Stockholms universitet

Survival Analysis and Failure Predictors of HDDs from the Backblaze Dataset

Sasha Tomic

Kandidatuppsats i matematisk statistik
Bachelor Thesis in Mathematical Statistics

Kandidatuppsats 2026:16
Matematisk statistik
September 2026

www.math.su.se

Matematisk statistik
Matematiska institutionen
Stockholms universitet
106 91 Stockholm

Mathematical Statistics
Stockholm University
Bachelor thesis **2026:16**
https://www.math.su.se

Survival Analysis and Failure Predictors of HDDs from the Backblaze Dataset

Sasha Tomic*

September 2026

Abstract

This thesis applies survival analysis to evaluate the reliability of data-centre hard disk drives (HDDs), using a large, left-truncated dataset of 14TB drives published by Backblaze and covering the period 2021–2025, in order to quantify how SMART error indicators, operating temperature, and power cycling are associated with a drive's remaining lifetime.

Diagnostics based on Schoenfeld residuals showed that the hazard ratios implied by a Cox proportional hazards model vary with drive age, so the proportional hazards assumption does not hold in these data. Because the Weibull distribution is the only accelerated failure time (AFT) family that also assumes proportional hazards, it was excluded on structural grounds; the analysis instead adopts a log-logistic AFT specification, which imposes no such restriction, with Weibull and log-normal fits reported throughout as robustness comparisons.

Reallocated and pending sectors (SMART 5 and 197) were associated with substantially reduced expected lifetimes, with estimated time ratios of approximately 0.49 and 0.15 respectively, and were the most consistent predictors across every distributional family and sensitivity check performed. Mean operating temperature showed a small positive association with lifetime, and drives manufactured by Western Digital survived longer than comparable Seagate and Toshiba units. The estimated log-logistic shape parameter implies a hazard that increases across the full observed age range for drives without SMART faults, but one that is already near its peak for drives reporting pending sectors.

*Postal address: Mathematical Statistics, Stockholm University, SE-106 91, Sweden. E-mail: sato7286@student.su.se. Supervisor: Jan-Olov Persson, Martina Favero.

Abstract

Denna uppsats tillämpar överlevnadsanalys (survival analysis) för att utvärdera tillförlitligheten hos hårddiskar i datacenter. Genom att använda ett storskaligt och vänstertrunkerat datamaterial från Backblaze för 14TB-hårddiskar (2021–2025) syftar studien till att kvantifiera hur SMART-fel, temperatur och antal systemstart påverkar enheternas återstående livslängd. Diagnostik med Schoenfeld-residualer visade att de riskkvoter (hazard ratios) som en Cox Proportional Hazards-modell implicerar varierar med hårddiskens ålder, varför antagandet om proportionella risker inte håller för detta datamaterial. Eftersom Weibull-fördelningen är den enda Accelerated Failure Time (AFT)-familjen som samtidigt antar proportionella risker uteslöts den av strukturella skäl. Analysen övergick istället till en log-logistisk AFT-modell, som inte kräver detta antagande; Weibull- och log-normala modeller redovisas genomgående som robusthetskontroller.

Resultaten visar att mekaniska avvikelser, mer specifikt omallokerade och väntande sektorer (SMART 5 och 197), fungerar som robusta varningsindikatorer som drastiskt komprimerar hårddiskens återstående livslängd, med skattade tidskvoter på cirka 0,49 respektive 0,15. Dessa effekter var även de mest stabila oavsett vilken fördelning eller känslighetsanalys som användes. Temperaturen visade en marginellt positiv effekt, vilket kan peka på överkylning i den klimatkontrollerade miljön. Modellen påvisade även statistiskt signifikanta skillnader i överlevnad mellan tillverkare, där enheter från Western Digital överlevde längre än jämförbara enheter från Seagate och Toshiba. Slutligen visar den skattade formparametern i den log-logistiska modellen att risken ökar under hela den observerade åldersperioden för hårddiskar utan SMART-fel, men redan närmar sig sitt maximum för hårddiskar med väntande sektorer.

Acknowledgements

I would like to express my gratitude to my thesis supervisors **Jan-Olov Persson** and **Martina Favero** for giving valuable feedback that helped iterate this report.

AI usage disclaimer: Various LLM-based tools were used in the creation of this report, primarily for code creation/debugging, and for refining the writing to be more formal.

Contents

1 Introduction

Hard Disk Drives (HDDs) are physical storage devices that store data magnetically on a spinning disk, using a small arm that moves rapidly across its surface to read and write data. Because of these moving mechanical parts, HDDs tend to wear down and fail more often than Solid State Drives (SSDs), which have no moving parts ([1]).

Despite this higher failure rate, HDDs remain the dominant storage medium in large-scale data centres, owing to their considerably lower cost per unit of storage ([2]). This combination of mechanical wear and large-scale deployment makes it valuable to understand which operating conditions and warning signs are associated with a drive nearing the end of its life: identifying such factors can help operators anticipate failures and reduce the risk of sudden, unrecoverable data loss.

This report addresses that question using five years (2021–2025) of daily operating logs published by the data storage company Backblaze for a large fleet of data-centre hard drives, described in full in Chapter 2.

1.1 Research Questions

The report aims to answer two questions:

- Can drive temperature, power-cycle count, and indicators of mechanical wear function as predictors of HDD failure – that is, are they systematically associated with a shorter or longer remaining drive lifespan?

- To what extent, and in which direction, does each of these factors affect a drive's failure risk?

The aim here is descriptive rather than predictive: the analysis quantifies associations between these factors and observed drive lifetimes, using the survival models introduced in Chapter 3, rather than producing forecasts for individual drives. Accordingly, no train/test split is performed, as no out-of-sample forecasting task is being evaluated.

2 Data Background

2.1 Overview and Layout

The data used in this report is the public dataset provided by the data storage and cloud backup company Backblaze ([3]). Backblaze runs data centers with a large variety of HDD models and publishes daily logs for each individual drive, spanning multiple years. The data is released quarterly, with one CSV file for each individual day in that quarter; each file contains one row per drive that was operational that day, so the raw data has a daily panel structure at the level of an individual drive. Modern HDDs have built-in tracking tools called S.M.A.R.T. (Self-Monitoring, Analysis, and Reporting Technology) which log many different aspects of the drive, like current temperature, power cycles, or various error counts.([5])

Due to the abundance of data and long time-frame, this dataset is well suited for statistical analysis; in this report we use a 5-year time window from 2021 to 2025. Over this timeframe, Backblaze logged a total of 417,288 unique hard disk drives, of which 18,628 experienced a terminal failure. A breakdown of the key variables in the CSV files is as follows:

- **Date** – the date the data was taken (daily).
- **Serial Number** – the unique identifier for each HDD.
- **Model** – the specific model of a HDD.
- **Failure** – a flag equal to 1 on the last day a given HDD was operational before failing, and 0 otherwise. A drive is removed from the daily logs once it has failed, so each drive contributes at most one failure to the dataset.
- **SMART variables** – a large selection of variables logging different aspects of the drive's functionality that day.

The raw daily CSV files contain 149 variables, the vast majority of which are manufacturer-specific SMART attributes. Because many of these attributes are proprietary or left blank by unsupported manufacturers, this study isolates a core subset of universally tracked variables to ensure parity across models. A short description of the key SMART attributes isolated for the study is:

- **SMART 9 (Power-On Hours)** – the cumulative time the drive has been powered on, used in Section 2.2 to construct the time scale of the survival analysis.
- **SMART 194 (Temperature)** – the average internal operating temperature of the drive that day (in Celsius).

- **SMART 12 (Power Cycle Count)** – the number of times the drive was shut down and spun back up.
- **SMART 5 & 197** – counts of reallocated and pending bad sectors, used as indicators of physical surface wear.

2.2 Variable Selection

The variables above were not chosen arbitrarily. Temperature (SMART 194) and power-on cycles (SMART 12) are both commonly cited as contributors to mechanical wear in hard drives, while reallocated and pending sector counts (SMART 5 and 197) are established indicators of physical degradation of the drive's storage medium. Backblaze itself reports using SMART 5 and 197 as two of the five SMART attributes it monitors specifically to anticipate drive failure, alongside SMART 187, 188 and 198, which are dropped here for lack of universal support across the three manufacturers in the study cohort ([11]).

