SCIENTIFIC CULTURE
www.sci-cult.com

SCIENTIFIC CULTURE, Vol. 9, No. 3, (2023), pp. 126-161
Open Access. Online & Print

DOI: 10.5281/zenodo.24321251

OUT-OF-SAMPLE COMPARISON OF BUDGETED DISK REPLACEMENT BY RISK DENSITY AND AGE ON ONE OPERATOR'S 2019 DRIVE STATISTICS

Luis Eduardo Muñoz Guerrero[1*]

[1]PhD, Universidad Tecnológica de Pereira, Facultad de Ingenierías, Pereira, Risaralda, Colombia. ORCID https://orcid.org/0000-0002-9414-6187

Received: 28/07/2023
Accepted: 28/08/2023

Corresponding Author: Luis Eduardo Muñoz Guerrero

ABSTRACT

Archives with a fixed replacement budget must choose which drives to replace. We compare a greedy risk-density policy with an age-based policy out of sample, treating budgeted replacement as a 0/1 knapsack problem. Backblaze drive statistics for two quarters of 2019 were used. The first quarter, with 115,231 distinct drives, estimated failure rates per model; 106,218 drives remained in the analysis population. The second quarter, sealed until the plan was frozen, gave the outcome. The budget was 5% of capacity. The greedy policy captured 2.18 failures per 1000 terabytes (TB), against 1.21 for the age policy: a difference of 0.97 failures per 1000 TB (95% interval 0.43 to 1.53), positive under the rule fixed in advance. Post hoc, without the half-terabyte or the six-terabyte class the difference is inconclusive. On first-quarter risks, the greedy value matched the exact optimum at four of five budgets and was within 0.003% at the fifth. On one operator's data, out-of-sample selection by estimated risk per terabyte captured more failures per terabyte than age-based replacement at the primary budget; post hoc, the advantage depends on the mix of drive classes and was not detected against ranking by model rate.

KEYWORDS: Knapsack Problem; Approximation Algorithm; Dynamic Programming; Integer Optimization; Reliability Engineering; Storage Systems; Disk Reliability; Archival Storage.

Copyright: © 2023. This is an open-access article distributed under the terms of the Creative Commons Attribution License. (https://cre-ativecommons.org/licenses/by/4.0/).

1. INTRODUCTION

Digital archives depend on the durability of their storage media, and hard disks remain a common medium for online and nearline storage. Replacing a disk before it fails costs money, and the budget for replacement is finite. An operator who can replace only part of a fleet in each period must decide which disks to replace.

Schroeder and Gibson (2007) analyze field-gathered disk replacement data from large production systems, and they report that replacement rates grow with age. Pinheiro et al. (2007) note that there is relatively little published work on the failure patterns of disk drives, and they report that models based on SMART parameters alone are unlikely to be useful for predicting individual drive failures. Prediction from internally monitored attributes is a problem of detecting rare events (Murray et al., 2005), and SMART-based analysis has been used to predict the need for a disk replacement ten to fifteen days in advance (Botezatu et al., 2016). The selection problem that follows a prediction, namely which disks to replace when the budget covers only part of the fleet, is the subject of this paper.

Replacement under a budget is a 0/1 knapsack problem: each drive has a risk of failure over the horizon and a replacement cost in whole terabytes, and the budget limits the total cost. Linear programming methods have been applied to discrete-variable extremum problems (Dantzig, 1957), and the linear relaxation of the knapsack problem gives a bound on the optimum. We compare a greedy risk-density policy with an exact dynamic programming over groups of identical drives, with an age-based policy and with a random control.

The study asks how much of the optimum a greedy risk-density policy attains, and whether, out of sample, it captures more failures per terabyte than an age-based policy in the following quarter. Its contributions are an out-of-sample design whose decision rule was fixed before the outcome was read and whose controls were specified in the analysis plan, although the implementation of one control (V5) was corrected after the first reading (see Methods); a comparison on one operator's fleet, decomposed by drive model and capacity; and a reproducible pipeline that generates every number in this manuscript from versioned results. We state the NP-completeness of the decision version and the one-half bound of the greedy rule for the problem studied, and we prove both here. The question is predictive and descriptive, and the analysis makes no causal claim.

The data are the publicly released Backblaze drive statistics for the first two quarters of 2019 (Backblaze, 2019). The study uses no human participants and no personal data.

2. METHODS

We used two quarters of daily drive-level records that Backblaze published for 2019 (Backblaze, 2019). The first quarter, from 1 January to 31 March, has 9,577,046 drive-day records for 115,231 drives of 49 models on 90 days, and 428 failures. Each record gives the date, the serial number, the model, the capacity in bytes, a failure indicator and the raw value of the SMART power-on hours attribute. The first quarter defines the population and estimates the failure rates. The second quarter, from 1 April to 30 June, provides the outcome only. Both archives were checked against the SHA-256 digests that were recorded when they were downloaded, which are listed in the data manifest of the repository. The second quarter was kept sealed: its loader refused to read it until the analysis plan had been committed under a version-control tag, and it logged every access.

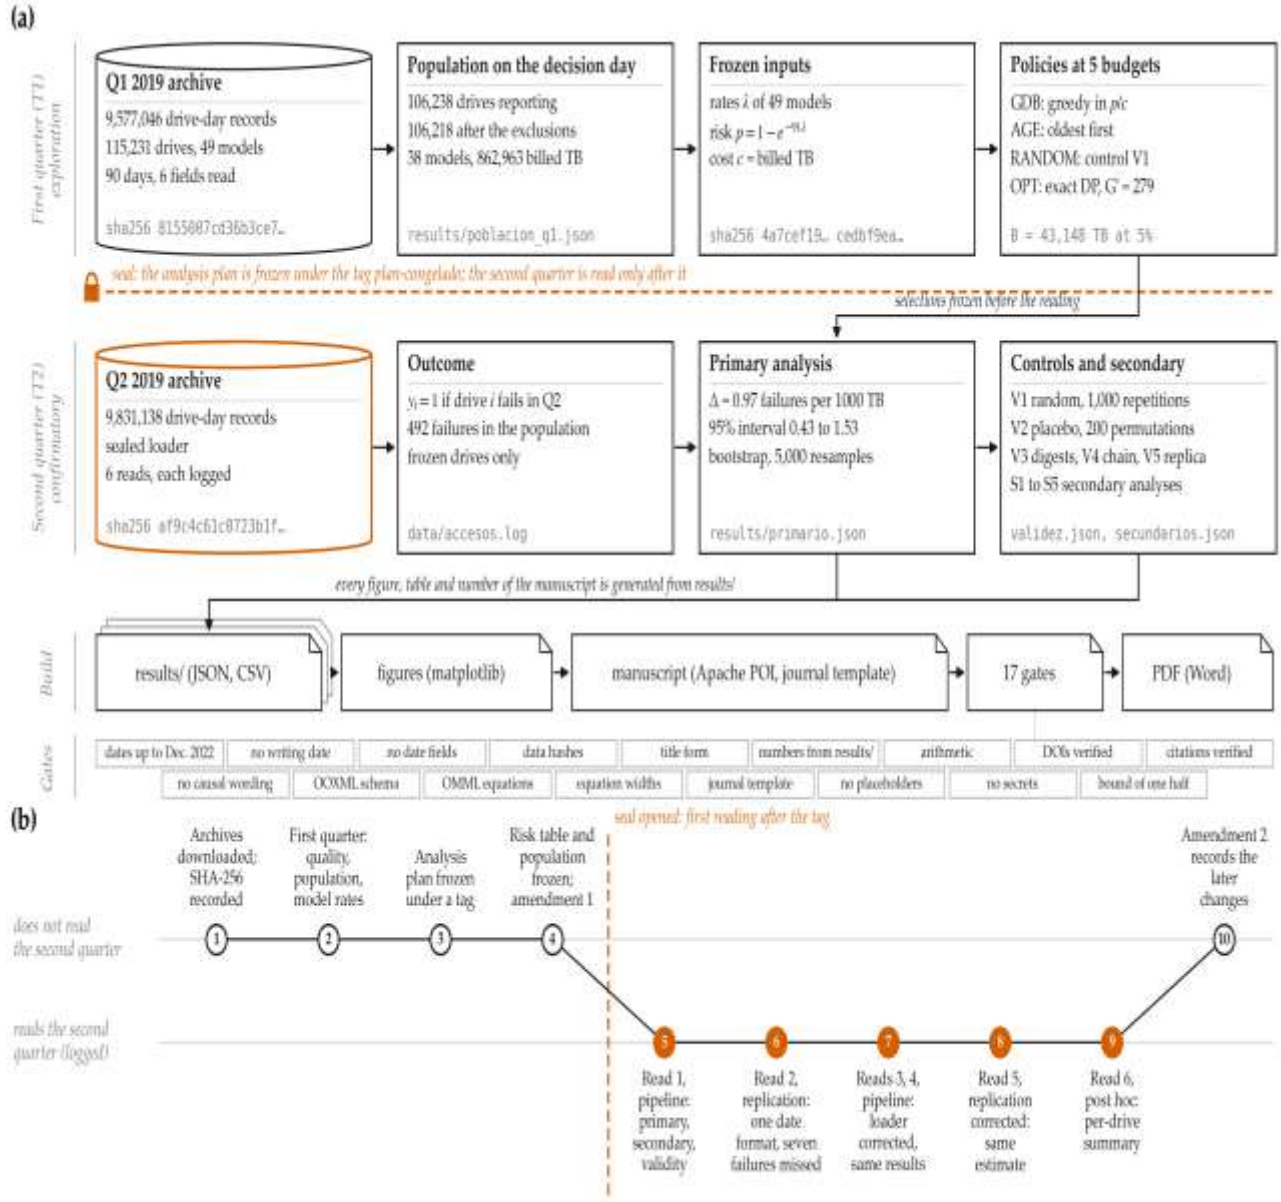


Figure 1. Pipeline and protocol. (a) Architecture, with the size of each input, the leading characters of the SHA-256 digests of the archives and of the frozen inputs, and, where applicable, the result file or the access log. The first quarter defines the population, the frozen risk table and the selections of the policies; the second quarter, sealed until the analysis plan was tagged (dashed line), gives the outcome. Every figure, table and number in this manuscript is generated from the versioned results. (b) The steps of the protocol in order, without dates, in two lanes: the steps that do not read the second quarter (open circles) and those that read it (filled circles). Each reading is logged, and the post hoc analyses use the per-drive summary derived in the last of them. Cylinders mark the two archives, the padlock the seal, and document symbols the generated files; the bottom strip names what each gate checks.

Figure 1a summarises the pipeline, and Figure 1b orders the steps of the protocol. The decision day is 31 March 2019. The population is the set of drives with a record that day, after excluding 20 drives that report no power-on hours. All of them belong to one model, DELLBOSS VD, and no drive on that day reports zero or negative capacity. This leaves 106,218 drives with 862,963 billed terabytes in total (physical capacity 862,084.6 TB), the sum of the whole-terabyte costs defined next. The replacement cost of drive i is its capacity in whole terabytes, ci = round(capacityi / 1012), so a drive of half a terabyte costs one TB. No drive costs zero: every model has one integer cost in the population, although 9 models report several raw capacity values in the first quarter. Age is the power-on hours of the drive on the decision day. The population includes two solid-state models, and neither policy replaced a drive of these models.

Table 1. Data, data quality and decision population (first quarter of 2019), and the records of the second quarter.

Quantity	Value
Drive-day records, first quarter	9,577,046
Distinct drives, first quarter	115,231
Models, first quarter	49
Days, first quarter	90
Failures recorded, first quarter	428
Models with no failure in the first quarter	30 of 49
Drives on the decision day (31 March)	106,238
Drives excluded (no power-on hours)	20
Analysis population	106,218
Models in the analysis population	38
Records without power-on hours, first quarter	1,534
Records with an invalid date or a repeated drive and day, first quarter	0 and 0
Population billed capacity (TB)	862,963
Records, second quarter (all inside the window of the plan)	9,831,138
Second-quarter records written year-month-day; month/day/year	9,720,605; 110,533
Population drives without any second-quarter record (post hoc count)	7
Population drives that failed in the second quarter (outcome)	492

Failure counts are drive-level events. The population excludes the drives without power-on hours, all of one model. No drive on the decision day reports zero capacity. The second quarter mixes two date forms, and the loader reads both (Methods).

Table I summarises the data and the population. The failure risk of drive i over a horizon of H = 91 days (the second quarter) is computed from the estimated first-quarter failure rate of its model m(i). Under a constant hazard per model, the risk is

$$p_i = 1 - \exp\left(-\hat{\lambda}_{m(i)}\, H\right)$$

Equation (1) gives the risk of drive i from the estimate of the rate λm of its model. Under the same model, the failures Fm of model m in the first quarter follow a Poisson law over its em drive-days, which count all of its records, including the 993 records with non-positive capacity, which carry no failures. The rate estimate and its exact 95% interval, Equation (2), are

$$F_m \sim \text{Poisson}(\lambda_m e_m), \qquad \hat{\lambda}_m = \frac{F_m}{e_m},$$

$$\lambda_m \in \frac{1}{2e_m}\left[\chi^2_{2F_m}(\beta/2),\ \chi^2_{2(F_m+1)}(1-\beta/2)\right]^{\text{I}}$$

n Equation (2), χ2d(γ) is the quantile of order γ of a chi-squared distribution with d degrees of freedom, β = 5% makes the interval a 95% interval, and the lower bound is zero when Fm = 0. With the failures of a selection as the count and its billed terabytes, in thousands, as the exposure, the same interval gives the exact Poisson intervals of the capture rates in the figures.

Table 2. Notation used in the formulation, the policies, the estimand, the controls and the post hoc analyses, with the field of the data that each symbol comes from.

Symbol	**Meaning**	**Domain**	**Data field**
i, k	Drive index; k is the first drive that does not fit	1 to n	serial_number (sorted)

n	Number of drives in the analysis population	positive integer	–
H	Horizon of the risk estimate, in days (primary value 91)	30, 91 or 182 days	–
λm	Constant daily failure rate of model m; with a hat, its estimate from the first quarter	non-negative real	model; failure (first quarter)
pi	Risk of drive i over the horizon H	0 to 1	–
ci	Replacement cost of drive i, in whole terabytes	positive integer	capacity_bytes
B	Budget, in whole terabytes	positive integer	–
f	Budget as a fraction of population capacity (primary value 0.05)	0 to 1	–
xi	Selection indicator of drive i	0 or 1	–
S	Replacement plan: the set of selected drives	subset of the drives	–
V(P), c(P)	Risk and cost of the greedy prefix P	non-negative real; positive integer	–
θ	Fraction of the first drive that does not fit, in the linear relaxation	0 to 1	–
LP	Value of the linear relaxation over the drives that fit	non-negative real	–
OPT, GDB, AGE, RANDOM	Exact, greedy risk-density, age-based and random policies, and the plan that each selects; in a ratio or an inequality, the expected failures Σ pi of that plan	plan: set of drives; value: non-negative real	smart_9_raw (power-on hours, for AGE)
D[b]	Best risk with capacity b in the dynamic programming	non-negative real	–
w, v	Cost and risk of an item in the dynamic programming	positive integer; non-negative real	–
G′	Number of 0/1 items after binary splitting of the group multiplicities	positive integer	model; capacity_bytes

yi	Outcome: 1 if drive i records a failure in the second quarter, 0 otherwise	0 or 1	failure (second quarter)
F(S), TB(S)	Failures and terabytes of plan S in the second quarter	non-negative integer; non-negative real	failure; capacity_bytes
Δ	Primary estimand: difference of the capture rates of GDB and AGE	real	–
Fm, em	Failures and drive-days of model m in the first quarter	non-negative integers	failure; date (first quarter)
χ2d(γ), Qγ	Quantile of order γ of a chi-squared distribution with d degrees of freedom; empirical quantile of order γ	non-negative real; real	–
β, ω, ε	One minus the level of the intervals, β = 5%; fraction of the AGE rate in the margin, ω = 20%; threshold of V2, ε = 5%	0 to 1	–
G, g, ng	Number of groups of identical drives; group index; drives in group g	positive integers	model; capacity_bytes

Values are those of the primary analysis unless stated otherwise. Data fields are the columns of the Backblaze records; a dash marks a symbol defined by the analysis.

Table II lists the symbols used in the formulation. A replacement plan is a set S of drives whose total cost does not exceed a budget B. The budget is a fraction f of the population capacity, B = floor(f • Σ ci). The primary analysis uses f = 0.05, which is 43,148 TB. The primary curve uses five budgets, 1%, 2%, 5%, 10%, 20% of population capacity, each rounded down in the same way; the unit-cost analysis (S5) sets every drive's cost to one, so its budget is 5% of the drive count, rounded down. The objective is the expected number of failures among the replaced drives over the horizon, which is the sum of their risks. This gives the 0/1 knapsack problem

$$\max \sum_{i=1}^{n} p_i\, x_i \quad \text{s.t.} \quad \sum_{i=1}^{n} c_i\, x_i \le B, \quad x_i \in \{0,1\} \tag{3}$$

Equation (3) states the problem. Its decision version is studied in Proposition 1.

