# Modelling HDD Failure Rates with Bayesian Inference

Sašo Stanovnik
sstanovnik@gmail.com

Mentor: prof. Erik Štrumbelj, PhD
erik.strumbelj@fri.uni-lj.si

Faculty of Computer and Information Science
University of Ljubljana, Slovenia

May 2016

# Problem

- Some hard drives fail earlier than others.
- An enormous amount of (unreliable) anecdotal evidence.
  - People don't know what they're talking about.
  - People make conclusions on sample sizes where $n = 2$.
- How long will a *typical* hard drive last?
- Are there differences between manufacturers?

# About the data

- We're using the Backblaze dataset[1].
- Consumer, not enterprise drives.
- Drives used in a datacenter environment.
  - $\sim$ 60 drives in a single enclosure[2].
  - Temperature and vibration are a factor.
- Daily SMART data for each disk in the datacenter.
  - 1087 days of data (from 2013Q2 to 2016Q1).
  - More than 65000 total drives.
- We only use SMART 9: Power–On Hours.

[1] https://www.backblaze.com/b2/hard-drive-test-data.html
[2] https://www.backblaze.com/b2/storage-pod.html

# Useful information

- HGST (Hitachi Global Storage Technologies) was founded when Hitachi acquired IBM's HDD branch.
- In 2012, HGST was acquired by WD and has been completely merged into WD in late 2015.
- The analysis merges HGST and Hitachi brandings[3].

[3] As reported by the drive's `model` attribute.

