When a Reference Goes Stale: Self-Diagnosis under Shared Sensor Drift

Brian Theory
AI-assisted research draft

13 November 2026

Classification: Technical note. Elementary finite-model results and a research prospectus; no priority claim. Internal AI review is distinct from external peer review.

Abstract

A calibrated observation mechanism need not remain calibrated. We study an actuator whose current reversal is to be diagnosed through two sensors sharing a drifting orientation bit. The environment is known, the target actuator bit is independently refreshed between tasks, and exact reference operations reveal the current sensor orientation at a stated price. Under these restrictive assumptions, routine observations provide no evidence about common orientation drift. We derive the exact diagnosis error as a function of reference age, a sharp maximum calibration interval for a prescribed mean-error bound at every deployed age, and the optimal long-run calibration cost over all history-dependent policies. Periodic calibration is sufficient because the observations relevant to scheduling are ancillary for the drifting orientation. We also derive an explicitly conditional extension for learning the drift rate and show why uncertain anchors, persistent targets, and health-informative sensor disagreement require different models. A separate coupled extension introduces hidden reference polarity and proves an exact reliability surcharge for a declared equal-drift maintenance model, including a certified regime with imperfect restoration and a counterexample to nominal-error reasoning. The results specialize elementary filtering and renewal arguments; unknown sensor polarity and anchor ambiguity already occur in the existing B-Theory corpus. The contribution is a fully reproducible contract for reference freshness, not a general theory of self-awareness or fault diagnosis.

1 Question, baseline, and contribution

Functional self-diagnosis requires evidence about a mechanism assigned to the agent. When the evidence channel changes, an apparent actuator fault may be a sensor-coordinate fault. An independently calibrated reference resolves that ambiguity only at the time and scope for which it is valid. The question here is therefore temporal: how often must an orientation reference be refreshed, and what does that requirement cost?

The actual baseline is stronger than a simple trusted-anchor model. *Causal Self-Diagnosis with Noisy Tests*, version 0.3 of 5 November 2026, explicitly includes unknown body and endpoint orientations, a simultaneous-flip ambiguity, and the insufficiency of a reference transmission without sensor anchoring, in Section 13.2 [1]. Its Section 14 also considers persistent preparation error and finite resources. Accordingly, neither an XOR ambiguity nor the need for independent orientation information is a new contribution here. *Sensor Decisions under Unknown Dependence*, version 0.2 of 9 November 2026, learns a stationary dependence law with known polarities and independent

deployment episodes [2]. The changed assumption in this note is an orientation state that keeps evolving after calibration.

Blind sensor calibration has a substantial independent history. Balzano and Nowak study gain and offset identification from structured signal ensembles; their Section 4 identifies unavoidable scaling conventions [3]. Those linear models do not establish our binary age formula, and our intentionally unstructured fresh targets remove the redundancy on which their positive results rely. The mathematical machinery used below is elementary Bayesian decision theory, a two-state Markov calculation, and cycle-cost accounting. We do not assert that the resulting formulas are historically new.

Reference timing also belongs to the established literature on information freshness. Sun and Cyr explicitly give the symmetric binary Markov aging probability in Section III-B, equations (11)–(12), and optimize nonlinear age penalties in discrete time in Section IV-C, Theorems 3–4 [6]. Their channel model includes positive service times and a sampling-rate constraint; our reference is instantaneous and has a per-use charge. We therefore do not present our renewal formula as a direct corollary of their theorem without changing its timing assumptions. The substantive connection is that a diagnosis error can itself be an age penalty. The underlying binary aging law and optimization of freshness are established antecedents, not new results of this note.

The specific outputs are: an explicit observational symmetry; an exact reference-age risk; an optimal renewal characterization under a stated information restriction; and failure cases identifying the assumptions that would have to change in a more realistic experiment. They support a technical addendum to the diagnostic research family. They do not yet justify a separate flagship research paper.

2 Operational contract

Time is indexed by deployment tasks $t = 0, 1, \ldots$. All logarithms in statistical bounds are natural. The physical and information variables are as follows.