We compare four policies. OPT is an exact dynamic programming over the G groups of identical drives (same model and capacity). Group g has ng drives of risk pg and cost cg, and OPT chooses how many of them, zg, to replace; its integer program, Equation (4), is

$$\max \sum_{g=1}^{G} p_g\, z_g \quad \text{s.t.} \quad \sum_{g=1}^{G} c_g\, z_g \le B, \quad z_g \in \{0, \dots, n_g\} \tag{4}$$

Equation (4) is equivalent to Equation (3), because any zg drives of a group have the same risk and cost. The dynamic programming works on 0/1 items obtained by binary splitting of each multiplicity ng into pieces s, as Equation (5) states:

$$n_g = 1 + 2 + \cdots + 2^{t_g - 1} + r_g, \qquad t_g = \lfloor \log_2(n_g + 1) \rfloor, \qquad 0 \le r_g < 2^{t_g}, \tag{5}$$

$$(w, v) = (s\, c_g,\ s\, p_g) \text{ for each piece } s, \qquad G' = \sum_{g=1}^{G} \lceil \log_2(n_g + 1) \rceil$$

Every number of drives from 0 to ng is a sum of distinct pieces of Equation (5), so the items represent every choice of zg. GDB considers the drives that fit the budget on their own, NB = {i : ci ≤ B}, in the order ɪɪ of decreasing risk density p/c, with ties broken by the drive index, which orders the serial numbers lexicographically as the first amendment specifies. It adds them in that order until the first drive k that does not fit, and it returns the better of this prefix P and the single drive i∗ with the highest risk in NB, the first by index among ties (Equation (6)):

$$j^* = \min\left\{ j : \sum_{l=1}^{j} c_{\pi(l)} > B \right\}, \qquad k = \pi(j^*), \qquad P = \{\pi(1), \dots, \pi(j^* - 1)\}, \tag{6}$$

$$i^* = \operatorname*{argmax}_{i \in N_B} p_i,$$

$$\mathrm{GDB} = \begin{cases} P & \text{if } V(P) \geq p_{i^*}, \\ \{i^*\} & \text{otherwise} \end{cases}$$

If every drive of NB fits, P = NB. AGE replaces the oldest drives first, in the order σ of decreasing power-on hours hi, with ties broken by the SHA-256 digest of the serial number; RANDOM scans a random order σ and serves as the negative control. Both scan the whole list and add every drive that still fits (Equation (7)):

$$x_{\sigma(j)} = \mathbf{1}\left[c_{\sigma(j)} \leq B - \sum\nolimits_{l<j} c_{\sigma(l)}\, x_{\sigma(l)} \right], \qquad j = 1, \dots, n \tag{7}$$

GDB stops at the first drive that does not fit, as the bound in Theorem 1 requires, whereas AGE and RANDOM continue; the two rules can therefore use different amounts of the budget, and in this population the difference is one terabyte (43,147 TB against 43,148 TB). A post hoc check that breaks the age ties by serial number gives the same number of failures. Table III compares the four policies and the three baselines of the post hoc analysis.

Table 3. Policies and baselines compared: order of the drives, ties, what happens at a drive that does not fit, guarantee and complexity.

Policy	Order of the drives	Ties	Drive that does not fit	Guarantee on the risk captured	Time	Memory
OPT	dynamic programming over groups of identical drives	within a group, the first drives by drive index	–	optimum	O(G′B)	O(B) for the value; G′(B + 1) bytes for the selection
GDB	risk per billed TB, p/c, decreasing	drive index	the scan stops; the better of the prefix and the best single drive is kept	at least OPT/2; OPT − GDB ≤ pk	O(n log n)	O(n)
AGE	power-on hours, decreasing	SHA-256 digest of the serial number	skipped; the scan continues	none	O(n log n)	O(n)
RANDOM	uniform random order	–	skipped; the scan continues	none (negative control)	O(n)	O(n)
Risk without cost (post hoc)	risk p, decreasing	drive index	two variants: the scan stops, or the drive is skipped	none	O(n log n)	O(n)
Model rate with an age tie-break (post hoc)	risk p, then power-on hours	SHA-256 digest	skipped	none	O(n log n)	O(n)
GDB with the oldest drives (post hoc)	as many drives of each group as GDB takes, oldest first	SHA-256 digest	–	as GDB: the same risk and cost	O(n log n)	O(n)

Complexity for n drives; G′ and B as in Proposition 2. GDB, AGE and the baselines sort the drives once and scan them once. The drive index is the position of the serial number in lexicographic order. Measured times are in the Results.

$$D[b] \leftarrow \max\{D[b],\ D[b-w]+v\}, \qquad b = B, \dots, w \tag{8}$$

Equation (8) is the recurrence of the dynamic programming over the items of Equation (5). After the first j items, the vector D holds Dj[b], the best risk with cost at most b among them, and the decision table K, read backwards from b = B, recovers the selection; ξl ∈ {0, 1} is the choice of item l, and Kj[b] = 0 for b < wj (Equation (9)):

$$D_j[b] = \max\left\{\sum_{l\le j} v_l\,\xi_l \,:\, \sum_{l\le j} w_l\,\xi_l \le b,\ \xi_l \in \{0,1\}\right\}, \tag{9}$$
$$K_j[b] = \mathbf{1}\left[\,D_{j-1}\left[b-w_j\right] + v_j > D_{j-1}[b]\,\right]$$

Theorem 1. Let NB = {i : ci ≤ B} be the drives that fit the budget, ordered by decreasing density pi/ci, with ties broken by index. Let P be the longest prefix of this order whose cost does not exceed B, with risk V(P), and let k be the first drive after P, with risk pk and cost ck. Let LP be the value of the linear relaxation of the problem over NB. Then LP ≤ V(P) + pk, GDB ≥ max{V(P), pk} ≥ LP/2 ≥ OPT/2, and OPT − GDB ≤ pk, where pk = 0 if the drives of NB fit together, that is, if their total cost is at most B.

Proof. If the drives of NB fit together, that is, if their total cost is at most B, the whole set is feasible and optimal, because all risks are non-negative, and GDB returns it. Otherwise k exists. The dual of the

linear relaxation, with a variable v for the budget and a variable μi for each bound xi ≤ 1, is Equation (10):

$$\min_{\nu \ge 0,\ \mu \ge 0} B\nu + \sum_{i \in N_B} \mu_i \quad \text{s.t.} \quad \nu c_i + \mu_i \ge p_i, \quad i \in N_B \tag{10}$$

Let v = pk/ck and μi = max(0, pi − vci) for i in NB. The pair (v, μ) is feasible for Equation (10), since vci + μi ≥ pi, and its value is Bv + Σ μi = V(P) + θpk, with θ as in Equation (11): the drives of P have pi/ci ≥ v, and the drives after P have μi = 0. The relaxation attains this value with the drives of P and the fraction θ of drive k, so the two values coincide; and OPT ≤ LP, because every feasible integer plan is feasible for the relaxation (Equation (11)):

$$\mathrm{LP} = V(P) + \theta\, p_k, \qquad \theta = \frac{B - c(P)}{c_k} < 1 \tag{11}$$

The best single drive that fits has risk at least pk, because k belongs to NB, and LP ≤ V(P) + pk, because θ < 1 and pk ≥ 0. Hence the chain of Equation (12) holds:

$$\mathrm{GDB} \ge \max\{V(P), p_k\} \ge \frac{V(P) + p_k}{2} \ge \frac{\mathrm{LP}}{2} \ge \frac{\mathrm{OPT}}{2} \tag{12}$$

Finally, OPT − GDB ≤ LP − max{V(P), pk} ≤ min{V(P), pk} ≤ pk. ■

Because OPT ≤ LP, Equation (11) and Equation (12) also bound the distance of GDB from the optimum without solving it (Equation (13)):

$$\frac{\mathrm{GDB}}{\mathrm{OPT}} \ge \frac{\mathrm{GDB}}{\mathrm{LP}} \ge \frac{V(P)}{V(P) + \theta\, p_k} \tag{13}$$

The quantities of Equation (13) are those that the greedy scan computes; the Results evaluate it at the primary budget.

Proposition 1. The decision version of the problem, which asks for a set of drives with cost at most B and risk at least T, is NP-complete, even when pi = ci for every drive.

Proof. Membership in NP is immediate, since a candidate set is checked in polynomial time. For hardness, let a1, ..., am be positive integers with sum 2T, which is an instance of PARTITION; if the sum is odd, the reduction outputs a fixed no-instance. Set pi = ci = ai, B = T and the risk threshold T. A set with cost at most T and risk at least T has cost exactly T, so it is a partition, and conversely. ■ Because the reduction keeps the integers of PARTITION, whose magnitudes may be exponential in their encoding length, it does not establish strong NP-completeness: the problem is only weakly NP-complete, and the dynamic programming of Proposition 2 is pseudo-polynomial in B. Risks are real numbers; membership in NP holds when they are represented with fixed precision, as in the implementation.

Proposition 2. If the drives are grouped into G groups of identical drives, with ng drives in group g, the grouped dynamic programming computes the optimal risk in time O(G′B) and in O(B) memory for the value, where G′ is the number of 0/1 items of Equation (5). Applied for b from B down to w to each 0/1 item of cost w and risk v, with D[b] = 0 at the start, the recurrence of Equation (8) keeps the invariant of Equation (9) after every item. Recovering the selection needs the decision table K of Equation (9), with G′(B + 1) entries, which the implementation stores as one byte per entry.

The outcome yi is 1 if the drive records a failure on any day of the second quarter, and 0 otherwise. A policy S captures F(S) = Σi∈S yi failures for TB(S) = Σi∈S ci terabytes, and its capture rate is F(S)/TB(S), which the tables and text report per 1000 TB. The primary estimand is the difference between the capture rates of GDB and AGE:

$$\Delta = \frac{F(\mathrm{GDB})}{\mathrm{TB}(\mathrm{GDB})} - \frac{F(\mathrm{AGE})}{\mathrm{TB}(\mathrm{AGE})} \tag{14}$$

Equation (14) defines the primary estimand. The estimand describes the failures observed among the drives each policy selected, per terabyte. It is not the benefit of replacing a drive, which would also require the failure rate of the replacement drives and the cost of a failure; the Discussion returns to this point.

The 95% interval is a percentile bootstrap with R = 5,000 resamples. In each resample r, the replaced drives of each policy, listed in the order of their index (the lexicographic order of serial numbers), are drawn with replacement, independently for the two policies, with one generator seeded with a fixed integer recorded in the analysis plan. With q(S) = F(S)/TB(S), the capture rate, and SG and SA, the

plans of GDB and AGE, the resampled differences and the interval, Equation (15), are

$$\Delta^{(r)} = q\big(S_G^{(r)}\big) - q\big(S_A^{(r)}\big), \qquad [L, U] = \big[\, Q_{\beta/2}\big(\Delta^{(r)}\big),\ Q_{1-\beta/2}\big(\Delta^{(r)}\big) \,\big] \tag{15}$$

where Qγ is the empirical quantile of order γ of the R differences and β = 5%. The decision rule of the plan compares the interval with zero and with the relevance margin δ, the fraction ω = 20% of the capture rate of AGE (Equation (16)):

$$\text{decision} = \begin{cases} \text{positive} & \text{if } L > 0 \text{ and } \Delta \geq \delta, \\ \text{informative null} & \text{if } L > -\delta \text{ and } U < \delta, \\ \text{inconclusive} & \text{otherwise,} \end{cases} \qquad \delta = \omega\, q(S_A) \tag{16}$$

The rule of Equation (16) does not require the lower bound to exceed the margin; the Results report that case separately. The controls are: V1, a random policy with the same budget, whose mean capture rate over R = 1,000 repetitions must match the population base rate q0 within two standard errors (Equation (17)); V2, a placebo that permutes the risk values across drives with a permutation ρ and lets GDB select again, over R = 200 permutations (Equation (18)); V3, a check that the frozen risk table and the frozen population reproduce the digests recorded before the second quarter was read; V4, the chain OPT ≥ GDB ≥ LP/2 ≥ OPT/2 at the primary budget and at five budgets in the post hoc analysis; and V5, an independent reimplementation written from the plan without the pipeline's code, which must reproduce the estimate within 10-9 times the larger of one and the reference value. V1 and V4 test the implementation, and V3 tests the frozen inputs. None of the controls tests whether the difference holds beyond this quarter. The statistics of V1 and V2 are

$$q_0 = \frac{\sum_i y_i}{\sum_i c_i}, \qquad q_{\text{RANDOM}} = \frac{1}{R}\sum_{r=1}^{R} q\left(S_{\text{RANDOM}}^{(r)}\right), \qquad \text{V1 passes if } |q_{\text{RANDOM}} - q_0| \leq \frac{2\, s_q}{\sqrt{R}} \tag{17}$$

$$\Delta^{(\rho)} = q\Big(\text{GDB}\big(p_\rho\big)\Big) - q(S_A), \qquad \text{V2 passes if } \frac{1}{R}\sum_{r=1}^{R} \mathbf{1}\left[\, \left|\Delta^{(\rho_r)}\right| \geq |\Delta| \,\right] < \varepsilon, \tag{18}$$

$$p_{\text{V2}} = \frac{\kappa + 1}{R + 1}$$

where qRANDOM and sq are the mean and the standard deviation of the R capture rates of RANDOM, ε = 5% is the threshold of the plan, and κ is the number of permutations whose difference reaches the observed one in absolute value. The add-one p-value of Equation (18) comes from a post hoc reproduction of V2 with a fixed seed; here κ = 0, so pV2 = 0.005.

Analyses that use the second quarter after its first reading are labelled post hoc. They do not replace the primary result, which keeps its rule. They were run on a per-drive summary of the second quarter that was derived once and cached, so that the quarter was not read again for each analysis. They cover the dependence between the two selections (a paired bootstrap), the dependence on individual models (bootstrap over models and leave-one-model-out), the composition by capacity class, real capacities, the uncertainty of the first-quarter rates (resampling and shrinkage), the tie-breaking rules, the drive-set comparison across horizons, baselines without density, censoring, and a cross-check of the exact optimum with a mixed-integer solver. Table IV lists the fixed settings of every analysis and whether the plan, its first amendment or a later analysis fixed them.

The paired bootstrap resamples with replacement the drives $\mathcal{U}$ selected by either policy, keeping for each drive whether GDB and AGE selected it, xGi and xAi; the model-block bootstrap resamples the models $\mathcal{M}$ that either policy selects, with the failures and terabytes of each selection within each model (Equation (19)):

$$\Delta_{\text{paired}}^{(r)} = \frac{\sum_{i \in \mathcal{U}^{(r)}} y_i\, x_i^G}{\sum_{i \in \mathcal{U}^{(r)}} c_i\, x_i^G} - \frac{\sum_{i \in \mathcal{U}^{(r)}} y_i\, x_i^A}{\sum_{i \in \mathcal{U}^{(r)}} c_i\, x_i^A}, \qquad \mathcal{U} = S_G \cup S_A, \tag{19}$$

$$\Delta_{\text{block}}^{(r)} = \frac{\sum_{m \in \mathcal{M}^{(r)}} F_{G,m}}{\sum_{m \in \mathcal{M}^{(r)}} \text{TB}_{G,m}} - \frac{\sum_{m \in \mathcal{M}^{(r)}} F_{A,m}}{\sum_{m \in \mathcal{M}^{(r)}} \text{TB}_{A,m}}$$

The leave-one-out analyses remove from the population $\mathcal{P}$ a subset $\mathcal{A}$, a model or a capacity class. Both policies select again under the same budget, and an analysis with an interval is classified by the rule of

Equation (16) with its own margin $\delta_{-\mathcal{A}}$ (Equation (20)):

$$\Delta_{-\mathcal{A}} = q\big(\mathrm{GDB}(\mathcal{P} \setminus \mathcal{A},\, B)\big) - q\big(\mathrm{AGE}(\mathcal{P} \setminus \mathcal{A},\, B)\big), \qquad \delta_{-\mathcal{A}} = \omega\, q\big(\mathrm{AGE}(\mathcal{P} \setminus \mathcal{A},\, B)\big) \tag{20}$$

The first-quarter rates are perturbed in two ways: the failures of each model are resampled, and the rates are shrunk towards the pooled rate λ0 of all the models of the first quarter by a gamma prior with shape α = 2 and mean λ0, whose posterior mean is λ(α)m. Each new selection S′ is compared with the primary selection S by the Jaccard index and by the drives that leave and enter it (Equation (21)):

$$F_m^* \sim \mathrm{Binomial}\big(e_m, \hat{\lambda}_m\big), \qquad \lambda_m^* = \frac{F_m^*}{e_m}, \qquad \lambda_m^{(\alpha)} = \frac{F_m + \alpha}{e_m + \alpha/\lambda_0}, \qquad \lambda_0 = \frac{\sum_m F_m}{\sum_m e_m}, \tag{21}$$

$$J(S,S') = \frac{|S \cap S'|}{|S \cup S'|}, \qquad \mathrm{out} = |S \setminus S'|, \qquad \mathrm{added} = |S' \setminus S|$$

Calibration compares, for each model m, the second-quarter failures Om of its drives $\mathcal{P}$ m with the failures Em expected from the first-quarter risks, against the central 95% range Im of a Poisson count with mean Em, whose quantile of order γ is QPoisE(γ), with β as in Equation (2). Model m lies outside the range if Om ∉ Im, and (O/E) is the ratio of the whole population (Equation (22)):

$$E_m = \sum_{i \in \mathcal{P}_m} p_i\,, \qquad I_m = \big[\, Q_{E_m}^{\mathrm{Pois}}(\beta/2),\; Q_{E_m}^{\mathrm{Pois}}(1-\beta/2)\,\big], \qquad (O/E)_{\mathcal{P}} = \frac{\sum_m O_m}{\sum_m E_m} \tag{22}$$

Discrimination is measured per drive by the area under the ROC curve of each ordering criterion, with mid-ranks rki for ties among the n drives, n1 of which failed in the second quarter and n0 did not (Equation (23)):

$$\mathrm{AUC} = \frac{\sum_{i:y_i=1} \mathrm{rk}_i - n_1(n_1+1)/2}{n_1\, n_0} \tag{23}$$

Table 4. Fixed settings of every analysis, and where each one was fixed.