# Exploration
## Disk lifetime

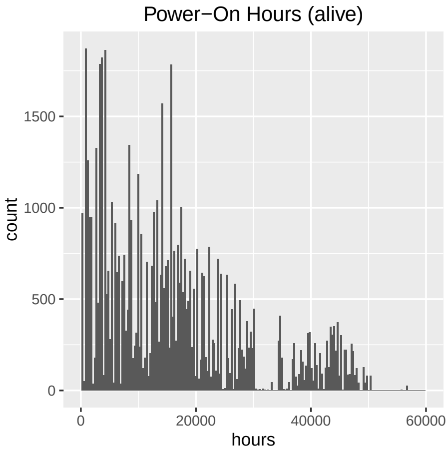


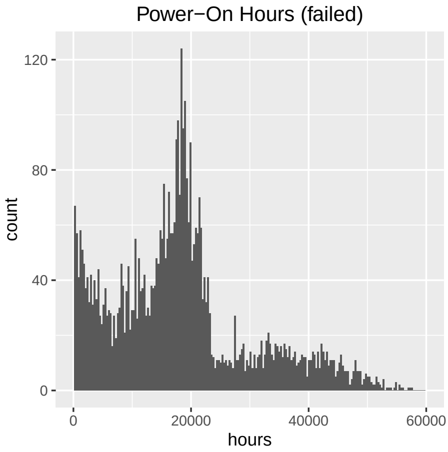

# Exploration

## Disk lifetime – manufacturers

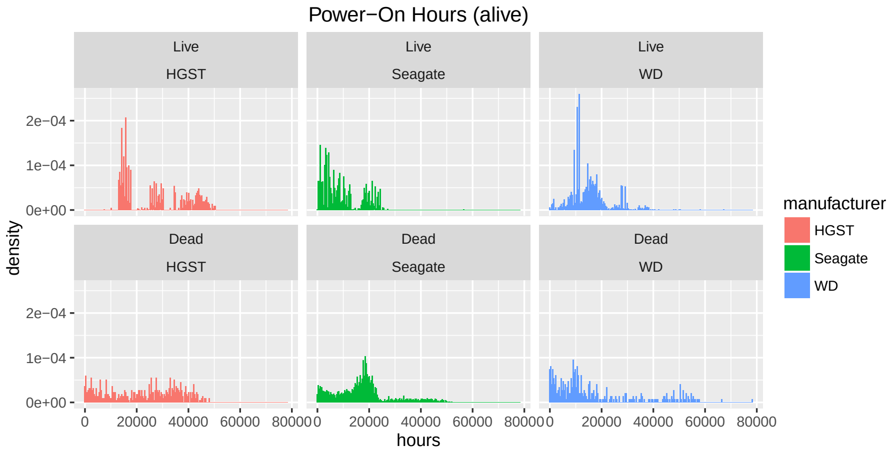

# Exploration

## Disk counts

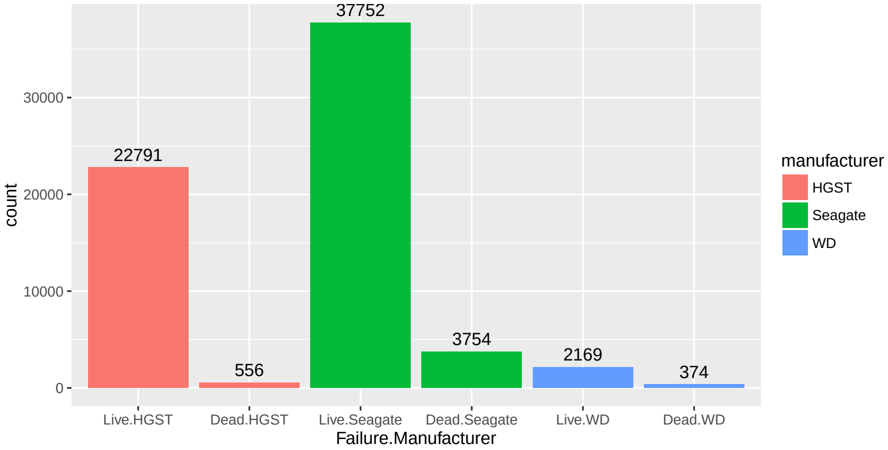

# About the Weibull distribution

- $x \sim Weibull(\alpha, \sigma)$

$$pdf(x) = \begin{cases} \frac{\alpha}{\sigma}\left(\frac{x}{\sigma}\right)^{\alpha-1} e^{-\left(\frac{x}{\sigma}\right)^{\alpha}} & x \geq 0 \\ 0 & x < 0 \end{cases}$$

- $\alpha, \sigma \in (0, \infty)$
- $x \in [0, \infty)$
- $\alpha$ has an important meaning
  - $alpha < 1$: early failure
  - $alpha = 1$: random failure
  - $alpha > 1$: age failure

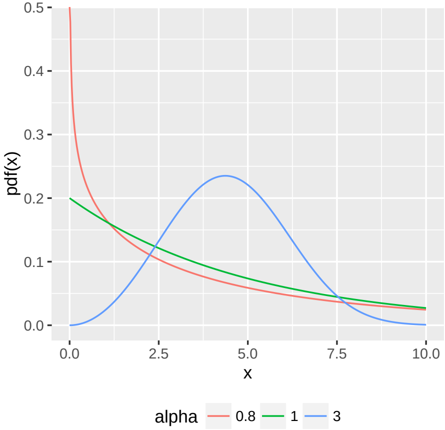

# Model

- The Weibull distribution is used for modelling failure rates.
- A single or a mixture of Weibull distributions?
  - Separate modelling of early–random–age failure characteristics.
    - Early failure $\alpha$ is $1 -$ `early_mortality_offset`.
    - Age failure $\alpha$ is $1 +$ `age_mortality_offset`.
- Non–informative priors, we have enough data.
- Right censoring[4] to include data about disks that are still alive.
- Separate modelling of different manufacturers.
  - The top 3 manufacturers: HGST, Seagate, WD.

[4] Picture it as including a probability that a certain disk has survived for as long as it has.

# Model

## Plate notation

Single Weibull

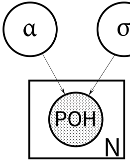


Three Weibulls

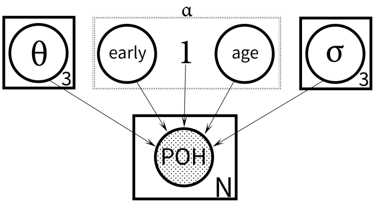


- Separate runs of the same model for different manufacturers.

# Model

## Convergence

- We ran 5000 iterations, 500 of which were warmup.
- All parameters across all models have $\sim 1200$ effective samples.
- The generated Power–On Hours have $\sim 4500$ effective samples
- $\hat{R}$ is always 1.
- Therefore, no reason to doubt convergence.

