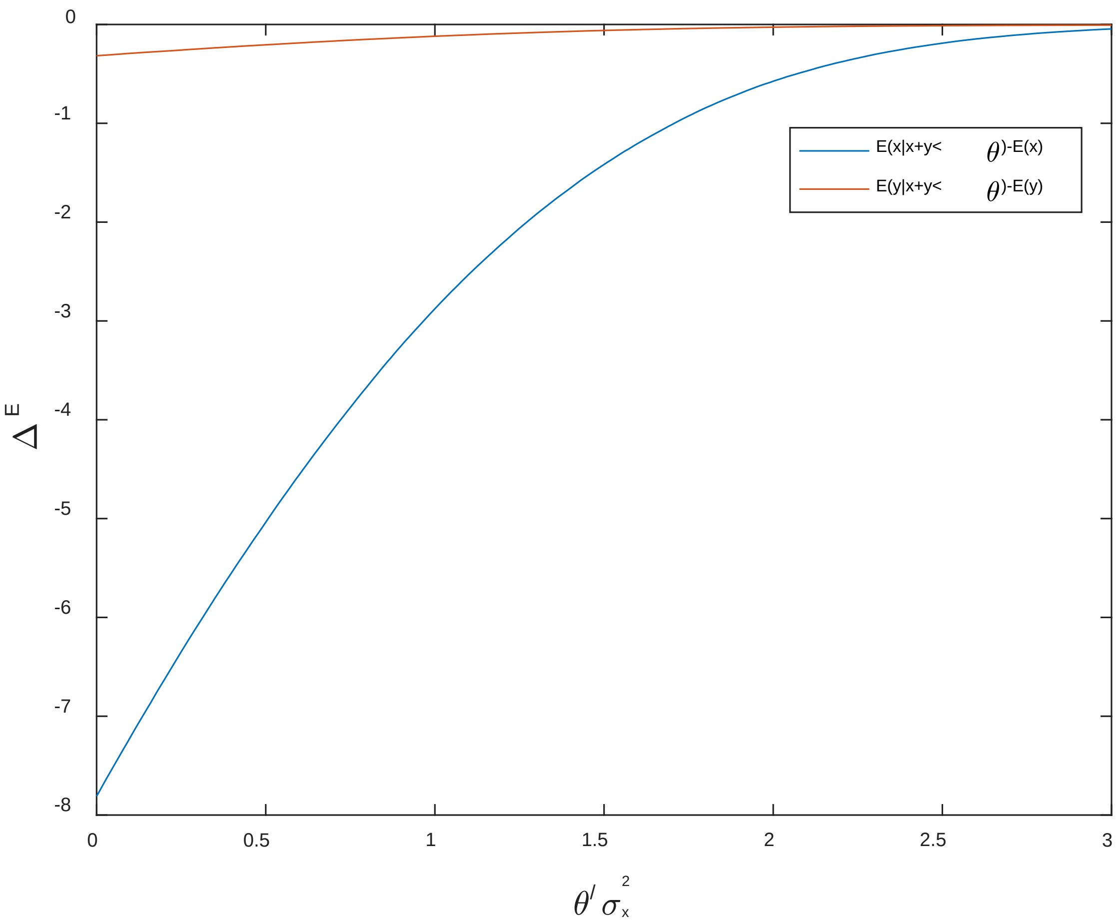

Priors for the Drake equation produce an oblong joint distribution

- The life and intelligence probability distributions move many orders of magnitude more than the others – tail-clipping observations have bigger effects on past great filter.
- Even very weak observations move them.