Setting	Value	Fixed in
Risk horizon H	91 days; 30 and 182 days in S3	plan
Budgets f	1%, 2%, 5%, 10%, 20% of the population capacity; primary 5%	plan
Cost ci and budget B	capacity in bytes / 1012, rounded; B = floor(f · Σ ci), 43,148 TB at 5%	plan
Interval of Δ	percentile bootstrap, 5,000 resamples of the selected drives of each policy, 95% (Equation (15))	plan
Order of the resamples	the drives of GDB, then those of AGE, each sorted by drive index, with one generator	amendment 1
Relevance margin	20% of the capture rate of AGE (Equation (16))	plan
Ties	drive index for GDB; SHA-256 digest of the serial number for AGE	amendment 1 (GDB); plan (AGE)
V1 and V2	1,000 random selections; 200 permutations of the risks, passing below 5% (Equation (17) and Equation (18))	plan
V5 tolerance	10-9 times the larger of one and the reference value; amendment 1 allows 10-6 for the interval, which was not needed	plan; amendment 1
Seeds	fixed integers: those of the bootstrap, RANDOM and the placebo are in the plan, the others in the code	plan; code
Cost measurement	1,000; 10,000; 30,000 and 106,218 drives; budgets of 0.5%, 5% and 10% of the capacity, rounded down; wall-clock time and peak memory, median of 5 repetitions after one warm-up; one thread	plan; changed after the first reading
First-quarter resampling	500 binomial resamples of the failures of each model (Equation (21))	post hoc

Setting	Value	Fixed in
Shrinkage of the rates	gamma prior with shape α = 2 and mean the pooled first-quarter rate (Equation (21))	post hoc
Random tie-breaks	200 random orders within groups of equal density	post hoc
Paired and model-block bootstraps	5,000 resamples of the drives in either selection, or of the 17 models (Equation (19))	post hoc
Intervals of the AUC	2,000 bootstrap resamples of the drives (Equation (23))	post hoc
Random instances of the guarantee	20,000 instances with at most 12 items, costs from 1 to 9, risks uniform on the open unit interval and budgets drawn uniformly	plan (the check); code (its settings)
Tight family of the bound	η = 10-4; B = 3, 10, 100 and 1,000	plan (a tight family); code (this one)
Cost of the other stages	median of 3 repetitions after one warm-up; one thread	not in the plan

The plan was tagged before the second quarter was read, and amendment 1 was written before that reading as well. Post hoc settings belong to analyses run after the first reading.

We measured wall-clock time and peak memory of the greedy policy, of the dynamic programming and of the grouping step that the dynamic programming needs, all from the same prepared inputs, on four subsamples of the population and at three budgets of the population billed capacity, rounded down, listed in Table IV; in a subsample of n of the N population drives, the budget is the population budget multiplied by n/N and rounded down (406 TB for 1,000 drives at 5%). Each time is a median over repetitions after one warm-up, with its interquartile range. Peak memory is measured in separate passes with tracemalloc, so that it does not affect the timing. The number of threads is fixed to one. The processor has eight cores and sixteen threads, the machine has 16 gibibytes (GiB) of memory, and the graphics processor is not used. The benchmark is not part of the confirmatory plan, and it describes this implementation rather than the algorithms in general.

From Table III and Proposition 2, the expected costs of the measured steps, Equation (24), are

$$T_{\mathrm{GDB}} = \Theta(n \log n), \qquad T_{\mathrm{DP}} = \Theta(G'B), \qquad M_{\mathrm{DP}} = G'(B+1)\ \text{bytes} + \Theta(B) \tag{24}$$

where n log n comes from sorting the drives and the bytes count the decision table of Equation (9). With the fraction f fixed, B grows in proportion to n. The measurements are read against these rates; no exponent is fitted.

The software is a pipeline in Python 3.11 with NumPy 1.24.1, SciPy 1.9.3 (whose mixed-integer solver, HiGHS, cross-checks the dynamic programming), pandas 1.5.2 and matplotlib 3.6.2. The manuscript is generated from the results by code, with Apache POI 5.2.3 on the journal template. Before it is exported to PDF, it passes verification gates for dates, numbers, claims, references, equations and the template. The analysis plan, its amendments and the frozen digests are version-controlled.

The analysis plan was committed under a tag before the second quarter was read. One amendment was written before that first reading; it fixes the tie-breaking rule, the order of resampling and the tolerance of the replication. A second record, written after the first reading, documents three changes made afterwards. First, the loader gained an explicit filter for the window of the plan and an explicit reading of the two date formats in the file, whereas the first version read every row without a date filter. The outputs of the two versions are the same, apart from a new report of the window: the secondary and validity files have identical checksums, and the primary result is unchanged. Second, the independent replication first read only one of the two formats; it left out seven failures of the population, and it was corrected and run again on the second quarter, which had already been read. Third, the post hoc analyses were added. The replication was written from the plan without the pipeline's code, but its first version was compared with the pipeline's output, so it is not a blind replication. The second quarter was read several times in total, each reading logged (Figure 1b), and the analyses of the Results use only the primary rules unless they are labelled post hoc. Table V lists what the plan fixed, what the first amendment added before the first reading, and what changed after it.

Table 5. What the analysis plan fixed, what the first amendment added before the second quarter was

read, and what changed afterwards.

Element	Analysis plan (tagged)	Amendment 1 (before the first reading)	After the first reading	Effect on the primary result
Question, data, population, risk, cost and budgets	fixed	–	–	none
Estimand, interval and decision rule	fixed	order of the resamples and of the selected drives	–	none
Ties in density and in age	drive index for GDB, without defining the index; SHA-256 digest for AGE	index defined as the lexicographic order of the serial numbers	–	none
Controls V1 to V5	fixed; V4 at every budget	tolerance of V5 for the interval	replication corrected after it missed one date form; V4 run at all five budgets, where before it had been run at the primary budget only	none: the same numbers
Loader of the second quarter	sealed until the tag	–	filter to the window and both date forms; identical outputs	none
Secondary analyses S1 to S5	fixed	–	post hoc checks added: the drive sets in S3 and the failure counts in S5	none
Cost measurement	three budgets, the smallest 4,315 TB; CPU and wall-clock time; a log-log fit with an interval	–	budgets rounded down, the smallest 4,314 TB; wall-clock time and memory only; no fitted exponent	none: not part of the confirmatory analysis
Analyses after the reading	not planned	–	post hoc analyses on a per-drive summary, labelled as such	they do not replace the primary result

The plan and the first amendment were committed under version control before the second quarter was read. The second amendment records the changes to the loader, the replication and the post hoc analyses; the other changes are recorded in the decisions file of the repository.

3. RESULTS

Population and data. The first quarter contains 9,577,046 drive-day records for 115,231 drives of 49 models over 90 days (Table I). On the decision day, 106,218 drives remained after the exclusions, with 862,963 billed TB. Of the 49 models, 30 recorded no failure in the first quarter, so their estimated risk is zero; these models hold 1,582 drives of the population, and 1 of them failed in the second quarter. In the second quarter, 492 population drives failed.

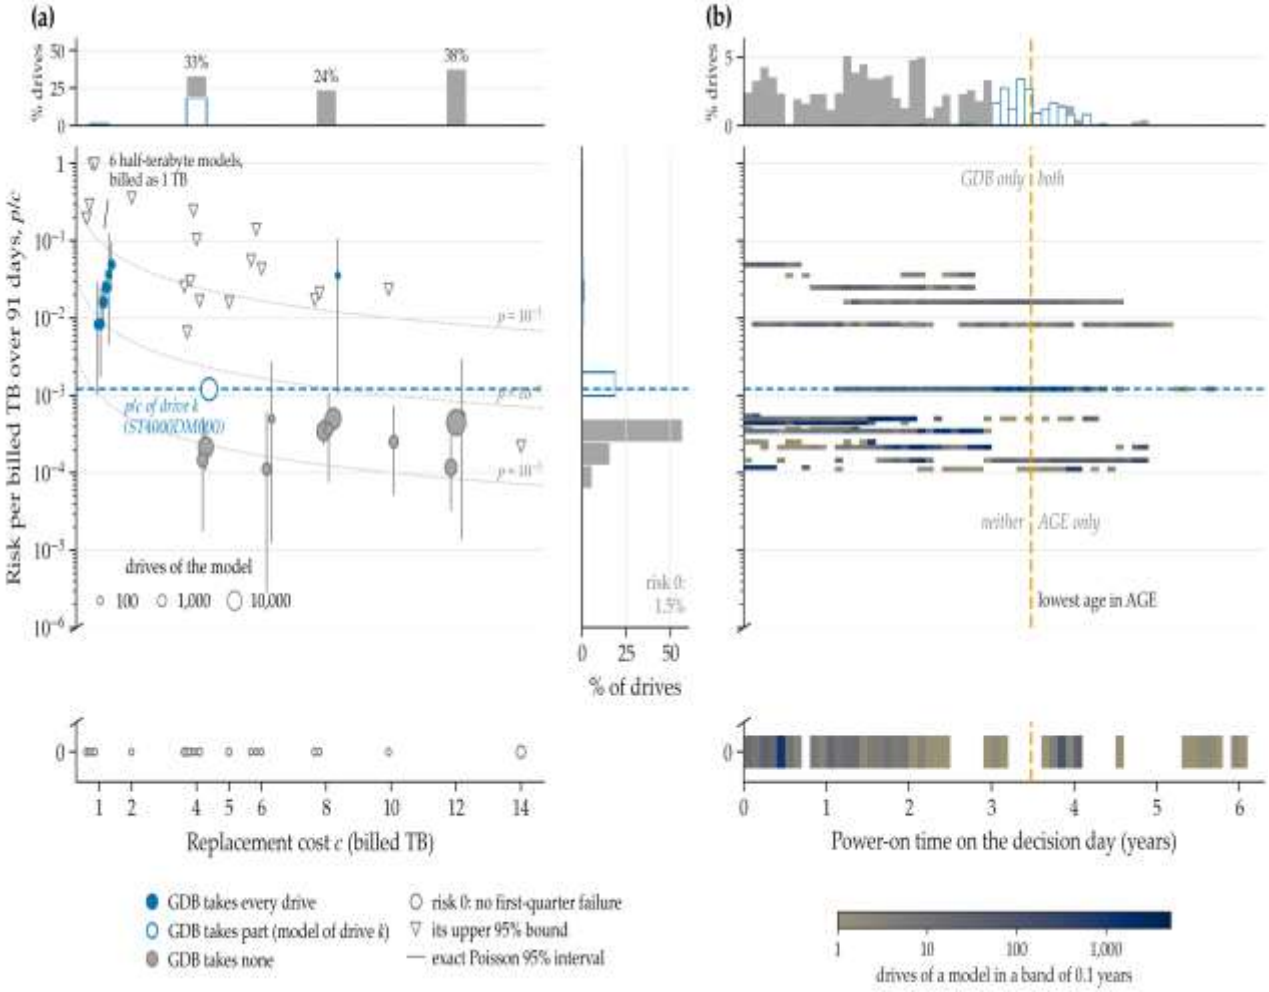


Figure 2. Risk by model and the two selection rules, first quarter only (no outcome is used). (a) Each point is a model; its area grows with the square root of its drives in the population. The horizontal position is the replacement cost, with models of equal cost spread apart. The vertical position is the risk per billed terabyte over the horizon, p/c, by which GDB ranks the drives, with the exact Poisson 95% interval of the model rate. The dashed line is the density of drive k, the first drive that does not fit at 5% of capacity: GDB takes every drive of the models above it, part of the model on it and none below it. Models without first-quarter failures have zero estimated risk and are drawn on the bottom row; triangles mark their upper bounds. Dotted curves join points of equal risk p (one in a thousand, a hundred and ten). The margins give the share of the population's drives in each band of p/c (right) and in each cost (top), split by what GDB takes of their model. (b) The drives of the population by power-on time on the decision day and by p/c: each model is a horizontal band, shaded by the number of its drives in each band of power-on time. The vertical line is the lowest power-on time among the drives that AGE selects at 5% of capacity; AGE takes all but 149 of the drives at least as old, which did not fit when its scan reached them. The labels name the selection that takes the drives of each region; the top margin is as in (a), by power-on time.

Risk, cost and age by model. Figure 2 shows how the greedy policy ranks the 38 models of the population. At the primary budget it takes every drive of the 7 models whose risk per billed terabyte exceeds that of drive k, part of ST4000DM000, the model of drive k, and no drive of the other 30. Of the 7 models taken in full, 6 are half-terabyte models, billed one terabyte each, so their risk per billed terabyte equals their risk. The first-quarter intervals are wide for small models: for 20 of the 30 models without selected drives, the upper bound of the risk per billed terabyte lies above that of drive k, and 18 of them recorded no failure in the first quarter. Figure 2b shows why the two selections overlap only in part: GDB takes the bands above the density of drive k, whole models at a time, whereas AGE takes the oldest drives across models, and it takes only part of 10 of the 12 models from which it selects drives.

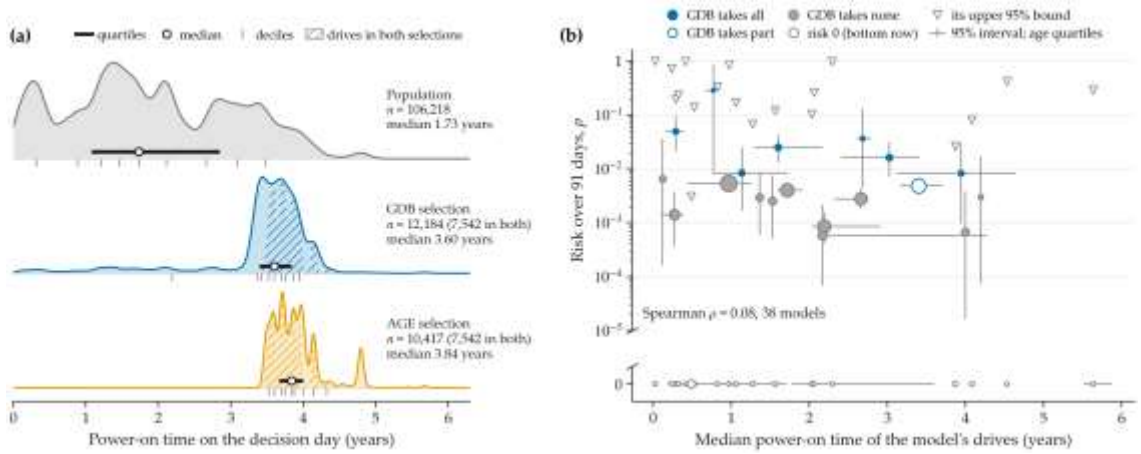


Figure 3. Age and risk, first quarter only (no outcome is used). (a) Distribution of the power-on time on the decision day (kernel density, each scaled to the same height) for the population and for the drives each policy selects at 5% of capacity, with the median, the interquartile range and the deciles (ticks) of each group; hatching marks the 7,542 drives that both policies select, on the scale of each selection. (b) Median power-on time of each model's drives against the model's risk over 91 days, with exact Poisson 95% intervals and the quartiles of its drives' power-on time (horizontal); markers as in Figure 2. Models with zero risk are drawn on the bottom row, and triangles mark their upper bounds.

