The Cost of Being Wrong Slowly

On the asymmetry between errors that announce themselves and errors that do not

A. Specimen · B. Placeholder

PAPER · PUBLISHED · VERSION v1.0 · 2026-08-01

SPECIMEN — FICTIONAL

This work does not exist. It is placeholder text written to exercise a typographic system, and its authors, venues, identifiers and findings are invented. Nothing in it should be cited, repeated, or believed.

ABSTRACT

Decision systems are usually judged by how often they are wrong. This specimen argues they should be judged by how long they take to notice. A system that errs frequently and corrects within a week is cheaper to run than one that errs rarely and takes a decade to admit it, even where the second has the better average. The argument is developed through six premises and one conditional result, and is offered as fictional placeholder text.

Part 0 — What this specimen is

This work does not exist. It is placeholder text written to exercise a typographic system: numbered premises, glosses, epistemic tags, callouts, footnotes, a table, a blockquote, and a reference list. Its arguments are constructed to be *shaped* like arguments rather than to be true, and no claim in it should be cited, repeated, or believed.[1]

What follows is nonetheless written straight. A specimen that hedges every sentence does not stress the typesetting the way real prose does — the line breaks fall differently, the rhythm of a hedged paragraph is not the rhythm of an argued one, and a system that looks composed on cautious text can still fall apart on confident text.

[1] Every reference in this specimen is fabricated, and deliberately so — no venue named here exists, and no digital object identifier is given anywhere in this work. A specimen that cited real sources would eventually be quoted as though it had.

The shape of the claim

The claim has three parts, and their modality differs. The first is a definition and is true by construction `[literal]`. The second is a claim about a class of dynamical systems and is exact only within a model that has not been built `[model-relative]`. The third transfers the result to human institutions and is, at present, a figure of speech doing structural work `[analogy]`. Readers who accept the first two and reject the third have understood the paper correctly.

Part 1 — The premises

Premise 1. Every decision system has a detection latency: the interval between committing an error and registering that an error was committed.

Latency is a property of the system, not of the error. The same mistake made inside two different institutions will surface at two different times, and the difference is structural rather than moral. Nothing here requires that anyone behaves badly.

Premise 2. Detection latency and error rate are independent parameters.
This is the load-bearing premise and the least intuitive. It is tempting to assume that careful systems are both accurate and quick to notice, and that sloppy ones are neither. But care spent on avoiding errors is care not spent on surfacing them, and the two budgets compete. A body that deliberates for eighteen months before acting has bought a lower error rate with a longer latency, and has not been told the price [defeasible].

Premise 3. The cost of an error grows with the time it remains uncorrected, and grows faster than linearly.
An error that persists is not merely an error that persists. Decisions are taken on top of it; the structures built on the mistake must themselves be unwound; and the people who acted in good faith on the false premise acquire a stake in its remaining true. This compounding is the whole of the argument.

Premise 4. A system cannot reliably measure its own detection latency.
The measurement requires knowing when the error began, which is available only after detection, and only through the same faculties that failed to detect it. Self-reported latency is therefore systematically understated — not through dishonesty, but because the clock is started when the alarm rings.

Premise 5. Openness lowers latency at a cost that is visible, and closure raises it at a cost that is not.
Admitting an outside challenge is expensive in a way everyone can see: it takes time, it is uncomfortable, and someone must answer. Refusing one is expensive in a way nobody can see, because the bill arrives later and is itemised under a different heading. Budgets respond to visible costs.

Premise 6. Where latency is unmeasured and closure is cheap, latency drifts upward.
Given Premises 4 and 5, no correcting pressure acts on latency. A quantity under no pressure and subject to a one-way cost gradient does not stay where it was put [model-relative].

A system optimised for being right is not the same as a system optimised for finding out that it is wrong. Where the two conflict, institutions reliably choose the first, because only the first produces evidence of virtue while the work is being done.