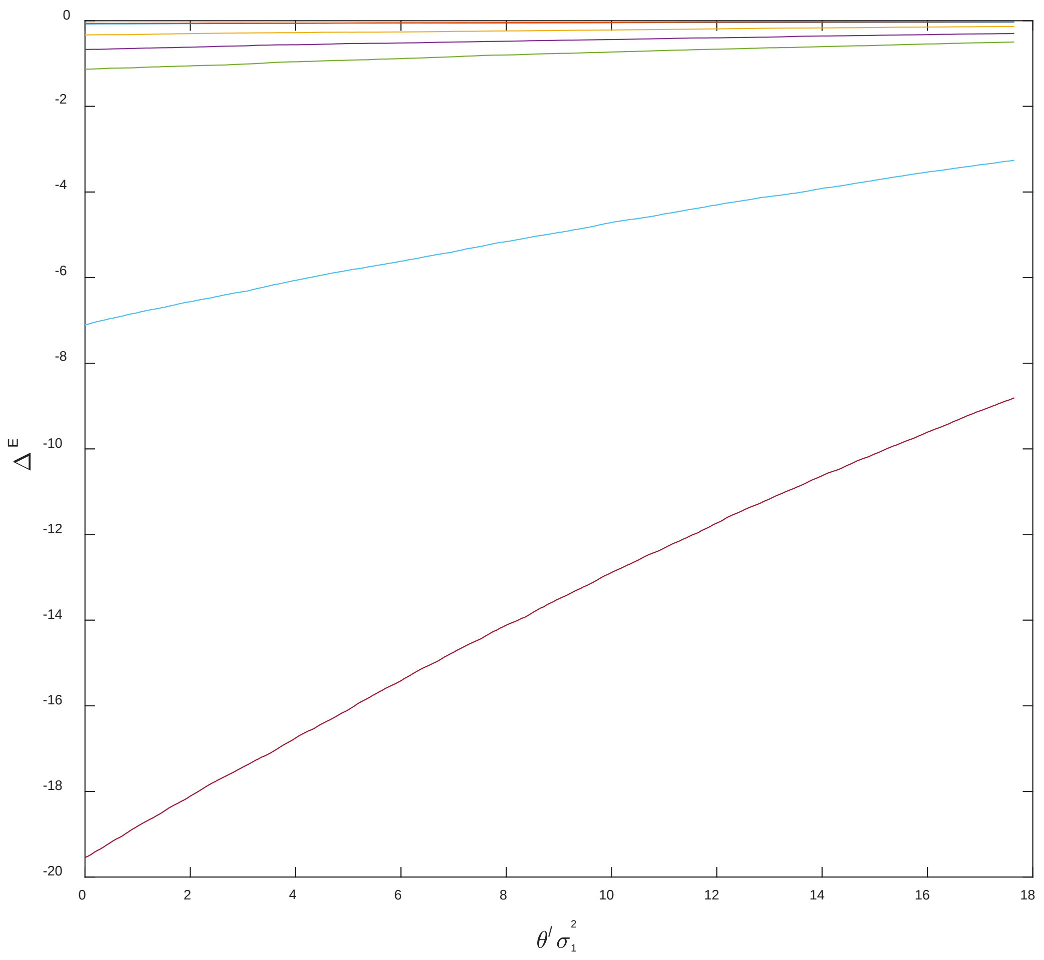

Modelling the Fermi observation

- Simple cut-off
  - 90% reduction for $N > N_{\text{threshold}}$ where $N_{\text{threshold}}$ is some large number.
  - 10% chance we are totally wrong about everything
- Failure to detect after sampling K stars out of N_{MW}:
  - $\Pr(\text{no detection}|N) = \left(1 - \left(\frac{N}{N_{MW}}\right)\right)^K$
- Observability within radius
  - If can see out to distance d, $\Pr(D_{closest} > d|N) = 1 - e^{-4\pi(N/N_{MW})d^3/3}$
- Ĝ search model
  - $\Pr(\text{no detection}|N) = 1 - P_{K3}(1 - (1 - P_{succ})^K)$
- Galactic settlement models
  - $\Pr(\text{no detection}|N) \approx e^{-N} + (1 - e^{-N})\left(1 - \frac{L^\alpha}{(\alpha+1)T^\alpha}\right)$
  - $\Pr(\text{no detection}|N) \approx e^{-\left(\frac{N}{L}\right)(T_{MW}-T)} + \left(1 - e^{-\left(\frac{N}{L}\right)(T_{MW}-T)}\right)\left(\frac{\alpha}{\alpha+1}\right)$
- Can adjust for miss probability

$$\Pr(N|\text{no detection}) = \frac{\Pr(\text{no detection}|N)\,P(N)}{P(\text{no detection})}$$