# Model

## Convergence: single Weibull

### Manufacturer subsets (poh_generated)

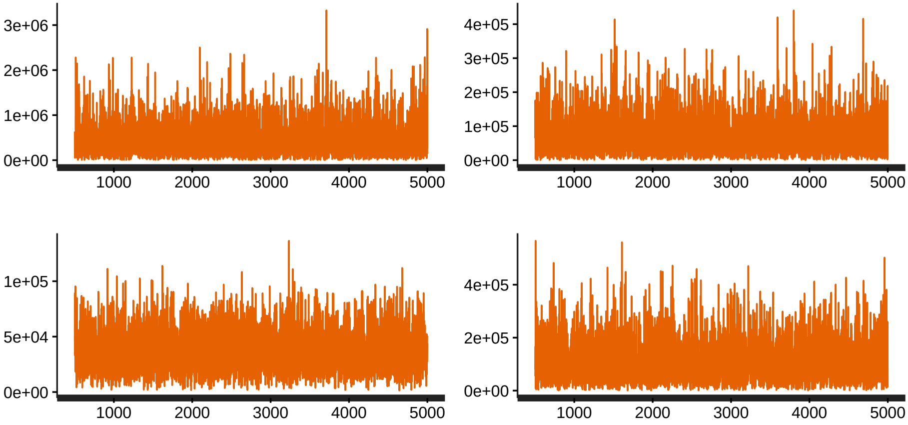


Note: other parameters also show no reason to doubt convergence, but there are too many to show.

# Model

## Convergence: triple Weibull

### Manufacturer subsets (poh_generated)

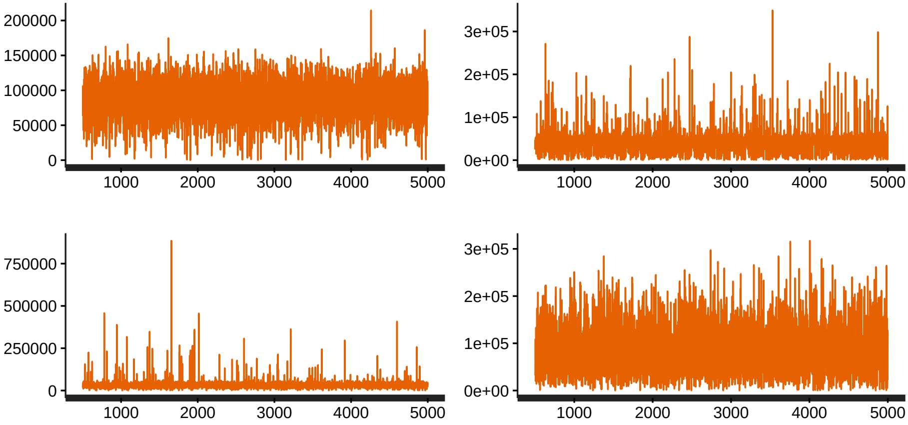


Note: other parameters also show no reason to doubt convergence, but there are too many to show.

# Results

## Single Weibull, manufacturers, $\alpha$

- Seagate appears to predominantly have age–related failures.
- HGST and WD are similar in this characteristic.

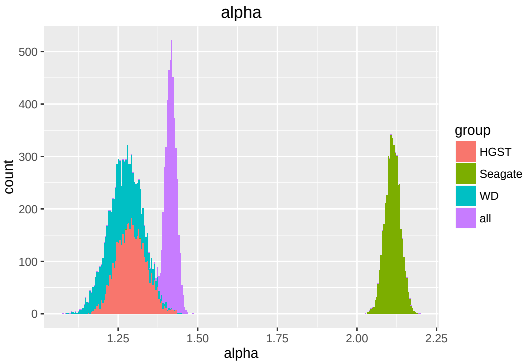

# Results

## Single Weibull, manufacturers, $\sigma$

- HGST drives have proportionately little failures (2% of all drives) and high lifetimes.
- The model likely overestimates based on limited data.
- Here, WD is estimated to generally have a higher lifetime than Seagate

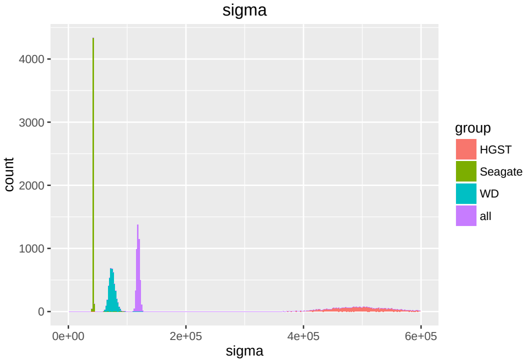

# Results

## Single Weibull, manufacturers, Predicted Power–On Hours

- HGST's prediction does not seem realistic.
- WD has a longer tail than Seagate.