SMART 9 is required in addition to these four: expressed in days, it gives each drive's cumulative operating age at the start and end of observation, which places every drive on the common left-truncated time scale $(t_{\text{entry}}, t_{\text{exit}}]$ used throughout Chapter 4, rather than treating the calendar date of each log entry as the survival time.

2.3 Study Cohort

Because the raw dataset is very large (hundreds of gigabytes uncompressed) some restriction of scope is necessary to keep the data manageable without discarding too much information. Since the average lifespan of a hard drive is around 4-7 years, a reasonably long observation window is needed to see a meaningful number of drives reach failure. ([4])

To balance data volume against timespan, the study is restricted to a 5-year window (2021–2025) and to three specific, high-volume drive models, each with a capacity of 14 Terabytes, to give an "apples-to-apples" comparison across manufacturers:

- Seagate: ST14000NM001G
- Toshiba: MG07ACA14TA
- Western Digital: WUH721414ALE6L4

Each model corresponds to a different manufacturer, so within this cohort "model" and "manufacturer" coincide; there is no case of two models from the same manufacturer. This 14TB cohort consists of 59,262 hard drives over the 5-year observation window, which is more than sufficient to observe a substantial number of failures.

2.4 Descriptive Statistics

Table 1 summarises the cohort by drive model. The three models differ noticeably both in scale (from 8,856 WDC drives to 39,159 Toshiba drives) and in reliability, with annualised failure rates ranging from 0.51% (WDC) to 1.40% (Seagate). Entry age – the drive's cumulative power-on time at the start of observation – has a median of under two months for all three models, though the Toshiba cohort includes a longer right tail, reaching entry ages of up to 823 days; this is returned to in Section 4.2, since it affects how far the Kaplan-Meier curves for different models can be meaningfully compared. Average operating temperature is broadly similar across models (31–34 °C), and reported SMART 5 and SMART 197 rates are roughly twice as high for Seagate drives as for the other two.

Table 1: Cohort description by drive model.

	Seagate	Toshiba	WDC	All models
Drives	11,247	39,159	8,856	59,262
Failures	715	1,864	214	2,793
Exposure (drive-years)	51,065	183,364	42,012	276,441
AFR (%)	1.40	1.02	0.51	1.01
Entry age, median (d)	48.7	55.0	24.0	48.5
Entry age, IQR (d)	13–99	8–171	8–77	8–138
Entry age, range (d)	0–122	0–823	0–100	0–823
Exit age, median (d)	1,686	1,854	1,849	1,849
Exit age, max (d)	1,947	2,648	1,924	2,648
Mean temperature (°C)	33.6	30.9	32.6	31.7
Temperature SD (°C)	3.8	4.0	3.7	4.1
Power cycles at entry, median	2	2	1	2
Power cycles at entry, range	0–20	1–164	1–208	0–208
SMART 5 > 0 (%)	7.72	3.06	2.70	3.89
SMART 197 > 0 (%)	7.58	4.44	4.20	5.00

The majority of drives never report a non-zero value for SMART 5 (3.89% of the cohort do) or for SMART 197 (5.00% do). Among those that do, raw values span several orders of magnitude: for SMART 5 the median non-zero reading is 8, the 90th percentile is 495, and the 99th percentile is 30,357; for SMART 197 the corresponding figures are 6, 367, and 13,535. This spread motivates treating both variables as binary indicators rather than using the raw counts directly. Important to note that how such errors are logged vary depending on HDD model, as each manufacturer has different HDD architecture. Creating a binary indicator thus enables parity between the models in our analysis.

2.5 Limitations

While the dataset is quite rich, it still has some shortfalls. There are no logs of any maintenance that might have been performed on specific drives, and there can be cases of technicians preemptively removing drives suspected to be at risk of failure based on elevated error counts, which would introduce informative right-censoring.

A further issue is one of measurement timing rather than target leakage in the predictive sense, since no forecasting task is performed in this report: SMART 5 and SMART 197 are frequently first logged in the days immediately preceding a failure, so an indicator built from a drive's entire observation window would in part be recording the failure event itself rather than a condition that preceded it. Section 4.1 addresses this by disregarding errors logged within a short window of failure.

A further limitation concerns the lack of standardization of SMART attributes across manufacturers. Only variables that are logged in a comparable way by all three manufacturers can be used in a joint model, which excludes many manufacturer-specific SMART fields that might otherwise carry useful information. In the present cohort each of the three models corresponds to a different manufacturer.

3 Theory

3.1 Survival Analysis Fundamentals

The foundation of survival analysis rests on specifying the distribution of a non-negative, continuous random variable T that represents the time elapsed from a well-defined starting point until the occurrence of an event (In our case HDD failure).

3.1.1 Probability Density and Survival Functions

Let T have a probability density function (Shortened to PDF) denoted by $f(t)$ and a cumulative distribution function (CDF) denoted by $F(t) = P(T \leq t)$. The CDF represents the probability that a failure occurs before time t.

In survival analysis however, we are primarily interested in what is known as the Survival Function, $S(t)$:

$$S(t) = P(T > t) = 1 - F(t) = \int_t^\infty f(u)du$$

.

The survival function describes the probability that the hardware survives longer than time t

3.1.2 The Hazard Function

Instead of modeling the unconditional probability of failure, survival analysis often focuses on the conditional failure rate. The Hazard Function, $h(t)$, represents the potential per unit time for an event to occur, given that the hardware has survived up to time t. It is defined as:

$$h(t) = \lim_{\Delta t \to 0} \frac{P(t \leq T < t + \Delta t \mid T \geq t)}{\Delta t}$$

By the definition of conditional probability and the density function, this can be rearranged as $h(t) = \frac{f(t)}{S(t)}$

3.1.3 Mathematical Relationship between $S(t)$ and $h(t)$

This section shows the direct link between the hazard and survival functions. Since $f(t) = -\frac{d}{dt}S(t)$, we can rewrite the hazard function as:

$$h(t) = -\frac{d}{dt} \ln S(t)$$

Integrating both sides from 0 to t now gives us the Cumulative Hazard Function, $H(t)$:

$$H(t) = \int_0^t h(u)du = -\ln S(t)$$

Exponentiating this gives us an important relationship used in maximum likelihood estimation for survival models ([7], ch. 2.1-2.6):

$$S(t) = \exp(-H(t)) = \exp\left(-\int_0^t h(u)du\right)$$

3.2 Information loss - Truncation and Right-Censoring

Standard regression techniques need the full exact value of the response variable for all subjects, however survival datasets like ours are inherently incomplete.

3.2.1 Right-Censoring

An observation is considered right-censored if an event (in our case HDD failure) has not occurred before the end of the observation period. Let T_i be the true unobserved failure time, and C_i be the censoring time, then we can only observe the random variable $Y_i = \min(T_i, C_i)$. We also define a censoring indicator δ_i, where $\delta_i = 1$ if the event is a true failure ($T_i \leq C_i$) and $\delta_i = 0$ if censored ($T_i > C_i$).

Assuming the censoring is independent of the failure (non-informative censoring), the likelihood function L for a sample of size n must include both exact failures and the partial information from survivors. A failed drive contributes its density $f(y_i)$, while a censored drive contributes its survival probability $S(y_i)$, in the end we get:

$$L = \prod_{i=1}^{n} f(y_i)^{\delta_i} S(y_i)^{1-\delta_i} = \prod_{i=1}^{n} h(y_i)^{\delta_i} S(y_i)$$

3.2.2 Left-Truncation

Left-truncation happens when drives enter the study after they have already been at risk for some time. In the Backblaze dataset, a drive might enter observation at an initial age $V_i > 0$. If the drive had failed prior to V_i, it would never have been recorded which means our inferences must be conditioned on survival up to V_i ([7], ch. 3.1-3.5).

3.2.3 Risk Set

To handle left-truncation and right-censoring simultaneously, we create the Risk Set, $R(t)$ which is the mathematical set of all individuals i who are under active observation and at risk of failure an instant just prior to time

t, where $R(t) = \{i : V_i < t \leq Y_i\}$. Accurate construction of $R(t)$ is an important requirement for calculating non-parametric estimators. ([7], ch. 4.6)

3.3 Non-Parametric estimators

Before imposing parametric assumptions on the baseline hazard $h_0(t)$, empirical distributions must be estimated.

