

Are There Manufacturer Differences in Hard-Drive Reliability?

Christoph Siemroth
University of Essex
CO4 3SQ Colchester, UK
christoph.siemroth@essex.ac.uk

Yeomyung Park
Sungkyunkwan University
03063 Seoul, Republic of Korea
byparkym@g.skku.edu

Abstract—Based on the Backblaze hard disk drive (HDD) dataset, we analyze whether the four major HDD manufacturers represented in the dataset—HGST, Seagate, Toshiba, Western Digital (WD)—show differences in short- to medium-term HDD failure rates. Using two different duration regression models, we find—holding constant drive age, capacity, form-factor, and drive temperature—that Toshiba’s failure rate is slightly above Seagate’s. HGST HDD failure rates are the lowest, about 41% of Seagate’s. WD HDD failure rates are significantly above HGST’s, but still only about 52% of Seagate’s. We also document the effects of age, capacity, temperature and drive location on failure rates.

Keywords: hard disk drive failure rates, HDD reliability, manufacturer differences

I. INTRODUCTION

The typical consumer uses only a handful of different hard disk drives (HDDs) in their lifetimes, not enough to find out by experience whether a specific HDD manufacturer is more reliable than another. Hard-drive manufacturers have a good idea about their own failure rates, given testing and returns, but they might know less about their competitors’ data. Large data storage and cloud service companies might be among the few with the necessary data to investigate manufacturer differences in reliability, since they observe the fate of thousands of HDDs. Here, we analyze the data from one of these data-centers.

The question of drive reliability is important, not just for data security but also for operating cost reasons: For data-centers and cloud storage providers with hundreds of thousands of disks, even a one percentage point drop in yearly failure rates translates into massive cost savings. This is especially true because among the major hardware parts of a server, HDDs are the most unreliable component [1].

In this article, we use the large Backblaze data-center dataset of daily HDD snapshots to investigate whether there are significant differences in the reliability and failure rates of HDDs between four major manufacturers, HGST, Seagate, Toshiba, Western Digital (WD).¹ Overall, we have more than 440,000 HDDs in the sample, combined with a runtime of more than 1.66 million years, allowing us to obtain very precise

¹As of 2025, three manufacturers in our sample (Seagate, Toshiba, WD) manufactured an estimated 99% of all HDDs worldwide [2], [3]. Hence, our study is very comprehensive in covering the meaningful manufacturers. HGST was a big manufacturer, but was acquired by WD in 2012 [4] and later stopped making its own HDDs.

estimates in the failure rates of HDDs by manufacturer, despite the low annual probability of failure of individual HDDs. Observing drives of several different manufacturers at the same company is an advantage, because the drives operate under the same conditions and procedures, thus limiting the number of confounding factors when comparing across manufacturers. Still, since drives of different manufacturers have different ages and installation years, the analysis is challenging and requires more than a comparison of annualized failure rates.

Using duration regression models, we can compare drives across manufacturers, but effectively holding drive age, capacity, form-factor, and drive temperature constant. We find that the HDDs from HGST have the lowest failure rate; about 41% of the failure rate of Seagate HDDs. WD HDDs have about 52% of the failure rate of Seagate HDDs, but it is significantly worse than HGST’s. The Toshiba HDD failure rate is about 107% of Seagate HDDs and hence the highest.

Beyond statistical significance, these manufacturer differences are of high practical importance: The HGST to Toshiba differences imply more than a doubling of the failure rate, thus more than doubling the time and cost spent on replacing drives. For interpretation, the manufacturer differences we estimate are based on a dataset spanning multiple model generations, and so they capture persistent manufacturer differences beyond specific models. In addition, since only a small share of HDDs in the Backblaze dataset were installed more than 10 years ago, our estimates capture short- to medium-term failure rates.

Our goal in this article is to estimate possible manufacturer differences, not to provide a comprehensive explanation for them. However, we can rule out that the manufacturer differences are due to differences in drive age, capacity, form-factor, or average drive temperature, as we statistically hold these constant. We also rule out that HDD positions (as measured by HDD cluster) in the data-center affect the manufacturer differences. Moreover, since WD acquired HGST in 2012 [4], one year prior to the start of our sample, it is possible that technology and procedures were shared between the two. This would explain why their failure rates are similar, compared to the other competitors’ failure rates.

Beyond manufacturer effects, we provide some evidence that drive age, HDD position/assignment in the data-center, drive temperature and capacity matter for the reliability of HDDs. Age displays manufacturer-specific patterns which

are not always monotone. Larger drive temperatures increase failure rates. And higher capacity drives tend to be more reliable. We also find evidence that reliability improvements by drive generation differ by manufacturer.

Since the Backblaze data is increasingly used in the academic literature, we also document some striking features about the data. First, a large share of HDDs is removed from the dataset without failure. Hence, “right-censoring”—when a drive’s failure date is not observed because the drive was still functioning in the most recent data snapshot—occurs not only for the large share of drives that have not yet failed, but also occurs for drives that were removed from the data in the past before failing. Second, HDD manufacturers are not evenly represented across time. In 2014, HGST was the most common manufacturer among newly installed HDDs, in 2017 Seagate was the most common manufacturer among newly installed HDDs, in 2023 it was Toshiba and in 2024 it was WD. This means there are differences in drive age by manufacturer, and serious analyses of reliability have to account for these age differences, given that drive age is a major determinant of drive failure rates. Third, the share of HDDs for which we do not observe an installation date or year—since they were in the sample from the very first day—is very small compared to the share of HDDs where the installation date can be inferred from the data. Hence, “left-censoring”—when a drive’s precise installation date and hence its survival duration is not observed—is not a major issue in this dataset.

An advantage of the Backblaze data over manufacturer returns data and consumer review data is that drives of all manufacturers are treated similarly, by the same technical staff and being used for the same purposes, since they are operated by the same company. Review and returns data from different consumers and companies, on the other hand, cannot rule out that those who prefer one manufacturer use the drives very differently from those who prefer other manufacturers. Hence, there are more confounding factors with returns and review data, in addition to possible selection problems with user reviews (e.g., negative experiences may be more likely to lead to a review than positive experiences).