Part 2 — The conditional result

Statement
If Premises 1 through 6 hold, **then** for any decision system that does not measure its own detection latency, expected total error cost rises over time even where the error *rate* is falling.

The result is conditional and the condition is strong. Premise 2 in particular is stated here without support and would need a model to earn `[defeasible]`. What the result offers is not a fact about the world but a place to look: if you wish to refute the argument, Premise 2 is where to aim, and the refutation would be surgical rather than rhetorical.

Why the falling error rate is a trap
The trap is that a falling error rate is *observable* and a rising latency is not. An institution watching its own dashboards sees improvement. It is not lying. It is measuring the quantity it can measure, and that quantity is genuinely getting better.

> *The most dangerous instrument is not the one that reads wrong. It is the one that reads right about the wrong thing, because no amount of care in reading it will reveal the error.*

A worked comparison

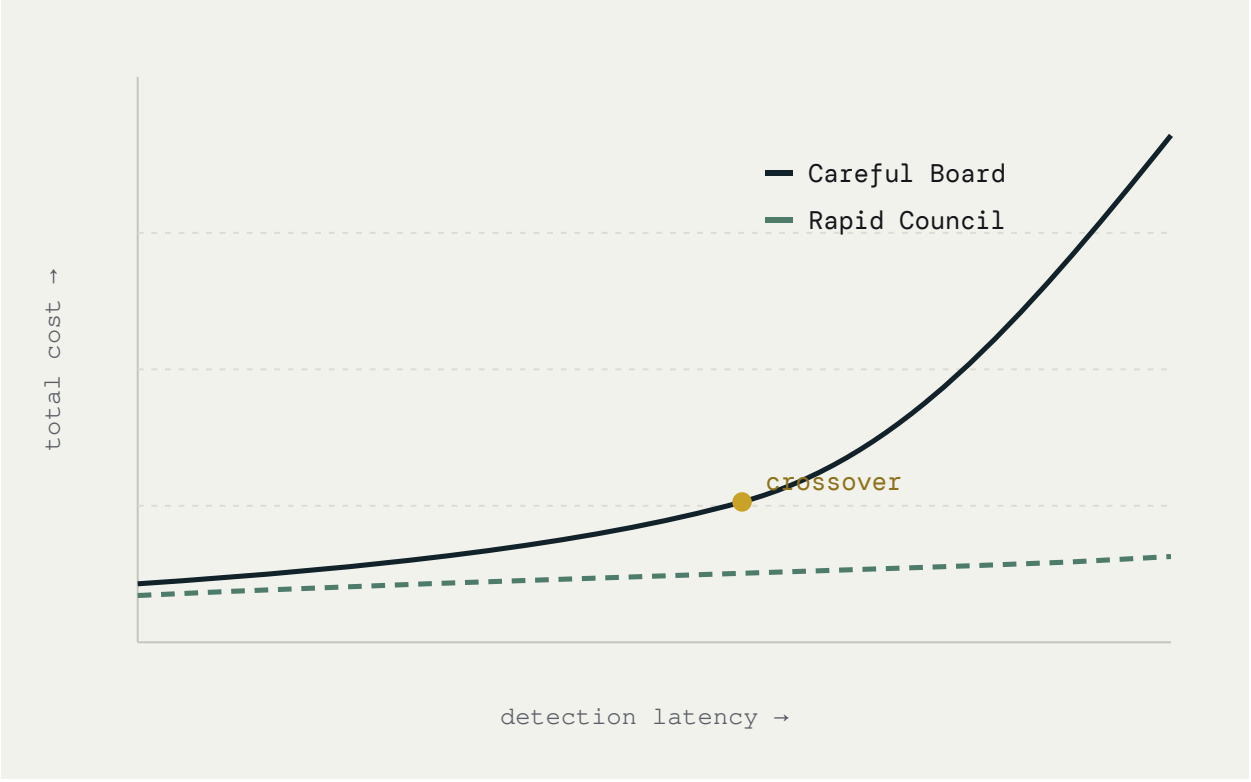

Figure 1 Total error cost against detection latency for two fictional bodies. The Careful Board errs less often but compounds for longer; the curves cross well before either body would notice. All values are invented.