- $B_t \in \{0, 1\}$ is the actuator reversal on task t. The sequence is independent and fair, and independent of all sensor variables. This is a supplied task ensemble, not a claim about the typical persistence of faults.
- $C_t \in \{0, 1\}$ is a common sensor-orientation state. Initially it is fair, and $C_{t+1} = C_t \oplus F_{t+1}$, where the F_t are independent Bernoulli(p) variables. Initially p is known and $0 < p < 1/2$.
- The environment contributes a known zero reversal. A command A_t produces the binary physical output $A_t \oplus B_t$. The two sensors return $Y_{it} = A_t \oplus B_t \oplus C_t \oplus N_{it}$, $i = 1, 2$.
- The N_{it} are independent Bernoulli(ϵ) variables, $0 \le \epsilon < 1/2$, independent over tasks and of everything else. Both sensors share C_t; independent readout noise does not imply independent calibration errors.
- Before issuing the current command or seeing its readings, the controller may pay $c \ge 0$ for an exact anchor $R_t = C_t$. This operation neither changes B_t, C_t nor advances the task clock. Its physical realization is an assumption.
- The controller then chooses A_t, receives both readings, and reports $\widehat{B}_t$. The loss is $\mathbf{1}\{\widehat{B}_t \neq B_t\} + c\mathbf{1}\{\text{anchor used}\}$. The next orientation transition happens only after the report.

Policies may depend on the complete retained action, reading, and anchor record, and on independent private coins. They cannot see B_t, C_t without anchoring, or future flips. Computation and record storage are unpriced; this note makes no computational-efficiency claim. The comparison is among policies with the same prior, known parameters, two mandatory readings per task, and report timing. A known-current-orientation oracle has a strictly stronger information interface and is labeled accordingly.

The target is the *current* actuator bit, not the original actuator state. There is no repair action in the core model. Calibrating the readout is not repairing the actuator. The self/world boundary and reference accuracy are assigned by the model; neither is inferred from a correlation.

Write $V_t = B_t \oplus C_t$ and $W_{it} = Y_{it} \oplus A_t$. The command cancels, so arbitrary adaptive commands cannot improve this task. If the last anchor occurred at $s \le t$, define its age as $k = t - s$. A task on which an anchor is purchased has age zero.

3 Ambiguity and what the second sensor can establish

Proposition 3.1 (Common-orientation ambiguity; elementary specialization)**.** *Without any anchored orientation, the complete routine transcript has the same law under* $(B_t, C_t)_{t\ge 0}$ *and* $(B_t \oplus 1, C_t \oplus 1)_{t\ge 0}$*. Under the supplied fair prior, every routine-only actuator estimate has Bayes error at least* $1/2$*. The pair of routine readings on each task is independent of the entire orientation trajectory, even though it carries information about* $B_t \oplus C_t$*.*

Proof. Flipping both coordinates leaves $B_t \oplus C_t$ unchanged and leaves the orientation-flip increments unchanged. The transition and observation kernels, and the prior, are invariant. Equal past transcripts induce identical action distributions; induction therefore preserves the equality for adaptive randomized policies. For any fixed orientation trajectory, the fresh bits $B_t \oplus C_t$ are independent fair bits. Integrating the known readout-noise channel yields a routine-transcript distribution independent of that trajectory, including adaptive commands after their known offsets are removed. Consequently C_t remains fair without an anchor. Given any routine observations, $B_t = V_t \oplus C_t$ is fair, proving the Bayes claim. □

This proposition is included to expose a contract, not to claim a new gauge principle. It is directly analogous to the already stated orientation ambiguity in B07. It also explains an empirical trap: agreement between sensors can establish a common reading while giving no evidence that their common coordinate system is correct.

With a fresh exact anchor, two equally accurate independent sensors have Bayes error $e = \epsilon$. To check this, on agreement the preferred bit is their common value; on disagreement the posterior is tied. Thus the error is

$$e = \epsilon^2 + \tfrac{1}{2}\{2\epsilon(1-\epsilon)\} = \epsilon. \tag{1}$$

Two readings do change conditional confidence: agreement error is

$$e_{\text{agree}} = \frac{\epsilon^2}{(1-\epsilon)^2 + \epsilon^2}, \qquad e_{\text{disagree}} = \tfrac{1}{2}. \tag{2}$$

Therefore the equality of *average* error with one sensor is not equivalence of the experiments for every decision problem. An odd number of independent readings would reduce e by ordinary majority voting, without removing uncertainty about C_t.

4 Exact reference aging

Theorem 4.1 (Age-dependent diagnosis risk; proved here as an elementary calculation)**.** *Under Section 2, the minimum expected classification error on a task whose latest exact anchor has age k is*

$$r_k = \frac{1-(1-2e)(1-2p)^k}{2}, \qquad e = \epsilon. \tag{3}$$

The statement holds conditional on any admissible past that schedules the task at that age. In particular, a history-triggered calibration rule cannot use previous routine data to obtain a lower risk at the same age.