II. RELATED WORK AND CONTRIBUTION

The recent literature on HDD failures has focused on predicting drive failures [5], often based on the individual drive’s SMART attributes [6]–[8] and using machine learning methods [9]–[13]. Hence, the focus in the literature is on reliable ahead-of-time prediction of failure during operation, rather than on quantifying the statistical factors that contribute to failures over a drive’s lifetime, as we do in this study. This existing literature does not investigate manufacturer differences, and indeed aims to find prediction methods that are valid across HDD models and manufacturers. This difference in research question also motivates differences in the methodology. If useful prediction of failure is the goal, then machine learning models might be best suited. But such models are not as well suited to quantify and interpret manufacturer differences, where more classical regression techniques are appropriate.

In addition to failure prediction, the literature also studies identification of slow failures, where HDDs are working but are much slower than normal [14].

Outside of the academic literature, Backblaze regularly posts descriptive analyzes of HDD failure rates on their blog.² In their posts, Backblaze reports HDD annualized failure rates by model, though not aggregated by manufacturer. Their methods differ substantially from ours. Backblaze compares the available drives at the time. In this article, we document that certain manufacturers are overrepresented in the first cohorts of HDDs, hence the drives of these manufacturers tend to be older. Thus, simply comparing drives across manufacturers is not like-for-like. To account for this, in our regression approach, we control for the start year of each drive, thus effectively comparing drives of similar age across manufacturers. We also control for drive form-factor, capacity, and we rule out that drive position in the data-center affects the manufacturer difference estimates.

Moreover, unlike annualized failure rate analyses, our duration model regressions take into account information contained in censoring events (when a drive exits the data without failure), thus using the information contained in the data to the fullest.³ Finally, the Backblaze analyses tend to focus on the data of the most recent quarter, whereas we use the full time range from the start in 2013 for a comprehensive analysis. Thus, our contribution is a revised methodology and an analysis of the full data sample.

Several studies use data from other data-centers. Researchers have investigated determinants of HDD failures in a large drive population at Google [16], but explicitly avoid reporting manufacturer differences “due to the proprietary nature of these data.” Hence, they leave a gap that we aim to fill. They investigate drive age, temperature, and workload as determinants of drive failure. Moreover, in Meta’s HDD fleet, increased age and lifetime cumulative workload both increase HDD failure rates [17]. Finally, an older study finds that, with much older drives, the interface (SCSI, SATA, FC) makes little difference to failure rates [18].

More recently, several studies investigate factors that affect failures among SSDs, which are increasingly used at scale in data-centers [19]–[22], but again without investigating manufacturer differences.

A separate strand of the literature takes a more statistical approach and investigates which distributions approximate the

²For example here for the first quarter of 2025: <https://www.backblaze.com/blog/backblaze-drive-stats-for-q1-2025/>.

³Duration models treat durations as a random variable with a distribution $f(t)$ to be estimated. If many drives fail at time t_1 , then $f(t_1)$ is estimated to be larger via maximum likelihood. If a drive is censored at time t_2 , then this contributes to the maximum likelihood function not only in that the drive has not failed for $t < t_2$ —as a annualized failure rate analysis would—but it anticipates that the true duration must be $t > t_2$. Thus, this observation enters the likelihood function as $1 - F(t_2)$ [15, eq. 20.23], where $F(\cdot)$ is the cumulative distribution function of the duration, thus making $t > t_2$ more likely. An annualized failure rate analysis, on the other hand, does not use the information contained in censoring. Given that there are many right-censored HDDs in the sample, it is important to explicitly model and properly take these partially observed outcomes into account.

lifetimes and failure rates of hard-drives [23], [24].

III. DATA

A. Data Source

We use the publicly available data from the Backblaze cloud storage provider [25]. Their data starts in 2013, since then documenting for every day and every HDD the drive status (failure or not), several SMART attributes, as well as its unique serial number and model number (from which we infer the manufacturer). We take as given Backblaze’s failure determination. In our analysis, we use all data available from Backblaze from 2013 (the earliest available) until and including the second quarter of 2025. The majority of HDDs in the dataset was added considerably after 2013, so we observe durations of more than 10 years for only 0.52% of HDDs in the dataset. Thus, the sample size at high drive age is considerably smaller than the sample size at low ages. Consequently, our estimates reflect short- to mid-term manufacturer differences in failure rates.

B. Data Preparation

We infer the HDD installation date from the first day that a unique serial number (i.e., an individual HDD) appears in the dataset. For the drives present on the very first day in the data, we do not observe the installation date, since these drives were likely in use before the first data snapshot was made. Hence, these drives are “left-censored”, i.e., the exact starting date is not observed. For this study, we exclude the drives with unknown installation date, as it is important to control for drive age (calculated from installation date) in the regressions. Only 21,195 drives have an unknown installation date (less than 5% of drives in the dataset).

There is a small number of SSDs in the dataset (often used as boot drives for the servers), which we exclude, as we are interested in differences of reliability within the same technology (HDDs). After excluding SSDs from the sample, there remained 4 Samsung HDDs in the dataset – not enough for a meaningful analysis, so we excluded these as well.

The drive capacity is incorrectly recorded for a few HDDs in the data. We correct these, because we use capacity (rounded up to terrabytes) as control variable in the regressions. We correct the errors by identifying all models for which HDDs do not all have the same capacity, and then manually correct the few deviations.

A large share of drives is removed by Backblaze from the dataset without failure. In personal communication with Backblaze, we learned this is because older drives of small capacity are replaced by newer ones with larger capacity to save space and upgrade to less power consumption per capacity, which is also documented online [26]. Overall, 146,943 drives—about 31% of the sample—are removed without failure. We retain these in the sample, treating them as “right-censored”, i.e., not knowing their failure date, but knowing they survived at least until the removal date. This is statistically equivalent to HDDs surviving until the most recent daily snapshot in the data, where we also do not observe the (future) failure date.

Overall, we observe for every unique HDD (serial number) whether it failed or not, its survival duration in months, as well as manufacturer, capacity, form-factor (inferred from model number), installation year, and daily temperatures. The sample consists of 443,156 HDDs, of which 442,998 have a positive duration.⁴ of these, We use up to 442,992 drives in the analysis, since the remaining 6 do not report temperatures.

C. Manufacturer representation over time

Figure 1 plots the number of installed HDDs by year and manufacturer. The figure shows that manufacturer shares are not evenly distributed across the years. Instead, HGST was the most common manufacturer for newly installed drives in 2014, whereas Seagate was the most common manufacturer in 2013 and from 2015-2020. Toshiba was the manufacturer with the most newly installed drives in 2021, 2023 and 2025. WD was responsible for the most drives installed in 2024. Overall, WD and Toshiba drives tend to be younger, as they are overrepresented in installations in more recent years, whereas HGST and Seagate drives tend to be older.