Two fictional bodies, each deciding one hundred matters a year.

Body	Errors per year	Median latency	Compounding factor	Annual error cost
The Rapid Council	12	3 weeks	1.1×	13.2 units
The Careful Board	4	40 months	6.5×	26.0 units
The Careful Board, year 10	3	61 months	9.4×	28.2 units

The Careful Board is three times more accurate and twice as expensive. By its tenth year it has improved its accuracy again and become *more* expensive still, because the latency drifted further than the error rate fell. Every number in this table is invented.[2]

> **The design objective is recovery time, not immortality. A system that cannot be captured is not available at any price; a system in which capture is visible, nameable, and cheap to reverse is available, and is cheaper than the pursuit of the first.**

[2] The compounding factors were chosen to make the table legible rather than to model anything. A real version of this table would need its units defined, which would take a section of its own.

Part 3 — Objections that survive

Three objections are strong enough to be worth stating, and none is answered here.

Latency is not one number. Different classes of error surface on different clocks within the same institution. Treating latency as a scalar is a modelling convenience that may not survive contact with a real body `[model-relative]`.

Some errors should not be corrected quickly. A system that reverses itself on every challenge has not lowered its latency; it has abolished its memory. There is a floor below which responsiveness becomes instability, and this specimen does not locate it.

The compounding factor is doing all the work. Premise 3 asserts faster-than-linear growth without saying how much faster. Choose a gentle enough curve and the result disappears. This is the honest weakness of the argument `[defeasible]`.[3]

[3] This is the objection the authors would raise first if the paper were someone else's, which is generally the test worth applying.

Part 4 — What would settle it

Figure 2 A kleroterion in the Stoa of Attalos, Athens. Citizens' tokens were slotted into the columns and a rank of them released at random, so that no faction could arrange who sat. A mechanism for making selection unarguable, cut in stone, two thousand three hundred years before the argument above. Photograph in the public domain.

A dynamical model in which latency and error rate are separately tunable, run to see whether the cost curves cross where the argument says they cross. Until such a model exists, the transfer from dynamical systems to human institutions is carried by the word "like", and the word "like" is not a mechanism [analogy].

The specimen therefore ends where a real paper of this shape would end: with the claim stated precisely enough to be attacked, the weakest premise named, and the experiment that would settle it described but not performed.[4]

[4] The structure — claim, premises, conditional result, surviving objections, the experiment not yet done — is the point of the specimen. It is the shape the canon's long-form papers are expected to take.

References

1. Placeholder, B. (2024). *Latency and the Illusion of Improvement*. Journal of Things That Do Not Exist, 12(3), 44–61.

2. Specimen, A. & Placeholder, B. (2025). *Six Premises About Nothing in Particular*. Proceedings of the Imaginary Society for Decision Structure, 7, 112–140.
3. Fictitious, C. (2023). *On Instruments That Read Right About the Wrong Thing*. Nowhere University Press.
4. Invented, D. (2026). *Against Scalar Latency*. Unpublished specimen manuscript.

Colophon

Set in **Young Serif** for display, **Literata** for text, and **DM Mono** for canon numbers, metadata and sidenotes. All three are licensed under the SIL Open Font License, vendored into the repository with their licences, and embedded here — the same files the website serves.

Typeset with Typst from the markdown source. The markdown is the work; this PDF is a build artifact and is never edited. Corrections go to the source and everything is rebuilt.

This is a **specimen**. It is fictional, it carries a prototype canon number, and it exists to test the typography of the canon rather than to say anything.

DG 001 · PAPER · v1.0 · 2026-08-01