3.3.1 The Kaplan-Meier Estimator

The Kaplan-Meier estimator calculates the conditional probability of survival without assuming an underlying parametric distribution.

Let $t_1 < t_2 < \cdots < t_k$ be the distinct, ordered failure times, then let d_j be the number of failures occurring exactly at time t_j, and let n_j be the cardinality of the risk set $R(t_j)$, then the conditional probability of surviving past t_j, given survival up to t_j, is estimated as $1 - \frac{d_j}{n_j}$. By the law of conditional probability we can get that the overall survival function is the product of these conditional probabilities:

$$\hat{S}(t) = \prod_{t_j \leq t} \left(1 - \frac{d_j}{n_j}\right).$$

Both right-censoring and left-truncation enter through $R(t_j)$, as defined in Section 3.2.3: a drive that was censored before t_j, or one that has not yet entered observation, does not contribute to n_j, while a drive that entered at $V_i < t_j$ is included in the risk set from V_i onward.

3.3.2 Variance of the Estimator and Confidence intervals

In order to be able to construct confidence intervals around the Kaplan-Meier curve we will need to estimate the variance of $\hat{S}(t)$. This can be done using Greenwood's Formula and gives us the variance:

$$\widehat{Var}(\hat{S}(t)) = \hat{S}(t)^2 \sum_{t_j \leq t} \frac{d_j}{n_j(n_j - d_j)}$$

Confidence intervals for $S(t)$ are constructed on the complementary log-log scale, which keeps the interval bounds inside $(0, 1)$. Writing $\theta(t) = \ln(-\ln S(t))$, the delta method gives

$$\widehat{\mathrm{Var}}[\theta(t)] \approx \frac{\widehat{\mathrm{Var}}[\hat{S}(t)]}{[\ln \hat{S}(t)]^2},$$

and a $(1-\alpha)$ confidence interval for $S(t)$ follows as

$$\left[\hat{S}(t)^{\,c},\ \hat{S}(t)^{\,1/c}\right], \qquad c = \exp\left(z_{1-\alpha/2}\sqrt{\widehat{\mathrm{Var}}[\theta(t)]}\right).$$

([7], ch. 4.2-4.3)

3.4 Cox Proportional Hazards Model

While non-parametric estimators capture baseline survival, controlling for multiple confounding variables requires a multivariate framework. Here we introduce the Cox Proportional Hazards Model, which is a semi-parametric model that estimates the relative risk of covariates without requiring a specific underlying failure distribution.

3.4.1 Model Formulation and Partial Likelihood

The Cox model assumes the hazard function is the product of an unspecified baseline hazard $h_0(t)$ and a parametric log-linear multiplier:

$$h(t \mid X) = h_0(t)\exp(\beta^T X)$$

Because the baseline hazard $h_0(t)$ is left completely unspecified, standard MLE is impossible so instead the parameters β are estimated using Cox's Partial Likelihood. At each ordered failure time t_i the probability that individual j fails out of all individuals in $R(t_i)$ is calculated. Because $h_0(t_i)$ appears in both the numerator and the denominator, it cancels out entirely:

$$L(\beta) = \prod_{i=1}^{k} \frac{\exp(\beta^T X_{(i)})}{\sum_{l \in R(t_i)} \exp(\beta^T X_l)}$$

Maximizing this partial likelihood provides the estimated coefficients $\hat{\beta}$. ([9] ch. 2.1-2.2)

3.4.2 The Proportional Hazards Assumption and Diagnostics

A consequence of this formulation is that the Hazard Ratio between any two identical subjects differing only by their covariate vectors X_A and X_B must be strictly time-invariant:

$$\mathrm{HR}_{A,B} = \frac{h_0(t)\exp(\beta^T X_A)}{h_0(t)\exp(\beta^T X_B)} = \exp(\beta^T (X_A - X_B))$$

To see whether this assumption holds, we analyze Schoenfeld residuals, which measure the difference between the observed covariate of the failed subject at time t_i and its expected value given the risk set. If the proportional hazards assumption is valid, the expected slope of the Schoenfeld residuals plotted against time is zero. A statistically significant non-zero slope suggests that the covariate's effect is time-dependent. In such cases, the Cox model is mathematically invalid.

Formally, for a subject i who fails at time t_i with covariate vector X_i, the Schoenfeld residual is defined as

$$r_i = X_i - \bar{X}(\hat{\beta}, t_i), \qquad \bar{X}(\hat{\beta}, t_i) = \frac{\sum_{l \in R(t_i)} X_l \exp(\hat{\beta}^T X_l)}{\sum_{l \in R(t_i)} \exp(\hat{\beta}^T X_l)}$$

where $\bar{X}(\hat{\beta}, t_i)$ is the risk-set average of the covariate, weighted by each subject's estimated relative risk. ([9] ch. 6.2)

3.5 Parametric Modelling with Accelerated Failure Time

While non-parametric estimators like Kaplan-Meier are useful for univariate exploration, they can't isolate the effect of specific covariates while controlling for others (e.g., assessing the impact of manufacturer origin while holding temperature constant). To achieve this, we employ multivariable parametric survival models.

Where a hazard-based model expresses the effect of a covariate as a multiplier of the instantaneous risk of failure, the AFT model expresses it instead as a multiplier of the time scale itself. For a physical system such as a hard drive this can be interpreted as a time ratio that can be read as a proportional change in expected lifetime, without reference to any baseline hazard. A further property of the framework, developed in Section 3.5.3, is that the requirement of proportional hazards between covariate profiles turns out to be a property of the chosen error distribution rather than of the AFT formulation itself, which means that unlike the Cox model, AFT models are not all bound by it.

3.5.1 The AFT Log-Linear Formulation

Unlike hazard-based models that assume covariates multiply the risk of failure, the AFT model assumes that the effect of a covariate is to stretch or compress the expected survival time of the subject. The framework models the natural logarithm of the survival time T as a linear function of the covariates:

$$\ln(T) = \beta_0 + \beta^T X + \sigma\epsilon$$

where β_0 is the intercept, σ is a scale parameter, and ϵ is an error term following a specified distribution.

3.5.2 The Acceleration Factor

By exponentiating the log-linear equation, we can express the failure time $T = \exp(\beta^T X) \cdot \exp(\sigma\epsilon)$. If we then let $T_0 = \exp(\beta_0 + \sigma\epsilon)$ represent the baseline failure time when all covariates are zero. The term $\exp(\beta_1 X_1 + \cdots + \beta_p X_p)$ acts as an Acceleration Factor.

If a covariate vector X yields a negative coefficient sum, the acceleration factor is < 1, which compresses time. Physically this means the system ages faster, shrinking its expected lifespan. If the sum is positive, the acceleration factor is > 1, meaning the lifespan is stretched.

Let $S_0(t)$ be the baseline survival function. The fundamental property of the AFT model is that the survival function for a subject with covariates X is simply the baseline survival function evaluated at a scaled time:

$$S(t \mid X) = S_0\big(t \cdot \exp(-\beta^T X)\big)$$

.

3.5.3 Choice of Error Distribution

To use the AFT model for inference we must specify the distribution of the error term ϵ. Three standard choices are considered: the Weibull, the log-logistic and the log-normal.

Under a Weibull error term, the model has scale parameter $\lambda(X) = \exp(\beta^T X)$ and shape parameter $\rho = 1/\sigma$, giving the survival and hazard functions

$$S(t \mid X) = \exp\left[-\left(\frac{t}{\lambda(X)}\right)^{\rho}\right], \qquad h(t \mid X) = \frac{\rho}{\lambda(X)}\left(\frac{t}{\lambda(X)}\right)^{\rho-1}.$$

A direct consequence of this form is that the hazard ratio between two covariate profiles X_A and X_B reduces to

$$\frac{h(t \mid X_A)}{h(t \mid X_B)} = \exp\big(-\rho\,\beta^T(X_A - X_B)\big),$$

which does not depend on t: for a one-unit change in a single covariate X_j, the implied hazard ratio is $\mathrm{HR}_j = \exp(-\rho\beta_j)$. The Weibull is in fact the only distribution for which the AFT and proportional hazards formulations coincide in this way ([8], sec. 2.3.4). It is therefore unavailable as a primary model whenever the proportional hazards assumption has been rejected, and is used in this thesis only as a comparator (Section 4.4.1).