IV. METHODS

We are using duration regression models, also called survival regression models, for the analysis. As discussed in the related literature section, merely comparing failure rates between manufactures ignores that the drives have different ages (and different capacities), potentially leading to inaccurate conclusions. Hence, we use a regression approach to control for these variables, thus effectively comparing drives of different manufacturers, holding these control variables constant. Moreover, duration models are designed for this kind of application where outcomes for many units are right-censored, i.e., the full duration is not observed. In these cases, we know the drive lasted at least as long as what we observe, but not when exactly it fails, and a model for statistical inference needs to account for that. This rules out linear ordinary least squares regression for this application.

Duration models are the appropriate tool for this analysis, because they ensure a “like-for-like” comparison in a sample where drives are not all the same age or capacity. A simple comparison of failure rates can be misleading: For instance, a manufacturer might appear more reliable simply because its drives are newer and have not yet worn out. Duration models correct for this by calculating the ‘hazard’—the immediate risk of failure—at every point in a drive’s life. This approach allows the model to use the information from every drive in the sample, no matter how old, including those that are still functioning and where the failure date is not yet known.

The Cox proportional hazard regression model is possibly the most popular duration model. It models the hazard or failure probability within a short time period around time t , conditional on surviving until time t and conditional on a

⁴The difference are HDDs that were installed on the last day of the sample, hence they have not (yet) survived a single day. The duration regression models we are using only make use of observations with positive duration.

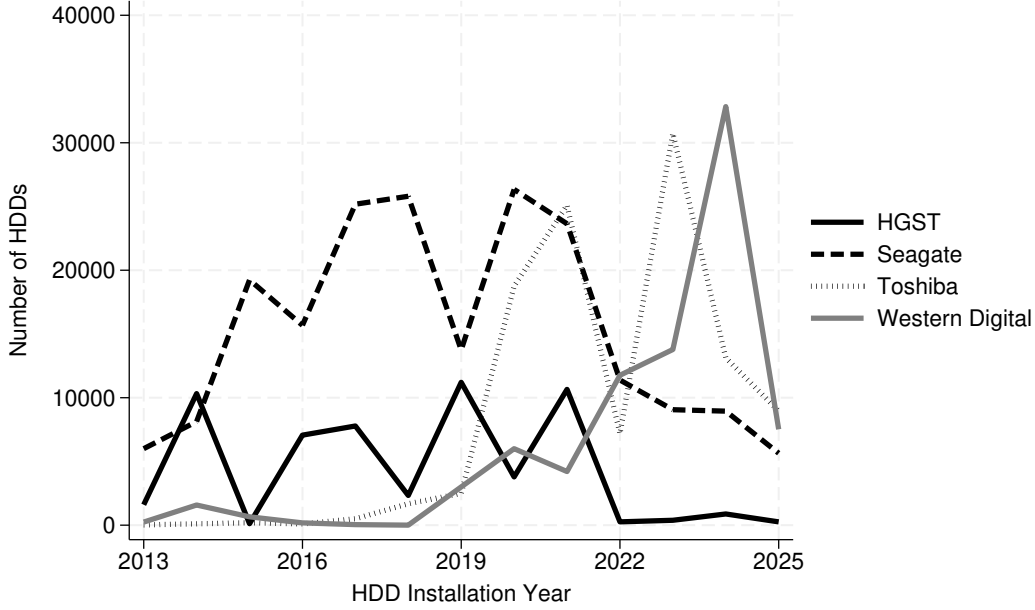


Fig. 1. Number of HDDs installed by manufacturer and year

vector of explanatory variables \mathbf{x} and a vector of estimated coefficients β , as

$$\lambda(t|\mathbf{x}, \beta) = \lambda_0(t) \exp(\mathbf{x}'\beta), \quad (1)$$

where $\lambda_0(t)$ is the baseline hazard function dependent on t but independent of \mathbf{x} . The strength of the Cox model is that the function $\lambda_0(t)$ does not have to be specified for estimation of β , hence there is no risk of misspecification of the hazard distribution. Consequently, the Cox semi-parametric model is lighter on distributional assumptions than a parametric model, its estimates are determined more by the data, and for this reason we choose it as our main model.

The coefficients of interest in this study are dummy variables denoting a specific manufacturer. Let x_1, x_2, x_3 be dummy variables equal to 1 if and only if the HDD is made by HGST, Toshiba, or WD, respectively, and 0 otherwise. Let $\beta_1, \beta_2, \beta_3$ be the corresponding coefficients. Seagate serves as the baseline category, with its drives being represented as $x_1 = x_2 = x_3 = 0$. Then the exponential of a manufacturer coefficient β_i represents the hazard ratio (HR) of that manufacturer and of the base category Seagate, i.e., the ratio of the hazard rate of the manufacturer relative to the Seagate hazard, at the same time and holding all other explanatory variables constant:

$$\text{HR} = \frac{\lambda_0(t) \exp(\beta_1 1 + \beta_2 x_2 + \dots)}{\lambda_0(t) \exp(\beta_1 0 + \beta_2 x_2 + \dots)} = \exp(\beta_1). \quad (2)$$

Consequently, estimation and interpretation of the coefficients in the Cox model do not require knowledge of the baseline hazard $\lambda_0(t)$. Throughout, we estimate robust standard errors for the statistical significance tests.

In addition to the semiparametric Cox model, we estimate the flexible parametric Weibull regression model, where the hazard is assumed to be

$$\lambda(t|\mathbf{x}, \beta) = \alpha t^{\alpha-1} \exp(\mathbf{x}'\beta), \quad (3)$$

where $\alpha \in \mathbb{R}, \beta$ are estimated. If $\alpha > 1$, then failure rates increase with survival time. If $\alpha < 1$, the failure rates decrease with time, and $\alpha = 1$ captures no duration dependence. An exponential regression model is a special case of the Weibull model with $\alpha = 1$, hence we do not show exponential regression estimates separately.⁵

We estimate these two different models to demonstrate that our findings do not depend on the model chosen or the functional form assumed. In each of these models, in addition to manufacturer dummy variables, we use as control variables:

- Installation year: one dummy variable for every year (2013-2025), ensuring drives of the same age/installation year are compared,
- Capacity: rounding capacity to the next highest terabyte, one dummy variable for every terabyte (1TB-24TB), ensuring that drives of the same capacity are compared,
- Form-factor: one dummy variable for 2.5 inch drives, ensuring that drives of the same size are compared.
- Average HDD temperature: Based on SMART attribute 194, which measures the temperature inside the HDD case in degrees Celsius, we calculate the average daily HDD temperature. This average temperature (centered

⁵As shown in Table I of the appendix, an exponential duration regression model and Gompertz duration regression model give us almost identical estimates of manufacturer failure rate differences. Thus, several regression models with varying distributional assumptions yield almost identical results, showing robustness in terms of model choice.

around 20°C) enters as a linear control variable, ensuring that drives with similar cooling solutions are compared.

