What has the Riemann Zeta Function got to do with Satellite Lifetimes

Matt Roughan <matthew.roughan@adelaide.edu.au>

The University of Adelaide

2025

What are we doing today?

This talk arises out of a weird conjunction:

- Satellite lifetimes
- Quantile-based distributions
- The Polylogarithm Distribution

Satellite Lifetime Motivation

- Satellite lifetimes are big business
  - more satellites being launched than ever before
  - StarLink has launched over 7,000 in less than 10 years
  - roughly $1 million (AU) per satellite
- Satellite lifetime controlled by
  - failures
  - fuel (loss from station keeping)
  - planned end-of-lifetime
- If the planned end-of-lifetime could be longer, a lot of money could be saved

Satellite Lifetime Data (Batthula *et al.*, 2022)

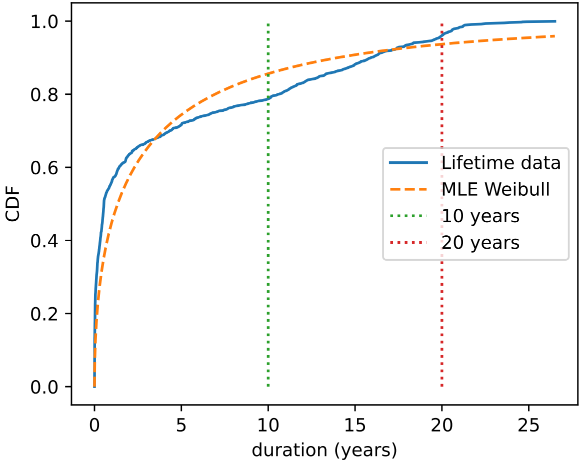

Satellite Lifetime Data (Batthula *et al.*, 2022)

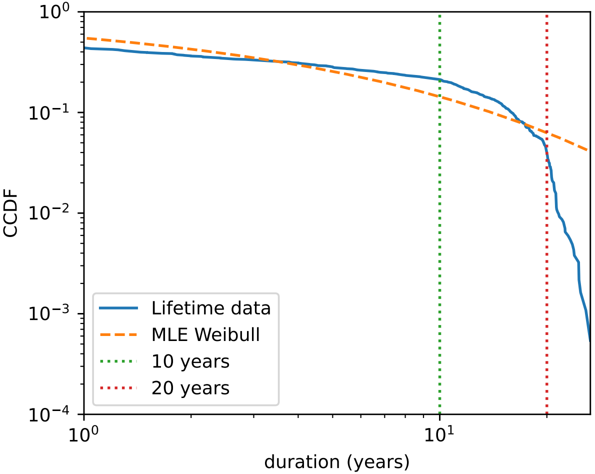

Satellite Lifetimes

- We see
  - somewhat heavy-tail lifetime, but
  - truncated around 10 years (planned lifetime)
  - try to reduce space junk (estimates of around 30,000 objects larger than 10cm)
- Seems like satellites might be being junked early
  - given more fuel, they could potentially live longer
- Weibull fit isn't doing great
  - Weibull and its generalisations are commonly used here
  - we want to hypothesize about the distribution if the artificial constraint were removed
  - we want to do model identification not just model fitting

Quantile-based distributions

Mostly when we teach distribution theory, we

- start with discrete RVs defined using PMFs, *e.g.*, geometric
- move onto continuous with PDFs
- note that these aren't always useful so we do CDFs

But we leave CDFs till late because they aren't intuitive???

Quantile-based distributions

Quantile-based distributions are an alternative

- Maybe it's more intuitive???
  - when we report distributions we often use quantiles, *e.g.*, median, IQR, box plots, ...
  - mean and variance of little use for heavy-tailed distributions
- We don't have to create new variants by generalising existing distribution (which is what most of the satellite lifetime work has done)
- These types of distributions have interesting properties
- They are easy to simulate from
- Easy to create classes of estimators

Quantile-based distributions

Roughly[1], the quantile function is

$$Q(p) = F^{-1}(p),$$

where $F(\cdot)$ is the Cumulative Distribution Function (CDF).
Notes:

- For convenience $Q(0)$ and $Q(1)$ are defined to be support
- Q is non-decreasing and left-continuous
- We'll put parameters in subscripts, *e.g.*, Q_s
- Q can easily add scale and location parameters, *i.e.*,

$$Q_{s,a,b}(u) = a + bQ_s(u)$$

- Simulation is by taking $X_i \sim U(0,1)$ and then

$$Y_s = Q_s(X_i)$$

[1] Actually, $Q(u) = \inf\{x \in \mathbb{R} \mid F(x) \geq u\}$ for $0 < u < 1$.