Under a log-logistic error term, the model has scale parameter $\alpha(X) = \exp(\beta^T X)$ and shape parameter $\gamma = 1/\sigma$, giving

$$S(t \mid X) = \left[1 + \left(\frac{t}{\alpha(X)}\right)^{\gamma}\right]^{-1}, \qquad h(t \mid X) = \frac{(\gamma/\alpha(X))\,(t/\alpha(X))^{\gamma-1}}{1 + (t/\alpha(X))^{\gamma}}.$$

This family imposes no such restriction on the hazard ratio, which instead varies with t and tends toward 1 as $t \to \infty$. Its structural assumption is of a different kind: rearranging the survival function gives

$$\ln\frac{1 - S(t \mid X)}{S(t \mid X)} = \gamma \ln t - \gamma \ln \alpha(X),$$

so the log-odds of failure by time t is linear in $\ln t$ with a slope γ common to all covariate profiles and an intercept that shifts with X.

The log-normal distribution, under which $\ln T$ is normally distributed with mean $\beta^T X$ and standard deviation σ, is likewise a pure AFT family with no proportional hazards restriction,

$$S(t \mid X) = \Phi\left(-\frac{\ln t - \beta^T X}{\sigma}\right).$$

3.5.4 Interpretation of the Log-Logistic Shape Parameter

As the primary model adopted in this thesis, the log-logistic AFT model's shape parameter γ merits closer interpretation. Its value determines the qualitative form of the hazard function:

- $\gamma \leq 1$: the hazard decreases monotonically from $t = 0$.
- $\gamma > 1$: the hazard rises from zero to a single maximum, located at

  $$t^\star(X) = \alpha(X)\,(\gamma - 1)^{1/\gamma},$$

  and declines thereafter.

Two aspects of this are relevant to how the model is interpreted in Chapter 5. First, γ is common to all covariate profiles, but the location of the peak, $t^\star(X)$, depends on X through $\alpha(X)$: different groups of drives are predicted to reach their point of maximum hazard at different ages, even though the shape of the hazard curve is the same for all of them. Second, whether the declining phase is observed at all in a given dataset depends on where $t^\star(X)$ falls relative to the range of ages actually reached. A fitted model with $\gamma > 1$ and $t^\star(X)$ beyond the oldest drive age in the data describes a hazard that is increasing throughout the observed range; the possibility of a later decline is a property of the fitted distribution, not a claim being made about the data. ([8], ch. 2.2-2.3)

3.6 Model Diagnostic Tools

3.6.1 AIC

The relative quality of parametric models is evaluated using the Akaike Information Criterion, which estimates the relative information loss between models, which will be useful for checking which distribution fits AFT better. AIC can be formulated as:

$$\text{AIC} = 2k - 2\ln(\hat{L})$$

where k is the number of estimated parameters and $\hat{L}$ is the maximized value of the likelihood function for the model. The AIC introduces a penalty term for complexity ($2k$) to prevent over-fitting. The model that ends up yielding the minimum AIC is preferred. ([8], ch 8.7)

3.6.2 Concordance Index

To measure the predictive discrimination of the survival model, we calculate Harrell's Concordance Index (C-index) which is a non-parametric rank correlation metric. The C-index estimates the probability that, for any randomly selected pair of subjects, the subject predicted by the model to have a shorter lifespan actually failed first.

$$c = \frac{\text{Number of concordant pairs}}{\text{Total number of evaluable pairs}}$$

A pair (i, j) is evaluable only if it can be established which subject survived longer. A C-index of 0.5 corresponds to random discrimination and a value of 1.0 to perfect discrimination between all evaluable pairs. What constitutes an acceptable value depends on the application and the event rate in the data; no fixed threshold is adopted here, and the values reported in Chapter 5 are compared across model specifications rather than against an external benchmark. ([6])

3.6.3 Goodness-of-Fit Overlay

Beyond comparing candidate distributions by AIC, it is useful to check whether a fitted parametric model is consistent with the data in absolute terms, independent of any comparison to competing parametric families. A standard graphical diagnostic for this is to overlay the fitted parametric survival function $\hat{S}(t)$ against the non-parametric Kaplan-Meier estimate for the same data, together with the Kaplan-Meier confidence band. Because the Kaplan-Meier estimator makes no distributional assumption, it serves as an empirical benchmark: if the parametric curve tracks the Kaplan-Meier estimate closely and remains within its confidence band across the observed range, the parametric model is judged to capture the observed survival experience adequately. Systematic or growing departure from the band, by contrast, indicates that the assumed distribution does not fit the data well, independent of how it compares to other parametric families under AIC. ([10], p. 11-13, p. 19)

4 Statistical modelling

The goal of this chapter is to formally construct a survival analysis framework that identifies and quantifies the mechanical, environmental, and usage-based predictors of hard drive failure, accounting for the unique characteristics of observational datacenter data.

4.1 Preparing the Data

Before any model is fitted, several decisions have to be made about how the raw daily logs described in Chapter 2 are converted into the row-per-drive format used throughout this chapter.

4.1.1 Cohort Formatting

The raw data is a daily panel: one row per drive per day. For this thesis, each drive is instead represented by a single row, using entry and exit ages derived from SMART 9 (Section 2.2) and either an averaged (temperature) or cumulative (power cycles, SMART 5/197 status) summary of its covariates over the observation period. A more advanced model would allow these covariates to vary over time within a drive, which would in particular let temperature and power-cycle effects be estimated from their trajectory rather than a single summary value; such time-varying covariate models are considerably more computationally expensive, however, and are left outside the scope of this thesis (Section 6.2).

4.1.2 Binarising SMART 5 and SMART 197

Section 2.4 showed that non-zero SMART 5 and SMART 197 readings span several orders of magnitude, and that both the share of affected drives and the typical raw value differ by manufacturer. Given this, and given that manufacturers are not known to share a common internal scale for these raw counts, it is not reasonable to treat the raw values as though they carried a comparable, proportional relationship to risk across drives. Both variables are therefore converted into binary indicators (present/absent) for the main analysis, which trades away information about severity within an affected drive in exchange for comparability across manufacturers.

4.1.3 The Lead-in Window

As noted in Section 2.5, SMART 5 and SMART 197 are frequently first logged in the days immediately preceding a failure. Because the aim of this analysis is to characterise how these indicators relate to a drive's subsequent survival, an indicator that also captures errors logged on or near the day of failure would in part be describing the failure event itself rather than a condition that preceded it. To address this, an indicator is only set to 1 if the corresponding error was logged more than 7 days before the drive's last recorded day; errors within that final week are disregarded, whether or not the drive went on to fail. The choice of a 7-day window is examined in Section 4.4.3.

4.1.4 Temperature Summary

The temperature covariate used throughout this chapter is a single value per drive, averaged over its entire observation period. This risks masking short-lived temperature spikes or periods of unusually high operating temperature within an otherwise moderate average. To check whether this is a serious concern, Table 2 reports the fraction of drives that ever exceeded a series of thresholds.

Table 2: Thermal spike threshold analysis: share of drives ever reaching a given temperature.

Model	Total Drives	% Hit 50 °C	% Hit 60 °C	% Hit 70 °C
WDC	8,864	21.51	2.26	0.10
Toshiba	39,197	16.18	1.08	0.03
Seagate	11,255	29.87	2.15	0.00

Only a small fraction of drives ever reach potentially unsafe temperatures, and almost none reach a severe spike of 70 °C. Given how rare high-temperature excursions are in this climate-controlled setting, the mean temperature is used as the primary summary for the remainder of the analysis.

This choice is checked directly by refitting the primary log-logistic model with mean temperature replaced by an alternative summary – the maximum temperature observed over a drive's lifetime – and comparing the two by AIC (Table 3).

Table 3: Comparison of temperature summaries by model fit.

Temperature summary	AIC	ΔAIC	Time Ratio	Concordance
Mean temperature	57,798.4	109.6	1.035	0.801
Max temperature	57,688.8	0.0	1.026	0.816

The maximum-temperature summary fits marginally better than the mean ($\Delta\text{AIC} = 109.6$) and gives an almost identical time ratio (1.026 vs. 1.035), so the two summaries support the same substantive conclusion: a small, statistically significant, positive association between higher observed temperature and longer expected lifetime. Mean temperature is nonetheless retained as the primary summary throughout this thesis. Conceptually, the research questions concern how sustained operating conditions relate to a drive's cumulative wear (Section 3.5.4), which a typical operating temperature captures more directly than a single peak reading; a maximum-based summary is better suited to detecting acute thermal events than to characterising long-run operating stress. This is reinforced by Table 2: because high readings are rare in this climate-controlled setting, a maximum-based summary would mostly distinguish a small number of drives that ever crossed a given threshold, rather than reflect graded differences in typical conditions across the cohort. The marginal fit advantage of the maximum is noted in Section 6.2 as a candidate refinement rather than a concern for the reported conclusions.