Across models, age does not track risk in this fleet: the rank correlation between a model's median power-on time and its first-quarter rate is 0.08 over the 38 models (Figure 3b). The age policy takes the oldest drives, each with at least 3.47 years of power-on time (Figure 3a); Table VI gives the median and the quartiles of each group. The greedy selection is also old. Of its 12,184 drives, 10,314 are ST4000DM000 drives, taken in serial order within that model, and 7,300 of these are also in the age selection, against about 4,053 expected if serial order were unrelated to age. Table VII shows the choice by serial prefix (post hoc): the greedy policy takes every S30 and W30 drive and a third of the Z30 drives, whereas the age policy takes most of the S30 and W30 drives and a quarter of the Z30 drives, so within this model the serial order favours older drives.

Table 6. The population and each selection at 5% of capacity, before the outcome (first quarter only).

Characteristic	Population	GDB	AGE
Models represented	38	8	12
Physical capacity (TB)	862,084.6	42,221.8	43,030.5
Billed TB per physical TB	1.001	1.022	1.003
Power-on time, median (years)	1.73	3.60	3.84
Power-on time, first to third quartile (years)	1.08 to 2.85	3.40 to 3.84	3.67 to 4.00
Expected failures per 1000 drives over 91 days	4.08	7.25	4.04
Drives in models without first-quarter failures	1,582	0	160

Expected failures use the first-quarter risks; no outcome is used. The drives, billed capacity and outcomes of each selection are in Table VIII.

Table 7. Post hoc: ST4000DM000 drives by serial prefix, and the drives of each prefix that each policy replaced.

Serial prefix	Drives	GDB drives	GDB failures (Q2)	AGE drives	AGE failures (Q2)	Neither: drives	Neither: failures (Q2)
S30	5,324	5,324	27	4,206	18	0	0
W30	212	212	0	165	0	0	0
Z30	14,249	4,778	32	3,403	26	8,997	41

The prefix is the first three characters of the serial number; the three prefixes cover every drive of the model. GDB breaks ties by the lexicographic order of the serial numbers, so it reaches S30 and W30 before

Z30. Failures are second-quarter outcomes, counted after the first reading.

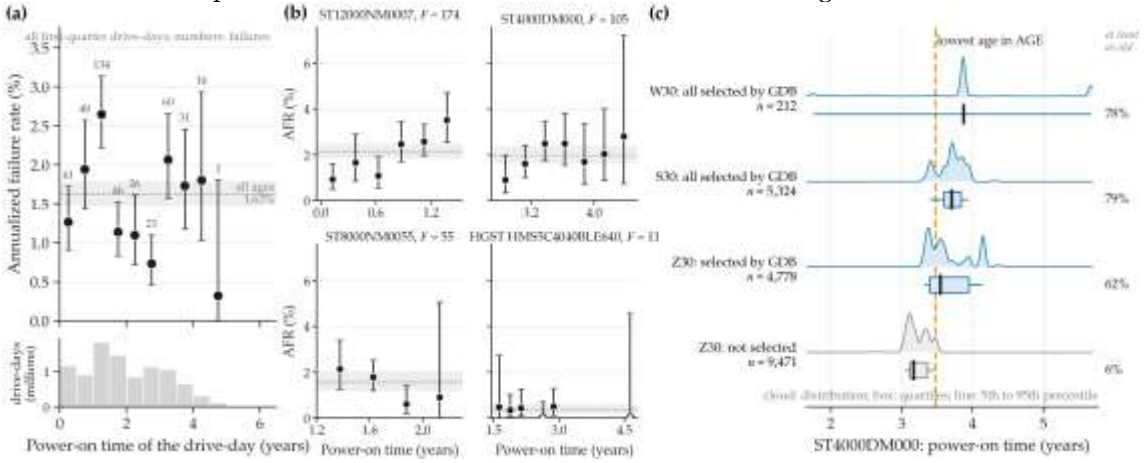


Figure 4. Age and failure in the first quarter (no outcome is used). (a) Annualized failure rate of all first-quarter drive-days with power-on hours, by the power-on time of the drive on that day, with exact Poisson 95% intervals; bands with fewer than 20,000 drive-days are not shown; the strip below gives the drive-days of each band and the numbers its failures; the dotted line and grey band give the rate of all bands together, with its exact interval. (b) The same within the four models with the most drives in the population, each over its own range of power-on time (F in each title); open markers have no failure. (c) Power-on time on the decision day of the drives of ST4000DM000, by serial prefix and by selection of GDB (box: quartiles; line: 5th to 95th percentile; bar: median), with the distribution of each group above its box; the dashed line is the lowest power-on time among the drives that AGE selects, and the right column the share of each group at least that old.

Age and failure. In the first quarter the failure rate does not rise steadily with power-on time (Figure 4a): it is 2.65% per year for drive-days at one to one and a half years (95% interval 2.22 to 3.14) and 0.73% at two and a half to three years (0.47 to 1.10). This curve mixes models of different ages, and within the four largest models the shape differs (Figure 4b): in ST12000NM0007, the largest, the rate grows from 0.94% to 3.51% over its first one and a half years. Figure 4c shows the serial order of Table VII in power-on time. The ST4000DM000 drives that GDB selects are older than those it leaves out: the median is 3.71 years in S30 and 3.88 in W30, all of them selected, and 3.55 among the selected Z30 drives, against 3.17 among the Z30 drives left out; 79% of the S30 drives and 62% of the selected Z30 drives are at least as old as the youngest drive in the age selection, against 6% of the Z30 drives left out.

Table VIII. Primary comparison at 5% of capacity.

Policy	Drives replaced	Capacity (TB)	Failures in Q2	Failures per 1000 TB	Expected failures from T1 risks
GDB (greedy)	12,184	43,147	94	2.18	88.36
AGE (oldest first)	10,417	43,148	52	1.21	42.08
OPT (exact on T1 risks; outcome post hoc)	12,184	43,147	94	2.18	–

Failures in Q2 are counted over the drives each policy replaced. Expected failures use the first-quarter risks over the 91-day horizon. The exact optimum selects the same numbers of drives and terabytes as GDB at this budget. Drives within a group of identical drives have the same risk and cost, so the optimum's failures depend on which of them are taken; the count shown is one such choice. The outcome is a post hoc count. The expected failures of OPT and the LP bound are in Table XI.

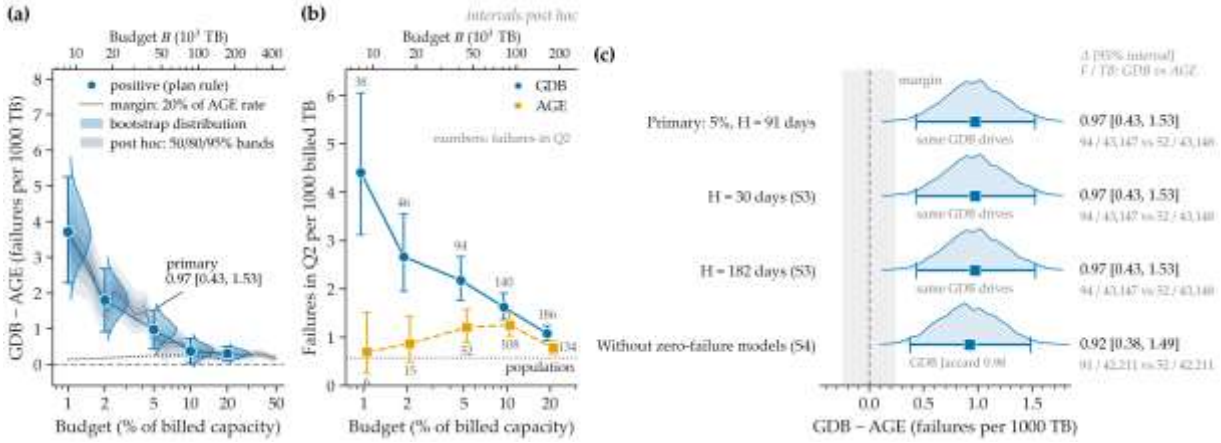


Figure 5. Primary result. (a) Difference in failures per 1000 TB between the greedy risk-density policy (GDB) and the age policy (AGE), with 95% bootstrap intervals, at five budgets; half-violins show each bootstrap distribution, reproduced post hoc with the seed of the plan, which gives the same intervals. Filled markers are differences that the rule of the plan classifies as positive, and the dotted line is the relevance margin of each budget (20% of the age-policy rate at the same budget). The fan gives the pointwise 50%, 80% and 95% bootstrap bands of the difference on a fine grid of budgets, with its estimate as a grey line (post hoc). (b) Capture rates of the two policies at the same budgets, with exact Poisson 95% intervals computed post hoc, and the population rate, which a random selection would capture in expectation; numbers are failures. The top axes give the budget in thousands of billed TB. (c) Differences for the horizon analyses (thirty, ninety-one and one hundred eighty-two days) and for excluding the models without first-quarter failures, with 95% intervals and bootstrap distributions (repeated post hoc with the seed of the plan; they reproduce the intervals). The other horizons select the same drives as the primary analysis; for the exclusion, the label gives the Jaccard index of its greedy selection with the primary one. Beneath each value: failures and billed TB of GDB and AGE. The grey band in (c) spans plus and minus the relevance margin (20% of the age-policy rate at 5% of capacity) around zero.

Primary result. At a budget of 5% of capacity (43,148 TB), the greedy policy replaced 12,184 drives, captured 94 failures in the second quarter and spent 43,147 TB, a capture rate of 2.18 failures per 1000 TB. The age policy replaced 10,417 drives, captured 52 failures and spent 43,148 TB, a rate of 1.21 per 1000 TB (Table VIII; Figure 5). The difference is 0.97 failures per 1000 TB, with a 95% bootstrap interval from 0.43 to 1.53. The greedy rate is 81% above the age rate, the lower bound is above zero, and the difference exceeds the relevance margin of 0.24 failures per 1000 TB (20% of the age rate). The pre-specified decision is therefore positive. Two qualifications apply, and both are reported below: the two selections share 7,542 drives, which the primary bootstrap treats as independent, and, post hoc, the difference depends on the mix of drive classes.

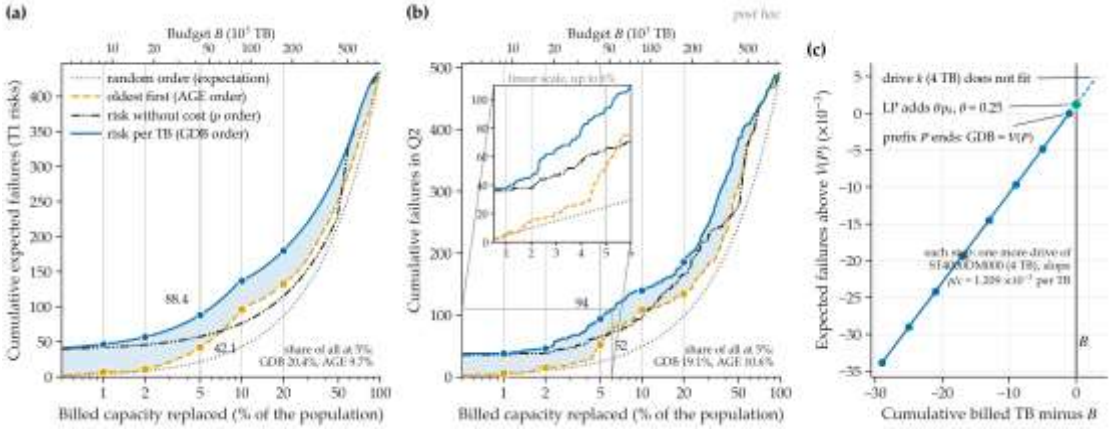


Figure 6. Capture along the orders of the two policies. (a) Cumulative expected failures from the first-quarter risks and (b) cumulative second-quarter failures (post hoc) against the billed capacity replaced, taking the drives in the order of GDB (risk density, solid) or of AGE (power-on hours, dashed); the dash-dotted curve takes them by risk without cost, a post hoc baseline. The dotted curve is the expectation of a

random order, and the shading marks where the order of density is ahead of the order of age. Vertical lines mark the budgets of the primary curve (the thicker one, the primary budget), and the markers give the selection of each policy at those budgets, with their values at the primary budget; notes give each selection's share of all the failures at that budget, and the top axes the budget in thousands of TB. The inset in (b) shows the first six per cent on a linear scale. (c) The end of the greedy prefix at the primary budget, with expected failures above V(P): the prefix P stops 1 TB short of B, the next drive k costs 4 TB and does not fit, and the linear relaxation LP adds the fraction θ = 0.25 of its risk; each step is one more drive of the model of drive k.

Capture at every budget. Figure 6 extends the comparison from the five budgets to the whole range, with the capture curves of Equation (25) in the density order π of Equation (6), which here covers every drive:

$$C_\pi(u) = \sum_{j\,:\,\sum_{l\le j} c_{\pi(l)} \le u} p_{\pi(j)}\,, \qquad C_{\mathrm{rand}}(u) = u\,\frac{\sum_i p_i}{\sum_i c_i}, \tag{25}$$

$$u_{1/2}(\pi) = \min\left\{u\colon C_\pi(u) \ge 1/2 \sum\nolimits_i p_i\right\}$$

In Equation (25), Cπ(u) counts the expected failures among the first u billed terabytes in the order π, the observed curve uses yi in place of pi, Crand(u) is the expectation of a random order, and u1/2(π) is the capacity that holds half of the expected failures; Cσ(u) and u1/2(σ) are the same quantities in the order σ of AGE (Equation (7)). At the primary budget, the greedy selection holds 20.4% of the expected failures of the population, against 9.7% for the age selection, and half of the expected failures lie within 29.1% of the capacity in the order of density, against 41.2% (Figure 6a). Ranking by risk without cost, a post hoc baseline, holds 13.2% of the expected failures within the primary budget and reaches half of them at 46.5% of the capacity. Post hoc, the second-quarter failures follow the same pattern. The greedy selection captured 19.1% of them with 5% of the capacity, and the age selection 10.6%. At the smallest budget, 1% of the capacity, the age selection captured 6 failures, close to the 4.9 expected from a random selection of the same capacity, against 38 for the greedy selection. The order of density captures at least as many failures as the order of age at every capacity below 70%; above it the two curves are close, and the order of age is ahead at some capacities (Figure 6b). At the primary budget the greedy prefix leaves 1 TB of the budget unused, and drive k, of 4 TB, does not fit. The bound OPT − GDB ≤ pk of Theorem 1 is then the risk of one drive, 0.0048, against 88.36 expected failures in the prefix, so GDB, OPT and LP coincide at the scale of the fleet (Figure 6c); by Equation (13), 1 − GDB/OPT ≤ θpk/LP = 0.0014%, without solving OPT. Per drive and without costs, ranking by first-quarter risk separates the drives that failed in the second quarter from the others with an area under the ROC curve of 0.65, against 0.48 for power-on hours, whose interval includes the value of chance (Table IX; post hoc).

Table 9. Post hoc: how well each ordering criterion separates the drives that failed in the second quarter, per drive.

Criterion	Area under the ROC curve	95% interval
Risk p (rate of the model, first quarter)	0.653	0.632 to 0.675
Risk per billed TB, p/c	0.611	0.590 to 0.632
Power-on hours on the decision day	0.484	0.461 to 0.508

Population of 106,218 drives, 492 of which failed in the second quarter. Ties take mid-ranks; the interval comes from 2,000 bootstrap resamples of the drives. An area of 0.5 is no better than chance. The area is per drive and does not weight drives by their cost; capture per terabyte is in Figure 6.

Shared and unshared drives (post hoc). The two selections share 7,542 drives, with 29,442 TB and 45 failures (Figure 7). Both selections spend almost the same capacity (43,147 and 43,148 TB), so the shared drives add the same failures and the same terabytes to both capture rates, and the difference comes from the drives that only one policy selected: 4,642 drives with 13,705 TB and 49 failures for the greedy policy (3.58 per 1000 TB), and 2,875 drives with 13,706 TB and 7 failures for the age policy (0.51 per 1000 TB). Of the drives that only the greedy policy selected, 1,625 are half-terabyte drives. The overlap depends on the budget (Figure 7b): the Jaccard index of the two selections is 0.03 at 1% of capacity, 0.50 at the primary budget and 0.77 at 10%, and it falls again to 0.42 at 20%.

Equation (26) splits the difference exactly into the drives that only one policy selected and the drives

that both selected:

$$\Delta = \left[\frac{F(S_G \setminus S_A)}{\mathrm{TB}(S_G)} - \frac{F(S_A \setminus S_G)}{\mathrm{TB}(S_A)}\right] + F(S_G \cap S_A)\left[\frac{1}{\mathrm{TB}(S_G)} - \frac{1}{\mathrm{TB}(S_A)}\right] \quad (26)$$

The two selections differ by one terabyte, so the second term of Equation (26) that of the shared drives, is 0.00002 failures per 1000 TB and the first term gives 0.97 (post hoc).