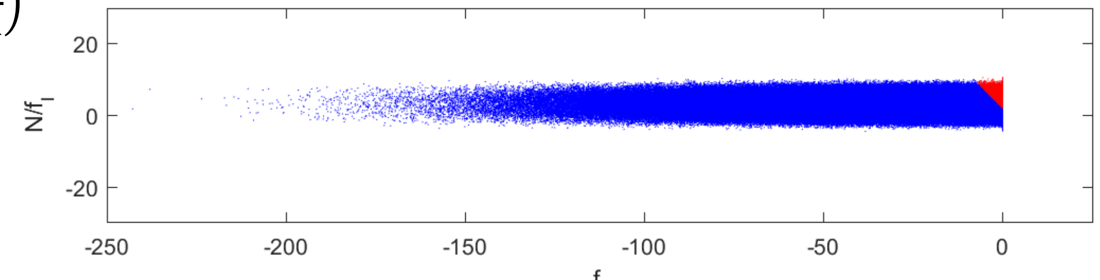

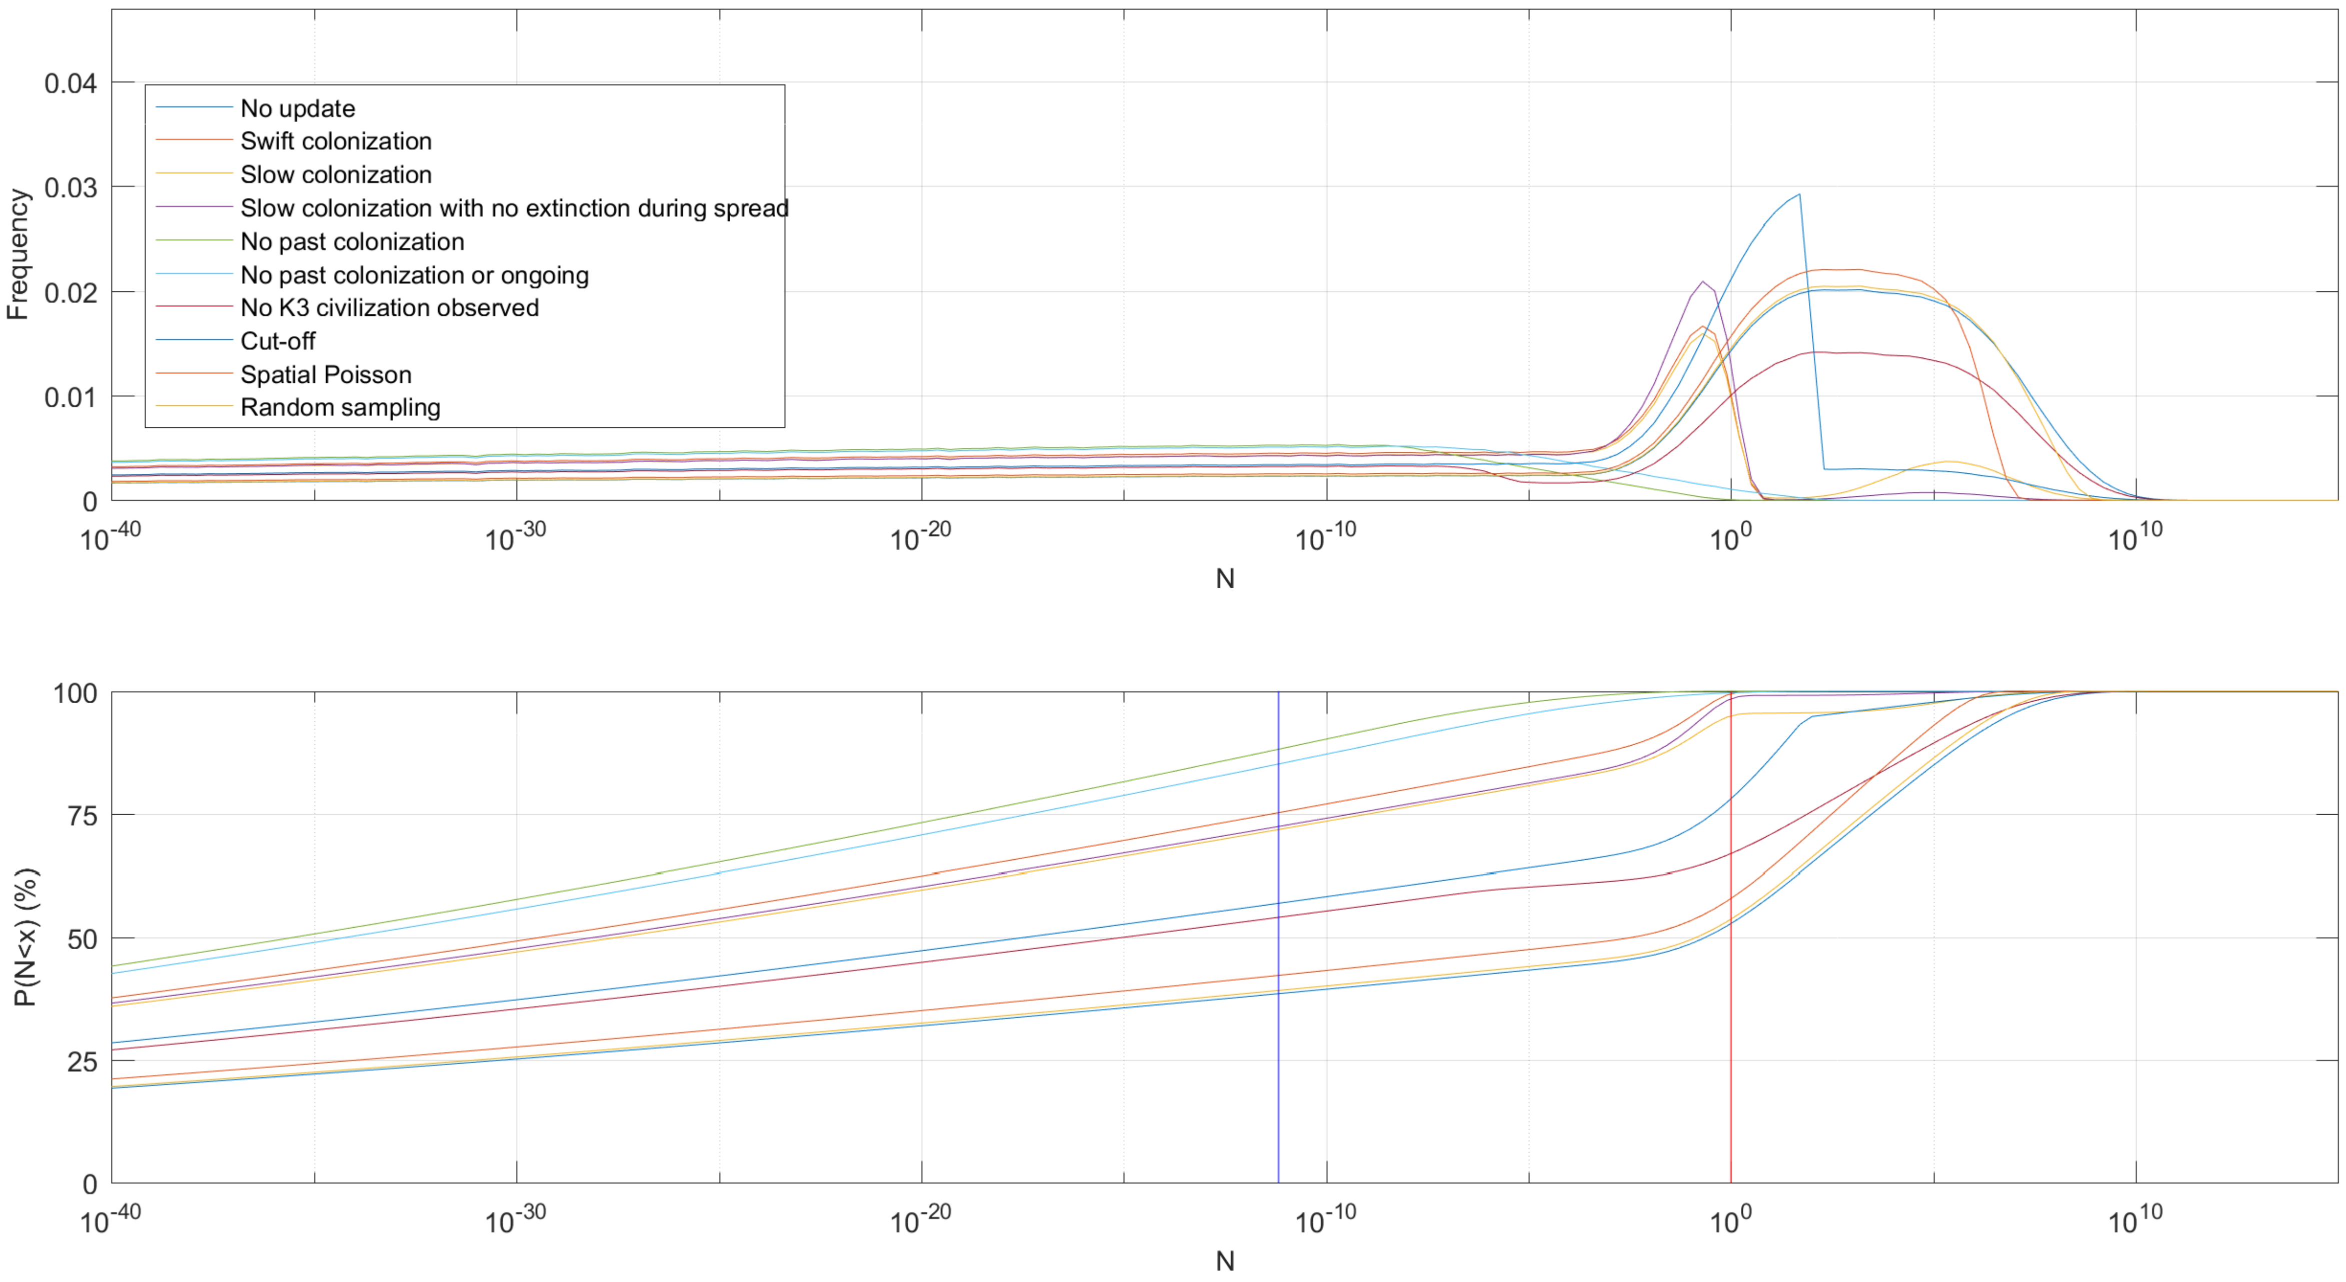

No update
Swift colonization
Slow colonization
Slow colonization with no extinction during spread
No past colonization
No past colonization or ongoing
No K3 civilization observed
Cut-off
Spatial Poisson
Random sampling
Frequency
0.04
0.03
0.02
0.01
0
10^{-40}
10^{-30}
10^{-20}
10^{-10}
10^{0}
10^{10}
N
P(N<x) (%)
100
75
50
25
0
10^{-40}
10^{-30}
10^{-20}
10^{-10}
10^{0}
10^{10}
N

Table 1. Comparison of conditioned credence distributions.

Update	Mean N	Median N	$\Pr[N < 1]$	$\Pr[N < 10^{-10}]$	Median f_l	Median L
No update	2.7×10^7	0.32	0.52	0.38	0.64	1×10^6
Random sampling	2.5×10^6	0.19	0.53	0.39	0.09	8.6×10^5