4.2 Non-Parametric Survival

Figure 1 shows the Kaplan-Meier estimate for the pooled 14TB cohort. Survival declines gradually through roughly the first 1,900 days of drive age, after which the curve steepens noticeably. This later steepening should be read with caution: as noted in Section 2.4, drives beyond roughly 1,900 days are drawn almost entirely from the Toshiba model, since the Seagate and WDC cohorts do not include drives old enough to be observed that far out. The apparent acceleration in the pooled curve is therefore substantially a statement about a shrinking, model-specific subset of the risk set, not a pooled effect observed across all three manufacturers.

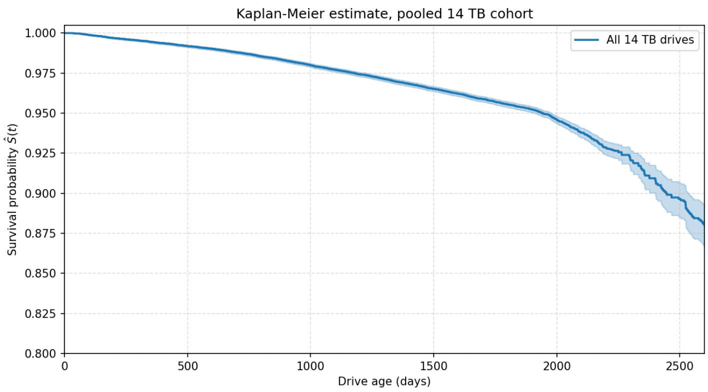

Figure 1: Kaplan-Meier estimate, pooled 14TB cohort. Shaded band is the 95% confidence interval.

Splitting the estimate by drive model (Figure 2) confirms this: WDC maintains the highest survival probability throughout the range it is observed in, Seagate the lowest, and Toshiba lies between the two while alone extending out to the full 2,648-day maximum observed age. Table 4 reports survival probabilities and 95% confidence intervals at three fixed ages; the intervals for the three models do not overlap at any of the three ages shown, and a likelihood-ratio test comparing the three models, adjusted for the remaining covariates, confirms the difference is not attributable to chance ($\chi^2 = 64.77$ on 2 df, $p = 8.6 \times 10^{-15}$).

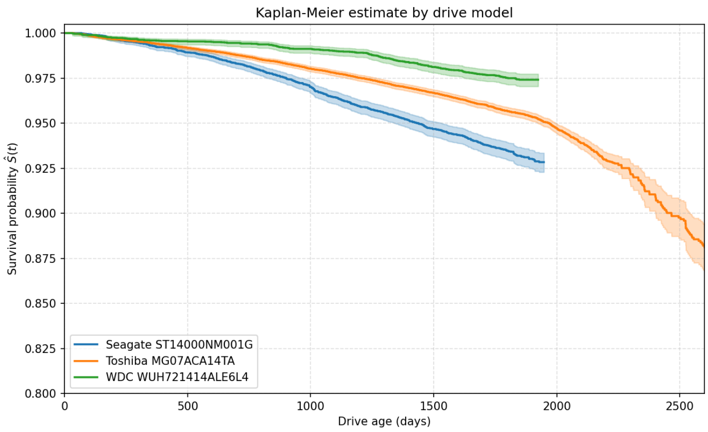

Figure 2: Kaplan-Meier estimate by drive model. Shaded bands are 95% confidence intervals.

Table 4: Survival probability by drive model at fixed ages.

	$\hat{S}(t)$ [95% CI]		
Model	$t = 1000$d	$t = 1500$d	$t = 2000$d
Seagate	0.970 [0.967, 0.973]	0.947 [0.942, 0.951]	0.928 [0.923, 0.934]
Toshiba	0.980 [0.979, 0.982]	0.967 [0.965, 0.969]	0.947 [0.945, 0.950]
WDC	0.991 [0.989, 0.993]	0.981 [0.978, 0.984]	0.974 [0.970, 0.977]

The separation by SMART status is considerably larger than the separation by model. Figures 3 and 4, and Table 5, show survival by reallocated- and pending-sector status respectively. Drives with no reallocated sectors have an estimated median survival that is not reached within the observation window; drives with reallocated sectors present have a median survival of 2,111 days. The same pattern holds for pending sectors, with a median survival of 1,954 days once present, against a median not reached in their absence.

While these non-parametric estimates already establish that drive model and SMART status are associated with clear differences in survival, they cannot isolate the effect of one covariate while holding the others fixed – a drive with reallocated sectors is not otherwise a random draw from the cohort, and may also differ systematically in temperature, power-cycle history, or manufacturer. Section 4.3 addresses this using the Cox proportional hazards model, before Section 4.4 moves to the parametric AFT specification.

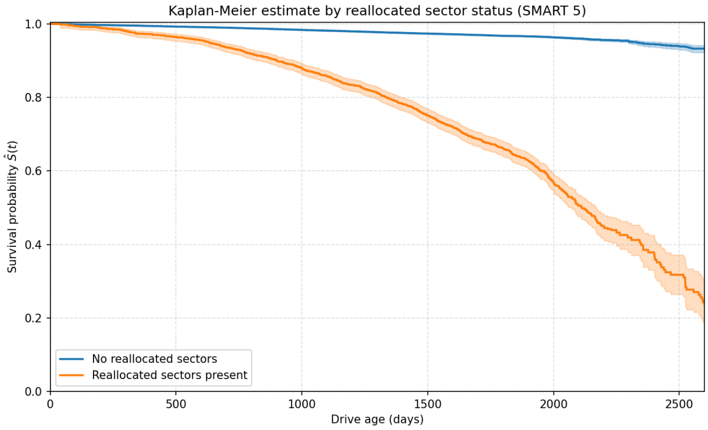

Figure 3: Kaplan-Meier estimate by reallocated-sector status (SMART 5).

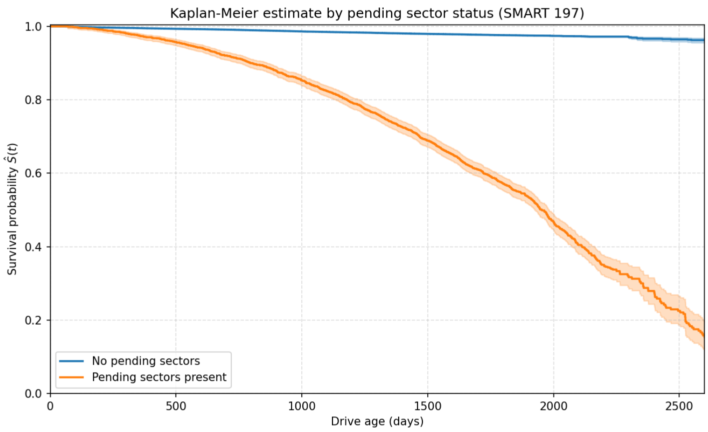

Figure 4: Kaplan-Meier estimate by pending-sector status (SMART 197).

Table 5: Survival probability by SMART status at fixed ages.

	$\hat{S}(t)$		
Group	$t = 1000$d	$t = 1500$d	$t = 2000$d
No reallocated sectors	0.984	0.974	0.964
Reallocated sectors present	0.880	0.750	0.570
No pending sectors	0.986	0.980	0.974
Pending sectors present	0.853	0.689	0.468

4.3 Testing the Proportional Hazards Assumption

Section 4.2 established, using non-parametric estimators, that reallocated sectors, pending sectors and drive model are all associated with clear differences in survival. To quantify these associations while controlling for the remaining covariates simultaneously, the natural next step is the Cox Proportional Hazards model, which sidesteps the need to specify the underlying distribution of drive lifetimes. Its validity, however, depends on the proportional hazards assumption holding.

To test this assumption, the model parameters are first estimated, and the Schoenfeld residuals are then calculated: the difference, at each failure time, between the observed covariate value of the drive that failed and the covariate value expected from the risk set at that moment. If proportional hazards holds, these residuals are independent of time; a statistical correlation test between the residuals and ranked failure times is used to check this.

