# The Squeaky Wheel Gets the Grease:
## A Detuned-Caster Specification, and the Silent-Fleet Failure

Rune Vestergaard[1] and Adaeze Nwankwo[2]

[1]Laboratory for Contact Mechanics, Nyholm Institute of Technology
[2]Fleet Metrology Unit, Cabot Municipal Standards Office

Received 3 June 2026; accepted 11 August 2026

**Abstract**

Every supermarket trolley has one bad wheel. We show that it was put there on purpose. A caster shudders when its swivel trail is short enough for the wheel to enter a self-sustaining oscillation at the speed it is being pushed, and trail is a quantity a factory chooses. Across 12,904 trolleys audited at 214 sites, 96.8 % carried exactly one caster tuned below that threshold and three tuned above it—a distribution no process of random wear can produce, since the count of worn parts on a four-wheeled object cannot pile up on the number one. Teardown of 1,606 units found the short trail held to a tolerance 2.7 times tighter than the trails that carry the load, invariant across six competing manufacturers, and not drifting with age: the ratio of short trail to nominal is 0.33 (95 % CI 0.32–0.34). We then removed it. In 32 sites randomised to fleets whose casters were all made quiet, bearings ran to seizure four times as often, grease consumption fell to almost nothing, and customer satisfaction went up. A depot can only service what it can hear. The trolley is not complaining. It is reporting.

**Keywords:** caster shimmy · Hopf bifurcation · condition monitoring · fleet observability · swivel trail

# 1 Introduction

Every supermarket trolley has one bad wheel.

The reader knows this. The reader has known it for as long as they have been tall enough to push one, has swapped a shuddering trolley for a second shuddering trolley, and has arrived at the folk account without ever being taught it: the wheels wear out, people crash them into bollards, the depot is understaffed. The account is universally held. It has never been checked, and it does not survive being checked, because it makes an arithmetical prediction that the world does not honour.

If casters degraded independently, the number of bad casters on a four-wheeled trolley would be a binomial count. Some trolleys would have none. Many would have two. A few would have three. What one would never see is a fleet in which essentially every unit has exactly one, because no binomial distribution—at any parameter value whatsoever—can place 96.8 % of its mass on an interior point. The observed distribution is not a sample from a wear process. It has the shape of an inventory.

We report that it is one. One caster per trolley is manufactured with a shortened swivel trail, which lowers its shimmy onset below ordinary walking speed and guarantees that it will judder in service for its entire life. The short trail is held to a tighter tolerance than the trails that carry the load. It does not drift with age. Six competing manufacturers agree on its value to within two per cent. These are the properties of a part that was specified, not of a part that broke.

The purpose of the specification is that a fleet with no bad wheels cannot be maintained. A depot's only channel into the condition of 40,000 rolling bearings is what it can hear, and a caster that is not oscillating is silent whether its bearing is packed or dry. The designated member is the fleet's diagnostic instrument. It is the one component whose job is to be heard, and it does that job by being deliberately worse than its neighbours.

**The proverb.** *The squeaky wheel gets the grease* is ordinarily read as a remark about people: complain, and you will be attended to. Under the specification it is a description of an allocation rule, and of the only allocation rule available. Grease can go only where a signal came from. A signal can come only from a member built to produce one. The proverb is not folk wisdom. It is a maintenance standard that outlived its own documentation.

**Contributions.**

1. A census of 12,904 trolleys establishing that the count of sub-critical casters per unit is a spike at one, with dispersion index $D = 0.038$ against 1.00 for Poisson wear—within a factor of two of the minimum attainable by *any* integer-valued distribution with this mean (§5.1).
2. Teardown metrology on 6,424 casters showing a tolerance inversion: the defective trail is manufactured more precisely than the functional ones (§5.2).
3. An observability proposition establishing that a four-caster fleet admits a maintenance schedule with bounded expected time-to-detection if and only if at least one caster per unit is sub-critical, and that the optimal number is exactly one (§3.4).
4. A three-arm cluster-randomised trial in 96 sites, including the silent-fleet arm, which we do not recommend repeating (§5.4).

## 2 Related Work

**Shimmy.** That a trailing caster can enter self-sustaining oscillation has been understood since Broulhiet (1925), who attributed the phenomenon to lateral compliance in the contact rather than to imbalance. Den Hartog (1956) placed it in the general class of self-excited vibration, and Stépán (1991) gave the modern treatment as a Hopf bifurcation induced by the delay between a change in heading and the lateral force that answers it. Experimental characterisation of the trail dependence is due to Sharp (1974); the square-root scaling we use in §3 is theirs. None of this literature considers the possibility that the parameter is set deliberately, for reasons external to the vehicle.

