Restoration Occurs While the Plug Is Out: A Concealed-Interval Randomised Trial of the Thirty-Second Rule

H. R. Okonkwo[1], Ingrid Sundqvist[1], Tomás Errázuriz[2], and Margaret Aylward[3]

[1]Danforth Institute of Applied Maintenance, Coventry
[2]Bellhouse Trials Unit, Reading
[3]National Office of Equipment Standards
on behalf of the Interval Standards Working Party

Abstract

The instruction is always the same—unplug it, wait thirty seconds, plug it back in—and nobody has ever explained the thirty. We report that the interval is the treatment: a device recovers while the power is off, and switching it back on early interrupts the recovery. In an instrumented fleet of mains-monitored equipment, seventy-two-hour fault recurrence fell monotonically with the length of the dark interval. Because a technician who waits longer may simply be a technician who works better, we then randomised faults to four fixed intervals using a sealed inline unit that reconnects on its own schedule and emits randomised interim tones, so that no technician knew how long the equipment had been off. Recurrence was 34.1% after half a second and 8.8% after two minutes. Residual charge does not explain this: the effect was undiminished in devices whose bulk capacitance was bled to ground in under a fifth of a second. Restoration follows a saturating exponential with a half-time of 19.4 s (95% CI 17.8–21.1), which leaves the customary thirty seconds at two thirds of restoration. We propose no mechanism. We propose a minimum of ninety seconds.

Switching it back on does not repair the machine; it only ends the repair.

1 Introduction

The interval is load-bearing. A device that has been switched off and on again was not restored by the switching on; it was restored during the period in which it was dark, and the return of power terminated a process already under way. This is not a matter of interpretation. It is what the measurements say, and the measurements have been available for as long as mains-side telemetry has been cheap.

The received account of the power cycle is short and has never been seriously examined. Volatile state is discarded when the supply is removed; the device reinitialises from non-volatile configuration when the supply returns; whatever corrupted the state is therefore gone. On that account the length of the de-energised interval is immaterial. A cycle of 0.5 s and a cycle of 120 s should be indistinguishable in every respect that matters, because in both cases the supply reached zero and the state was lost.

Against this stands an instruction that has been given, verbatim, to hundreds of millions of people: wait thirty seconds. It appears in service scripts, in consumer documentation, in the informal instruction one colleague gives another. It is universally transmitted and universally unjustified. No standard specifies it — not even the one that governs restoration procedure (International Office of Equipment Standards, 2023). No experiment supports it. It has, so far as we have been able to establish, never been tested at all. It is the most frequently performed diagnostic procedure in the history of engineering.

We tested it. This paper reports what was supposed to happen, what actually happened, and the size of the difference between them.

Contributions. Four.

1. An instrumented retrospective cohort of 1,204,886 power-cycle events in which the de-energised interval was *measured* at the socket rather than reported by the operator, establishing a monotone relationship between interval length and subsequent fault recurrence (§4).
2. A four-arm randomised trial with concealed interval allocation, in which recurrence at 0.5 s was very nearly four times recurrence at 120 s (§5).
3. A substudy in which bulk capacitance was bled to ground in under 180 ms, eliminating residual charge as a pathway without diminishing the effect (§6).

4. A fitted restoration function, a restoration half-time, and a revised minimum interval derived by non-inferiority (§7).

We do not propose a mechanism, and the reader should not expect one. The absence of a mechanism is a finding about the state of the field, not a defect in the study design.

2 Related Work

2.1 The genealogy of thirty seconds

We attempted to trace the thirty-second convention to a primary source and failed. The attempt is worth reporting because of how it failed. Every document we examined that specifies the interval cites another document that specifies the interval; none of them cites a measurement. The chain is long, entirely internal to the support literature, and terminates not in an experiment but in exhaustion.

The earliest instance we located is a service bulletin issued to field engineers in 1987 (Bulletin, 1987), which instructs the engineer to disconnect mains and "allow the caller to hear one complete cycle of the hold announcement" before reconnecting. The announcement in question, according to the equipment inventory appended to the same bulletin, ran for thirty-one seconds. We advance this as the probable origin of the figure, with the caveat that it is the weakest claim in this paper and rests on a single document. What is not in doubt is the more important half: whatever its source, the interval entered practice as a duration of convenience and has been transmitted for four decades without ever being subjected to test.