Every covariate (temperature, SMART 5, SMART 197, power-on cycles, and the manufacturer/model indicators) shows a statistically significant correlation with time at $p < 0.05$. All covariates except the Western Digital model indicator have $p < 0.00005$; the WDC indicator has $p = 0.0466$. The proportional hazards assumption is therefore rejected: the hazard ratios implied by the fitted coefficients are not constant over drive age, but change as drives grow older.

This rejection is a statement about the hazard *ratio*, not about the hazard itself. The Cox model leaves the baseline hazard $h_0(t)$ completely unspecified, and therefore already permits a hazard that rises with drive age; what the model requires is that the ratio between the hazards of any two covariate profiles stay constant over time. The test above shows that this requirement does not hold in these data.

This has a direct consequence for which parametric alternative can be used. As established in Section 3.5.3, the Weibull distribution is the only distribution that is simultaneously an accelerated failure time model and a proportional hazards model: under a Weibull error term the implied hazard ratio is $\exp(-\rho\beta_j)$, which does not depend on t. Fitting a Weibull AFT model would therefore not relax the assumption just rejected – it would retain it, while additionally imposing a parametric form on the baseline hazard, making the specification strictly more restrictive than the Cox model already rejected above. Section 4.4 instead uses the log-logistic AFT model, in which the hazard ratio between two covariate profiles is free to vary with age.

4.4 Accelerated Failure Time Models

As established previously, the log-logistic AFT model is used as the primary specification. This section reports the model-selection diagnostics that support that choice; the fitted coefficients and their interpretation are presented in Chapter 5.

4.4.1 Model Selection

The three candidate error distributions introduced in Section 3.5.3 were fitted to the design established in Section 4.1, using the same six covariates and the same left-truncated time scale. All three specifications converge to an in-sample concordance of 0.801; Table 6 compares them by log-likelihood and AIC.

Table 6: AFT model selection: log-likelihood, AIC and shape parameter by distribution.

Distribution	Log-likelihood	Parameters	AIC	ΔAIC	Shape estimate
Weibull	−28,891.19	8	57,798.38	0.0	$\hat{\rho} = 1.432$
Log-logistic	−29,048.01	8	58,112.01	313.6	$\hat{\gamma} = 1.526$
Log-normal	−29,420.88	8	58,857.76	1,059.4	$\hat{\sigma} = 1.429$

The Weibull specification attains the lowest AIC. This is not the operative criterion here: AIC measures marginal distributional fit and is silent on the regression structure a family imposes, and the Weibull's advantage on this criterion is obtained precisely under the constant-hazard-ratio restriction rejected in Section 4.3. Restricting attention to the two families that do not carry that restriction, the log-logistic fits the data substantially better than the log-normal (ΔAIC = 745.7), and it is adopted as the primary specification for Chapter 5. The Weibull fit is retained as a robustness check rather than discarded outright: Section 5.2 shows that refitting the model under all three distributions preserves the sign and approximate ranking of every covariate, with the reallocated-sectors estimate the one exception showing a more than incidental difference between families.

4.4.2 Goodness of Fit

Earlier we introduced the goodness-of-fit overlay as a standard graphical diagnostic: a fitted parametric survival curve is checked against the non-parametric Kaplan-Meier estimate and its confidence band, which together serve as an empirical benchmark that carries no distributional assumption of its own. Figure 5 applies this to a subset of 9,977 pristine Seagate drives (186 failures) with no reallocated or pending sector history, overlaying the fitted log-normal and log-logistic survival functions against the Kaplan-Meier estimate for that subset. Both parametric curves track the empirical estimate closely and remain within its confidence band across the observed age range, and are visually near-indistinguishable from one another (log-normal AIC = 4,613.4, log-logistic AIC = 4,616.2 on this subset); the preference for the log-logistic rests on the structural argument of Section 4.3, not on a difference in fit visible here.

Collecting the specification choices made above, the model fitted throughout the remainder of this thesis is:

$$\begin{aligned}\ln(T_i) = \beta_0 + \beta_1 X_{\text{temp},i} + \beta_2 X_{\text{realloc},i} + \beta_3 X_{\text{pend},i} \\ + \beta_4 X_{\text{Toshiba},i} + \beta_5 X_{\text{WDC},i} + \beta_6 X_{\text{pwr},i} + \sigma\epsilon_i,\end{aligned} \tag{1}$$

where X_{realloc}, X_{pend}, X_{Toshiba} and X_{WDC} are binary indicators (Seagate is the reference category), and ϵ_i follows a standard logistic distribution under the log-logistic specification.

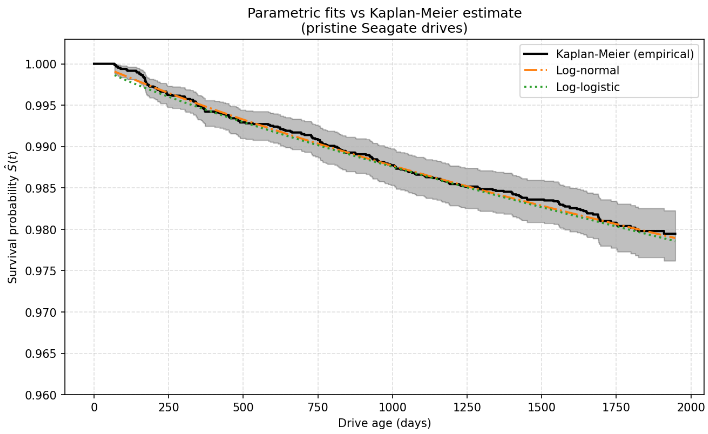

Figure 5: Fitted log-normal and log-logistic survival functions against the Kaplan-Meier estimate, pristine Seagate drives (n = 9,977, 186 failures). Shaded band is the Kaplan-Meier 95% confidence interval.

4.4.3 Sensitivity Analysis

We previously introduced a 7-day lead-in window, disregarding SMART 5 and SMART 197 readings logged within the final week before a drive's last recorded day. Because this window length is a modelling choice rather than a value given by the data, its effect on the estimated time ratios is checked directly: the log-logistic and Weibull models are refit under 1-day, 7-day and 30-day lead-in windows, on an identical cohort of 59,262 drives common to all three, so that only the covariate definition changes between them.

Figures 6 and 7 show the result. For both SMART covariates and both distributions, the estimated time ratio changes only modestly across the full range of lead-in lengths: under the log-logistic model, the pending-sectors time ratio moves from 0.129 at a 1-day window to 0.163 at 30 days, and the reallocated-sectors ratio moves from 0.504 to 0.508; the Weibull model shows a closely comparable pattern (0.123 to 0.162, and 0.546 to 0.556). The association between these SMART indicators and reduced survival is therefore not an artefact of errors logged immediately before failure as it is already present when the final month of a drive's observation is disregarded entirely, and so the 7-day window used throughout the rest of this thesis is retained as a reasonable middle ground.

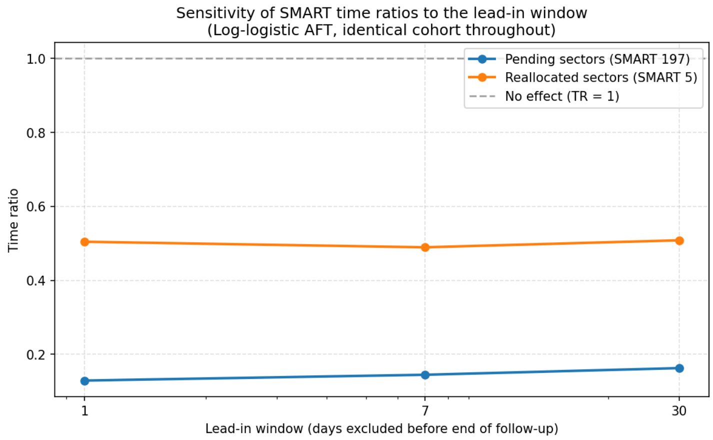

Figure 6: Sensitivity of SMART time ratios to the lead-in window, log-logistic AFT, identical cohort across all three windows.