Consequently, we specify for the regressions:

$$\begin{aligned} \mathbf{x}'\boldsymbol{\beta} = & \beta_1\text{HGST} + \beta_2\text{Toshiba} + \beta_3\text{WD} + \beta_4\text{Install-2014} \\ & + \dots + \beta_{15}\text{Install-2025} + \beta_{16}\text{Capacity-2TB} \\ & + \dots + \beta_{29}\text{Capacity-24TB} + \beta_{30}\text{Form-Factor-2.5} \\ & + \beta_{31}\text{Average-HDD-Temperature.} \end{aligned} \quad (4)$$

Therefore, the baseline category is a Seagate HDD installed in 2013, with 1TB capacity, form-factor 3.5 inch and twenty degree Celsius average daily drive temperature. The hazard ratios with respect to this base category are represented by coefficient vector $\boldsymbol{\beta}$, as described in (2).

V. RESULTS

A. Manufacturer differences in failure rates

First, we plot a non-parametric estimate of the hazard (failure) rate by manufacturer and drive age in Figure 2. Since the number of HDDs in the sample at higher ages is small, the graph is only plotted for ages where the manufacturer had at least 1,000 HDDs at risk, giving us more reliable estimates of the failure rate. As the figure shows, the length of the lines differs drastically by manufacturer: WD HDDs were added at scale only later by Backblaze, so there are fewer than 1,000 WD HDDs at age 60 months. At the same time, HGST drops down to below 1,000 HDDs at risk only at about 115 months of age.

In the range where we can compare the failure rate across all manufacturers—from months 0 to about 60—the failure rates are relatively stable across ages and the ranking between manufacturers is also quite stable. According to Figure 2, HGST HDDs have the lowest failure rates at most ages, closely followed by WD HDDs, which do a bit better at over 40 months of age. Future data and hence longer HDDs durations will be needed to see whether WD continues to have the lowest failure rates at higher ages. Seagate HDD failure rates start out as the highest, but are overtaken by Toshiba’s after about 50 months. Toshiba HDD monthly failure rates more than quadruple at age 60 months compared to early ages, a spike that no other manufacturer shows at any age in the sample. Seagate HDD failure rates are quite stable over time, whereas HGST HDD failure rates are a bit more volatile, but nevertheless largely maintain the lowest failure rates.

Regarding the level of HDD failure rates, for up to 50 months of HDD age, no manufacturer’s failure rate exceeds 2 out of 1,000 HDDs within a month or (annualized) 2.4% within a year.⁶ Between 50 and 100 months, Toshiba’s failure rates jump to between 5 and 6 out of 1,000 HDDs within a month, while the failure rates of the other manufacturers tend to stay at 2.5 out of 1,000 HDDs per month or below.

Second, we get to the main statistical analysis in Table I. This table displays results of the two regression models

⁶The values along the lines in Figure 2 represent a weighted average of the monthly failure rate, including up to 5 months before and 5 months after, with more weight given to closer months. For details, see the caption of Figure 2.

TABLE I
HDD MANUFACTURER DIFFERENCES IN FAILURE RATES

	(1) Cox HR	(2) Weibull HR
HGST	0.395*** (0.009)	0.411*** (0.009)
Toshiba	1.070** (0.033)	1.073** (0.033)
WD	0.514*** (0.021)	0.520*** (0.021)
Average HDD Temperature	1.021*** (0.002)	1.021*** (0.002)
Control Installation Year	Yes	Yes
Control Capacity	Yes	Yes
Control Form-Factor	Yes	Yes
P-value of test HGST = WD	0.000	0.000
Observations	442,992	442,992

Note: Failure rate ratios/hazard ratios (HR) of various manufacturers relative to manufacturer Seagate, holding constant drive age (installation year), capacity, form factor, average drive temperature. A smaller HR is better. Robust standard errors are displayed below the HRs. ***Significantly different from HR = 1 at the 1% level; **significant at the 5% level; *significant at the 10% level.

described in the methods section. The numbers for HGST, Toshiba, and WD represent the estimated ratio of failure rates (i.e., hazard ratios, HR) of these manufacturers with Seagate, while holding constant the installation year (drive age), the drive capacity, the drive form factor, and average drive temperature. See equation (2) and the explanation above for the precise interpretation. Thus, Table I compares the failure rates across manufacturers for comparable drives.

The estimates of both the Cox regression model and the Weibull regression model are on the same scale (i.e., both are hazard ratios in an exponential specification), and the estimated values are very similar (i.e., within 2 percentage points), demonstrating robustness of the results with respect to the regression model chosen. Hence, in the following we won’t discuss the estimates of the two models in Table I separately.

Table I shows that HGST has the lowest failure rate at about 39%-41% of Seagate’s failure rate. This difference in failure rates is statistically significant at any conventional confidence level ($p < .001$). Thus, HGST has the lowest HDD failure rates among the four manufacturers. Next, WD has the second lowest failure rate, about 51%-52% of Seagate’s failure rate. Again, this difference to Seagate is highly significant. Moreover, the significance test at the bottom of Table I shows that the hazard ratios (and hence failure rates) for HGST and WD are significantly different ($p < .001$). Hence, HGST’s failure rate is significantly lower than WD’s.

Toshiba’s hazard ratio with Seagate is about 1.07, so its failure rates are estimated to be about 7% larger than Seagate’s. The difference between these two manufacturers is significantly different at the 5% confidence level, not as clear-cut as the differences between other manufacturer pairs.