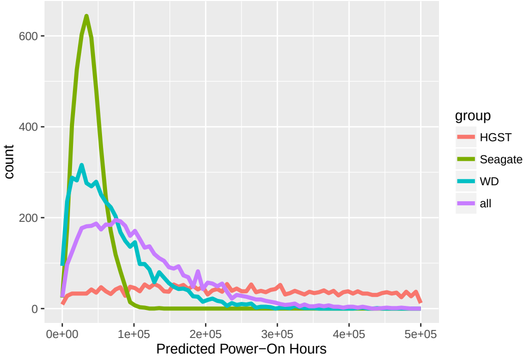

# Results

## Triple Weibull, manufacturers, early failure offsets

- This mostly tends to 0 (to achieve $\alpha = 1$).
- HGST is different here in that it shows more confidence in (some) early mortality.

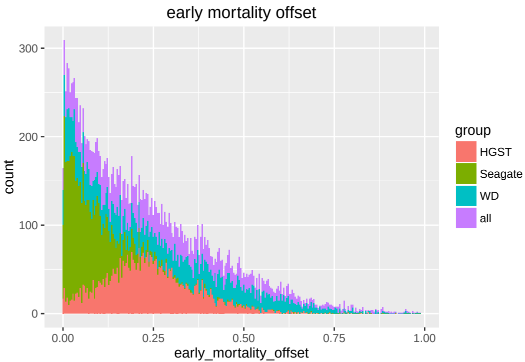

# Results

## Triple Weibull, manufacturers, age failure offsets

- WD trends most heavily towards age mortalities, followed by HGST.
- All drives combined appear to fail randomly, likely because of the mixture of different manufacturers.

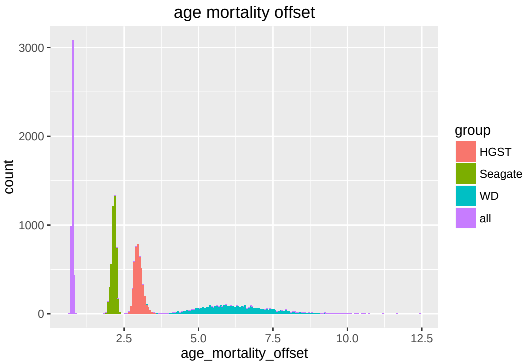

# Results

## Triple Weibull, manufacturers, sigmas

- The parameter for early failures is the same across all drives, but Seagate is very confident in very early failures.
- Random failures occur latest for WD. HGST has a very wide spread and all drives combined fail early if they fail randomly.
- Age–related failures occur latest for HGST, then WD, then Seagate.

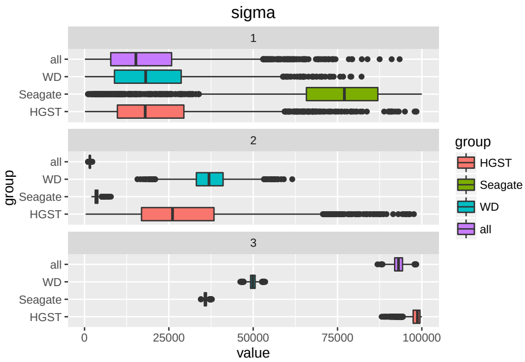

# Results

## Triple Weibull, manufacturers, mixture ratios

- Early failures are fairly infrequent for all manufacturers.
- WD has most random failures (at the expense of age–related failures), followed by Seagate.