Quantile-based distributions – a question?

Does anyone have a reference for the following:

$$\mathbb{E}\left[h(Y)\right] = \int_0^1 h\big(Q_Y(p)\big)\,dp,$$

Quantile-based distributions examples

Distribution	**Parameters**	**Quantile Function**	**Support**
Uniform(a, b)	$a < b$	$Q(p) = a + p(b - a)$	$[a, b]$
Exponential(λ)	$\lambda > 0$	$Q(p) = -\frac{1}{\lambda}\ln(1 - p)$	$[0, \infty)$
Normal(μ, σ)	$\sigma > 0$	$Q(p) = \mu + \sigma\,\Phi^{-1}(p)$	$\mathbb{R}$
Logistic(μ, s)	$s > 0$	$Q(p) = \mu + s\,\ln\left(\frac{p}{1-p}\right)$	$\mathbb{R}$

The Tukey-λ distribution

$$Q(p;\lambda) = \begin{cases} \frac{1}{\lambda}\left[p^\lambda - (1-p)^\lambda\right], & \text{if } \lambda \neq 0, \\ \log\left(\frac{p}{1-p}\right), & \text{if } \lambda = 0, \end{cases}$$

where λ is the shape parameter.

It is unusual because:

- $\lambda > 0$, the distribution has finite support; and
- $\lambda \leq 0$, the distribution has infinite support.

Notes:

- The distribution is symmetric about 0
- It interpolates across distributions: *e.g.*,
  - $\lambda = -1$: approximately Cauchy (heavy tailed)
  - $\lambda = 0$: logistic
  - $\lambda = 0.14$: approximately normal
  - $\lambda = 1$: uniform
- We don't have closed forms for many of its bits and pieces for all λ, but the numerical calculations aren't hard

The Tukey-λ distribution

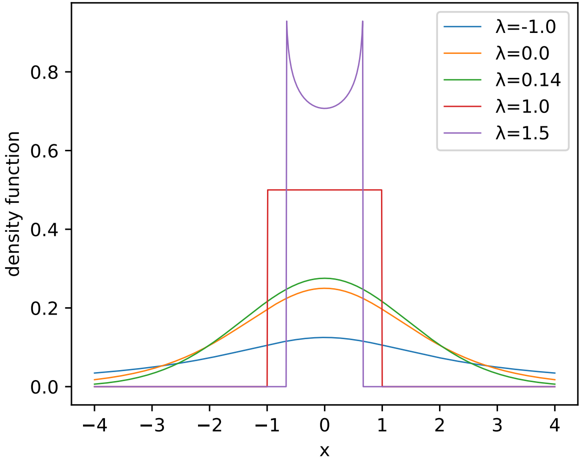

Alternatives?

- There aren't a lot of other distribution families defined by their quantile (though we often have closed forms for the quantile after the fact)
- In particular the Tukey-λ is symmetric — is there a one-sided variant (other than trivial functional transformation)

The Polylogarithm Distribution (PLD)

So let's define a one-sided (non-negative) distribution

- Exponential is a good starting point

$$Q_\lambda(p) = -\ln(1-p)/\lambda$$

- It's natural to generalise the log function with the polylogarithm[2]

$$Q_s(p) = \mathrm{Li}_s(p) = \sum_{k=1}^{\infty} \frac{p^k}{k^s}, \qquad (1)$$

where $0 \leq p < 1$ and $s \in \mathbb{R}$.

- When $s = 1$ we get the power series for the exponential distribution with $\lambda = 1$ (as above)

[2]We can extend the polylogarithm to complex arguments and parameters and outside of the range of the sum's convergence through analytic continuation, but we won't need that here.