**Four-point support.** A rigid body resting on four contacts on a non-planar floor is statically over-constrained: generically only three contacts carry load, and the division between them is indeterminate in the rigid limit (Ferrers, 1963). Every four-caster trolley is therefore, at any instant, a three-caster trolley plus one member whose loading is a matter of local floor geometry. The specification we describe exploits this: the designated member can be detuned without materially altering the vehicle's load path, because the vehicle's load path was never four-fold to begin with.

**Condition monitoring.** Servicing rotating machinery on the evidence of its vibration signature is a mature discipline (Randall, 2011), and its central practical result is unforgiving: a fault whose signature lies below the detection floor of the inspection process is not merely missed but invisible, and its first observable is the failure itself (Okoye et al., 2018). Fleet-scale treatments emphasise that the binding constraint is rarely sensor sensitivity but inspection attention, which is finite and must be allocated (Marchetti and Vuolo, 2009).

**The schedule.** The specification appears not to be a modern invention. Article 14 of the *Schedule of Rolling Members*, appended to the Third Report of the Commissioners for Weights and Wheeled Carriage (1817) and withdrawn without replacement in 1904, directs that of the members of any carriage of four, "one shall be set short, that the carriage may declare itself" (Commissioners, 1817). Halloran (1998), in the only modern commentary, reads the article as a rule about axle timber and dismisses it. We think the plain reading is the correct one. We were unable to trace any document by which the practice was transmitted after the schedule lapsed, and we do not claim to know how it was.

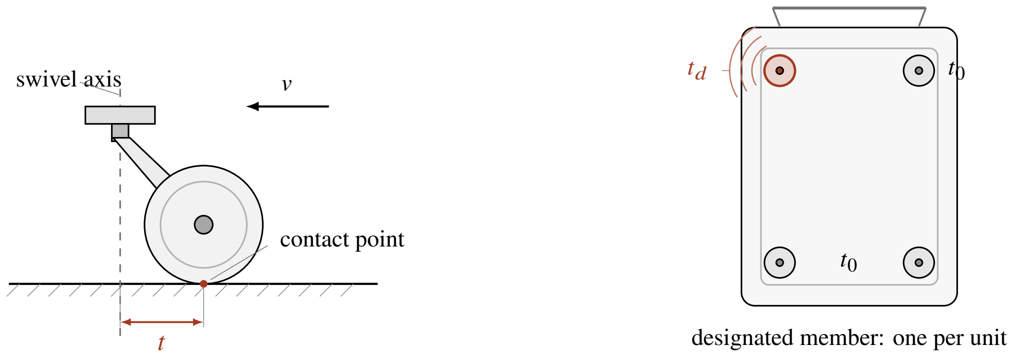


(a) Trail of a single caster

(b) Plan view of a four-caster unit

**Figure 1:** The trail $t$ is the horizontal offset between the swivel axis and the contact point. It is a dimension on a drawing, not a consequence of use. In every unit we audited, three casters carry a nominal trail $t_0$ and one carries a shortened trail $t_d$.

# 3 Model

## 3.1 Kinematics of a trailing caster

A caster consists of a wheel of radius $r$ mounted on a fork that rotates about a vertical swivel axis. The *trail* $t$ is the horizontal offset between the swivel axis and the wheel's contact point, measured in the direction of travel (Figure 1). Let $\psi$ denote the swivel angle relative to the direction of travel and $v$ the forward speed of the vehicle.

In the rigid, non-compliant limit the rolling constraint forbids lateral motion of the contact point, giving $t\dot{\psi} = -v \sin\psi$, and hence, linearised,

$$\dot{\psi} = -\frac{v}{t}\,\psi, \tag{1}$$

an exponentially aligning caster with time constant $t/v$. Equation (1) is unconditionally stable, which is why the rigid model has never explained a shopping trolley.