Hence, in this large sample of HDDs, with ages mostly under 10 years, we see a clear ranking: HGST has the lowest




Fig. 2. Hazard rate by manufacturer and drive age: Number of HDDs failed this month / number of HDDs at risk and not censored this month; discrete by month (non-parametric estimate using Epanechnikov kernel with bandwidth 5 months, making the graph a moving average). The graph only displays durations where the manufacturer has at least 1000 HDDs at risk.

failure rate, followed by WD, both of which have about half or less of Seagate’s failure rate. After this, it is close between Seagate and Toshiba, with Toshiba having slightly higher failure rates. The maximum difference between HGST and Toshiba implies more than a doubling of the drive failure rate. This suggests that the time and money spent on replacing faulty drives is more than doubled when comparing these two manufacturers. The results from the regressions in Table I are in line with the non-parametric values in Figure 2.

B. The effect of drive temperatures on failure rates

According to Table I, average drive temperatures have a hazard ratio of 1.021; a statistically significant effect on drive failure rates. The estimated hazard ratio of 1.021 indicates that an additional degree in the average daily temperature increases failure rates by 2.1% (not percentage points). Due to compounding, an additional 10 degrees of mean temperature is predicted to increase the failure rate by $1.021^{10} = 1.231 \approx 23.1\%$.⁷ Thus, suppose a drive at 40°C average temperature has a 2% failure rate. Then according to our estimates, cooling the drive by 10°C is predicted to reduce the failure rate to $2\% \times 1/1.231 \approx 1.625\%$.

C. Does accounting for position in the data-center change the results?

Clearly, conditions such as the quality of the maintenance staff, or the load on the drives might affect HDD failure rates.

⁷The inner 90% of the average temperatures in our data range from 22.5°C to just above 40°C, so a 10°C change is comfortably within the observed sample temperature ranges.

All of these factors depend on the position and assignment of a drive in the data-center: A server dedicated to a heavy user might see higher failure rates than another server for an inactive client. Backblaze has added position variables to their drive data only from quarter 2 and 3 of 2023. Beforehand, it is unknown where a drive was installed.⁸

In this section, we re-estimate the Cox model manufacturer differences while accounting for the cluster that an HDD is part of. In Backblaze’s terminology [26], a cluster is a collection of vaults sharing common resources (such as networking equipment), and a vault is a collection of 20 servers. Each server uses dozens of HDDs. Overall, there are 7 different clusters in the dataset. If some clusters see heavier loads, have higher temperatures or less capable maintenance staff, then by controlling for clusters in the regression, we can remove and disentangle these effects from manufacturer differences.⁹ The downside of controlling for drive position is that our sample size is reduced by about 90,000 drives, because the cluster variable is not available for drives that dropped out of the sample before this variable was reported.

Table II shows the estimates. In column 1, the Cox regression includes dummy variables for each of the clusters as control variables. The Cox regression in column 2 uses the

⁸In order to include as many drives as possible in this analysis, we use the later recorded position variable values even if a drive was installed at an earlier time when these variables were not yet reported. Thus, in this analysis we only lose drives that dropped out of the sample before the position variable was reported in Q2/3 of 2023.

⁹This is important if some clusters have proportionally more drives of a certain manufacturer: If a cluster sees heavier loads and has (say) more Seagate drives, then it would disadvantage Seagate drives.

TABLE II

COMPARING HDD MANUFACTURER DIFFERENCES WITH AND WITHOUT DRIVE POSITION CONTROLS (SMALLER SAMPLE)

	(1) Cox HR	(2) Cox HR Without
HGST	0.814*** (0.033)	0.833*** (0.033)
Toshiba	0.977 (0.048)	1.029 (0.049)
WD	0.442*** (0.029)	0.471*** (0.030)
Cluster 20	1.287*** (0.114)	
Cluster 31	1.229*** (0.064)	
Cluster 40	0.736*** (0.046)	
Cluster 50	1.154** (0.073)	
Cluster 52	0.511** (0.144)	
Cluster 60	0.489* (0.203)	
Control Installation Year	Yes	Yes
Control Capacity	Yes	Yes
Control Form-Factor	Yes	Yes
Control Avg HDD Temperature	Yes	Yes
P-value of test HGST = WD	0.000	0.000
Observations	355,429	355,429

Note: Failure rate ratios/hazard ratios (HR) of various manufacturers relative to manufacturer Seagate, holding constant drive age (installation year), capacity, form factor, average drive temperature, and drive position/cluster (only in column 1, cluster 0 is the omitted category). A smaller HR is better. Robust standard errors are displayed below the HRs. ***Significantly different from HR = 1 at the 1% level; **significant at the 5% level; *significant at the 10% level.

same observations as column 1, but omits the cluster variable controls. Comparing the manufacturer hazard ratio estimates from column 1 and 2 in Table II, there is little difference, and certainly nothing that changes the ranking among manufacturers. Hence, in our case, HDD position/assignment to cluster does not change the manufacturer HDD reliability estimates a lot. However, the estimates of manufacturer differences in Table II (column 2) are somewhat different from Table I, which is solely due to the change in sample: Table II uses about 355,000 HDDs, whereas Table I uses the full sample of 442,000 HDDs. Consequently, the conclusions from Table I are the more credible ones: While not controlling for HDD position, we have shown this variable does not change the estimates noticeably, and Table I uses the larger sample.

D. Drive failure rates by age, capacity, and position

Beyond manufacturer effects, our dataset allows to make some inferences regarding other factors' influence on HDD reliability and failure rates. As seen in Figure 2, failure rates clearly can change by drive age, and in addition, these age-specific patterns appear to differ by manufacturer. While Toshiba HDD failure rates change a lot by drive age, Seagate failure rates are more constant. Taking Toshiba as an example,

TABLE III

THE EFFECT OF CAPACITY ON FAILURE RATES

	(1) Cox HR	(2) Weibull HR
HGST	0.409*** (0.009)	0.418*** (0.009)
Toshiba	0.863*** (0.018)	0.875*** (0.018)
WD	0.595*** (0.021)	0.602*** (0.021)
Capacity in TB (linear)	0.966*** (0.005)	0.961*** (0.005)
Control Installation Year	Yes	Yes
Control Form-Factor	Yes	Yes
Control Avg HDD Temperature	Yes	Yes
P-value of test HGST = WD	0.000	0.000
Observations	442,992	442,992

Note: Failure rate ratios/hazard ratios (HR) of various manufacturers relative to manufacturer Seagate, holding constant drive age (installation year), capacity (with a linear term), form factor, average drive temperature. A smaller HR is better. Robust standard errors are displayed below the HRs. ***Significantly different from HR = 1 at the 1% level; **significant at the 5% level; *significant at the 10% level.

failure rates can easily change by a factor of 5 over the course of the HDD lifetime. Hence, age is an important determinant of HDD failure rates, though it is not a simple linear or monotone effect, especially when considering HGST in Figure 2.