Polylogarithms

- Already used in some distributions, *e.g.*,
  - Zeta distribution characteristic and probability generating functions
  - Fermi–Dirac and Bose-Einstein integrals
- Has a strong relationship to other important functions, *e.g.*,
  - The Riemann zeta function

$$\mathrm{Li}_s(1) = \zeta(s), \text{ for } Re(s) > 1$$

  - Bernoulli polynomials
  - Polygamma functions
- Has lots of known properties

"...almost all the formulas relating to it, have something of the fantastical in them, as if this function alone among all others possessed a sense of humor.", Don Zagier, 2007.

PLD: Properties

It has properties similar to the Tukey-λ

- It has both finite and infinite support
  - For $s > 1$ the distribution has finite support
  - For $0 < s \leq 1$ it has infinite support but finite expectation
  - For $s \leq 0$ the expectation is infinite
- For certain parameters we get other distributions
  - for $s \geq 8$, it closely approximates the uniform distribution,
  - for $s \simeq 1.6$, it approximates the (non-negative) triangular distribution,
  - for $s = 1.0$, it is exactly the exponential distribution,
  - for $s = 0.0$, it is exactly an inverse beta distribution, and
  - for large negative s, it approximates a generalised extreme value distribution with infinite mean.

PLD: CDFs

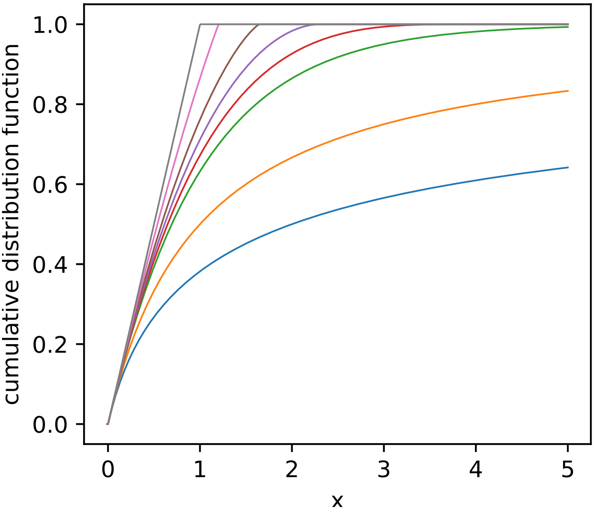

Calculated using a Golden Section search (280 μseconds).

PLD: PDFs

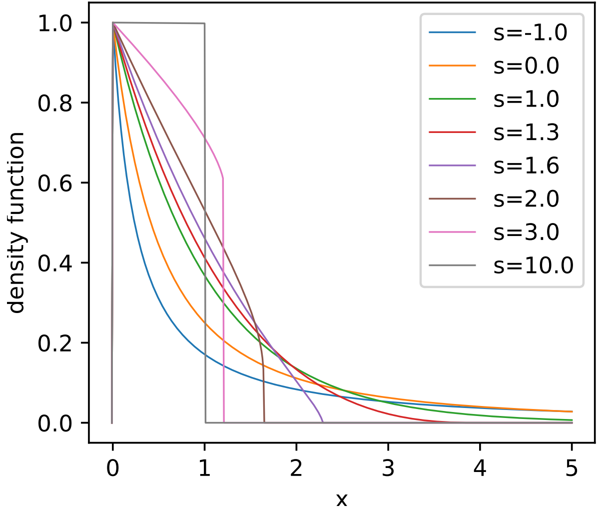

Calculated from CDF and Quantile derivative

$$\frac{d}{dz}\mathrm{Li}_s(z) = \mathrm{Li}_{s-1}(z)/z$$

PLD: Mean and Variance

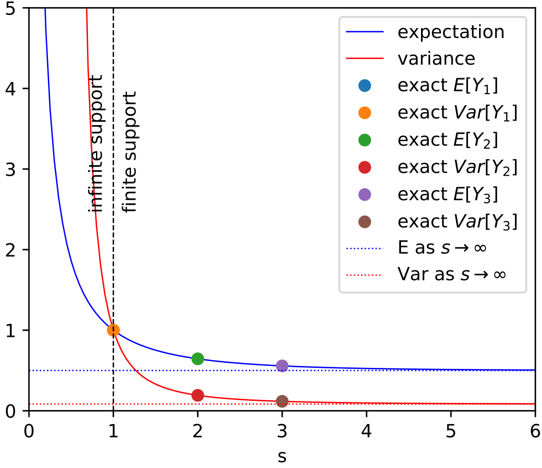

PLD: Support

$$\begin{aligned} \mathrm{Li}_s(0) &= 0 \\ \lim_{z\to 1} \mathrm{Li}_s(z) &= \begin{cases} \zeta(s), & \text{for } s > 1, \\ \infty, & \text{otherwise,} \end{cases} \end{aligned}$$