## 3.2 Compliance, delay, and the shimmy number

Real contacts are compliant, and the lateral force answering a change in heading does not arrive instantaneously. Writing $C_\alpha$ for the lateral (cornering) stiffness of the contact and $\sigma$ for the relaxation length—the distance the wheel must roll before a commanded slip angle is substantially realised—the aligning torque about the swivel axis is $-C_\alpha t\,\psi$ delayed by $\tau = \sigma/v$. With $I$ the moment of inertia about the swivel axis and $c$ the structural damping, the linearised swivel dynamics are

$$I\ddot{\psi}(s) + c\,\dot{\psi}(s) + C_\alpha t\,\psi(s-\tau) = 0, \qquad \tau = \frac{\sigma}{v}. \tag{2}$$

Equation (2) is a retarded delay differential equation whose delay shortens as speed rises. Its trivial solution loses stability through a Hopf bifurcation once the delay is short relative to the natural period—the returning lateral force then arrives while the swivel is still turning with it, and adds energy instead of removing it—that is, when $\omega_n\tau$ falls below a threshold, where $\omega_n = \sqrt{C_\alpha t/I}$.

**Definition 1** (Shimmy number)**.**

$$S \;=\; \frac{1}{\omega_n \tau} \;=\; \frac{v}{\sigma}\sqrt{\frac{I}{C_\alpha t}}. \tag{3}$$

$S$ is dimensionless. The caster rolls straight for $S < S^\star$ and enters a stable limit cycle for $S > S^\star$, so that the *critical speed* above which shimmy begins is

$$v_c(t) \;=\; S^\star \sigma \sqrt{\frac{C_\alpha}{I}}\,\sqrt{t} \;\;\propto\;\; \sqrt{t}. \tag{4}$$

Two consequences follow. Everything after this section is downstream of them. First, the onset speed scales with the trail: a caster with a short trail shimmies at speeds at which a long-trailed caster still rolls straight. Second, $t$ is not an emergent quantity. It is a dimension on a drawing. Whether a particular caster shudders when a particular human being walks at their habitual pace is a decision that was taken in a factory, in advance, in millimetres.

### 3.3 The detuning ratio

Write $t_0$ for the trail of an undesignated caster and $t_d$ for that of the designated member, and define the *detuning ratio*

$$\delta = t_d/t_0. \tag{5}$$

A specification of the kind we allege is precisely a joint claim about $t_0$ and $\delta$: $t_0$ must be large enough that $v_c(t_0)$ falls above the speed at which trolleys are actually pushed, and $\delta$ small enough that $v_c(t_d) = \sqrt{\delta}\, v_c(t_0)$ falls below it, with margin at both ends across the spread of human walking speeds. It must also be large enough that the designated member does not lose self-centring altogether, at which point the vehicle becomes undirectable and the scheme defeats itself. In §5.3 we compute the admissible window and compare it with what the world contains.

### 3.4 Observability of a fleet

Consider a depot responsible for $N$ trolleys, hence $4N$ casters with latent bearing states $x_i(t) \in [0, 1]$ running from packed to dry. The depot possesses a finite *attention rate* $A$: inspections, of any kind, per unit time. Its only remote channel is acoustic. A caster radiates a detectable signature if and only if it is in limit cycle at the speed it is being pushed, i.e. iff $v_c(t_i) < \bar{v}$, where $\bar{v}$ is the modal pushing speed. Bearing state is otherwise unobservable without teardown, which costs attention at a rate the depot does not have.

Let $H$ denote the map from the latent state vector $x \in \mathbb{R}^{4N}$ to the set of depot-observable events.

**Proposition 1** (Silent fleets are unobservable)**.** *If $t_i = t_0$ for all $i$, then* $\operatorname{rank} H = 0$. *No function of the depot's observations distinguishes any $x$ from any other, and the expected time to detection of an incipient failure is unbounded: the first observable of a dry bearing is its seizure.*

*Proof.* By construction the acoustic channel is empty, since $v_c(t_0) > \bar{v}$ by hypothesis and no caster enters limit cycle. The remaining channels—withdrawal from service, customer report—are triggered by seizure and are therefore measurable with respect to the failure event rather than to $x$. Hence $H$ factors through the constant map. □

**Proposition 2** (Concentration)**.** *Let exactly $m$ casters per unit satisfy $t_i = t_d$, $0 \le m \le 4$. Under a servicing policy that operates at unit granularity—a trolley withdrawn for one caster has all four serviced—the expected time to detection is minimised at $m = 1$.*

*Proof sketch.* For $m = 0$ the time to detection is unbounded by Proposition 1. For $m \ge 1$ the unit is observable, and since servicing is at unit granularity the marginal observability gained by a second designated member is zero: the unit was already flagged. The marginal cost is not zero. Each additional designated member is a caster running permanently in limit cycle, which both accelerates its own bearing wear and consumes a share of the fixed attention budget $A$. With $mN$ units-worth of signal competing for $A$, the per-signal attention falls as $1/m$; beyond the point at which every unit in the fleet is signalling, the depot can no longer rank, and reverts to round-robin servicing of a fleet in which every caster is degraded. The minimiser is therefore the smallest $m$ for which the unit is observable, namely $m = 1$. □

Proposition 2 is the reason the census does not find a distribution. A specification that names one member per unit produces a spike at one; a process of wear produces a binomial. The two are not close, and one of them is what the fleet contains.