In addition, the effect of HDD position in the data-center matters. While accounting for HDD position does not change estimates of HDD manufacturer differences (see Table II), there are still large differences between clusters in terms of hazard rates. Specifically, the cluster-specific estimates of the regression in column 1, Table II suggest that cluster 20 with the highest failure rates almost has $1.287/0.489 \approx 2.63$ times the failure rates of cluster 60 with the lowest failure rates – for HDDs with the same manufacturer, capacity, age, form-factor, and temperature. This difference is even larger if we do not hold temperature constant (not shown in the table). So loads and other location-specific factors matter a lot in practice.

Drive capacity not only represents more storage volume, but also technical progress and updated manufacturing processes, since higher capacity drives are newer. To see whether higher capacity drives are more reliable, we re-estimate the Cox model from Table I, but include capacity as a linear term rather than a series of dummy variables as that table does. The estimates are displayed in Table III. As before, this regression also holds age, manufacturer, and form-factor constant. The hazard ratio of capacity in terrabytes is estimated to be 0.966, and the difference to HR = 1 is highly statistically significant ($p < .001$). Hence, in this linear approximation, a drive with 1TB higher capacity is predicted to have a 3.4% (not percentage points¹⁰) lower failure rate. So, based on this data, modern drives are not only becoming larger in terms of storage volumes, but also more reliable.

¹⁰Hence, suppose a 1TB drive has an annual failure rate of 2%. Then a 2TB drive would be predicted to have a $2\% \times 0.966$ annual failure rate.

There are at least three classes of possible explanations for this estimated improvement in reliability for higher capacity (i.e., newer) drives. First, improvements in the manufacturing process itself. Manufacturing processes improve over time as manufacturers gain experience and iterate. Thus, modern high-capacity drives benefit from improved servo technology, precision head fabrication techniques (e.g., WD switched from a dry hole process to a Damascene process, arguing it improved quality [27]), and quality control methodologies. Second, improvements in the HDD technology itself are likely a cause for the improvement. For example, during the sample window enterprise drives increasingly switched from air to helium-filled enclosures. Due to its lower density, helium reduces aerodynamic drag on the HDD platters, thereby decreasing mechanical stress, heat generation, and spindle bearing wear compared to conventional air-filled drives [28], [29]. Other changes in HDD technology helped increase capacity, but are less likely to explain increased reliability (e.g., shingled magnetic recording uses overlapping data tracks which increase complexity). Third, a theoretical possibility is that the requirements of HDD customers (increasingly data-centers, as workstations and retail consumers switch to SSDs) have shifted towards more reliability and a higher willingness to pay for it over time. Given the changed demand, the manufacturers may have switched to higher quality parts and larger investments in manufacturing.

Finally, whether the HDD form-factor—3.5 inch or 2.5 inch—matters for failure rates cannot be conclusively answered with the Backblaze data, because only about 0.6% of the HDDs have a small form factor. It would require a different kind of setting and dataset to answer this question.

E. Failure rates across drive generations and manufacturers

In this section, we ask whether the gradient at which drives become more reliable over time/over the generations is different by manufacturer. Unlike in the previous section, we focus here not on capacity as a proxy for drive generation, but directly focus on installation year. While it is possible that older generation drives are installed in a given year, on the whole we expect drives to get more modern with advanced installation years. In these regressions, we will not control separately for drive capacity, as drive capacity and drive generation are highly correlated. Put differently, we want to measure the effect of newer drives on reliability, but if we were to hold capacity constant, we would not pick up the full effect.

Thus, in Table IV, we estimate regressions as in Table I, except that (1) we do not control for drive capacity, (2) we let the effect of installation year be linear rather than nonparametric via dummy variables, and (3) we allow the effect of installation year be different by manufacturer, which we achieve by interacting the installation year linear variable with the manufacturer dummy variables.

For the purposes of this analysis, we re-centered the Installation Year variable to 2013, the first year of the sample. To interpret the coefficients in Table IV, note that the

TABLE IV
DIFFERENCES IN MANUFACTURER GENERATIONAL TRENDS

	(1) Cox HR	(2) Weibull HR
HGST	0.091*** (0.005)	0.097*** (0.005)
Toshiba	1.450*** (0.117)	1.643*** (0.127)
WD	1.076 (0.085)	1.033 (0.076)
Installation Year	0.905*** (0.004)	0.884*** (0.003)
HGST×Installation Year	1.420*** (0.012)	1.412*** (0.012)
Toshiba×Installation Year	0.944*** (0.010)	0.927*** (0.010)
WD×Installation Year	0.917*** (0.010)	0.916*** (0.009)
Control Form-Factor	Yes	Yes
Control Avg HDD Temperature	Yes	Yes
P-value Toshiba trend = WD trend	0.039	0.359
Observations	442,992	442,992

Note: Failure rate ratios/hazard ratios (HR) holding constant form factor. Installation Year is centered on 2013, the first sample year. A smaller HR is better. Robust standard errors are displayed below the HRs. ***Significantly different from HR = 1 at the 1% level; **significant at the 5% level; *significant at the 10% level.

manufacturer dummy variables are the hazard ratios of that manufacturer relative to Seagate *at installation year 2013*, and relative to the assumed linear time trend. As we are interested in changes over time rather than manufacturer differences in a specific year, the coefficients of the manufacturer dummies are not of much interest in this table. The estimated coefficient on Installation Year is 0.905 or 0.884, depending on the regression model. Thus, Seagate drives get more reliable with each installation year, approximately by 10%-11%, which is quite large in magnitude. This installation year effect does not capture the same thing as the capacity effect of the previous section (which instead measured the reliability effect of an additional 1TB of capacity of a drive installed in the same year), so these estimates are not in conflict.