Hence

$$\mathrm{supp}(X_s) = \begin{cases} [0, \zeta(s)], & \text{for } s > 1, \\ [0, \infty)\,, & \text{otherwise.} \end{cases}$$

PLD: Mode and Median

- The mode is always at 0.0
  This can be derived from the fact that second derivative of the quantile function

  $$Q''_s(p) = \frac{\mathrm{Li}_{s-2}(p) - \mathrm{Li}_{s-1}(p)}{p^2},$$

  is positive for all real s and $0 \leq p$ because the polylogarithm function is decreasing with respect to s for $0 < p < 1$.
- The median is (trivially) $Q_s(0.5)$

PLD: Expectation

$$
\begin{aligned}
\mathbb{E}[Y_s] &= \int_0^1 \mathrm{Li}_s(x)dx \\
&= \lim_{T\to 1} \sum_{k=1}^{\infty} \frac{1}{k^s} \int_0^T x^k dx \\
&= \lim_{T\to 1} \sum_{k=1}^{\infty} \frac{T^{k+1}}{k^s(k+1)} \\
&= \begin{cases} \sum_{k=1}^{\infty} \frac{1}{k^s(k+1)}, & \text{for } s > 0, \\ \infty, & \text{otherwise.} \end{cases}
\end{aligned} \tag{2}
$$

PLD: Expectation (Convergence Bound)

The sum is bounded by the Riemann zeta function

$$\begin{aligned}\zeta(s+1)-1 &= \sum_{k=1}^{\infty}\frac{1}{(k+1)^{s+1}} \\ &\leq \sum_{k=1}^{\infty}\frac{1}{k^{s}(k+1)} \\ &\leq \sum_{k=1}^{\infty}\frac{1}{k^{s+1}} \\ &= \zeta(s+1). \end{aligned} \tag{3}$$

Where we know that the $\zeta(s+1)$ is finite when $Re(s) > 0$

PLD: Expectation (Part II)

When s is close to zero the series converges very slowly ☹

However, we use integration by parts with $\frac{d}{dz}\mathrm{Li}_s(z) = \mathrm{Li}_{s-1}(z)/z$ to get

$$\mathbb{E}[X_s] = \zeta(s+1) - \mathbb{E}[X_{s+1}],$$

so that we can calculate in a faster converging region

PLD: Higher Moments

- Tricky (numerically) but doable
  - Care about singularity at $z = 1$ needed for $s \leq 1$
  - Need to use an alternative series expansion around $z = 1$ that uses the Riemann zeta function
- Can prove bounds on finiteness of moments
  - Interval $[1 - 1/n, 1 - 1/(n+1))$ has moments up to n all finite, and the $(n+1)$th moment infinite.

PLD: Fitting

- One of the aims here is that fitting the parameter replace model identification
- We should use a fitting process matched to the definition
- We are working in the context of lifetime estimations, *i.e.*, non-negative distributions (*e.g.*, assume support starts at 0 and we don't need a location parameter)
- For Tukey-λ folks use a probability-plot-correlation-coefficient (PPCC) chart, but its visual and doesn't allow for a scale parameter
- Instead use L-statistics, which are a linear combination of order stats