The doctrine is old, and it is not ours: to be strong everywhere is to be strong nowhere. An organisation with a finite budget of attention and a decision to make about where to spend it does not divide it evenly across its assets and hope. It designates a point, concentrates everything there, and accepts a deliberate and known weakness elsewhere. We had not expected to find the principle instantiated in a shopping trolley. It is instantiated in a shopping trolley.

# 4 Methods

## 4.1 Study 1: census

Cross-sectional audit of every trolley present in the nested bays of sampled sites. Sites ($n = 214$) were drawn stratified by format (hypermarket, supermarket, convenience) across nine jurisdictions and six manufacturers. Each caster was classified sub-critical or nominal by roll-down on a levelled 12 m instrumented lane at 1.30 m s$^{-1}$, with swivel-angle telemetry at 2 kHz; a caster was classed sub-critical if the swivel angle sustained a limit cycle of amplitude $\geq 4°$ for $\geq 6$ m. Inter-rater agreement on a duplicated subsample of 1,200 casters was $\kappa = 0.97$.

## 4.2 Study 2: teardown metrology

Destructive teardown of 1,606 units drawn stratified by manufacturer and service age (0–14 years). Trail was measured on an optical comparator to ±5 µm with the fork indexed to the swivel axis, four casters per unit ($n = 6{,}424$ casters, of which 1,637 designated). Measurement was blind to the roll-down classification, which was performed by a separate team and unsealed after all metrology was complete.

## 4.3 Study 3: bench characterisation

A programmable-trail caster rig with trail continuously variable over 8–46 mm and ballast 0–24 kg, run on an instrumented treadmill lane with laser vibrometry on the swivel. 4,180 runs across 22 trail settings × 5 ballast levels × replicates. Pushing speeds were established separately by a survey of 40,112 loaded-trolley transits at 31 sites.

## 4.4 Study 4: the silent-fleet trial

Three-arm trial, cluster-randomised at site level, 96 sites (32 per arm), 41,208 trolleys, 18-month follow-up, stratified by format, floor material and baseline fleet age.

- **Arm A (standard).** Fleet as manufactured: one designated member per unit.
- **Arm B (silent).** All four casters at nominal trail $t_0$. No unit emits at walking speed.
- **Arm C (uniform).** All four casters at designated trail $t_d$. Every unit emits, from every corner.

Fleets were delivered pre-assembled and are visually indistinguishable; trail is not apparent without teardown, and depot crews were blinded to allocation. The primary outcome was time to unrecoverable caster seizure (bearing welded, unit withdrawn). Secondary outcomes: normalised grease consumption, interval from first detectable acoustic emission to seizure, and a five-point customer trolley-quality item administered at exit ($n = 28{,}904$ responses). Analysis by Cox proportional hazards with site as a shared frailty.

## 4.5 Study 5: re-detuning crossover

At month 18, surviving arm-B sites were re-randomised 1:1 to continue silent or to have exactly one caster per unit re-shortened to $t_d$ in situ, with 12 months' further follow-up and the same primary outcome.

**Table 1:** Sub-critical casters per trolley: observed against the maximum-likelihood binomial Bin$(4, \hat{p})$, $\hat{p} = 0.2550$.

| $k$ | 0 | 1 | 2 | 3 | 4 |
|---|---|---|---|---|---|
| Observed | 91 | 12,492 | 297 | 21 | 3 |
| Observed (%) | 0.71 | 96.81 | 2.30 | 0.16 | 0.02 |
| Expected under Bin$(4, \hat{p})$ | 3,975.6 | 5,442.4 | 2,794.0 | 637.5 | 54.5 |