Next, we can consider the interaction terms, which capture a difference of the time gradient relative to Seagate. Both the Toshiba and WD interaction terms are significantly less than 1, indicating that the improvement over time (i.e., by installation year) is better than for Seagate. To obtain the gradient for WD, we need to multiply the interaction term coefficient with the Installation Year coefficient. Thus, WD drive reliability improves by an estimated $1 - 0.905 \times 0.917 = 0.17$, i.e., about 17% (not percentage points) per installation year, in this linear approximation. And it being a linear approximation, we should be hesitant extrapolating the effect beyond the sample period. Toshiba's improvement by installation year are very similar to WD's (not significantly different at the 5% significance level in column 2), as its similar estimated interaction coefficients and the second-last line in the table indicate. Hence, Toshiba, too, sees large improvements in reliability by drive genera-

tion. Interestingly, HGST’s gradient over time is worse than Seagate’s, as the estimated coefficient for HGST’s interaction term is significantly larger than 1. We need to keep in mind, however, that HGST was mostly installed in the first few years of the sample and later stopped making new HDDs, so any later installed HDDs are effectively older generation ones. So the shorter sample for HGST combined with a possible bias—where later HGST drive installations had to be older generation drives—explains the worse improvement gradient over time. Consequently, we do not consider the comparison between HGST and the others “like for like” when it comes to improvements over time, though there is no issue with comparing reliability holding installation year constant, as we did in previous sections.

VI. WHEN DOES IT PAY TO BUY THE MORE RELIABLE HDD?

In this section, we briefly derive a relationship between the acceptable purchase price premium $P_{\text{premium}} = P_{\text{reliable}} - P_{\text{base}}$ of a more reliable drive over a base drive and the reduction in the annualized failure rate of the drive, $\Delta = \text{AFR}_{\text{base}} - \text{AFR}_{\text{reliable}}$. This relationship can then be used to answer whether buying the more reliable drive is justified.

We start with the Total Cost TC of a drive over its lifetime L (in years), which is the sum of its purchase price P and the expected cost due to failures over the lifetime:

$$TC = P + \text{AFR} \times C \times L,$$

where AFR is the annual failure rate and C is the cost of drive failure¹¹.

From a business perspective, reliability has the benefit of lower expected replacement costs and the possible downside of a higher purchase price. The total cost of the reliable drive is less than the total cost of the base drive if and only if $TC_{\text{reliable}} \leq TC_{\text{base}}$, and rearranging for the price premium:

$$P_{\text{premium}} \leq \Delta \times C \times L. \quad (5)$$

Clearly, with longer lifetimes L , reliability differences have a longer time to pay off, justifying larger initial price premia. And with lower replacement costs C , only smaller initial price premia are justified. Thus, inequality (5) gives us a simple tool to trade off procurement costs and better reliability.¹²

To use this expression in an example, suppose the lifetime of a drive (after which it will be replaced) is $L = 10$ years and the cost of replacement is $C = \$100$, then the AFR difference from Seagate (around 2%, Figure 2) to HGST is $\Delta = 0.02(1 - 0.41) = 0.0118$, using our regression estimate of HGST having a failure rate that is 41% of Seagate’s. The maximum acceptable price premium for an HGST drive over Seagate then would be $\$100 \times 0.0118 \times 10 = \11.8 . And while this number may seem small, suppose the drives actually

have the same price, so choosing the less reliable one for the drive fleet costs over \$3m more in terms of total cost in the case of Backblaze with 300k active drives just due to more replacements.

VII. CONCLUSION

We provide estimates of manufacturer differences in HDD reliability. When comparing drives of the same age, capacity, and form-factor, HGST HDDs are the most reliable, whereas Toshiba is the least reliable, closely followed by Seagate. The HGST failure rate is less than half of Seagate’s, whereas WD’s failure rate is about half of Seagate’s. We can rule out that these manufacturer differences are due to different positions (as measured by cluster) in the data-center, due to differing temperatures, due to different ages, different capacities, or different form-factors, as we account for these variables in the statistical analysis.

Our findings should be interpreted and used with care. First, we have no data that consistently tracks drive read/write workloads across drives of different manufacturers, because different manufacturers use different SMART attributes and implement them differently, thus making workload comparisons across manufacturers difficult. Hence, we cannot directly confirm that drives of all manufacturers of the same age have similar workloads. Our cluster-location analysis indicates that accounting for drive position (which also determines tasks and workload) does not meaningfully change our manufacturer differences estimates, which is reassuring. But absent more direct workload data, our analysis is based on the informed assumption that workloads do not drastically differ between manufacturers. Second, Backblaze, like other data-centers, largely uses enterprise drives. Hence, our conclusions apply to enterprise drives, and manufacturer differences might be different among consumer drives. Additionally, the dataset contains HDD models from over 10 years ago, and manufacturer differences among the newest models (which did not yet have time to prove themselves) might be different. That said, we see large manufacturer differences in a sample spanning several model generations, so the manufacturer differences we obtained cannot be explained by single models or single model generations.

Our results are of interest to data-centers and cloud storage providers, where even small differences in failure rates translate into large costs due to their scale. In Backblaze’s case of over 300,000 active HDDs (which is still small compared to Amazon, Microsoft, Google, etc.), a doubling of the annual failure rate from about 1% (Toshiba, Figure 2) to 2% (Seagate) requires replacing an additional 3,000 HDDs per year. These failed HDDs may be under warranty, depending on the contract, so no additional hardware cost might apply. Yet the staff time required for replacement is not covered by warranty. Consequently, if salaries for data-center technicians increase faster than price differences in HDD models, then it may be increasingly beneficial to shift to more expensive but more durable HDDs.

¹¹Assuming warranty, this is the labor and logistics cost of replacement plus expected costs due to downtime.

¹²Clearly, the model behind (5) could still be further extended by allowing replacement costs to differ by manufacturer (relevant if not the entire lifetime is covered by warranty) or by introducing discounting of future costs.

ACKNOWLEDGMENT

We would like to thank Stephanie Doyle from Backblaze, who patiently answered our questions about the data and operations at Backblaze. We thank Backblaze for publicly sharing the data. We are grateful to Youngik Eom, Hyungmin Cho, and Panagiotis Kanellopoulos for their helpful comments and advice.