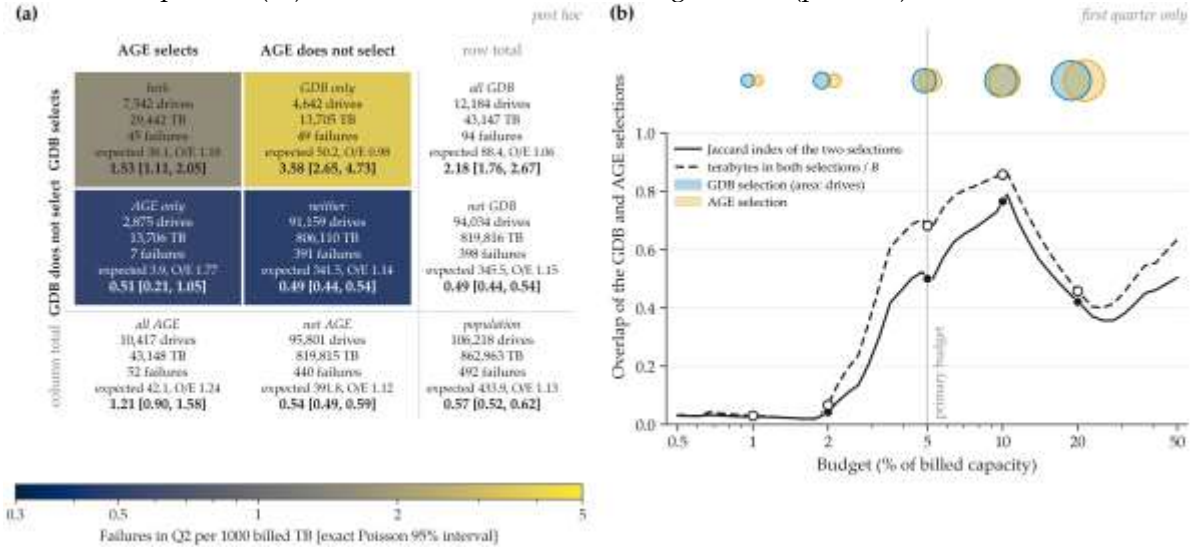


Figure 7. Shared and unshared drives. (a) Post hoc: each cell crosses selection by GDB with selection by AGE at 5% of capacity and gives the drives, the billed terabytes and the second-quarter failures, with the failures per 1000 billed TB and their exact Poisson 95% interval; the colour encodes that rate. Each cell also gives the failures expected from the first-quarter risks and the ratio of observed to expected failures (O/E). The margins give the totals of each selection, of the drives outside it and of the population. (b) First quarter only: the overlap of the two selections across budgets, as the Jaccard index of the two sets of drives (solid) and the share of the budget spent on drives that both select (dashed); markers give the five budgets of the plan. Above them, Euler diagrams at those budgets: circle areas are proportional to the drives of each selection (one scale for all) and the overlap to the shared drives.

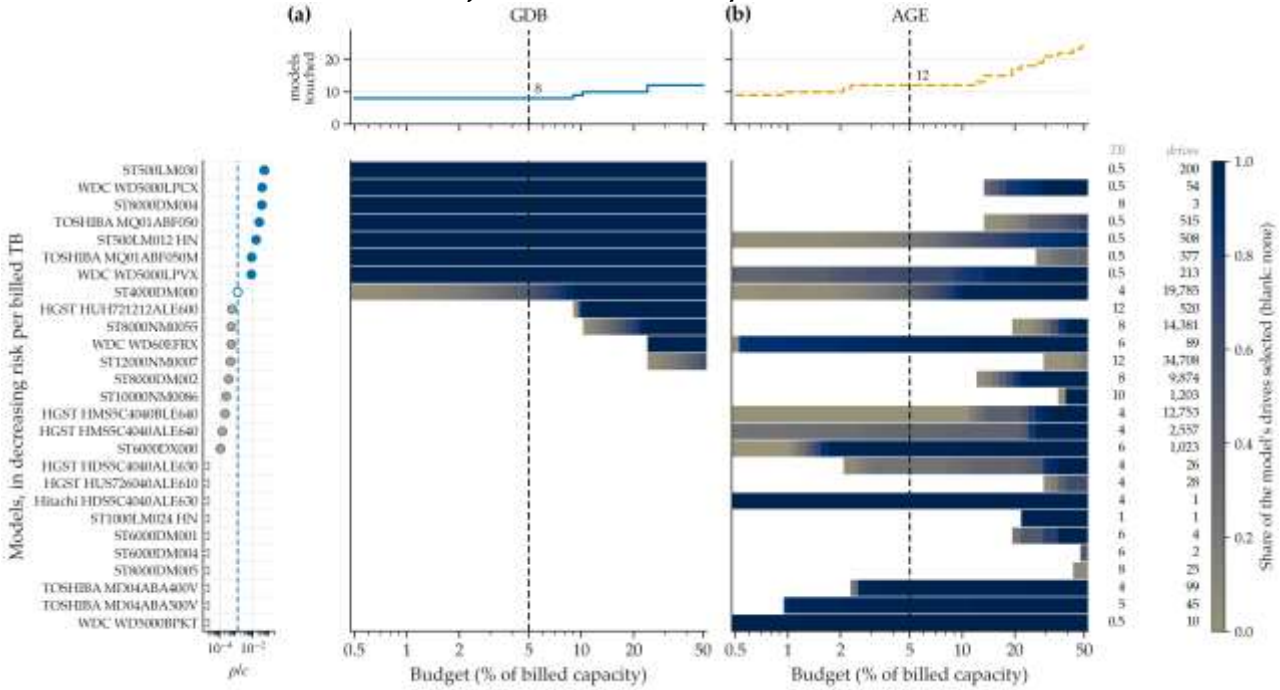


Figure 8. How each policy fills the budget, model by model (first quarter only; no outcome is used). Share of the drives of each model that GDB (a) and AGE (b) select at budgets from 0.5% to 50% of billed capacity,

with the models in decreasing order of risk per billed terabyte; blank cells mean no drive selected, and the dashed line marks the primary budget. Only the models that either policy selects at some budget are shown. The left column gives p/c, which sets the order (markers as in Figure 2; triangles: zero risk; dashed: drive k); the top strips count the models touched at each budget, and the right columns give each model's capacity class and drives.

Figure 8 extends the comparison of the two selections to every budget of the grid. GDB fills the budget model by model, in decreasing order of risk per billed terabyte, and at every budget it takes at most one model in part. AGE takes drives from many models at once: at the primary budget it selects drives of 12 models, 10 of them only in part, whereas GDB selects drives of 8 models and takes 1 of them in part.

Table 10. Validity controls, consistency checks and sensitivity analyses.

Analysis	Setting	Estimate	Decision
V1 random control	1,000 repetitions, same budget	mean 0.572 per 1000 TB; base 0.570	pass
V2 placebo	200 permutations of the risks	0 reach the observed difference; add-one pV2 = 0.005 (post hoc)	pass
V3 frozen inputs	digests before the second quarter was read	identical	pass
V4 consistency chain	OPT ≥ GDB ≥ LP/2 ≥ OPT/2 at five budgets, the primary one among them (post hoc)	holds at every budget	pass
V5 independent reimplementation	written from the plan without the pipeline's code (not blind; see Results)	same numbers of drives and failures, difference, interval and base rate	pass
S2 optimum gap	GDB/OPT on first-quarter risks, five budgets	lowest 0.99997	descriptive
S3 horizon	thirty, ninety-one and one hundred eighty-two days	0.97 in each case; same drives selected (post hoc check)	descriptive
S4 models without first-quarter failures excluded	104,636 drives, budget recomputed	0.92 [0.38; 1.49] per 1000 TB	descriptive
S5 unit cost	budget 5,310 drives	GDB 64 failures (12.05 per 1000 drives); AGE 18 (3.39)	descriptive (counts post hoc)

Intervals are 95% bootstrap percentile intervals. Decisions follow the rule of the analysis plan where the plan defines a margin. S2 uses first-quarter risks only. S3 to S5 are pre-specified secondary analyses in the plan; the plan defines no margin for them, so they are reported as descriptive, without a decision rule; the drive-set comparison in S3 and the failure counts in S5 are post hoc checks. V3 and V4 are consistency checks, not validity tests. The placebo percentiles are in the Results.

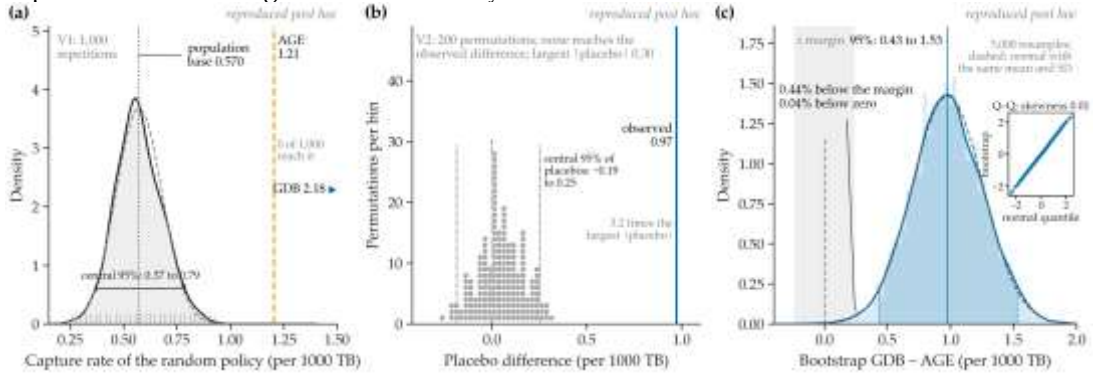


Figure 9. Distributions behind the controls and the primary interval (post hoc reproduction). They are recomputed with the seeds of the analysis plan from the per-drive summary of the second quarter, derived after its first reading, and they reproduce the stored summaries. (a) Capture rate of the random policy over 1,000 repetitions (V1), as a kernel density with a rug of the observed rates (repetitions with the same rate share a tick), with the population base rate (dotted) and the rate of AGE (dashed); the rate of GDB lies beyond the axis. Dashed: the normal with the same mean and standard deviation; the bracket spans the central 95%. (b) Placebo difference over 200 permutations of the risks (V2), one point per permutation, with the observed difference; dotted lines bound the central 95% of the placebos. (c) Bootstrap distribution of the primary difference over 5,000 resamples (histogram and kernel density), with its central 95% shaded between dotted lines, the observed difference (solid line) and zero (dashed); the grey band spans plus and minus the relevance margin; the dashed curve is the matching normal, and the inset a normal quantile plot of the standardized resamples, with their skewness.

Controls. The random control captured 0.572 failures per 1000 TB on average over 1,000 repetitions, against a population base rate of 0.570 (Table X). The difference of 0.002 is below two standard errors (0.007 per 1000 TB). In none of the 200 permutations of the risk values did a placebo reach a difference as large as the observed one. The post hoc placebo distribution, with the same seed, has a median of 0.05 per 1000 TB and 5th and 95th percentiles of −0.14 and 0.23. In a post hoc reproduction with the same seeds (Figure 9), no random repetition reached the capture rate of the age policy, and the largest placebo difference in absolute value was 0.30 per 1000 TB. The frozen risk table and the frozen population reproduced their digests before the second quarter was read. The chain OPT ≥ GDB ≥ LP/2 ≥ OPT/2 held at all five budgets within a floating-point tolerance of 10-9, and the greedy value was at least 0.99997 of the optimum on the first-quarter risks at each of them. Both policies respected the budget. An independent reimplementation, written from the analysis plan without the pipeline's code (not blind: its first version was compared with the pipeline's output; see Methods), reproduced the primary estimate: 12,184 and 10,417 drives, 94 and 52 failures, and the same difference, interval and base rate.

Budget and sensitivity. The difference was positive at all five budgets, from 1% to 20% of capacity (Figure 5a). It was largest at 1% (3.71 per 1000 TB) and smallest at 20% (0.30). On a fine grid of 53 budgets from 0.5% to 50% of capacity (post hoc), the difference stays between 0.17 and 8.35 per 1000 TB without falling steadily, and its pointwise 95% bootstrap band, from 1,000 resamples at each budget, excludes zero at every budget of the grid (Figure 5a). The lower bound of the interval exceeded the relevance margin at 1%, 2%, 5%. Estimating risk over thirty or one hundred eighty-two days left the estimate unchanged (Table X); a post hoc check found that the same drives were selected. Excluding the models without failures in the first quarter left the difference positive (0.92 per 1000 TB, interval 0.38 to 1.49). With unit costs (S5, each drive costs one unit; the counts are post hoc), the greedy policy captured 64 failures among 5,310 drives, and the age policy 18.

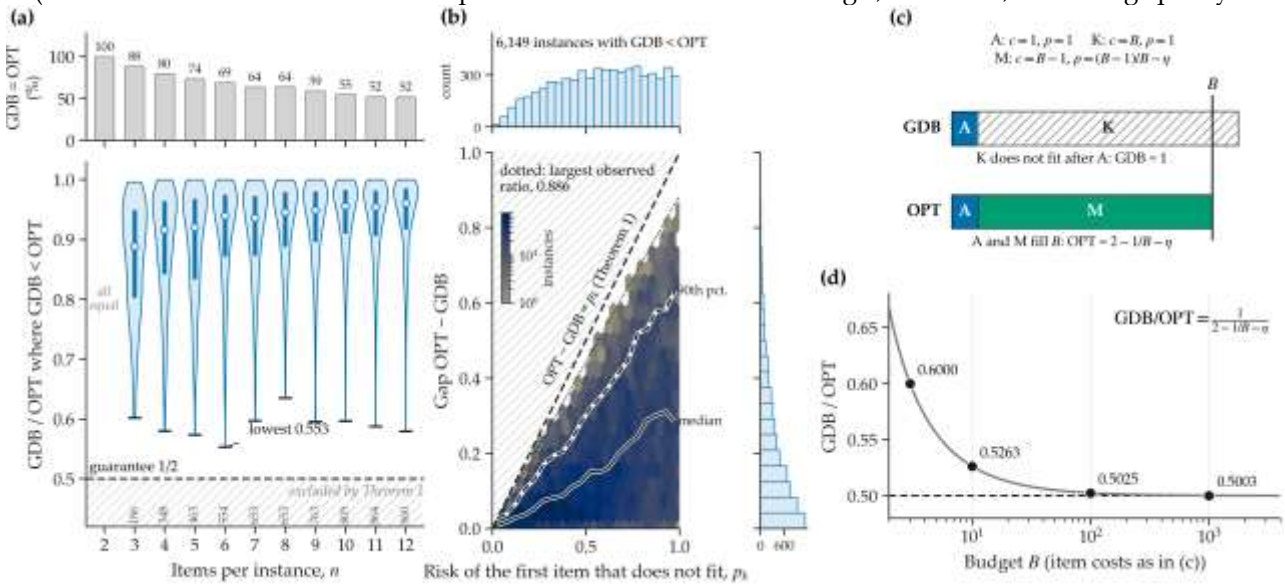


Figure 10. Guarantee on random instances and the tight family. (a) Random instances with at most 12 items, by number of items: share of instances in which the greedy value equals the optimum, found by

exhaustive search (top), and, over the instances in which it is below, the distribution of the ratio of the greedy value to the optimum (violins, with quartiles, medians and the lowest ratio, and the number of such instances at the foot of each); the dashed line marks the one-half guarantee. The exhaustive check over these instances (Results) shows that no ratio falls below the bound; it is not evidence that the bound is tight. (b) Gap OPT − GDB against the risk p_k of the first item that does not fit, over the instances in which the greedy value is below the optimum (hexagonal bins, logarithmic colour scale); the dashed diagonal is the bound of Theorem 1, the hatched region above it is excluded by the theorem, and the dotted line is the largest observed ratio; the solid and dashed curves are the median and the 90th percentile of the gap in bands of p_k one twentieth wide, and the margins give the distributions of p_k (top) and of the gap (right). (c) The family of three items whose ratio approaches one half, drawn with B = 10. (d) Its ratio against the budget, with the four budgets computed.