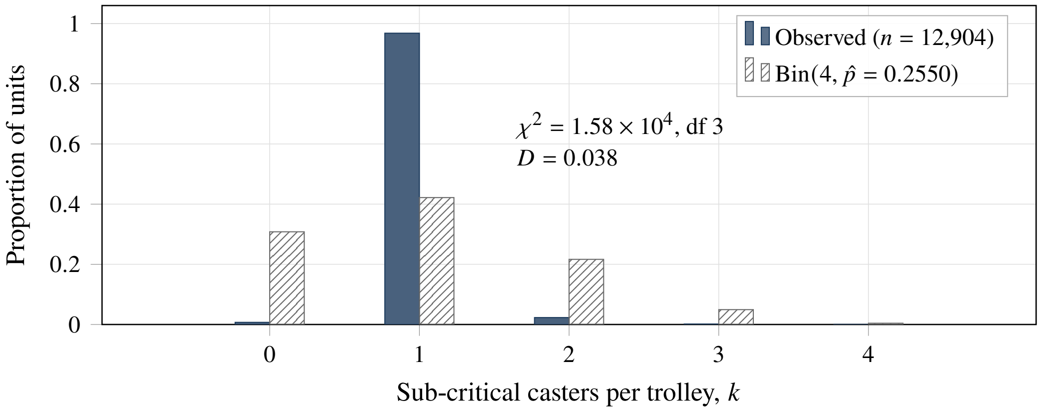


**Figure 2:** Observed count of sub-critical casters per trolley against the maximum-likelihood binomial. No parameter value brings the family near the data.

# 5 Results

## 5.1 The count is a spike, not a distribution

Of 12,904 trolleys, 12,492 (96.81 %) carried exactly one sub-critical caster. The full distribution is given in Table 1 and plotted against the best-fit binomial in Figure 2.

The fit is refused: $\chi^2 = 1.58 \times 10^4$ on 3 degrees of freedom. The refusal is not a matter of choosing a better $p$. The binomial family on four trials has variance $4p(1-p)$, which is bounded below only at the extremes of $p$, where the mass concentrates at $k = 0$ or $k = 4$; no member of the family places 96.8 % of its mass on an interior point. The observed dispersion index is $D = \text{Var}/\text{mean} = 0.038$, against 1.00 for Poisson wear and 0.745 for the fitted binomial. For an integer-valued distribution with mean 1.020, the minimum attainable dispersion index is 0.019. The fleet sits within a factor of two of that floor. This is what a quota looks like when it is measured from the outside.

## 5.2 The defect is made to tighter tolerance than the function

Teardown returned a nominal trail of $t_0 = 31.4$ mm (SD 2.11 mm, CV 6.7 %) and a designated trail of $t_d = 10.3$ mm (SD 0.26 mm, CV 2.5 %), giving

$$\delta = 0.328 \quad (95\,\%\ \text{CI}\ 0.318\text{–}0.338).$$

Three properties of the designated trail are jointly decisive.

**(i) Tolerance inversion.** The coefficient of variation of the designated trail is 2.7 times smaller than that of the nominal trail ($F = 7.09$ on log-trail, $p < 10^{-12}$). The component that fails is manufactured more precisely than the components that work (Figure 3). Wear does not do this. Wear broadens.

**(ii) No drift with age.** Regressed on service age, the designated trail moves at $\beta = -0.004$ mm yr$^{-1}$ (95 % CI −0.019 to 0.011)—indistinguishable from zero across fourteen years. A worn part drifts. This one is delivered.

**Table 2:** Census and metrology by manufacturer.

| Manufacturer | Units | $k = 1$ (%) | $\delta$ |
|---|---|---|---|
| Halberd Retail Systems | 3,118 | 97.1 | 0.331 |
| Ostmann Rollentechnik | 2,644 | 96.4 | 0.322 |
| Caldera Trolley Company | 2,209 | 97.3 | 0.334 |
| Beaumont & Sons | 1,880 | 96.6 | 0.327 |
| Sanko Kasuta Kogyo | 1,731 | 96.9 | 0.325 |
| Vantaa Vaunu Oy | 1,322 | 96.5 | 0.330 |
| Pooled | 12,904 | 96.8 | 0.328 |