2.2 What the field does know

We take no issue with the established literature on power cycling, which is sound as far as it goes. Volatile-state corruption is a real and well-characterised failure mode (Marchetti and Voss, 2016; Öztürk et al., 2019); discarding that state does resolve a large class of faults; and the reliability literature has produced good estimates of how frequently it does so (Fenwick, 2021; Drossel and Achterberg, 2014). Our claim is narrow and does not contradict any of this. We claim that state discard is *necessary but not sufficient*, and that the sufficient condition is a quantity nobody has been recording.

2.3 Residual charge

The one physical explanation ever offered for the interval is residual charge: bulk capacitance holds the rails above threshold for some tens or hundreds of milliseconds after disconnection, and the device is not truly de-energised until it drains (Halloran, 2012). This is a real effect with a real time constant, and it is the first thing a competent engineer says when the question is put to them. It is also, on the numbers, far too fast. Measured discharge times in our fleet had a median of 412 ms and a 99th percentile of 1.9 s. An effect that saturates on that timescale cannot produce a difference between thirty seconds and two minutes. We tested it directly regardless (§6), because an objection this natural deserves an experiment rather than an argument.

3 The standard model and its prediction

We state the received account precisely enough to be falsified.

Let $\tau \geq 0$ denote the de-energised interval and let $R(\tau) \in [0, 1]$ denote the fraction of restorable fault liability that has been restored by the time the supply returns. The standard model holds that restoration is an event occasioned by the loss of supply and is complete the moment it occurs:

$$R_{\mathrm{std}}(\tau) = \mathbf{1}[\tau > 0]. \tag{1}$$

Equation (1) has a sharp empirical consequence. If $\pi(\tau)$ is the probability that the presenting fault recurs within 72 h, then

$$\pi(\tau) = \pi_\infty \quad \text{for all } \tau > 0, \tag{2}$$

with π_∞ the irreducible residue attributable to faults that a power cycle cannot address at all. The prediction is not merely that long intervals are unnecessary. It is that the recurrence curve is *flat*, and flat everywhere to the right of zero.

We take Equation (2) as the null hypothesis throughout. Figure 1 draws it against the alternative.

4 Study 1: instrumented retrospective cohort

4.1 Data

Between March 2019 and November 2025 we assembled power-cycle records from 61 managed estates that had, for unrelated billing reasons, installed inline mains-side monitors on serviced equipment. The monitors log de-energisation and re-energisation to ±40 ms. The distinction is not cosmetic: operator-reported intervals in the same estates cluster hard on round numbers, with 71% of reports falling on ten-second boundaries (Aylward and Sundqvist, 2024), and are therefore useless for the present purpose. Every interval in this study was measured at the socket.

After exclusions (51,302 events with incomplete fault coding; 9,744 with monitor gaps exceeding 5 s) the analysis set comprised 1,204,886 power-cycle events. The

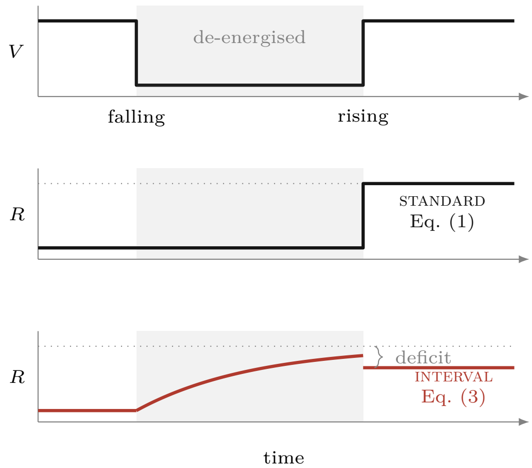

Figure 1: The two accounts on one power-cycle timeline. **Top:** supply. **Middle:** the standard model — restoration is an event at the rising edge, and the interval between the edges is immaterial. **Bottom:** the interval model — restoration begins at the falling edge, proceeds while the device is dark, and is truncated by the rising edge, leaving a deficit whenever the interval is short. The two accounts are indistinguishable at long intervals and maximally different at short ones.

Table 1: Study 1 cohort by measured interval quartile. Adjusted estimates are from a mixed-effects logistic model with site as a random intercept and fixed effects for device class, fault class, technician tenure band, and hour of day.