Proof. The parity $D_k = F_{s+1} \oplus \cdots \oplus F_{s+k}$ has sign expectation

$$\mathbb{E}(-1)^{D_k} = \prod_{j=1}^{k} \mathbb{E}(-1)^{F_{s+j}} = (1-2p)^k.$$

Hence its probability of one is $q_k = [1-(1-2p)^k]/2$. Future flips are independent of the most recent anchor value and of all routine data. In particular, deciding not to anchor on intervening tasks does not select unusually stable orientations: the decision uses only information independent of those flips. The event defining a given age may depend on earlier anchor values, but the most recent anchor restarts the independent flip process.

For any current reading pair w, let $z(w) = \mathbb{P}(V_t = 1 \mid w) - \mathbb{P}(V_t = 0 \mid w)$. Its sign specifies the optimal estimate of V_t. Mixing over the orientation parity multiplies this posterior contrast by $(1-2q_k)$, up to the known sign of C_s. Since $q_k \le 1/2$, the sign does not reverse. Conditional error is $[1-(1-2q_k)|z(w)|]/2$. Averaging and using $\mathbb{E}|z(W_t)| = 1-2e$ proves (3). This proof also covers disagreement, where $z = 0$. □

Corollary 4.2 (Maximum interval for a mean-error requirement)**.** *If $e \le \alpha < 1/2$, a policy that must have mean classification error at most α at every deployed age can use only ages*

$$k \le k_{\max} := \left\lfloor \frac{\log((1-2\alpha)/(1-2e))}{\log(1-2p)} \right\rfloor. \tag{4}$$

Periodic calibration every L tasks satisfies the requirement exactly when $L \le k_{\max} + 1$.

Proof. The risk requirement is equivalent to $(1-2p)^k \ge (1-2\alpha)/(1-2e)$. Taking logarithms reverses the inequality because $\log(1-2p) < 0$. The ages in a length-L cycle are $0, \ldots, L-1$ and the risk increases with age. □

This is a guarantee about average error over the two readings and the supplied target ensemble. It is not a bound on every realized posterior error: disagreement has posterior error $1/2$ even at age zero. When $\alpha < e$ the fixed observation interface is insufficient. When $\alpha \ge 1/2$ no calibration is required for this criterion.

5 Optimal calibration timing

For an integer $L \ge 1$, periodic calibration has average cost

$$\begin{aligned} g(L) &= \frac{c}{L} + \frac{1}{L}\sum_{k=0}^{L-1} r_k \\ &= \frac{1}{2} + \frac{1}{L}\left[c - \frac{(1-2e)\{1-(1-2p)^L\}}{4p}\right]. \end{aligned} \tag{5}$$

Never calibrating has error $1/2$ and zero calibration cost. Let

$$g_* := \min\left\{\frac{1}{2}, \inf_{L\geq 1} g(L)\right\}. \tag{6}$$

For a permitted policy π, write $I_t = 1$ if it buys an anchor at task t and $I_t = 0$ otherwise. Its expected upper-average loss is

$$J(\pi) = \limsup_{T\to\infty} \frac{1}{T}\,\mathbb{E}_\pi\left[\sum_{t=0}^{T-1}\left(\mathbf{1}\{\widehat{B}_t \neq B_t\} + cI_t\right)\right].$$

The expectation is taken before the upper limit. Replacing the upper limit by a lower limit of the expected averages gives the same optimal value: the proof gives a finite-horizon lower bound for every policy and an attaining policy whose expected averages converge.

Theorem 5.1 (No benefit from ancillary scheduling)**.** *For $0 < p < 1/2$ the infimum of $J(\pi)$ over all policies in Section 2 equals g_*. If $c < (1-2e)/(4p)$, some finite renewal period attains the value. If $c \geq (1-2e)/(4p)$, never calibrating is optimal.*

Proof. At each task, conditioning on the scheduler's available information before the current readings gives the minimum age cost in Theorem 4.1. Every reporting rule has conditional classification risk at least that minimum, with equality for a Bayes report. By the tower property, expected total classification cost is therefore at least the expectation of the sum of these age costs. This does not condition on the entire realized future schedule, which could contain information from future anchors. For a finite run, group the age costs pathwise into consecutive calibration segments. Every segment of length L, including the final truncated one, has cost $c + \sum_{k=0}^{L-1} r_k = Lg(L) \geq Lg_*$. Any initial segment before the first anchor has cost $L/2 \geq Lg_*$. The inequality holds for every sequence of segment lengths, hence also after expectation over randomized scheduling. Every finite-horizon expected average is therefore at least g_*, even if a policy has no limiting average.