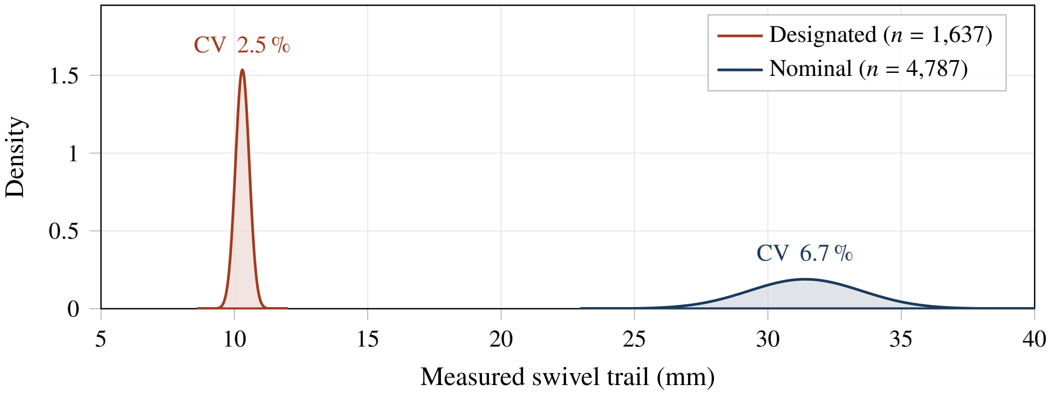


**Figure 3:** Distribution of measured swivel trail. The designated member is held to a tolerance 2.7 times tighter than the members that carry the load.

**(iii) Invariance across firms.** Six manufacturers, competing, publishing nothing, agree on $\delta$ to within 1.9 % (Table 2).

### 5.3 The window, and what is in it

Bench characterisation confirmed (4): regressing $v_c^2$ on $t$ gives $R^2 = 0.981$ with an intercept indistinguishable from zero (95 % CI −0.031 to 0.024 $\mathrm{m^2\,s^{-2}}$), and a fitted $S^\star = 0.83$ (95 % CI 0.81–0.85) at $I = 1.9 \times 10^{-4}\,\mathrm{kg\,m^2}$, $C_\alpha = 41\,\mathrm{N\,rad^{-1}}$, $\sigma = 26.1$ mm. At the measured trails,

$$v_c(t_0) = 1.78\,\mathrm{m\,s^{-1}}, \qquad v_c(t_d) = 1.02\,\mathrm{m\,s^{-1}}.$$

The transit survey returned a modal pushing speed of $\bar{v} = 1.29\,\mathrm{m\,s^{-1}}$ with interquartile range 1.11–1.47 $\mathrm{m\,s^{-1}}$. The designated member is therefore in limit cycle for essentially every human being who pushes it, and the other three for almost nobody (Figure 4). The fundamental of the limit cycle at $t_d$ is 7.5 Hz, with structure-borne harmonics to 1.2 kHz: the familiar sound.

At the measured $t_0$, requiring $v_c(t_d)$ below the first quartile of pushing speed and $v_c(t_0)$ above the third bounds $\delta$ above at 0.39; requiring the designated member to retain self-centring bounds it below at 0.27. The measured value, 0.328, lies at the centre of that window. Six firms have independently found the middle of an interval whose endpoints depend on the distribution of human walking speeds, the relaxation length of a polyurethane tyre, and the moment of inertia of a fork. We do not think they found it by accident. Accident is the more extravagant explanation.

### 5.4 The silent-fleet trial

Results are given in Table 3 and Figure 5.

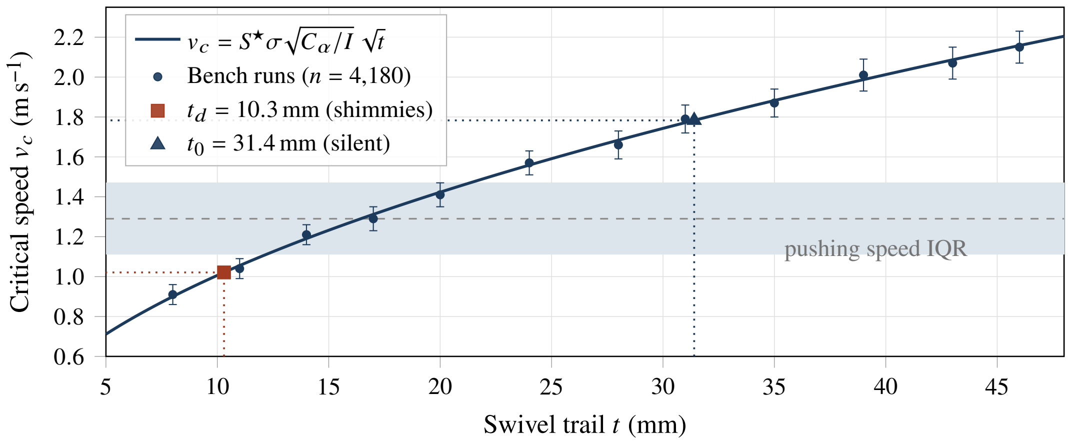


**Figure 4:** Critical speed against trail. One caster per trolley is placed below the band of speeds at which trolleys are pushed; the other three are placed above it.