PLD: Fitting

- Match ratio of IQR to median to the data (it's insensitive to scale), *i.e.*, choose $\hat{s}$ to match

$$\frac{IQR}{m} = \frac{Q_s(0.75) - Q_s(0.25)}{Q_s(0.5)}.$$

  The function is strictly increasing, so the search for $\hat{s}$ is almost trivial.
- Then match scale parameter by taking $Q_{b,s} = bQ_s$ with

$$\hat{b} = \frac{\hat{m}}{Q_{\hat{s}}(0.5)}.$$

- Generalise to α and $1 - \alpha$ percentiles

PLD: Fit on simulated data

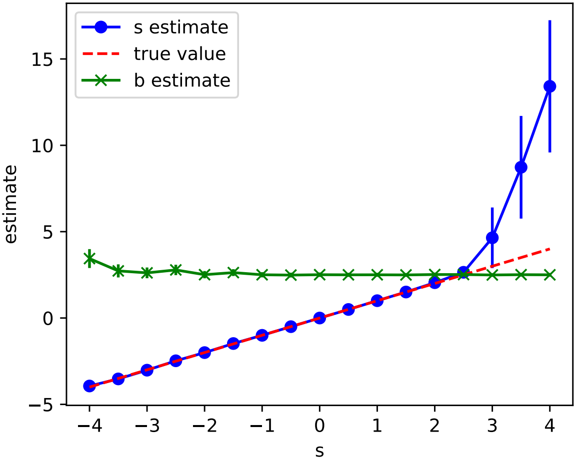

Positive bias for large s, but the distributions here are nearly identical (uniform).

Back to Satellite Data

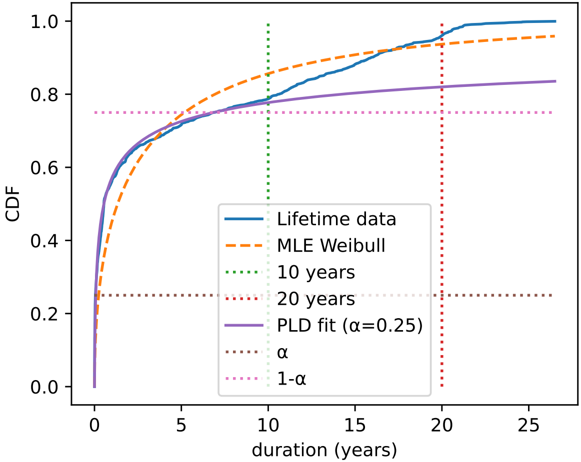

Back to Satellite Data

- Fit using $\alpha = 0.25$ resulting in
  - shape $\hat{s} \simeq -1.88$ (so heavy-tailed, infinite mean)
  - scale $\hat{b} \simeq 38.05$ (median 197 days)
- Much better fit to the body
  - that's what we want here
  - we want to model normal behaviour, because tail is artificially truncated

What does it all mean?

- Planned fuel depletion is killing satellites before the hardware wears out
- Extrapolation suggests that, given a satellite survives until 10 years, its chance of surviving an additional 10 years would be over 90%. An additional 10 years and it only drops to 86%.
- If sufficient fuel were available, satellite lifetimes might be easily doubled or even tripled
  - current fuel accounts for around 10–25% of satellite mass at launch

Conclusion

- Quantile-defined distributions have some interesting properties
- Polylogarithm is an ideal function to use in the quantile definition because (i) so much is known about it, (ii) it is just well-behaved enough, without being boring.
- Satellite lifetimes could be greatly extended with comparatively little cost

Extra Slides

A Lemma

$$\mathbb{E}\left[h(Y)\right] = \int_0^1 h\big(Q(p)\big)\, dp$$

Proof.

Take random variable $X \sim U(0,1)$, and note that (where it exists) the expectation of a function $g(x)$ of a random variable is given by

$$\mathbb{E}\left[g(X)\right] = \int g(x) dF_X.$$

Then note that we can create a random variable Y with quantile function $Q(p)$ by taking $Y = Q(X)$ and hence when the integral exists

$$\mathbb{E}\left[h(Y)\right] = \int h\big(Q(x)\big) dF_X = \int_0^1 h\big(Q(p)\big) dp.$$

□

PDF from Quantile

The density function $f(x)$ can be computed by noting

$$\left.\frac{dF^{-1}}{dp}\right|_{F^{-1}(p)=x} = 1/f(x).$$

Where a search is needed to find the point where $F^{-1}(p) = x$. Then the derivative can be derived using the derivative of the polylogarithm function

$$f(x) = \left.\frac{p}{\mathrm{Li}_{s-1}(p)}\right|_{\mathrm{Li}_s(p)=x}.$$