Optimality and guarantee. On the first-quarter risks, the greedy value was within 0.003% of the exact optimum at all five budgets, and it equalled the optimum at four of them (Table X, S2). Table XI gives, for each budget, the greedy value, the optimum from the dynamic programming and from a mixed-integer solver (HiGHS, through SciPy; a post hoc check), the linear relaxation and the risk of the first drive that does not fit. The dynamic programming and the solver agree to six decimals at four budgets. At 10% of capacity the solver returned a lower value; the dynamic programming solution is feasible and better, so the solver stopped short of the optimum with its default tolerance, and the dynamic programming method is taken as exact. The gap OPT − GDB is bounded by the risk of the first drive that does not fit (Theorem 1). At every budget it was within this bound, and the largest ratio of the gap to that risk, over the random instances, was 0.886 (Figure 10b). The random instances follow the settings of Table IV; the 19,126 instances with a positive optimum were kept. On these instances the lowest ratio of the greedy value to the optimum was 0.553, above the guaranteed one-half (Figure 10a), and the greedy value equalled the optimum in 68% of them, with median ratio 1.00 and fifth percentile 0.84. The one-half bound is approached by a family of three items (Figure 10c and d): item A has cost one and risk one, item K has cost B and risk one, and item M has cost B − 1 and risk (B − 1)/B − η, with η = 10-4. The greedy rule takes A and is then stopped by K, which does not fit after A; the optimum takes A and M, which fill the budget (Equation (27)):

$$\mathrm{GDB} = 1, \qquad \mathrm{OPT} = 2 - \frac{1}{B} - \eta, \qquad \frac{\mathrm{GDB}}{\mathrm{OPT}} = \frac{1}{2 - 1/B - \eta} \to \frac{1}{2} \quad \text{as } B \to \infty,\ \eta \to 0 \tag{27}$$

Equation (27) gives the ratios 0.6000, 0.5263, 0.5025 and 0.5003 for budgets of 3, 10, 100 and 1,000, so the bound is tight in the limit.

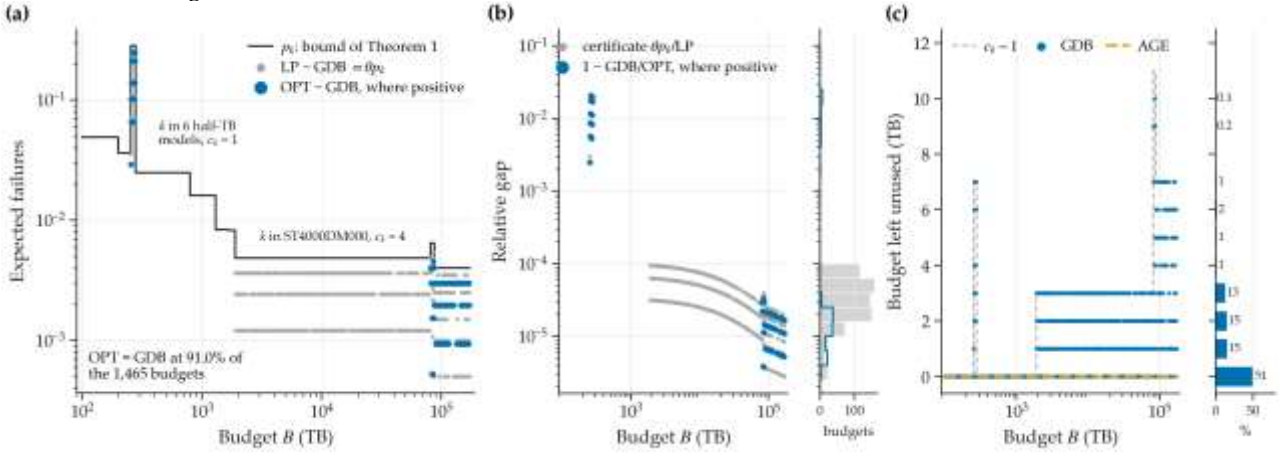


Figure 11. The guarantee of Theorem 1 on the fleet, first quarter only (no outcome is used), at 1,465 budgets from 100 TB to 20% of capacity. (a) Gap OPT − GDB where it is positive, below LP − GDB = θp_k and the risk p_k of the first drive that does not fit: OPT − GDB ≤ LP − GDB ≤ p_k, the chain of Theorem 1; the labels name the model of drive k on the longest stretches of the bound. (b) Relative gap 1 − GDB/OPT where it is positive, and the certificate θp_k/LP of Equation (13), which bounds it; the margin counts budgets per band (grey: certificate; blue: gap). (c) Budget left unused by GDB, which stops at drive k, and by AGE, which continues the scan, with the bound c_k − 1; the margin gives the percentage of budgets at which GDB leaves each amount unused. OPT comes from one dynamic programming over the groups of identical drives, whose

vector D[b] gives the optimum at every budget b.

Figure 11 extends Table XI from five budgets to 1,465 budgets from 100 TB to 20% of capacity, with one dynamic programming whose vector D[b] gives the optimum at every budget. The greedy value equals the optimum at 91.0% of them. Where it does not, the gap stays below pk, at most 0.87 of it, as Theorem 1 requires. The relative gap reaches 2.1% and exceeds 0.1% only at the 16 budgets from 255 to 277 TB; above 1,000 TB it is at most 0.0033% (Figure 11a, b). The stopping rule leaves at most 10 TB of the budget unused, whereas AGE, which continues the scan, uses the whole budget at every budget computed (Figure 11c).

Table XI. Greedy value, exact optimum, linear relaxation and the bound of Theorem 1 at the five budgets (first-quarter risks).

Budget	B (TB)	GDB	OPT, dynamic programming	OPT, HiGHS	LP	OPT − GDB	pk
1%	8,629	46.621728	46.621728	46.621728	46.624146	0	0.004837
2%	17,259	57.059740	57.059740	57.059740	57.059740	0	0.004837
5%	43,148	88.364104	88.364104	88.364104	88.365313	0	0.004837
10%	86,296	137.004015	137.008014	137.004015	137.008860	0.003999	0.006461
20%	172,592	180.184034	180.184034	180.184034	180.184534	0	0.003999

Values are expected failures over the horizon from the first-quarter risks; no outcome is used. A difference OPT − GDB below 10-9 is shown as 0. HiGHS ran with its default tolerance, as a post hoc check. Every row satisfies OPT ≤ LP, GDB ≥ LP/2 and OPT − GDB ≤ pk.

Table XII. Post hoc: drives replaced at the primary budget, by real capacity class.

Capacity class (TB, real)	Drives in population	GDB drives	GDB failures	GDB billed TB	AGE drives	AGE failures	AGE billed TB
0.5 (billed as one TB)	1,909	1,867	35	1,867	252	5	252
4	35,293	10,314	59	41,256	9,025	44	36,100
5	45	0	0	0	44	0	220
6	1,123	0	0	0	1,096	3	6,576
8	25,304	3	0	24	0	0	0
Other classes	42,544	0	0	0	0	0	0

Failures and drives are counted in the second quarter for the drives each policy replaced. Billed TB use the integer cost of the plan. Classes with no drive replaced by either policy are included in Other classes.

Composition by drive class (post hoc). The difference is concentrated in the half-terabyte class. The six half-terabyte models that the greedy policy replaces in full, billed as one TB each, contribute 35 of the 94 failures of the greedy policy and 5 of the 52 of the age policy (Table XII). The greedy policy replaces every one of their 1,867 drives, whereas the age policy replaces 252 drives of the whole nine-model class, some of them of WDC WD5000BPKT, a model that the greedy policy does not select. The first-quarter estimates identify this class as high-risk: in the second quarter its population failed at 18.3 failures per 1000 billed TB, against 0.53 for the rest of the population. Outside this class the greedy rate is higher: the greedy policy captured 59 failures in

41,280 TB (1.43 per 1000 TB) and the age policy 47 in 42,896 TB (1.10). No interval is computed for these subsets. Removing the class and selecting again with the same budget of 43,148 TB gives a difference of 0.32 per 1000 TB, interval −0.14 to 0.81, which the rule does not classify as positive. This is not the difference of the two rates above, because both selections change. The billing convention does not change the primary difference: with real capacities, it is 0.97 per 1000 TB, interval 0.43 to 1.52.

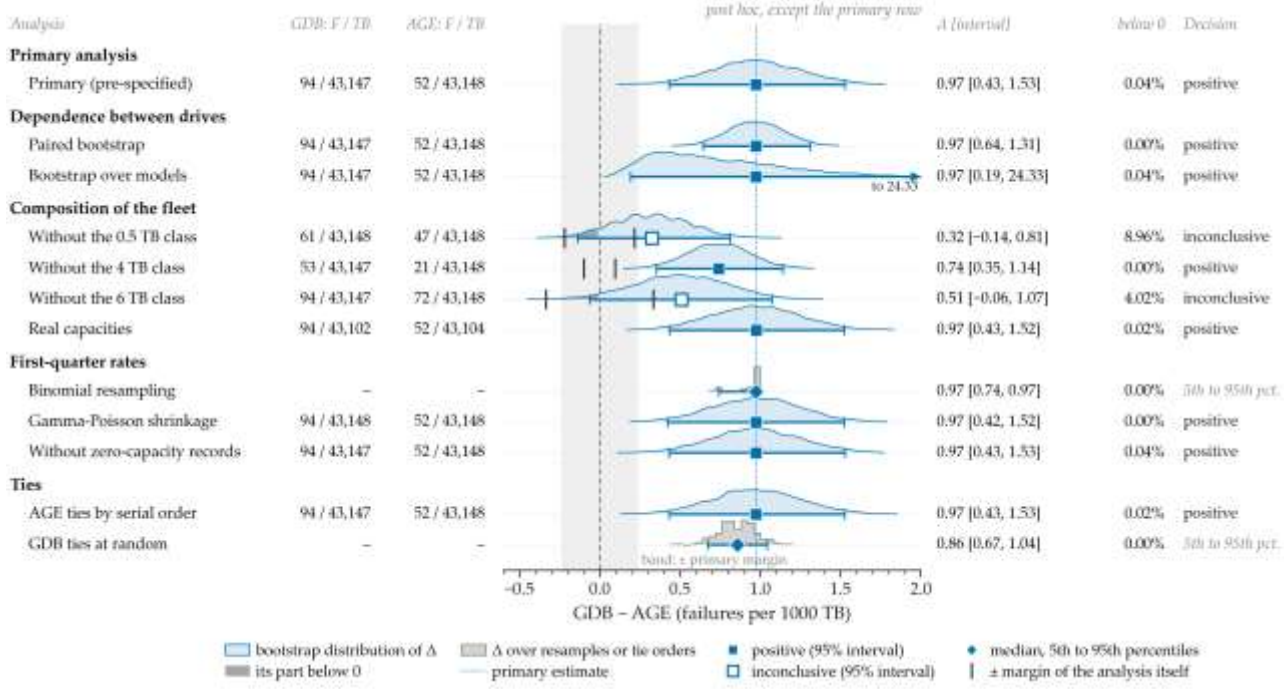


Figure 12. Sensitivity of the primary difference (post hoc, except the primary row). Each row changes one element of the primary analysis and keeps its rule. The paired bootstrap (Equation (19)) resamples the 15,059 drives selected by either policy, and the model-block bootstrap the 17 models. The class rows remove every drive of one real capacity class and select again with the same budget of 43,148 TB. Real capacities use the capacity in terabytes with decimals for the cost and the denominator. The first-quarter rows resample each model's failures (500 binomial resamples), shrink the rates with a gamma-Poisson prior with alpha equal to two, or remove the records with non-positive capacity from the counts. The last two rows break the ties of AGE by serial order and those of GDB at random (200 orders). Squares are estimates with 95% intervals, filled when the rule of the plan classifies the result as positive and open when it classifies it as inconclusive; diamonds are medians with the 5th to 95th percentiles. The grey band spans plus and minus the relevance margin of the primary analysis; each row is classified with the margin of its own analysis, 20% of its AGE rate, marked by short bars in the class rows. The dotted line is the primary estimate. Above each row is its distribution: the bootstrap of the difference (5,000 resamples, repeated post hoc with the seeds of each analysis, which reproduce every interval; the part below zero in dark grey) or, in the percentile rows, a histogram over the resamples or tie orders, each scaled to its own height. The columns give the failures and billed TB of each selection, the estimate and interval, the share below zero and the decision. The upper bound of the model-block interval, 24.33, lies beyond the axis.

Robustness (post hoc). Figure 12 shows the analyses that test the primary difference, and Figure 13 the difference without each model. The paired bootstrap, which accounts for shared drives, gives 0.64 to 1.31 per 1000 TB, narrower than the primary interval, so the primary interval is conservative with respect to this dependence. The model-block bootstrap is very wide, because a large part of the greedy selection comes from a few models: its interval runs from 0.19 to 24.33 per 1000 TB, and its lower bound is below the margin (0.24), so its positive decision rests on the rule's wording. Removing any single model keeps the point estimate positive; the smallest, 0.51 per 1000 TB, comes from removing ST6000DX000, which holds 1,023 of the 1,123 drives of the six-terabyte class and gives the same difference as removing the whole class. It is also the only removal whose interval, computed with the bootstrap of the plan, includes zero: −0.05 to 1.10 per 1000 TB. Removing the whole class of half-terabyte drives, or the six-terabyte class (0.51 per 1000 TB, interval −0.06 to 1.07), makes the result inconclusive, whereas removing the four-terabyte class leaves it positive (0.74, 0.35 to 1.14). Shrinking the first-quarter

rates with a gamma-Poisson prior gives 0.97 per 1000 TB (interval 0.42 to 1.52), and counting only the records with positive capacity gives 0.97 (0.43 to 1.53). Resampling the first-quarter failures changes the selection in most resamples; the 5th and 95th percentiles of the difference over the resamples are 0.74 and 0.97 per 1000 TB, and the difference is at least the margin in 100% of them. Breaking the ties of the age policy by serial order gives the same failures and the same budget, and an interval that agrees to two decimals. Random tie-breaks within equal densities put the observed difference at the 88th percentile of 200 random orders. That is compatible with chance: the serial order is not shown to favour the greedy policy, and the sign does not change. The selected drives themselves are compared in Table XIII: the two other horizons and the rates without zero-capacity records select exactly the same drives, and the shrunken rates leave out 9 drives and add 46. Random tie-breaks replace a median of 4,938 of the 12,184 drives of GDB with other drives of the same groups of identical drives, so the risk and the cost of the selection do not change. With five seeds, the lower bound of the primary interval ranges from 0.41 to 0.44 and the upper bound from 1.49 to 1.53 (post hoc), so the Monte Carlo error is in the second decimal.

Table XIII. Stability of the selected drives under each variant of the analysis (first-quarter data only).

Variant	Policy	Drives selected	Jaccard index with the primary selection	Primary drives left out	Drives added
First-quarter resampling (500 resamples)	GDB	varies	0.833, 0.999 and 1.000 (percentiles)	none in 33.2% of the resamples	–
Shrunken first-quarter rates	GDB	12,221	0.9955	9	46
Rates without zero-capacity records	GDB	12,184	1	0	0
Horizon of thirty days	GDB	12,184	1	0	0
Horizon of one hundred eighty-two days	GDB	12,184	1	0	0
Random tie-breaks within groups (200 orders)	GDB	12,184	0.419, 0.423 and 0.428 (percentiles)	4,880, 4,938 and 4,990 (percentiles)	as many as left out
AGE ties by serial order	AGE	10,417	0.9983	9	9
Real capacities	GDB	12,404	0.9823	0	220
Real capacities	AGE	10,438	0.9980	0	21

The Jaccard index is the number of drives in both selections divided by the number in either; 1 means the same drives. Percentiles are the 5th, the median and the 95th over the resamples or orders. These comparisons use no outcome, because a selection depends on the first quarter only.

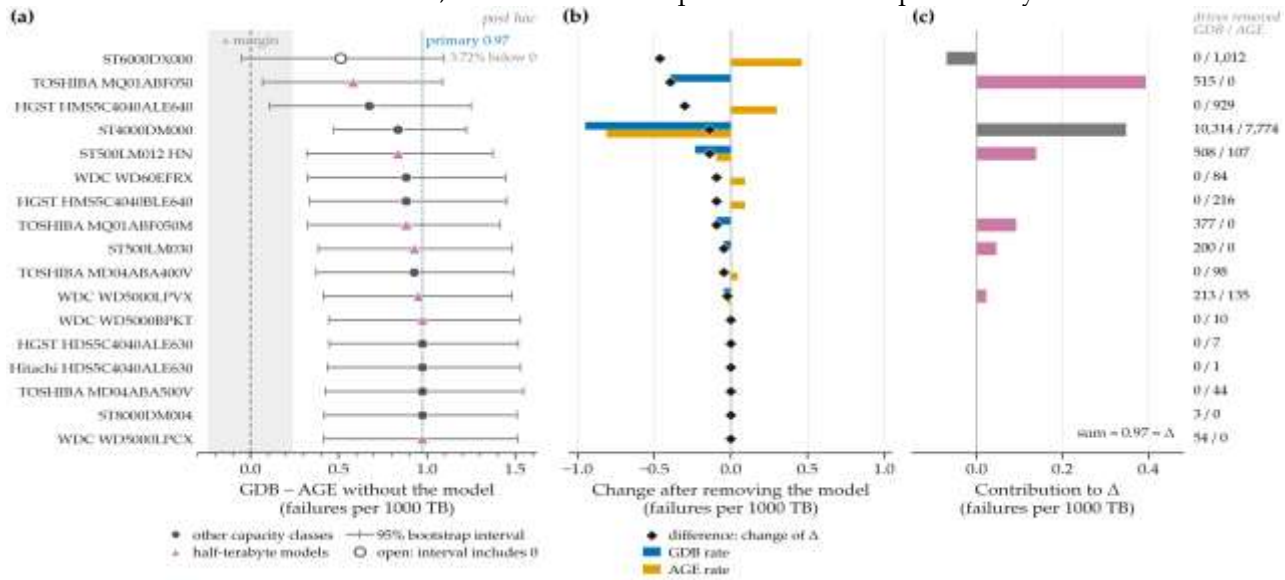


Figure 13. Difference without each model (post hoc). (a) GDB minus AGE after removing all drives of one of the 17 models that either policy selects, with both policies selecting again under the same budget, with its 95% bootstrap interval (the procedure and the seed of the plan); an open marker is an interval that includes zero, with the share of resamples below zero beside it. Triangles mark the half-terabyte models. The grey band spans plus and minus the relevance margin, and the dotted line is the primary difference. (b) The change of the capture rate of each policy after the removal; their difference is the change of the difference in (a). (c) The contribution of each model to the primary difference: the failures of the model in each selection divided by the billed terabytes of that selection, GDB minus AGE, which sum to the difference; the right column gives the drives of the model that each policy had selected.