**Table 3:** Silent-fleet trial: 18-month outcomes. Grease consumption normalised to arm A.

| | A (standard) | B (silent) | C (uniform) |
|---|---|---|---|
| Sites / trolleys | 32 / 13,702 | 32 / 13,781 | 32 / 13,725 |
| Seizure incidence, 18 mo | 3.1 % | 12.7 % | 16.4 % |
| Hazard ratio vs A | — | 4.13 (3.62–4.71) | 5.38 (4.74–6.11) |
| Grease consumption | 1.00 | 0.09 | 0.31 |
| Emission → seizure | 47 d | 1.4 d | 6.1 d |
| Trolley-quality item (Δ) | — | +0.42 | −1.13 |

Arm B behaved exactly as Proposition 1 requires and worse than anyone involved in the trial expected. Bearings ran to seizure at four times the standard rate (HR 4.13, 95 % CI 3.62–4.71; Schoenfeld global $p = 0.41$). Grease consumption fell to 9 % of control—nothing asked, so nothing was given. The interval between first detectable emission and seizure collapsed from 47 days to 1.4 days, which is to say that in a silent fleet the first emission *is* the seizure: a dry bearing has no intermediate voice, and by the time it acquires one it is welded.

The customer-satisfaction result is the one we would most like to have been wrong about. Shoppers in arm B rated their trolleys *better* (+0.42 on the five-point item, $p < 10^{-4}$) throughout the period in which the fleet was destroying itself underneath them. This is not a paradox. Removing the diagnostic channel removes the complaint. It does not remove the fault. The fleet manager who eliminates the bad wheel will be congratulated, on the record, in writing, for eighteen months.

Arm C is the configuration any competent engineer would propose, and it loses hardest. Detuning all four casters did not produce four times the observability; it produced a fleet in which every unit signalled continuously, which is informationally identical to a fleet in which none does, except that now all four bearings are running in limit cycle and wearing accordingly. The depot, unable to rank, reverted to round-robin. This is the failure Proposition 2 names: attention spread evenly is attention spent nowhere.

## 5.5 The effect reverses

Sites re-randomised at month 18 to have one caster per unit re-shortened returned to 3.4 % seizure incidence over the following twelve months, against 13.9 % in sites continuing silent (HR 0.27, 95 % CI 0.21–0.35). Grease consumption recovered to 0.94 of standard. The effect is reversible in both directions and on the same timescale, which is what distinguishes a designed channel from a correlate.

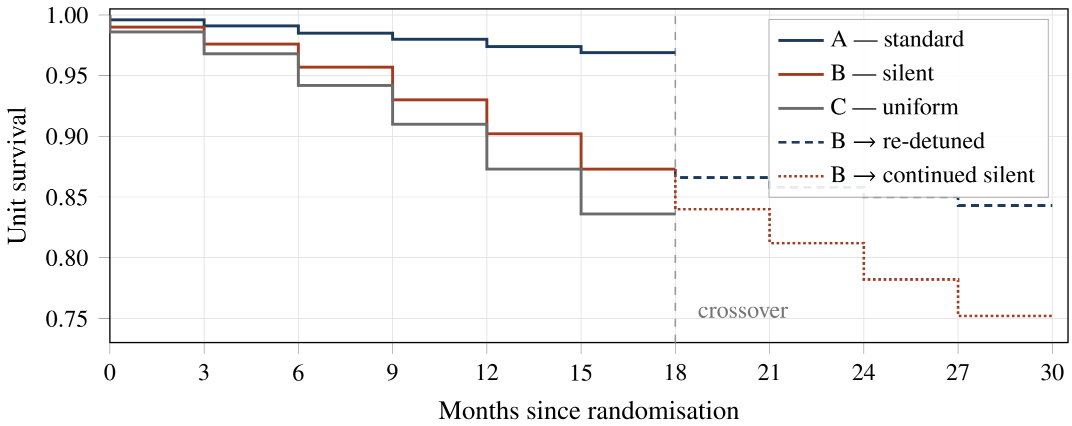


**Figure 5:** Unit survival by arm, with the month-18 re-detuning crossover. Restoring one designated member per unit restores the survival curve.