A fixed finite period attains $g(L)$ by direct summation, and the never-anchor policy attains $1/2$. Put $A = (1-2e)/(4p)$. If $c \geq A$, equation (5) gives $g(L) \geq 1/2$ for every finite L. If $c < A$, then for sufficiently large L, $c - A\{1-(1-2p)^L\} < 0$, so some $g(L) < 1/2$. Since $g(L) \to 1/2$, the infimum below $1/2$ must be attained in a finite set of periods. This supplies an attaining policy. □

The cycle argument is a specialization of ordinary average-cost control, not a general claim that periodic recalibration is best. It relies on all scheduler observations being uninformative about the unobserved orientation increments. A sensor-health alarm, a persistent actuator target, unequal orientation-dependent noise, or an anchor taking time invalidates that premise.

At $p = 0$, a single exact calibration followed by no further anchors gives limiting risk e at any finite initial charge. The formula dividing by p must be read by continuity, but "never anchor" and "anchor once" remain different policies. At $p = 1/2$, only age zero benefits: $g(L) = 1/2 + (c + e - 1/2)/L$. At $e = 1/2$ the readings carry no actuator information even with calibration. At $c = 0$ anchoring every task is optimal. These cases make clear why a frequency statement needs its parameter domain.

6 Learning a drift rate and deciding when to abstain

6.1 A finite training contract

Suppose p is unknown but constant. For an integer $m \geq 1$ and $0 < \delta < 1$, perform exact reference observations at $m+1$ consecutive task boundaries, yielding m observed independent flip indicators.

With their mean $\widehat{p}$, set

$$p_U = \min\left\{\frac{1}{2}, \widehat{p} + \sqrt{\frac{\log(1/\delta)}{2m}}\right\}. \tag{7}$$

Hoeffding's one-sided bound [4, Theorem 2] implies $\mathbb{P}(p \le p_U) \ge 1 - \delta$. For fixed k, r_k is increasing in p on $[0, 1/2]$. When $p_U < 1/2$, a frozen period satisfying Corollary 4.2 at p_U therefore meets the mean-error requirement at every deployed age on that one good-training event. At the clipping boundary $p_U = 1/2$, the logarithmic formula is not used: if $e \le \alpha < 1/2$, choose period one, since age zero has risk e and every positive age has risk $1/2$ at the upper-bound model. If $\alpha < e$, the fixed observation interface cannot meet the requirement. The training anchors cost $(m+1)c$; their cost is not hidden in deployment performance. This is a standard concentration wrapper, not an optimal sample-complexity theorem.

The per-task mean guarantee is conditional on the parameter being covered. It does not say that no deployment error occurs. A finite number of correct calibrations cannot rule out later changes in p, arbitrary adversarial polarity drift, or a defective reference. Time-varying rate estimation and sparse-anchor experimental design remain separate problems.

6.2 Abstention uses conditional risk

Suppose an abstention costs $a \in (0, 1/2)$ and an incorrect report costs one. For the symmetric sensor pair, disagreement must be rejected because its conditional error is $1/2$. On agreement the conditional error at age k is

$$r_k^{\text{agree}} = \frac{1 - (1 - 2e_{\text{agree}})(1 - 2p)^k}{2}. \tag{8}$$

The optimal terminal rule accepts agreement exactly when $r_k^{\text{agree}} \le a$. It is incorrect to threshold the unconditional r_k and claim a per-report guarantee. The coverage on an accepted age is $(1-\epsilon)^2 + \epsilon^2$; it is zero at rejected ages. Optimizing calibration jointly with selective coverage is an extension with a new cycle reward, not a consequence of (5) without recomputation.

Abstention from attribution need not prohibit action. If all currently compatible explanations permit the same shutdown or reference-restoration action, it can be selected without deciding which actuator label is true. The general common-action principle is already explicit in B07, Section 7.1; no new theorem is claimed for it here.

7 Reproducible checks and worked example

The supplied `code/reproduce.py` uses rational arithmetic to enumerate the two sensor noises, current target, and orientation parity. It checks (3) at four values each of p and ϵ and ages zero through twelve: 208 checks. It separately enumerates all binary calibration schedules for horizons one through eight and compares them with an age-state Bellman recurrence. These checks test the timing convention and implementation; the proofs establish the unlimited-age and all-policy statements.

For $p = 1/50$, $e = 1/10$, and $c = 1/2$, equation (3) gives $r_0 = 0.1$, $r_{10} \approx 0.234067$, and $r_{20} \approx 0.323199$. A mean-error requirement of 0.2 permits maximum age seven and hence period at most eight. The plotted renewal curve and tabulated values are evaluated from exact rational expressions before conversion to decimal. A finite search through period 200 is identified as such in the output; its minimum is corroborative, not a substitute for the global cycle theorem.