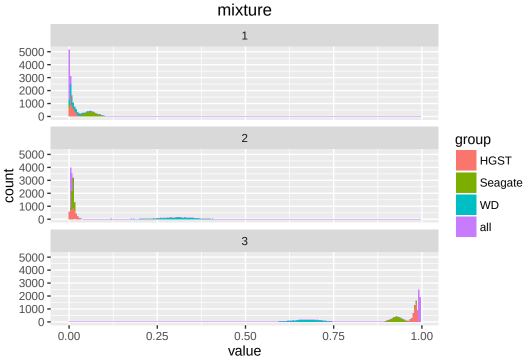

# Results

## Triple Weibull, manufacturers, Predicted Power–On Hours

- HGST stands out as the most reliable.
- WD is better than Seagate, but has a spike in early failures.
- Modelling all drives together does not give useful information as different manufacturers are blended together.

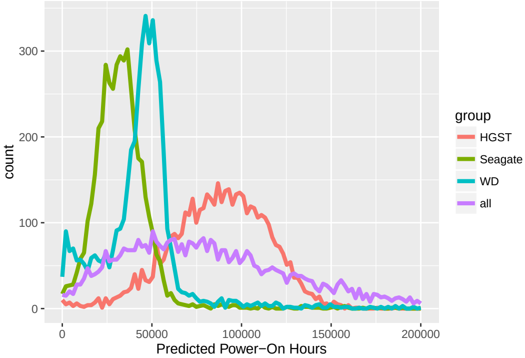

# Model quality estimation

## PSIS leave–one–out approximation using `loo`

| | elpd.estimate | elpd.se |
|---|---|---|
| Single Weibull | -62159 | 811.48 |
| Triple Weibull | -61977 | 808.79 |

- The triple Weibull model outperforms the single model, but is within the error tolerance.

# Interpretation

## Proportion of failed drivers by year

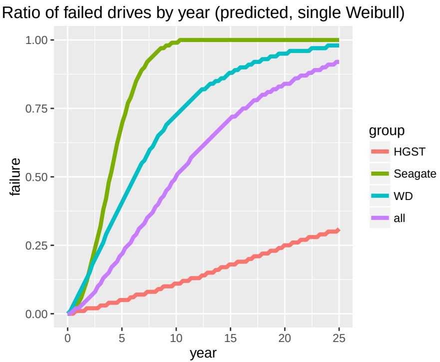


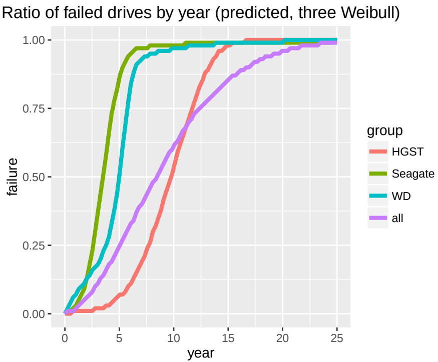


- The three Weibull model is more rational.
- HGST is most reliable, followed by WD, then Seagate.
- WD has more early failures.

# Interpretation

## Hazard function

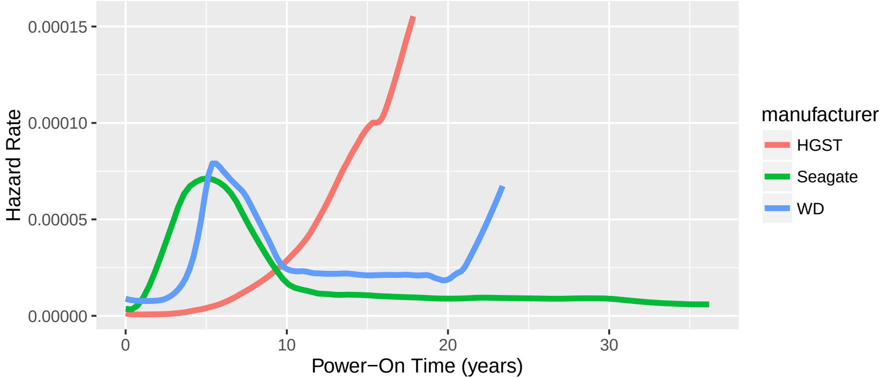


1. Seagate and WD have similar hazards. WD's hazard rises later, but then experiences more random failures.
2. HGST fails primarily with age.

# Conclusions

- The mixture of Weibulls is more effective than a single one.
  - We can model detailed disk characteristics.
- There are differences between manufacturers.
  - HGST drives appear most reliable, followed by WD, then Seagate.
  - Early failures are characteristic for WD.
- We need to remember that we have a (relatively) short time window of observation.
  - Also the environment in which these drives operate.
- Think about all these interpretations for yourself, anectodal "evidence" is bad.