Removing a model changes both selections (Figure 13b). Without ST6000DX000, which only the age policy selects, the age rate rises by 0.46 per 1000 TB, as AGE selects other drives in place of its 1,012 drives of that model; without HGST HMS5C4040ALE640 it rises by 0.30. Neither model adds to the difference in the primary selections: their contributions are −0.07 and 0.00 per 1000 TB (Figure 13c). The contributions sum to the difference, and the largest come from TOSHIBA MQ01ABF050 (0.39) and ST4000DM000 (0.35).

Table XIV. Baselines without density (post hoc, except the GDB and AGE rows).

Policy or baseline	Drives	Failures (Q2)	Per 1000 TB	Per 1000 drives	GDB minus baseline, per 1000 TB (95% interval)	Decision
GDB (density)	12,184	94	2.18	7.72	reference	–
Risk without cost, stopping at the first drive that does not fit	5,308	64	1.48	12.06	0.70 (0.13 to 1.26)	positive
Risk without cost, skipping drives that do not fit	5,309	64	1.48	12.06	0.70 (0.12 to 1.25)	positive
Model rate, age as tie-break, skipping	5,309	74	1.72	13.94	0.46 (−0.13 to 1.06)	inconclusive
GDB with the oldest drives within each group	12,184	103	2.39	8.45	−0.21 (−0.84 to 0.43)	inconclusive
AGE (oldest first)	10,417	52	1.21	4.99	0.97 (0.43 to 1.53), GDB minus AGE	positive

Each baseline uses the same budget as GDB. Per 1000 TB is the comparison metric; per 1000 drives is shown because the baselines replace different numbers of drives. Intervals are bootstrap percentile intervals. The GDB and AGE rows repeat the primary comparison; all other rows are post hoc.

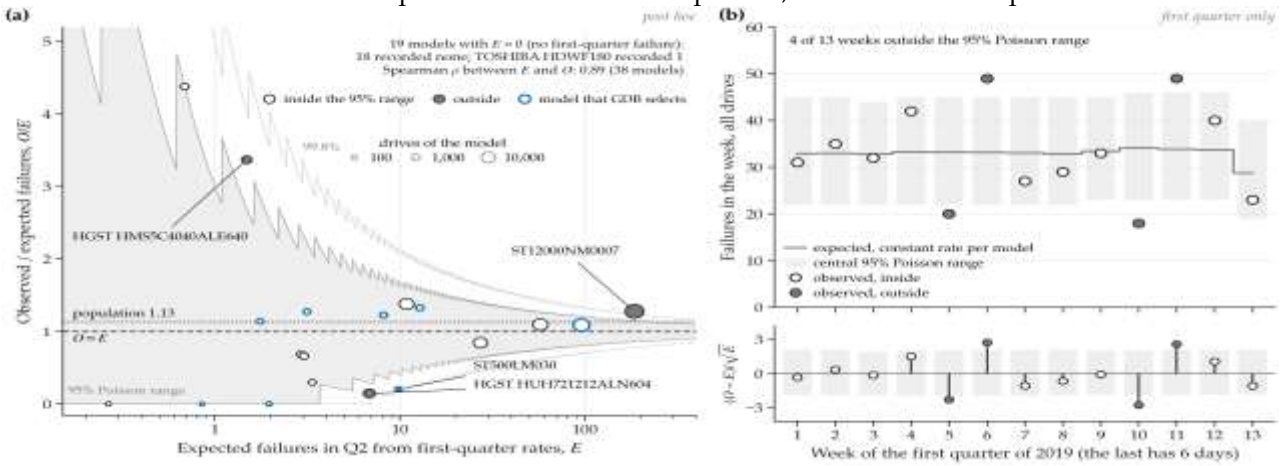


Figure 14. Calibration of the risk estimates. (a) By model (post hoc), as a funnel plot. Ratio of observed to expected second-quarter failures, O/E, against the failures E expected from the first-quarter rates, for the 19 of the 38 models of the population with positive expected failures; marker area grows with the square root of the model's drives. The grey region is the central 95% range of a Poisson count whose mean is the expected

value (Equation (22)), and the dotted curves the central 99.8% range; both ignore the uncertainty of the first-quarter rates. Filled markers and labels mark the models outside the 95% range, and blue edges the models GDB selects. The dashed line is O = E and the dotted line the population ratio; the note gives the Spearman correlation of E and O. The 19 models without first-quarter failures have zero expected failures and are not drawn; 18 of them recorded none, as the note in the panel says. (b) Constant rates within the first quarter (no outcome is used): failures of all drives in each week against those expected from the constant rate of each model (line), with the central 95% range of a Poisson count with that mean (grey); filled markers lie outside it. The strip below gives the Pearson residuals, (O − E)/√E, with the same Poisson range. The last week has six days.

Baselines and calibration (post hoc). In this sample, ranking drives by first-quarter risk without the cost selected fewer drives for the same budget and captured 64 failures against 94; the difference is 0.70 per 1000 TB (95% interval 0.13 to 1.26; Table XIV). Against a ranking by model rate with an age tie-break, the greedy policy is ahead by 0.46 per 1000 TB, but the interval includes zero, so the density rule is not statistically distinguishable from this simpler ranking in this sample. Within each group, keeping the oldest drives gave 103 failures against 94 for the greedy choice, a difference that the interval does not separate from zero, so the data neither show nor exclude that the age of a drive carries information within a group. The first-quarter risks gave expected counts somewhat below the second-quarter failures: 88.36 expected for the greedy policy against 94 observed (ratio 1.06), and 42.08 against 52 for the age policy (ratio 1.24, interval 0.92 to 1.62). Across the whole population, the expected count is 433.9 against 492 observed (ratio 1.13, 95% interval 1.04 to 1.24, holding the first-quarter risks fixed). The first-quarter rates rank the models in a similar order to the second quarter: their rank correlation with the second-quarter rates across 38 models is 0.67. Individual models are less well calibrated: of the 19 models with a positive expected count, 4 fall outside the range of Poisson variation around it (Figure 14), 2 above it and 2 below, and of the 19 models without first-quarter failures, whose expected count is zero, 1 recorded a failure. ST500LM030, for example, was expected to fail 9.8 times among its 200 drives, and 2 of them failed. The constant rate per model behind Equation (1) does not hold exactly within the first quarter itself: the weekly failures of the fleet vary more than constant rates allow, with 4 of the 13 weeks, and 2 of the 7 blocks of two weeks (the last of six days), outside the central 95% range of a Poisson count with the expected mean, where 5% would fall outside by chance (Figure 14b). The analysis does not separate changes of the hazard from changes in when failures are recorded.

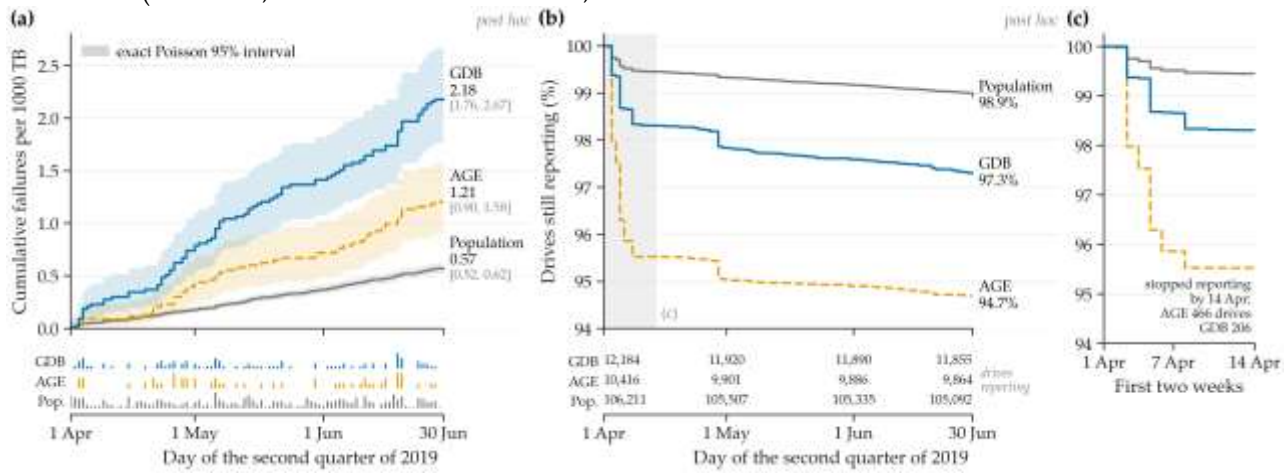


Figure 15. Failures and censoring over the second quarter (post hoc). (a) Cumulative failures per 1000 billed TB among the drives each policy selected and in the population, dating each failure by the drive's last record, with exact Poisson 95% intervals (shaded) and the final value with its interval at the right; the strip below marks the failures of each day, with the height of each mark proportional to the day's count within its row. (b) Percentage of the drives of each selection, and of the population, that still report on each day; the grey strip marks the first two weeks, shown in (c). The table below (b) gives the drives still reporting on the first day of each month and on the last.

Censoring (post hoc). Some drives stopped reporting before the end of the second quarter. Of the greedy policy's 12,184 drives, 329 had their last record before 30 June, and 94 of those had failed. For the age policy the numbers are 552 of 10,417, with 500 drives that stopped without a failure. Per 1000 observed drive-days, the failure rates are 0.087 for the greedy policy and 0.058 for the age policy. Censoring

falls more heavily on the age policy. If its censored drives had failed, its count would be higher and the difference would shrink; per 1000 TB, the two capture rates (2.18 and 1.21) would be equal if 42 of its 500 censored drives had failed. The binomial test

$$\mathrm{rep}_S(t) = \frac{1}{|S|}\sum_{i\in S} \mathbf{1}\,[\tau_i \geq t], \qquad F_S(t) = \sum_{i\in S} y_i\ \mathbf{1}[\tau_i \leq t] \tag{28}$$

In Equation (28), τi is the last day with a record of drive i in the second quarter: a failure is dated on that day, and a drive is still reporting on day t if τi ≥ t; τi = 0 for a drive without any record in the second quarter. The gap between the capture rates opens in April and widens through the quarter: at the end of April, May and June the greedy selection had captured 0.74, 1.41 and 2.18 failures per 1000 TB, and the age selection 0.39, 0.72 and 1.21 (Figure 15a). Most of the drives that stopped reporting did so in April: at the end of April, 95.0% of the age selection and 97.8% of the greedy selection still reported, and at the end of June, 94.7% and 97.3% (Figure 15b). Table XV gives the models of the drives that stopped reporting without a failure. Of the 500 such drives of the age policy, 461 stopped in the first week of April, and 492 belong to five models: for example, all 44 drives of TOSHIBA MD04ABA500V that it selected stopped, and 81 of the 84 drives of WDC WD60EFRX. Of the 235 such drives of the greedy policy, 193 also stopped in the first week, and 223 are ST4000DM000 drives.

conditional on the total failures is not valid here, because the selections share drives. Figure 15 shows when the failures and the losses of drives occurred, with the curves of Equation (28):

Table XV. Post hoc: drives that stopped reporting without a failure, by model, for each selection.

Model	GDB: stopped	GDB: share of its drives of the model	AGE: stopped	AGE: share of its drives of the model	Median day of the last record
ST4000DM000	223	2.2%	212	2.7%	4
ST6000DX000	0	–	135	13.3%	2
WDC WD60EFRX	0	–	81	96.4%	2.5
TOSHIBA MD04ABA500V	0	–	44	100.0%	4
HGST HMS5C4040ALE640	0	–	20	2.2%	2
ST500LM012 HN	5	1.0%	5	4.7%	2
WDC WD5000LPVX	3	1.4%	3	2.2%	5
Other models (3)	4	–	0	–	–
All models	235	1.9%	500	4.8%	–

A drive stopped reporting if its last second-quarter record is before 30 June; drives that failed are not counted, and the one drive of AGE without any second-quarter record is not included. Models with no drive that stopped are not listed, and the share is over the drives of that model that the policy selected. Day 1 is 1 April 2019. The median day is over the population drives of the model that stopped.

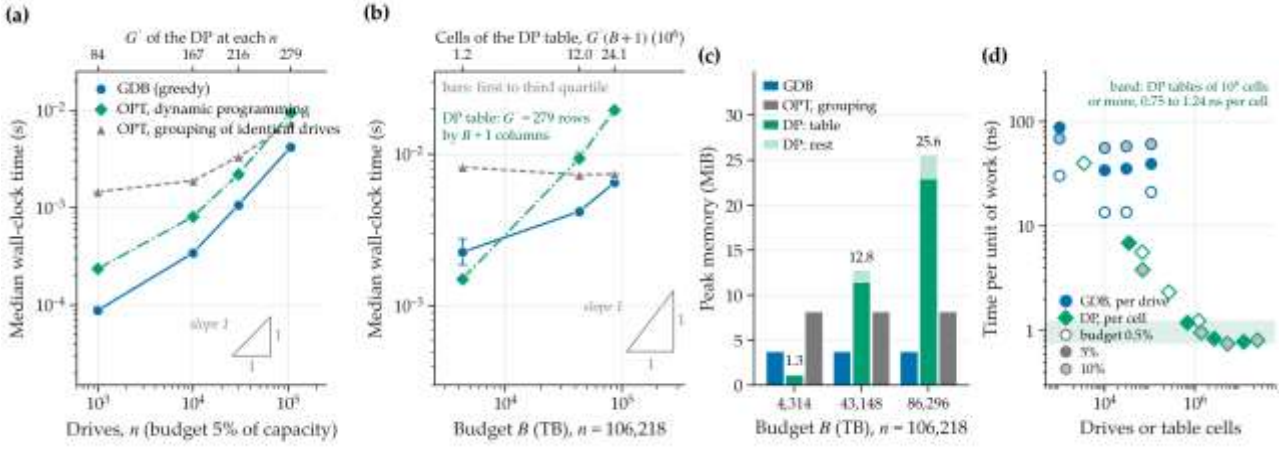


Figure 16. Measured cost at the 5% budget against the number of drives (a), and at three budgets for the full population: wall-clock time (b) and peak memory (c), whose bar for the DP splits into its decision table, G'(B + 1) bytes, and the rest. Median wall-clock times; error bars span the first and third quartiles around the

median. Each panel shows GDB, the dynamic programming (DP) of OPT and, separately, the grouping of identical drives that precedes it. All are measured from the same prepared inputs. In (a) and (b), the right triangles drawn in thin grey lines mark a slope of one on the logarithmic axes, that is, cost proportional to the number of drives or to the budget. The top axes give G' at each size (a) and the cells G'(B + 1) of the DP table at each budget (b). (d) Median time per unit of work in the twelve measured combinations: per drive for GDB and per table cell for the DP (fill: budget); a ratio of measurements, not a fitted model, flat where time is proportional to the work.

Table XVI. Cost and peak memory at the 5% budget, one thread.

Drives	GDB time (ms)	DP time (ms)	Grouping time (ms)	GDB memory (MiB)	DP memory (MiB)	Grouping memory (MiB)
1,000	0.09	0.24	1.45	0.04	0.05	0.08
10,000	0.34	0.81	1.90	0.38	0.78	0.65
30,000	1.06	2.20	3.33	1.07	2.89	2.19
106,218	4.19	9.40	7.31	3.76	12.81	8.20

Medians of five repetitions, with the same prepared inputs for both methods. Groups of identical drives: 38; G′ = 279. The table of G′(B + 1) = 279 × 43,149 one-byte entries accounts for 11.48 MiB of the 12.81 MiB peak of the dynamic programming; the remaining 1.33 MiB is not itemized.

Table XVII. Measured cost of further stages of the pipeline, for the full population and the primary budget, one thread.