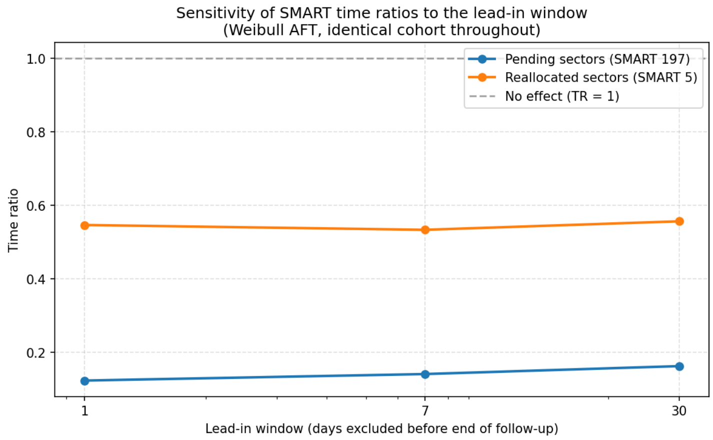

Figure 7: Sensitivity of SMART time ratios to the lead-in window, Weibull AFT, identical cohort across all three windows.

A caveat applies to this result, we see that the AIC rises with the length of the lead-in window (log-logistic: 56,927.2 at 1 day, 58,112.0 at 7 days, 59,690.3 at 30 days) even though the same 59,262 drives are used throughout: a longer window recodes more genuine pre-failure SMART events as absent, which weakens rather than strengthens the fitted association. This is an expected consequence of discarding signal for the sake of the timing argument, not evidence against the model.

5 Results

Section 4.2 already established, via Kaplan-Meier estimation, that drive model and SMART status are associated with clear differences in survival. This chapter reports the adjusted, quantitative estimates from the log-logistic AFT model selected in Section 4.4, together with the answers to the two research questions posed in Chapter 1.

5.1 The Study Cohort

The composition of the study cohort – drive counts, failure counts, exposure, annualised failure rates, and covariate ranges by model – is described in full in Section 2.4 and is not repeated here. In summary: the cohort comprises 59,262 drives across three 14TB models, contributing 276,441 drive-years of exposure and 2,793 observed failures, for an overall annualised failure rate of 1.01%, ranging from 0.51% (WDC) to 1.40% (Seagate).

5.2 Estimated Time Ratios

Table 7 reports the time ratios estimated by the primary log-logistic AFT model, together with the Weibull estimates retained as a robustness comparison. A time ratio below 1 corresponds to a shorter expected lifetime relative to the reference category (Seagate ST14000NM001G, no reallocated or pending sectors); a time ratio above 1 corresponds to a longer one.

Table 7: Estimated time ratios, log-logistic AFT (primary) and Weibull AFT (comparator).

	Log-logistic		Weibull	
Covariate	TR	95% CI	TR	95% CI
Pending sectors	0.145	[0.132, 0.158]	0.140	[0.128, 0.154]
Reallocated sectors	0.489	[0.450, 0.531]	0.533	[0.496, 0.573]
Power cycles at entry	0.988	[0.982, 0.995]	0.988	[0.983, 0.994]
Mean temperature	1.040	[1.033, 1.046]	1.035	[1.029, 1.041]
Toshiba model	1.065	[0.998, 1.136]	1.067	[1.006, 1.132]
WDC model	1.498	[1.360, 1.651]	1.487	[1.354, 1.633]

All six covariates carry the same sign under both distributions, and every covariate except the Toshiba model indicator is significant at $p < 0.001$ under both. The Toshiba indicator itself has $p = 0.056$ under the log-logistic model and $p = 0.031$ under the Weibull model – significant at the conventional 5% threshold in one specification and not quite in the other, so it should be read as a borderline effect rather than a clearly established one.

The two largest effects by a wide margin are the SMART indicators. Pending sectors are associated with an expected lifetime roughly one-seventh that of an otherwise identical drive; reallocated sectors, with roughly half. Both are consistent with the non-parametric comparison in Section 4.2. Temperature is statistically significant but small in magnitude: each additional degree Celsius of mean operating temperature is associated with a 4.0% increase in expected lifetime (the feasability of this result is discussed in section 6). Power cycles at study entry show a small per-cycle effect (a 1.2% reduction in expected lifetime per additional cycle) that compounds to a substantial one over a wider range: a drive entering observation with 100 power cycles recorded has an expected lifetime only around 31% that of an otherwise identical drive with none ($0.988^{100} \approx 0.308$). Because the median drive in this cohort enters with only 2 recorded power cycles, this 100-cycle figure describes a contrast between drives near the upper tail of the distribution rather than a typical within-cohort comparison as the per-cycle time ratio in Table 7 is the estimate that should be quoted without further qualification. Relative to the Seagate reference, WDC drives show an expected lifetime roughly 50% longer.

Robustness across distributional families. We previously fitted all three candidate distributions to the same design. Every covariate keeps its sign across all three, and the ranking of effect sizes is essentially unchanged. The one departure from close agreement is reallocated sectors, where the log-normal estimate (TR = 0.385) is noticeably more extreme than the Weibull and log-logistic estimates (0.533 and 0.489); the SMART 197, power-cycle, temperature and WDC estimates all agree to within 0.02–0.06 across all three families. The Toshiba effect is also the least stable of the six: significant at the 5% level under the Weibull specification only ($p = 0.031$ vs. 0.056 and 0.221 under log-logistic and log-normal respectively). Taken together, the substantive conclusions – large, highly significant effects for both SMART indicators; a smaller but real manufacturer difference favouring WDC; modest associations with temperature and power cycles – do not depend on which of the three distributions is used, though the exact magnitude of the reallocated-sectors effect and the significance of the Toshiba effect are somewhat sensitive to that choice.

Concordance. The log-logistic model achieved an in-sample concordance of 0.801. Five-fold cross-validation gives a mean concordance of 0.8007 (SD 0.0078), indicating no meaningful overfitting; the Weibull and log-normal models show the same pattern (in-sample 0.8012 and 0.8014; cross-validated 0.8009 and 0.8008 respectively).

The shape parameter. The estimated log-logistic shape parameter is $\hat{\gamma} = 1.526$, placing the fitted hazard in its increasing phase (Section 3.5.4) for a drive without either SMART indicator across essentially its entire observed lifetime: evaluated at the cohort mean temperature and median power cycles at entry, the fitted hazard for a pristine Seagate drive peaks at approximately 13,200 days – roughly five times the oldest drive age actually observed (2,648 days). For a drive reporting pending sectors, the same calculation places the turning point at roughly 1,900 days, inside the observed range. The practical implication is that the hazard ratio between a drive with pending sectors and an otherwise identical drive without is not constant with age, but attenuates as both drives grow older.

6 Discussion

6.1 Interpreting the Coefficients

When reviewing the time ratios reported in Section 5.2, a few values may stand out as counter-intuitive. Most notably, mean operating temperature carries a positive time ratio (1.040 per degree Celsius): higher observed temperature is associated with a slightly longer, not shorter, expected drive lifetime. The effect is small relative to the SMART indicators, and was of similar magnitude and significance under every specification examined, so it is unlikely to be a modelling artefact. One possible interpretation is that within this tightly controlled data-centre environment, drives may be operating below their ideal temperature range, such that a slightly warmer operating point is mildly beneficial rather than harmful. This interpretation should be treated with ample caution: the covariate was observed only over a narrow range (cohort mean 31.7 °C, SD 4.1), and the association is with mean, sustained temperature rather than with short-term fluctuations or interactions with drive age, either of which could plausibly behave differently. A causal claim about overcooling can not be provided by this observational dataset. A further consideration is that mean temperature is computed over each drive's observation window, so unlike the power-cycle covariate it is not measured at a point prior to the outcome. Drives that failed early contribute correspondingly short averaging windows, and any systematic difference in operating conditions between short- and long-lived drives would be absorbed into this covariate; the positive association should therefore be read alongside this possibility as well as the overcooling interpretation.

Power cycles at study entry show a modest negative association with expected lifetime: a time ratio of 0.988 per additional cycle, compounding to approximately 0.31 over 100 cycles. Two points of interpretation apply. First, this covariate measures the number of cycles recorded before a drive entered the observation window, not a cumulative count that grows with the drive's own observed lifetime; the association therefore reflects handling and commissioning history prior to entry, rather than an accumulating stress that grows over the period actually being modelled, which is a materially different claim from "more power cycles shorten drive life" in general. Second, because the median drive in this cohort enters with only 2 recorded cycles, the 100-cycle contrast is a comparison between drives that are unusual for this cohort rather than a typical one; the per-cycle estimate is the figure that generalises. Backblaze's own reporting on this SMART attribute reaches a similarly inconclusive position, noting a difference in average power-cycle count between failed and operational drives without being able to rule out confounding by drive age ([11]).