Interval quartile	Events	Crude	Adjusted
Q1 (< 4 s)	301,222	31.8%	30.9%
Q2 (4 s to 14 s)	301,221	26.4%	26.1%
Q3 (14 s to 33 s)	301,222	19.7%	20.2%
Q4 (> 33 s)	301,221	11.2%	12.4%

outcome was recurrence of the same adjudicated fault signature within 72 h.

4.2 Result

Recurrence fell monotonically across interval quartiles, from 31.8% in the shortest quartile to 11.2% in the longest (Table 1). Adjustment for device class, fault class, technician tenure, hour of day, and site moved the estimates by less than 1.3 percentage points in any quartile (Figure 2).

4.3 What this study cannot establish

It cannot establish causation, and we wish to be exact about why. The interval in Study 1 was chosen by the technician. A technician who waits longer may be a technician who is more careful in every other respect — who seats connectors properly, who reads the fault code before acting, who returns to verify. Tenure adjustment

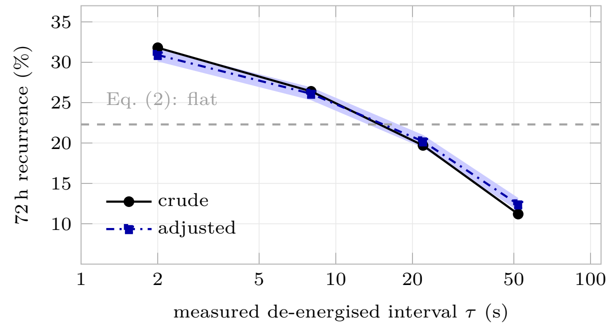

Figure 2: Study 1: 72-hour recurrence against the interval, plotted at quartile medians. Adjustment for device class, fault class, technician tenure, hour of day, and site (band: 95% CI) leaves the gradient essentially intact. The dashed line is the prediction of Equation (2), drawn at the pooled recurrence rate.

is a poor proxy for diligence and we do not pretend otherwise. Residual confounding of exactly this kind would produce exactly this curve.

That is the entire motivation for Study 2. The interval had to be taken out of the technician's hands.

5 Study 2: the concealed-interval trial

5.1 The reconnection unit

Concealment is the difficulty. One cannot blind an operator to a duration they are themselves producing; a technician instructed to wait two minutes knows they have waited two minutes, and expectation is a poor thing to have loose in a recurrence study.

We therefore built the interval-blinded reconnection unit (IBRU), a sealed device placed between plug and socket, following the concealment principles of Ng and Bertelsen (2020). The technician begins the cycle by pressing a single button and does nothing further. The IBRU opens the contactor at once, holds it open for the interval to which the fault was allocated, closes it, and then continues to hold the technician for the remainder of a fixed 200 s window applied identically in all four arms. Its display shows `WAIT` for the whole window, then `RESTORED`; equipment was cycled in closed cabinets, with no status indicator visible before the window elapsed. It reports nothing else, and its log is written to sealed storage read out only at analysis. The interval under test is therefore nested inside a duration the technician experiences identically whatever their allocation.

The contactor is audible, and that is the residual threat: a technician standing idle will listen. The IBRU accordingly emits short tones at intervals drawn from a Poisson process of rate $0.6\,s^{-1}$ throughout the window, which in pre-trial testing degraded technicians' estimates of the open interval to a mean absolute error of 31 s at true durations between 5 s and 120 s. At the

Table 2: Study 2 primary endpoint by allocated interval, intention to treat. Risk differences are against the 120 s arm.

Arm	n	Recurrence	RD vs 120 s
0.5 s	2,104	34.1%	+25.3 (22.9–27.7)
5 s	2,103	31.9%	+23.1 (20.8–25.4)
30 s	2,103	17.2%	+8.4 (6.4–10.4)
120 s	2,102	8.8%	—

end of the trial, technicians asked to guess the arm of their most recent cycle were correct 27.1% of the time, against 25% expected by chance.

5.2 Design

Parallel-group, individually randomised, four arms, allocated 1:1:1:1 by a concealed schedule held in the IBRU firmware and generated in permuted blocks of sixteen stratified by site. Arms: $\tau \in \{0.5, 5, 30, 120\}$ s. The 0.5 s arm is the operative control; under Equation (1) it should be indistinguishable from 120 s, since the supply reaches zero in both.