Stage	Median time (s)	First to third quartile (s)	Peak memory (MiB)
Reading the first-quarter archive (9,577,046 records)	27.5	27.5 to 27.7	1,334.6
Population and frozen risk table	0.851	0.848 to 0.860	255.9
AGE at the primary budget	0.026	0.025 to 0.026	1.7
RANDOM, one selection	0.009	0.009 to 0.009	0.5
Bootstrap of the primary difference (5,000 resamples)	0.402	0.395 to 0.410	0.6
V1: 1,000 random selections	8.62	8.61 to 8.74	0.5
V2: 200 permutations of the risks	1.34	1.34 to 1.35	5.0
First-quarter resampling, 500 resamples (post hoc)	5.74	5.72 to 5.78	13.1
Manuscript: content, equations and the DOCX file	1.71	1.71 to 1.72	–

Medians of 3 repetitions after one warm-up. Peak memory is measured with tracemalloc in separate passes and counts the memory allocated through Python and NumPy; it is not measured for the manuscript, which runs in its own process. GDB, the dynamic programming and the grouping are in Table XVI. The bootstrap, V1, V2 and the resampling are timed on the per-drive summary of the second quarter derived after its first reading, so their timing is post hoc. The second quarter is not read again to time it, because that would be a new reading of the confirmatory data. Stages not listed, such as the guarantee check on random instances or the intervals of the AUC, were not timed.

Cost. Table XVI and Figure 16 give the measured times. At 106,218 drives and a budget of 5%, the median time was 4.19 ms for the greedy policy, 9.40 ms for the dynamic programming and 7.31 ms for grouping the identical drives. The greedy policy was faster than the exact method including its grouping step at every size and budget we measured. With four sizes and three budgets, the timings do not support a fitted exponent, and we report none. The greedy policy needs 3.76 MiB at the full size, rising with the number of drives from 0.04 MiB at 1,000 drives, against 12.81 MiB for the dynamic programming at the full size. These times describe this implementation in Python, and they reflect the implementation as much as the algorithms. Table XVII gives further stages for the full population. Reading the first-quarter archive took 27.5 s, with a peak of 1,334.6 MiB, and the random control V1, the slowest of the analysis stages measured, 8.62 s, against 4.19 ms for the greedy policy: the selection itself is a small part of the cost of the study.

Exploratory (post hoc). Figure 17 shows the failure rate per billed terabyte of every model that either policy replaced. Five of the six half-terabyte models have the five highest rates among the models the greedy policy replaced; the sixth, WDC WD5000LPCX, recorded no failure. For that class as a whole, the rates per billed terabyte are close for the two policies: 18.7 for the greedy policy and 19.8 for the age policy. ST4000DM000, the model that both policies replaced most, has 1.43 failures per 1000 TB under the greedy policy and 1.41 under the age

policy, so for this model the rates are nearly equal.

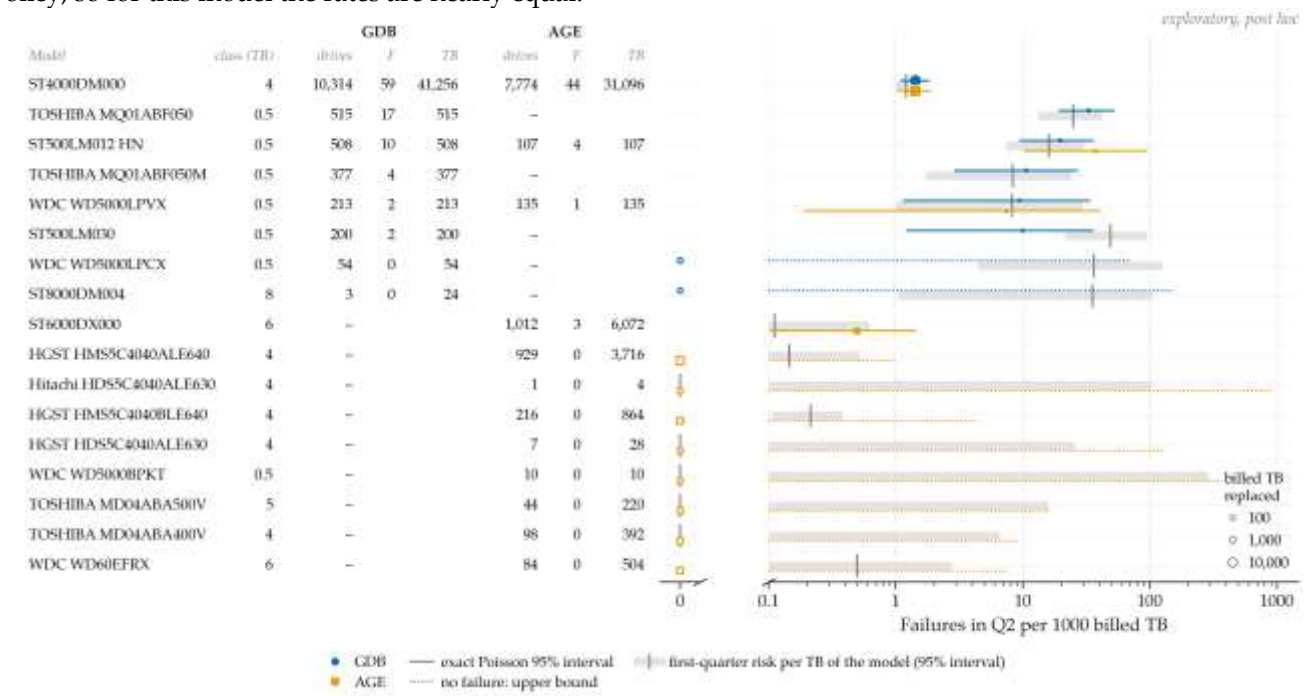


Figure 17. Exploratory (post hoc). Failures in the second quarter per 1000 billed TB for each model that either policy replaced, with exact Poisson 95% intervals; marker area grows with the billed terabytes replaced. Zero rates are drawn on the left strip, with their upper bounds dotted. Behind each row, a grey bar and tick give the first-quarter expectation for the model, its risk per billed TB over the horizon, per 1000 TB, with the interval of its rate (tick at zero without first-quarter failures). The table gives each model's capacity class and, per policy, the drives, failures (F) and billed TB. The figure tests no hypothesis.

4. DISCUSSION

The main finding is that, on these data, ranking drives by estimated failure risk per terabyte captured more failures out of sample than replacing the oldest drives, at the primary budget under the rule fixed in advance, and at the other budgets of the primary curve. Post hoc, the finding depends on the composition of the fleet. Most of the difference comes from six half-terabyte models, which the first-quarter estimates identify as high-risk and which the greedy rule replaces in full. Without the half-terabyte class the point estimate stays positive, but its interval includes zero, and the rule classifies the result as inconclusive. We therefore present the primary result as positive under its rule and, post hoc, as dependent on the mix of models.

Post hoc, the density rule is not statistically distinguishable from a simpler ranking by model rate with an age tie-break. This comparison is the one that bears on whether the knapsack structure contributes. The advantage over ranking by risk alone is positive, but its interval is wide, and the advantage over ranking by model rate is not detected. What the data support is that model-level rates predict second-quarter rates across models, and that the policies differ in the models they select; for ST4000DM000, the model both policies replaced most, the two rates are nearly equal. This fits the field evidence that age and model are hard to separate: Pinheiro et al. (2007) report that the higher baseline failure rate of drives three and four years old is more strongly influenced by the underlying reliability of the particular models in that vintage than by disk drive aging effects, and that drive model factors are strongly mixed into their age data. In this fleet, a model's median age does not track its first-quarter rate (Figure 3b). The guarantees of Theorem 1 concern the optimum for the risk estimates they are given; they do not measure out-of-sample performance, and the closeness of the greedy value to the optimum should not be read as such a measure.

The data come from one operator's fleet during the first half of 2019 (two quarters, from January to June), and the policies were judged on the failures of the following quarter. The failures captured per terabyte are not the loss of stored objects. That loss depends on replication or erasure coding, on how quickly a failed drive is rebuilt, and on whether failures are correlated across replicas. Latent sector errors, which go undetected until the corresponding disk sectors are accessed (Bairavasundaram et al.,

2007), are outside the scope of this analysis, which counts drive-level failures only. Rosenthal (2010) argues that storage demands continually push beyond the capabilities of systems implementable at an affordable price. A replacement budget is one such limit; the present analysis measures one side of that budget and does not measure preservation outcomes.

Table XVIII compares the data, aims and methods of the closest studies we read with those of this study. The horizon of 91 days is much longer than the lead time of ten to fifteen days that SMART-based analysis achieved in Botezatu et al. (2016). The decision here is a quarterly budget decision rather than a last-minute one, and this analysis does not test whether its results transfer between the two problems. The sensitivity to the horizon is also limited: in a post hoc check, the selected drives do not change for thirty or one hundred eighty-two days, so the horizon analysis is not a test of robustness, and it is reported as such.

Table XVIII. The closest studies we read, and this one: data, aim, method and what each reports.

Study	Data	Aim	Method	What is reported
Pinheiro et al. (2007)	More than one hundred thousand disk drives deployed across Google's services	Failure statistics and the parameters that correlate with failures	Failure rates and their correlation with SMART parameters and other factors	Annualized failure rates; models based on SMART parameters alone are unlikely to predict individual failures
Schroeder and Gibson (2007)	Field replacement data from high-performance computing and internet services sites; about one hundred thousand disks, some for an entire lifetime of five years	Field replacement compared with the MTTF of the data sheets	Analysis of disk replacement records	Annual replacement rates typically above one percent, two to four percent common and up to thirteen percent; rates grow with age
Murray et al. (2005)	Attributes monitored internally by individual drives	Predicting hard-drive failure	Multiple-instance naive Bayes, support vector machines, clustering and nonparametric tests	Failure-prediction performance at low false-alarm rates
Botezatu et al. (2016)	More than thirty thousand disks from two manufacturers, monitored over seventeen months	Predicting disk replacement	SMART parameters selected by statistical techniques, then a predictor	The need for replacement predicted ten to fifteen days in advance; separately, an accuracy of up to ninety-eight percent
This study	Backblaze, two quarters of 2019: 106,218 drives of 38 models	Which drives to replace under a fixed budget	0/1 knapsack: greedy risk density and the exact optimum, against an age policy, judged on the next quarter	Second-quarter failures captured per 1000 billed TB

Each entry for a published study paraphrases a passage of that study recorded in the repository; the list covers the studies we read and is not a systematic review.

Limitations. Table XIX lists the threats to the validity of the result, the analysis that examined each one and what it found. Four are not examined by any analysis here: the blindness of the replication, the SMART attributes other than the power-on hours, the restriction to one operator and one period, and an estimand that is not the benefit of replacing a drive; the conclusions are therefore limited to rates per model. The constant hazard per model is examined only within the first quarter, where the weekly failures vary more than it allows. Censoring is described but not corrected: drives that stop reporting count as having no failure, and the censoring analysis gives the failures that would equalize the two rates, not a bound on the effect. The post hoc analyses use the second quarter after its first reading, so they inform the design of a next study rather than confirm this one.

Table XIX. Threats to validity, how each was examined, and what was found.

Threat	How it was examined	What was found	Where
The result depends on a few drive classes or models	removing each class and each model; bootstrap over models	inconclusive without the half-terabyte or the six-terabyte class; the lower bound of the model-block interval is below the margin	Table XII; Figure 12; Figure 13
The two selections share drives	paired bootstrap	a narrower interval than the primary one	Figure 12
Drives stop reporting before the end of the quarter	described, not corrected: such drives count as having no failure; drives that stopped, by model and over time; failures that would equalize the rates	heavier on AGE, mostly in the first week of April; 42 of its 500 censored drives would have to fail, which is not a bound on the effect	Figure 15; Table XV
First-quarter rates are estimates, some from few drive-days	binomial resampling; gamma-Poisson shrinkage; stability of the selection	a positive difference in every resample; the same selection in a third of them	Figure 12; Table XIII
Choice of the horizon	horizons of thirty and one hundred eighty-two days	the same drives are selected, so this is not a test of robustness	Figure 5c; Table XIII
Ties, and a serial order linked to age	random ties; AGE ties by serial order; the oldest drives within each group	the observed difference at the 88th percentile of random ties; in ST4000DM000 the serial order favours older drives	Figure 12; Table VII; Table XIV
Half-terabyte drives billed as one terabyte	real capacities	the same difference to two decimals	Figure 12
Miscalibrated risk estimates	expected and observed failures, by policy and by model	ratio 1.13 in the population; 4 of 19 models outside the Poisson range	Table VIII; Figure 14
Analyses chosen after the outcome was read	labelled post hoc; second amendment	the primary result keeps its rule	Figure 1b; Table V
Constant hazard per model	weekly and two-week failures of the first quarter against the constant rates	4 of 13 weeks and 2 of 7 two-week blocks outside the central 95% Poisson range: more variation than constant rates allow	Figure 14b
The replication was not blind	not examined: it was written from the plan but compared with the pipeline after its first version	–	Methods
SMART attributes other than the power-on hours	not examined: no SMART-based predictor was evaluated	–	Discussion
One operator and two quarters of 2019	not examined	–	Discussion
The estimand is not the benefit of replacing a drive	not examined: it needs the failure rate of new drives and the cost of a failure	–	Methods

Where gives the element of the paper that reports each check.

5. CONCLUSIONS

On one operator's 2019 drive statistics, a greedy rule that ranks drives by estimated failure risk per terabyte captured more failures per terabyte than an age-based rule out of sample at the primary budget, under the rule fixed in advance. Post hoc, the advantage depends on drive classes: without the

high-risk half-terabyte class, or without the six-terabyte class, the difference is inconclusive. Post hoc, against a simpler ranking by model rate, the density rule shows no detectable advantage. The greedy value was close to the exact optimum on the estimated risks, Theorem 1 guarantees at least one half of the optimum, and, in a post hoc check, an independent solver agreed with the dynamic programming at the primary budget. Generalizing beyond this fleet requires data from other operators and periods, and a predictor based on SMART attributes evaluated with the same protocol.

DECLARATIONS

Data availability: The data are published by Backblaze (Backblaze, 2019). No licence is declared for them here, because the terms of the 2019 version of the data page were not checked for this study. The data manifest of the repository lists the SHA-256 digests of the two archives. The analysis code, the analysis plan with its two amendments, and the results are kept under version control and are available from the corresponding author.

Ethics: The study involved no human participants and used only publicly released, de-identified drive-level data.

Author contributions: (CRediT) Luis Eduardo Muñoz Guerrero: conceptualization, methodology, software, validation, formal analysis, investigation, data curation, writing – original draft, writing – review and editing.

Funding: This research received no external funding.

Competing interests. The author declares no competing interests.

REFERENCES

BACKBLAZE, Inc. (2019) Backblaze Drive Stats: Hard Drive Reliability Test Data, 2019 Q1 and Q2. Dataset. Available at https://www.backblaze.com/cloud-storage/resources/hard-drive-test-data.

BAIRAVASUNDARAM, L. N., GOODSON, G. R., PASUPATHY, S. and SCHINDLER, J. (2007) An analysis of latent sector errors in disk drives. In Proceedings of the 2007 ACM SIGMETRICS International Conference on Measurement and Modeling of Computer Systems, pp. 289–300. doi:10.1145/1254882.1254917.

BOTEZATU, M. M., GIURGIU, I., BOGOJESKA, J. and WIESMANN, D. (2016) Predicting disk replacement towards reliable data centers. In Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, pp. 39–48. doi:10.1145/2939672.2939699.

DANTZIG, G. B. (1957) Discrete-variable extremum problems. Operations Research, 5(2), pp. 266–288. doi:10.1287/opre.5.2.266.

MURRAY, J. F., HUGHES, G. F. and KREUTZ-DELGADO, K. (2005) Machine learning methods for predicting failures in hard drives: a multiple-instance application. Journal of Machine Learning Research, 6(27), pp. 783–816. Available at https://www.jmlr.org/papers/v6/murray05a.html.

PINHEIRO, E., WEBER, W.-D. and BARROSO, L. A. (2007) Failure trends in a large disk drive population. In Proceedings of the 5th USENIX Conference on File and Storage Technologies (FAST '07), pp. 17–29. Available at https://www.usenix.org/legacy/event/fast07/tech/full_papers/pinheiro/pinheiro.pdf.

ROSENTHAL, D. S. H. (2010) Bit preservation: a solved problem? International Journal of Digital Curation, 5(1), pp. 134–148. doi:10.2218/ijdc.v5i1.148.

SCHROEDER, B. and GIBSON, G. A. (2007) Disk failures in the real world: what does an MTTF of 1,000,000 hours mean to you? In Proceedings of the 5th USENIX Conference on File and Storage Technologies (FAST '07), pp. 1–16. Available at https://www.usenix.org/legacy/events/fast07/tech/schroeder/schroeder.pdf.