Spatial Poisson	7.8×10^4	0.0048	0.57	0.42	3.1×10^{-6}	4.5×10^5
No K3 civilization observed	1.9×10^7	1.2×10^{-15}	0.66	0.54	4×10^{-19}	9×10^5
Settlement update	0.072	8.1×10^{-35}	0.996	0.85	3×10^{-38}	1×10^6

Original mean $\log(f_L) = -19.97$
Posterior mean $\log(f_L) = -28.49$

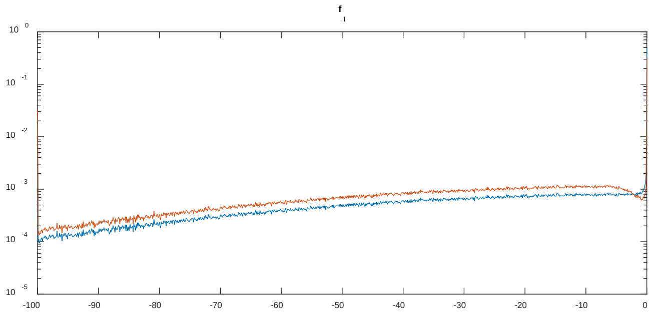

Original mean $\log(L) = 5.50$
Posterior mean $\log(L) = 5.02$

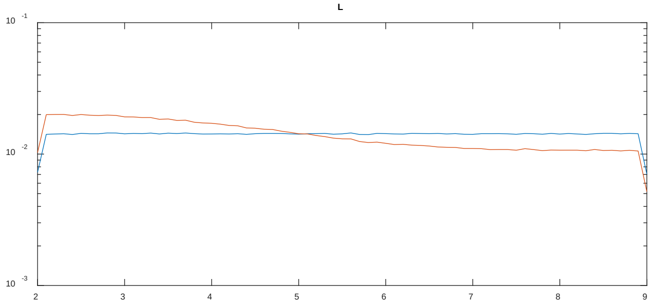

Conclusion 2: the great filter is likely in the past

- Given the priors and the Fermi observation, the default guess should be that the low-probability term(s) are in the past.
- The conclusion can be changed if:
  - We reduce the uncertainty of past terms to less than 7 orders of magnitude
  - The distributions have weird shapes
- Note that a past great filter does *not* imply our safety
  - (The stars just don't foretell our doom)

Summary

- The Fermi question is not a paradox: it just looks like one if one is overconfident in how well we know the Drake equation parameters.
- Doing a distribution model shows that even existing literature allows for a substantial probability of very little life, and a more cautious prior gives a significant probability for rare life.
- The Fermi observation makes the most uncertain priors move strongly, reinforcing the rare life guess and an early great filter.
- Getting even a little bit more information can update our belief state a lot!