REFERENCES

- [1] K. V. Vishwanath and N. Nagappan, "Characterizing cloud computing hardware reliability," in *Proceedings of the 1st ACM symposium on Cloud computing*, 2010, pp. 193–204.
- [2] Market Reports World, "Hard disk drives market size, share, growth, and industry analysis," <https://www.marketreportsworld.com/market-reports/hard-disk-drives-market-14720196>, 2025.
- [3] Forbes, "C1q 2025 hdd industry update," <https://www.forbes.com/sites/tomcoughlin/2025/05/03/c1q-2025-hdd-industry-update/>, 2025.
- [4] —, "A very long engagement – Western Digital finally acquires Hitachi's HDD business," <https://www.forbes.com/sites/tomcoughlin/2012/03/09/a-very-long-engagement-western-digital-finally-acquires-hitachis-hdd-business/>, 2012.
- [5] Y. Wang, E. W. Ma, T. W. Chow, and K.-L. Tsui, "A two-step parametric method for failure prediction in hard disk drives," *IEEE Transactions on Industrial Informatics*, vol. 10, no. 1, pp. 419–430, 2013.
- [6] M. M. Botezatu, I. Giurgiu, J. Bogojeska, and D. Wiesmann, "Predicting disk replacement towards reliable data centers," in *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 2016, pp. 39–48.
- [7] C. A. Rincón, J.-F. Pâris, R. Vilalta, A. M. Cheng, and D. D. Long, "Disk failure prediction in heterogeneous environments," in *2017 International Symposium on Performance Evaluation of Computer and Telecommunication Systems (SPECTS)*. IEEE, 2017, pp. 1–7.
- [8] E. Barelli and E. Ottaviani, "Unsupervised anomaly detection for hard drives," in *PHM Society European Conference*, vol. 6, no. 1, 2021, pp. 7–7.
- [9] N. Aussel, S. Jaulin, G. Gandon, Y. Petetin, E. Fazli, and S. Chabridon, "Predictive models of hard drive failures based on operational data," in *2017 16th IEEE international conference on machine learning and applications (ICMLA)*. IEEE, 2017, pp. 619–625.
- [10] J. Xiao, Z. Xiong, S. Wu, Y. Yi, H. Jin, and K. Hu, "Disk failure prediction in data centers via online learning," in *Proceedings of the 47th International Conference on Parallel Processing*, 2018, pp. 1–10.
- [11] S. Lu, B. Luo, T. Patel, Y. Yao, D. Tiwari, and W. Shi, "Making disk failure predictions SMARTer!" in *18th USENIX conference on file and storage technologies (FAST 20)*, 2020, pp. 151–167.
- [12] J. Ahmed and R. C. Green II, "Cost aware lstm model for predicting hard disk drive failures based on extremely imbalanced smart sensors data," *Engineering Applications of Artificial Intelligence*, vol. 127, p. 107339, 2024.
- [13] G. Han, Y. Wang, Q. Hai, Z. Yang, and S. Peng, "Adversarial domain generalization network for new enabled hard disk drive failure prediction," *IEEE Transactions on Magnetics*, 2025.
- [14] R. Lu, E. Xu, Y. Zhang, F. Zhu, Z. Zhu, M. Wang, Z. Zhu, G. Xue, J. Shu, M. Li *et al.*, "Perseus: A Fail-Slow detection framework for cloud storage systems," in *21st USENIX Conference on File and Storage Technologies (FAST 23)*, 2023, pp. 49–64.
- [15] J. M. Wooldridge, *Econometric analysis of cross section and panel data*. MIT press, 2010.
- [16] E. Pinheiro, W.-D. Weber, and L. A. Barroso, "Failure trends in a large disk drive population," in *FAST '07: 5th USENIX Conference on File and Storage Technologies*, 2007.
- [17] Z. Miller, O. Medaiyese, M. Ravi, A. Beatty, and F. Lin, "Hard disk drive failure analysis and prediction: an industry view," in *2023 53rd Annual IEEE/IFIP International Conference on Dependable Systems and Networks-Supplemental Volume (DSN-S)*. IEEE, 2023, pp. 21–27.
- [18] B. Schroeder and G. A. Gibson, "Understanding disk failure rates: What does an mttf of 1,000,000 hours mean to you?" *ACM Transactions on Storage*, vol. 3, no. 3, pp. 8–31, 2007.
- [19] I. Narayanan, D. Wang, M. Jeon, B. Sharma, L. Caulfield, A. Sivasubramaniam, B. Cutler, J. Liu, B. Khessib, and K. Vaid, "Ssd failures in datacenters: What? when? and why?" in *Proceedings of the 9th ACM International on Systems and Storage Conference*, ser. SYSTOR '16. New York, NY, USA: Association for Computing Machinery, 2016. [Online]. Available: <https://doi.org/10.1145/2928275.2928278>
- [20] E. Xu, M. Zheng, F. Qin, Y. Xu, and J. Wu, "Lessons and actions: What we learned from 10k SSD-related storage system failures," in *2019 USENIX Annual Technical Conference (USENIX ATC 19)*, 2019, pp. 961–976.
- [21] S. Maneas, K. Mahdavian, T. Emami, and B. Schroeder, "A study of SSD reliability in large scale enterprise storage deployments," in *18th USENIX Conference on File and Storage Technologies (FAST 20)*, 2020, pp. 137–149.
- [22] S. Han, P. P. Lee, F. Xu, Y. Liu, C. He, and J. Liu, "An in-depth study of correlated failures in production SSD-based data centers," in *19th USENIX Conference on File and Storage Technologies (FAST 21)*, 2021, pp. 417–429.
- [23] Z.-S. Ye, M. Xie, and L.-C. Tang, "Reliability evaluation of hard disk drive failures based on counting processes," *Reliability Engineering & System Safety*, vol. 109, pp. 110–118, 2013.
- [24] S. S. Arslan and E. Zeydan, "On the distribution modeling of heavy-tailed disk failure lifetime in big data centers," *IEEE Transactions on Reliability*, vol. 70, no. 2, pp. 507–524, 2020.
- [25] Backblaze, "Hard drive data and stats," <https://www.backblaze.com/cloud-storage/resources/hard-drive-test-data>, 2025.
- [26] —, "How backblaze scales our storage cloud," <https://www.backblaze.com/blog/how-backblaze-scales-our-storage-cloud/>, 2024.
- [27] WD, "Next-generation technologies for a new decade of big data," <https://innovation.wdc.com/downloads/HDD-Capacity-Tech-Brief.pdf>, 2017.
- [28] J. Yang, C. P. Tan, and E. H. Ong, "Heat transfer enhancement by helium gas filled in a hard disk drive enclosure," in *ASME International Mechanical Engineering Congress and Exposition*, vol. 44281, 2010, pp. 155–159.
- [29] A. Aoyagi, B. Stipe, R. Wood, S. Campbell, and X. Che, "Helium sealed hard disk drive," *Journal of Magnetism and Magnetic Materials*, vol. 564, p. 170146, 2022.