# Addendum

# AD 1: stan code for the single Weibull model

```
data {
  int<lower=0> n;
  int poh[n];
  int <lower=0, upper=1> failure_bools[n];
}

parameters {
  real<lower=0.0> alpha;
  real<lower=0.0> sigma;
}

model {
  for (i in 1:n) {
    if (failure_bools[i] == 1) {
      poh[i] ~ weibull(alpha, sigma);
    } else {
      // right censoring
      increment_log_prob(
          weibull_ccdf_log(poh[i], alpha, sigma)
      );
    }
  }
}
```

```
generated quantities {
  real poh_generated;
  real log_likelihood[n];

  poh_generated <- weibull_rng(alpha, sigma);
  for (i in 1:n) {
    if (failure_bools[i] == 1) {
      log_likelihood[i] <-
          weibull_log(poh[i], alpha, sigma);
    } else {
      log_likelihood[i] <-
          weibull_ccdf_log(poh[i], alpha, sigma);
    }
  }
}
```

# AD 2: stan code for the triple Weibull model (1/3)

```
data {
  int<lower=0> n;
  int poh[n];
  int <lower=0, upper=1> failure_bools[n];

  real<lower=0.0> prior_gamma_alpha0;
  real<lower=0.0> prior_gamma_beta0;
}

parameters {
  simplex[3] theta;
  real<lower=0.0, upper=1.0> early_mortality_offset;
  real<lower=0.0, upper=100> age_mortality_offset;
  real<lower=0.0, upper=100000> sigma[3];
}
```

# AD 2: stan code for the triple Weibull model (2/3)

```
model {
  real temp[3];

  sigma ~ gamma(prior_gamma_alpha0, prior_gamma_beta0);

  for (i in 1:n) {
    if (failure_bools[i] == 1) {
      temp[1] <- log(theta[1]) + weibull_log(poh[i], 1 - early_mortality_offset, sigma[1]);
      temp[2] <- log(theta[2]) + weibull_log(poh[i], 1   /* random mortality */, sigma[2]);
      temp[3] <- log(theta[3]) + weibull_log(poh[i], 1 + age_mortality_offset, sigma[3]);
    } else {
      // right censoring
      temp[1] <- log(theta[1]) + weibull_ccdf_log(poh[i], 1 - early_mortality_offset, sigma[1]);
      temp[2] <- log(theta[2]) + weibull_ccdf_log(poh[i], 1   /* random mortality */, sigma[2]);
      temp[3] <- log(theta[3]) + weibull_ccdf_log(poh[i], 1 + age_mortality_offset, sigma[3]);
    }

    increment_log_prob(log_sum_exp(temp));
  }
}
```

# AD 2: stan code for the triple Weibull model (3/3)

```
generated quantities {
  real poh_generated;
  int cat;
  real log_likelihood[n];
  real temp[3];

  cat <- categorical_rng(theta);
  if (cat == 1) {
    poh_generated <- weibull_rng(1 - early_mortality_offset, sigma[1]);
  } else if (cat == 2) {
    poh_generated <- weibull_rng(1, sigma[2]);
  } else {
    poh_generated <- weibull_rng(1 + age_mortality_offset, sigma[3]);
  }

  for (i in 1:n) {
    if (failure_bools[i] == 1) {
      temp[1] <- log(theta[1]) + weibull_log(poh[i], 1 - early_mortality_offset, sigma[1]);
      temp[2] <- log(theta[2]) + weibull_log(poh[i], 1  /* random mortality */, sigma[2]);
      temp[3] <- log(theta[3]) + weibull_log(poh[i], 1 + age_mortality_offset, sigma[3]);
    } else {
      temp[1] <- log(theta[1]) + weibull_ccdf_log(poh[i], 1 - early_mortality_offset, sigma[1]);
      temp[2] <- log(theta[2]) + weibull_ccdf_log(poh[i], 1  /* random mortality */, sigma[2]);
      temp[3] <- log(theta[3]) + weibull_ccdf_log(poh[i], 1 + age_mortality_offset, sigma[3]);
    }
    log_likelihood[i] <- log_sum_exp(temp);
  }
}
```