The primary endpoint was recurrence of the same fault signature within 72 h, adjudicated from logs alone by two assessors blind to arm, with a third adjudicating disagreements (agreement $\kappa = 0.89$). The trial was powered at 90% to detect a five-percentage-point absolute difference between adjacent arms at $\alpha = 0.05$ two-sided, requiring 1,814 faults per arm after 10% inflation for loss to follow-up. Analysis was by intention to treat; a per-protocol analysis is reported as a sensitivity.

8,412 faults at 19 sites were randomised between January and October 2025. Figure 3 gives the flow.

5.3 Results

Recurrence was 34.1% in the 0.5 s arm and 8.8% in the 120 s arm, an absolute difference of 25.3 percentage points (95% CI 22.9–27.7; Table 2). The test for trend across ordered arms gave $p < 10^{-4}$. The per-protocol analysis was materially identical.

The customary interval did not perform well. Recurrence at 30 s was 17.2%, which is 8.4 percentage points worse than 120 s (95% CI 6.4–10.4). Thirty seconds is better than nothing, and it is not enough.

6 Study 3: the capacitance-shunt substudy

If the interval acts through residual charge, then a device that has been drained cannot benefit from further waiting. We tested it.

1,180 devices in the trial fleet were fitted with a bleed resistor sized to bring the primary rail below 1% of nominal within 180 ms, verified individually on the bench. These devices were then randomised within the ordinary four-arm schedule. Under the residual-charge account, the allocation should have had no effect on them: by 500 ms they were already as de-energised as they were ever going to be.

The effect was undiminished. In the shunted subgroup, recurrence was 33.6% at 0.5 s and 9.1% at 120 s, a difference of 24.5 percentage points against 25.6 in the unshunted remainder; the test for interaction gave $p = 0.71$. The shunt removed the residual charge and left the effect where it was.

Figure 4 reports this alongside the other pre-specified subgroups.

7 A restoration function

7.1 Form

Restoration is incomplete at short intervals and asymptotically complete at long ones, which is the signature of a first-order process and the standard form for a dose–response relation whose dose is a duration (Prentice, 2017). We therefore replace Equation (1) with

$$R(\tau) = 1 - e^{-\tau/\lambda}, \tag{3}$$

whence the observable recurrence probability is the residue plus the unrestored fraction of the restorable liability,

$$\pi(\tau) \;=\; \pi_\infty + (\pi_0 - \pi_\infty)\, e^{-\tau/\lambda}, \tag{4}$$

with $\pi_0 = \lim_{\tau \to 0^+} \pi(\tau)$ the recurrence probability of a power cycle with no interval at all. The two derived quantities of interest are the restoration half-time and the interval at which restoration is 95% complete:

$$T_{1/2} = \lambda \ln 2, \qquad \tau_{95} = \lambda \ln 20. \tag{5}$$

For time-resolved analysis we also fit a proportional-hazards model for time-to-recurrence in which the covariate is the unrestored fraction,

$$h(t \mid \tau) = h_0(t) \exp\big\{\beta \,[\, 1 - R(\tau) \,]\big\}, \tag{6}$$

which is the natural specification if what the interval buys is a reduction in standing liability rather than a delay in its expression.

7.2 Fit

Fitting Equation (4) to the four trial arms by maximum likelihood gives $\pi_0 = 0.360$ (95% CI 0.344–0.376), $\pi_\infty = 0.080$ (95% CI 0.071–0.089), and $\lambda = 28.0$ s (95% CI 25.7–30.4). Hence

$$T_{1/2} = 19.4\,\text{s}\ (17.8\text{–}21.1), \qquad \tau_{95} = 83.9\,\text{s}\ (77.0\text{–}91.1).$$