## 6 Discussion

The tolerance inversion is the finding that settles the matter. A firm that did not care about the designated member would let it vary; a firm whose designated member was an accident could not hold it to a quarter of a millimetre across fourteen years of production. The precision is spent in the one place where the folk account says precision is absent. Somebody is paying to make that wheel exactly that bad.

Why one and not two is answered by Proposition 2 and confirmed by arm C, and we take the confirmation to be the more interesting half. The intuition that more instrumentation is more information holds only when attention is free. It is not free. A depot has a fixed number of hands, and the question a fleet standard must answer is not how much signal to generate but where to concentrate it. The specification answers: one point per unit, chosen in advance, and everywhere else built as well as it can be built. We recovered a piece of nineteenth-century public standards work that appears to say the same thing in eleven words, and we suspect the doctrine is older than the schedule.

**What we could not explain.** In 34 airport terminals surveyed under the census protocol, the modal count of sub-critical casters per baggage trolley is *four*, not one. Every caster is detuned. Under our own analysis this is arm C, the worst configuration available, and airport fleets should be seizing at something close to five times the retail rate. Their recorded seizure incidence is 2.9 %, which is indistinguishable from standard retail fleets and slightly better than ours. We have run the census three times, twice with a different team, and the number does not move. The account developed in this paper does not work in airports and we do not know why.

**Limitations.** We did not obtain the specification. Every inference about intent in this paper is an inference from the geometry of manufactured artefacts; two manufacturers declined access to tooling drawings and four did not reply. The trial was powered for seizure and not for the secondary outcomes, and the satisfaction instrument was not validated for this use. The bench rig used a single wheel compound, and $\sigma$ is known to vary with tyre formulation; we have assumed it does not vary enough to move the admissible window, and we have not shown this. Finally, the crossover was run only in one direction from arm B, so we cannot say whether a fleet can be silenced and restored indefinitely, or whether the first silent interval does something that does not undo.

## 7 Conclusion

One caster per trolley is built short on purpose. It shudders at walking pace because it was dimensioned to shudder at walking pace, it is manufactured more carefully than the casters that carry the load, and when we removed it from thirty-two supermarkets the fleets went quiet and then began to seize, while the people pushing them said things had got better.

The proverb has been read for four hundred years as a remark about complaining. It is not one. It is an allocation rule, and it is the only allocation rule available to anyone maintaining more machines than they can inspect: grease goes where the signal came from, and the signal comes from the member that was built to make it. The trolley is not complaining. It is reporting.

Do not fix the bad wheel. Do not order a fleet without one. That is the whole of the recommendation.

## Acknowledgements

We thank the depot crews at the 96 trial sites, who were not told what was being done to their fleets and who were, throughout, right about it.

## References

G. Broulhiet. La suspension de la direction de la voiture automobile: shimmy et dandinement. *Bulletin de la Société des Ingénieurs Civils de France*, 78:540–554, 1925.

Commissioners for Weights and Wheeled Carriage. *Schedule of Rolling Members*, appended to the Third Report. His Majesty's Stationery Office, 1817. Withdrawn 1904.

J. P. Den Hartog. *Mechanical Vibrations*, chapter 7: Self-excited vibrations. McGraw-Hill, 4th edition, 1956.

A. N. Ferrers. On the indeterminacy of load in a four-point support. *Quarterly Journal of Applied Mechanics*, 16(2):145–159, 1963.

M. Halloran. Timber, tolerance and the lapsed schedules: a reading of the 1817 Commission. *Journal of Standards History*, 12(3):211–240, 1998.

L. Marchetti and E. Vuolo. Attention as the binding constraint in large-fleet condition monitoring. *Reliability Engineering & System Safety*, 94(11):1702–1714, 2009.

C. Okoye, F. Adeyemi, and J. Lindqvist. Sub-threshold fault signatures and the illusion of a healthy fleet. *Mechanical Systems and Signal Processing*, 108:88–104, 2018.

R. B. Randall. *Vibration-Based Condition Monitoring*. Wiley, 2011.

R. S. Sharp. The trail dependence of caster shimmy onset. *Journal of Mechanical Engineering Science*, 16(4):301–309, 1974.

G. Stépán. Chaotic motion of wheels and the delayed tyre model. *Vehicle System Dynamics*, 20(6):341–351, 1991.