The clearest and most robust predictors of reduced lifetime remain the two SMART indicators. Reallocated sectors (SMART 5) are associated with a time ratio of 0.489, compressing expected lifetime to roughly half; pending sectors (SMART 197) are associated with a time ratio of 0.145, a substantially larger effect. Both pertain to physical defects on the drive's storage medium, but the mechanisms they signal differ: SMART 5 indicates that a bad sector was identified and successfully reallocated to a healthy area of the disk, meaning the drive has already recovered from the specific defect being counted, whereas SMART 197 indicates a sector that is currently unreadable and awaiting reallocation, signaling a more acute, unresolved condition. This distinction is consistent with the substantially larger effect estimated for pending sectors throughout Chapters 4 and 5, and with the more extreme non-parametric separation shown in Figure 4 relative to Figure 3.

The log-logistic shape parameter offers a structural explanation for why the Cox proportional hazards assumption was rejected, beyond simply having established that it was. Because the log-logistic hazard rises to a maximum before declining, and because the location of that maximum, $t^\star(X) = \alpha(X)(\gamma-1)^{1/\gamma}$, scales with a drive's covariate profile, drives with and without SMART faults are not simply subject to the same hazard shape at different heights – they are, in effect, at different stages of that shape at any given age. Section 5.2 illustrated this concretely: for a drive without SMART faults, the fitted turning point falls far beyond any age observed in this cohort, so its hazard is rising throughout; for a drive with pending sectors, the turning point falls inside the observed range, so its hazard is closer to levelling off. The resulting hazard ratio between the two groups is consequently largest for younger drives and attenuates as both groups age. This is exactly the kind of age-dependence a constant hazard ratio cannot express, and offers a physically plausible account of the Schoenfeld-residual rejection in Section 4.3: a drive that develops pending sectors early and nonetheless survives is not, on this account, equivalent in risk to one that develops them only recently, which is a distinction the Cox and Weibull specifications are structurally unable to draw.

The manufacturer/model contrasts are the largest effects in the model after the two SMART indicators: WDC drives show an expected lifetime roughly 50% longer than the Seagate reference, with Toshiba intermediate and only borderline distinguishable from Seagate. As with any observational comparison between manufacturers, numerous unobserved factors could plausibly contribute to this difference beyond intrinsic build quality, such as differences in typical workload, rack placement, firmware, or batch-level manufacturing variation are all candidates this dataset cannot separate from a manufacturer effect. We also previously noted a structural limitation specific to this cohort: because each of the three models corresponds to a distinct manufacturer, the model and manufacturer covariates are perfectly confounded here, and the reported effect should be read as a joint contrast rather than attributed to manufacturer identity specifically.

6.2 Limitations and Future Work

Several limitations should be acknowledged. First, the data are observational rather than experimental; the estimated effects should be interpreted as associations, not as strict causal effects. This is particularly relevant for the manufacturer/model contrast discussed above, and for the temperature finding, where an observational association cannot rule out confounding by unmeasured operating conditions.

Second, the temperature covariate used throughout is a per-drive average over the full observation period (Section 4.1.4). A maximum-temperature summary fit marginally better by AIC without changing the substantive conclusion, and remains a candidate refinement for future work, particularly if paired with a more detailed model of temperature over time rather than a single lifetime summary, like with time-varying covariates.

Third, SMART attributes are not perfectly standardised across manufacturers (Section 2.5), which required restricting to a small set of comparatively robust attributes and converting SMART 5 and SMART 197 into binary indicators. This improved comparability across manufacturers at the cost of discarding information about the severity of an affected drive's degradation; Section 2.4 showed that raw non-zero values for these attributes range over several orders of magnitude, so two drives coded identically as "affected" may differ substantially in their underlying condition.

Fourth, the analysis relies on static, per-drive summaries of covariates measured either at study entry (power cycles) or averaged over the full observation window (temperature, SMART status), rather than on covariates that vary over time within a drive's own record. This is a reasonable simplification for the scope of this thesis, but it means that within-drive temporal dynamics – for instance, a SMART fault appearing shortly before an unrelated temperature spike – are not modelled explicitly.

Finally, right-censoring is treated as non-informative throughout. This is the standard assumption in survival analysis, but Section 2.5 already notes a plausible mechanism by which it could be violated in this setting: a technician preemptively retiring a drive with elevated error counts would remove that drive from observation for a reason directly related to its underlying risk.

These limitations point to natural extensions. Modelling SMART attributes and temperature as fully time-varying covariates, rather than static per-drive summaries, would address the second and fourth limitations directly, though at considerably greater computational cost than the specification used here. Incorporating frailty terms would be a further natural step: drives are installed in specific server pods and originate from specific manufacturing batches, so individual drive lifetimes are unlikely to be strictly independent, an assumption this thesis has maintained throughout. Extending the cohort to additional capacities, product families, or observation periods would also help establish how far the specific estimates reported here generalise beyond the three 14TB models examined.

6.3 Conclusion - Answering the Research Questions

Returning to the two questions posed in Section 1.1: temperature, power-cycle count, and the two SMART fault indicators all function as predictors of HDD failure in the sense intended by this thesis – each is associated with expected drive lifetime at $p < 0.001$ once the remaining covariates and drive model are controlled for, with only the Toshiba-vs-Seagate contrast falling short of conventional significance. The direction and magnitude of these associations differ substantially. Pending and reallocated sectors are associated with the largest effects by a wide margin, compressing expected lifetime to roughly one-seventh and one-half respectively, and remained the most consistent predictors across every distributional family, lead-in window, and robustness check performed. Power cycles at study entry and mean operating temperature both show statistically significant but small associations – a modest reduction and a modest increase in expected lifetime per unit, respectively – and drives from Western Digital showed substantially longer expected lifetimes than the Seagate reference.

As emphasised throughout, these are associations between groups of drives defined by shared covariate values, not predictions for any individual drive; no forecasting task was evaluated, and none of the findings above should be read as such.

References

[1] Emerson College Technology & Media, *Hard Drive Overview* (n.d.) https://support.emerson.edu/hc/en-us/articles/21709168771867-Hard-Drive-Overview#introduction-0-0

[2] Brad Warbiany, Western Digital, *The Long-Term Case for HDD Storage* (2025) https://blog.westerndigital.com/long-term-case-for-hdd-storage/

[3] Backblaze, *Hard Drive Reliability Test Data*[Dataset] (2026), https://www.backblaze.com/cloud-storage/resources/hard-drive-test-data

[4] Don Hall, Enterprise Storage Forum, *Life Expectancy of a Drive: HDD, SSD, and Flash* (2023), https://www.enterprisestorageforum.com/hardware/life-expectancy-of-a-drive/

[5] *Wikipedia*, (2026), Self-Monitoring, Analysis and Reporting Technology, https://en.wikipedia.org/wiki/Self-Monitoring,_Analysis_and_Reporting_Technology#Known_ATA_S.M.A.R.T._attributes

[6] Statistical Odds & Ends, *What is Harrell's C-index?* (2019), https://statisticaloddsandends.wordpress.com/2019/10/26/what-is-harrells-c-index/

[7] Klein, J. P., & Moeschberger, M. L. (2003). *Survival Analysis: Techniques for Censored and Truncated Data*. 2nd ed. Springer.

[8] Kalbfleisch, J. D., & Prentice, R. L. (2002). *The Statistical Analysis of Failure Time Data*. 2nd ed. Wiley.

[9] Therneau, T. M., & Grambsch, P. M. (2000). *Modeling Survival Data: Extending the Cox Model*. Springer.

[10] Latimer, N.R̃. (2011). *NICE DSU Technical Support Document 14: Survival analysis for economic evaluations alongside clinical trials – extrapolation with patient-level data*. Decision Support Unit, ScHARR, University of Sheffield.

[11] Klein, A. (2016). *What SMART Stats Tell Us About Hard Drives.* Backblaze Blog. https://www.backblaze.com/blog/what-smart-stats-indicate-hard-drive-failures/