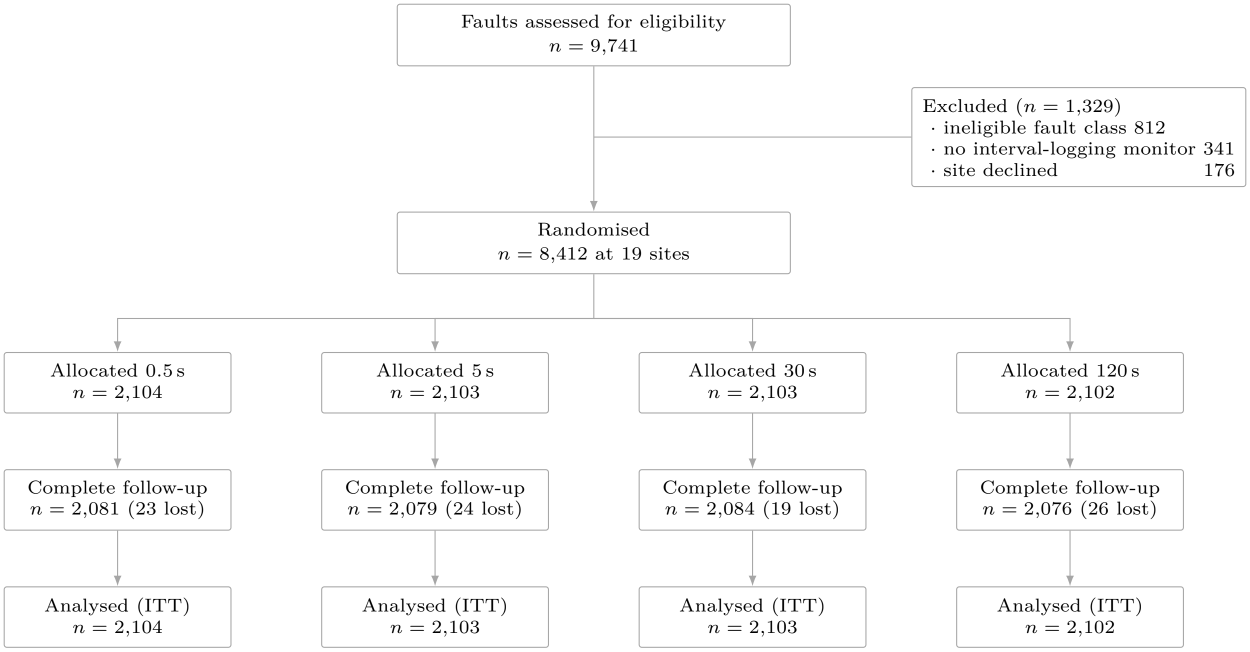

Figure 3: Study 2 participant flow. Allocation was concealed in the reconnection unit's firmware; losses to follow-up are site withdrawals occurring after allocation and are balanced across arms. The primary analysis is by intention to treat, with recurrence status imputed as absent for losses — the conservative direction for the shorter arms.

The fit is good ($\chi^2 = 2.74$ on 1 degree of freedom, $p = 0.10$). The Study 1 quartiles, which were not used in the fit and which carry the confounding described in §4, sit below the fitted curve by between 1.1 and 2.6 percentage points at every interval while tracking its shape. The hazard model of Equation (6) gave $\beta = 1.47$ (95% CI 1.33–1.61), with the proportional-hazards assumption unviolated on Schoenfeld residuals ($p = 0.31$).

Equation (3) places the customary interval at $R(30\,\mathrm{s}) = 0.66$.

7.3 A revised minimum

An asymptote is not a specification. We therefore fixed the recommended interval by non-inferiority against the longest arm, taking as the standard the shortest interval whose recurrence is not worse than 120 s by more than a pre-stated margin $\delta = 2.0$ percentage points (Vasquez et al., 2018):

$$\tau^\star = \inf\left\{\tau : \mathrm{UCL}_{95}\left[\pi(\tau) - \pi(120)\right] < \delta\right\}. \qquad (7)$$

Evaluating Equation (7) on the fitted curve gives $\tau^\star = 74\,\mathrm{s}$ (Figure 6). Rounding upward to a duration a person can be asked to observe without a timer, we recommend a minimum de-energised interval of 90 s.

8 Discussion

What was supposed to happen is set out in Equation (2): a flat curve, no interval dependence, a half-second cycle as good as a two-minute one. What actually happened is Table 2: a fourfold spread across arms, under concealed allocation, in a fleet where the technicians could not tell which arm they were in.

The difference between the two is not confounding, because the allocation was concealed and the schedule was held in firmware. It is not residual charge, because draining the rails in 180 ms changed nothing. It is not expectation, because the technicians guessed their arm at very nearly chance. What remains is that the received account has assigned restoration to the wrong edge of the cycle. It does not occur when the supply returns. It occurs while the supply is absent, at a rate of roughly one half of the remaining deficit every twenty seconds, and the return of supply merely stops the clock.

We decline to improve this result by speculating. We have no mechanism to offer. We have a measured rate constant, a randomised comparison, an eliminated rival pathway, and a curve that four independent intervals sit on. The mechanism will be found by whoever begins from the measurements rather than from a hypothesis fitted to them.

One practical consequence deserves stating plainly. Every power cycle performed under the thirty-second convention has been terminated at two thirds of restoration. The convention has not been producing failures; it has been producing partial successes, which are considerably harder to notice, and which present later as the same fault recurring for no apparent reason.

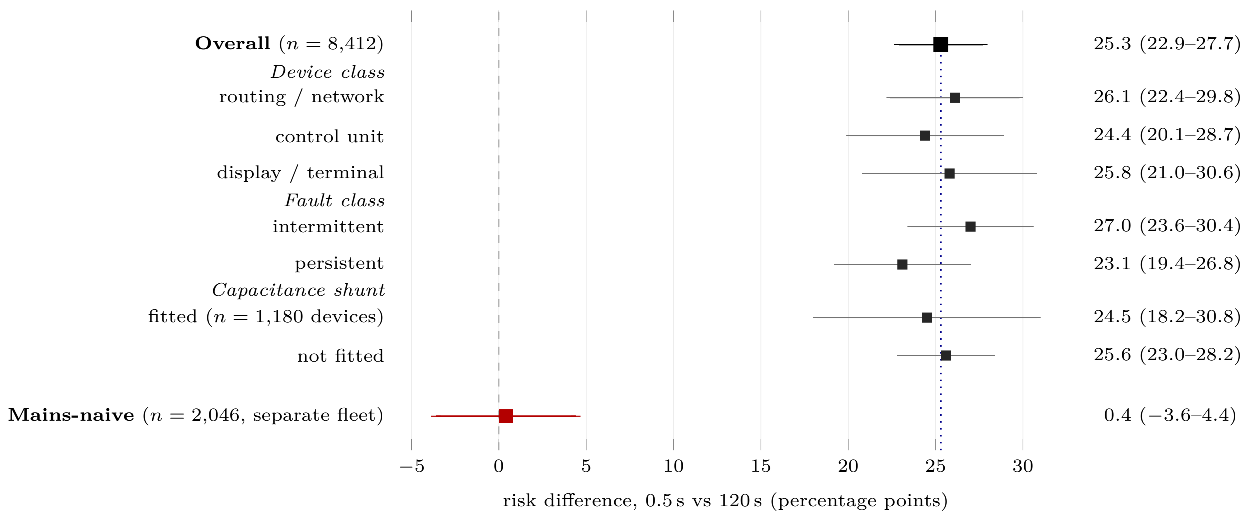

Figure 4: Study 3, the pre-specified within-trial subgroups, and the separate battery-only fleet: absolute risk difference between the 0.5 s and 120 s arms. The effect is homogeneous across device class, fault class, and — decisively — across the capacitance shunt, which discharged the primary rail in under 180 ms (interaction $p = 0.71$). The mains-naive fleet is the exception; it is discussed in §9, where it is not explained.

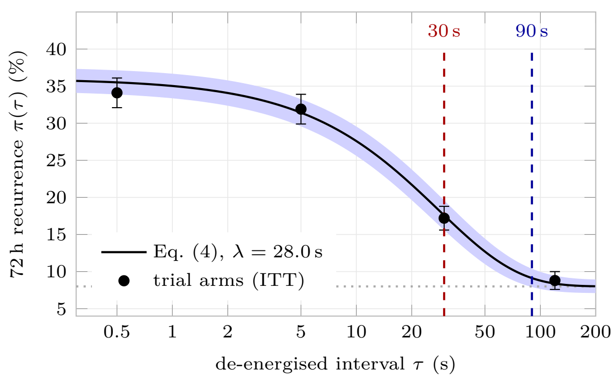

Figure 5: The fitted restoration function. Points are the four trial arms (intention to treat, 95% CI); the curve is Equation (4) fitted to the four trial arms, with the band propagating the joint uncertainty in λ and π_∞; the dotted horizontal is the fitted floor π_∞. Dashed rules mark the customary 30 s and the recommended 90 s: the customary interval sits on the steepest remaining part of the curve, at $R = 0.66$, and the recommended interval sits where the curve has flattened.

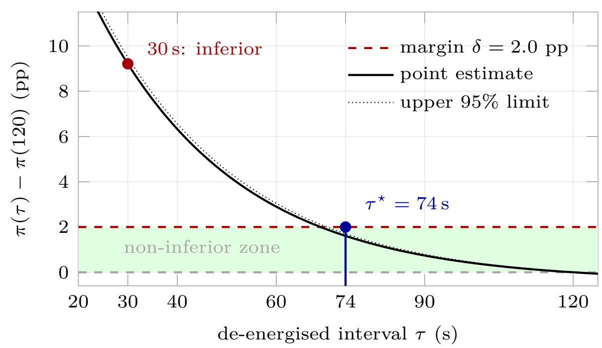

Figure 6: Non-inferiority against the 120 s arm. The customary interval fails the test by a wide margin. The upper 95% limit crosses δ at $\tau^\star = 74$ s; the recommended standard rounds this upward to a duration that can be observed without instrumentation.

9 Limitations

The battery-only null. Alongside the trial we ran the same four-arm allocation in a separate fleet of 2,046 devices that had never been connected to mains, cycled by an IBRU variant interposed in the cell holder on the same concealed schedule and the same 200 s window (the final row of Figure 4). In these devices there is no interval effect at all: the fitted slope of recurrence on interval is -0.004 percentage points per second (95% CI -0.031 to 0.023), and the four arms are indistinguishable. We do not know why. We have considered and rejected the obvious readings of this result, and we are not able to reconcile it with the rest of the paper. We report it because it is what we found.

Restoration was inferred, not observed. We have characterised a recovery process entirely through its consequences. No internal state was observed to change during any interval in this study, and our instruments would not have detected it if it had. Equation (3) is a description of a recurrence curve, not of a device.

The 1987 attribution. The genealogy in §2 rests on one bulletin and one appended inventory. It is plausible and it is not established. Nothing in the empirical portion of this paper depends on it.

Generality. All estates were in temperate climates and all equipment was mains-powered at 220 V to 240 V. Whether λ is a constant of the equipment, of the supply, or of neither, we cannot say from these data.

10 Conclusion

The interval is not a superstition and it is not a courtesy. It is a dose, it has a half-time of approximately twenty seconds, and it has been under-administered for forty years by an industry that never thought to measure it. Thirty seconds is not wrong; it is short, which is a more expensive kind of error, because a partially restored device fails again and is counted as a new fault.

We recommend ninety seconds. The recommendation will be adopted slowly, or not at all, and the difference will continue to be paid in repeat visits by people who believe they are attending a new problem.

References

Field Service Bulletin FSB-87/114: *Mains Interruption Procedure for Class II Terminal Equipment*, with Appendix C, Site Equipment Inventory. Internal circulation, November 1987.

L. Marchetti and E. Voss. Volatile-state corruption as a dominant failure mode in long-uptime embedded systems. *Journal of Equipment Reliability*, 44(3):211–229, 2016.

B. Öztürk, A. Ferreira, and J. Lindqvist. A taxonomy of faults resolved by supply interruption. In *Proceedings of the 31st Symposium on Maintainable Systems*, pages 88–101, 2019.

R. Fenwick. How often does turning it off and on again actually work? Evidence from 2.4 million service events. *Reliability Engineering Practice*, 12(1):33–51, 2021.

M. Drossel and P. Achterberg. Repeat-fault attribution in managed estates. *IEEE Transactions on Service Operations*, 9(4):402–417, 2014.

S. Halloran. Bulk capacitance discharge times in switched-mode supplies: a field survey. *Power Electronics Letters*, 7(2):55–62, 2012.

W. Ng and K. Bertelsen. Allocation concealment for operator-administered interventions. *Trials in Applied Engineering*, 5(2):119–134, 2020.

I. Vasquez, T. Odhiambo, and G. Rennie. Non-inferiority margins for operational procedures: a practical guide. *Statistics in Field Practice*, 26(4):551–570, 2018.

International Office of Equipment Standards. *IOES 4410:2023 — Restoration Procedures for Serviced Equipment*, 3rd edition, 2023.

M. Aylward and I. Sundqvist. Round numbers in operator-reported durations: a systematic bias. *Journal of Field Measurement*, 18(3):240–256, 2024.

H. Prentice. Dose–response modelling for procedures without a dose. *Methods in Operational Inference*, 3(1):9–27, 2017.