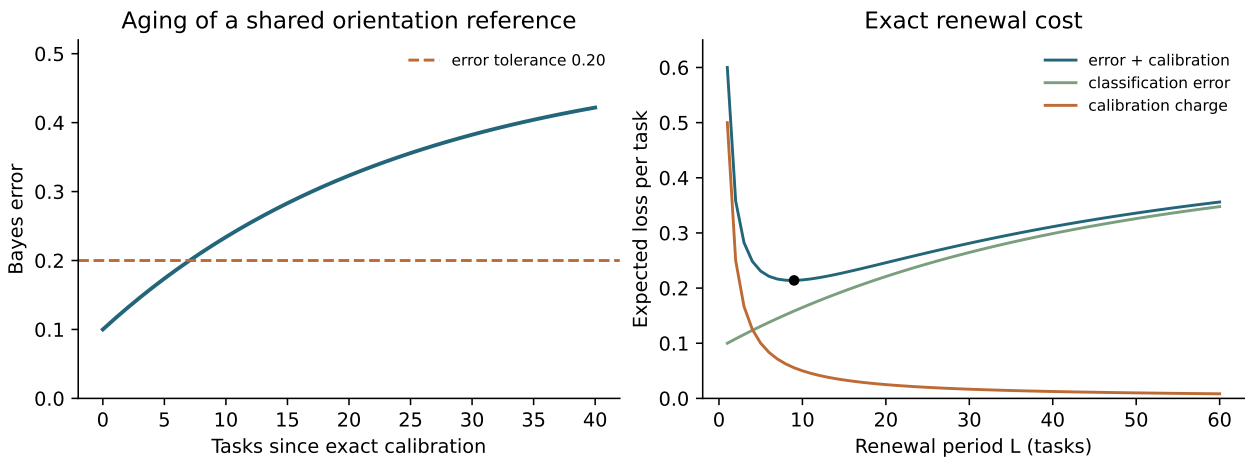


Figure 1: Original analytic plots for the declared example. Left: diagnosis risk as the exact anchor ages. Right: the renewal objective decomposed into error and calibration cost. The marked point is the best of periods 1–200. These are model calculations, not device measurements. Editable SVG and generating code are included.

8 Failure boundaries and unresolved investigations

A fixed actuator is a different problem. If $B_t = B$ forever and a noiseless anchored measurement identifies it once, later sensor drift does not erase the retained actuator label. Reusing formula (3) would incorrectly predict a need to relearn a bit already known. Fresh independent targets are the reason routine histories cannot track C_t here. A persistent but changing actuator requires joint filtering of actuator and sensor changes.

An imperfect reference is not exact renewal. If the latest anchor has independent orientation error η, a decoder using only that anchor has risk $[1 - (1 - 2e)(1 - 2\eta)(1 - 2p)^k]/2$. This follows by multiplying sign correlations. It is *not* the full-history optimum: earlier imperfect anchors can provide additional information. If a reference shares an unknown permanent orientation with the sensors, the initial symmetry persists and absolute diagnosis may remain impossible.

Repetition does not manufacture independence. Reusing one readout-noise bit gives one corrupted observation, irrespective of the number of times it is read. Repeated common-mode observations cannot replace an independent reference. Conversely, structured task distributions can make blind calibration possible, as in the linear signal assumptions of Balzano and Nowak [3]; the impossibility here is not universal.

General joint drift is unresolved. A useful next model has persistent B_t, separate sensor orientation chains, and imperfect calibrations of known reliability. Its posterior is a finite-state hidden Markov filter. One should compare exact finite-horizon belief-state control with periodic and alarm-triggered calibration under identical readout costs. No optimality theorem for those extensions is asserted. A second question is whether a lower bound on calibration expenditure survives a robust class of unknown flip rates, with the physical reference itself degrading. These are explicit research proposals.

9 Hidden reference polarity and a coupled maintenance theorem

This section is a new, deliberately narrow D03–D05 coupling. It changes the action and information contract; none of its formulas should silently replace the exact-anchor results above. A reference being physically operational is not the same as its polarity being known.

9.1 Filtering a reference that can itself drift

Let C, H be sensor and reference polarity bits, with independent increments $F \sim \text{Bernoulli}(p)$ and $G \sim \text{Bernoulli}(q)$ between ticks, $0 < p, q < 1/2$. A noiseless but potentially misoriented reference returns $Z = C \oplus H$. Initially C, H are independent fair bits. Routine data still concern independent fair targets and reveal no additional orientation information. Observing Z confines the joint belief to two states. Write $m = \mathbb{E}[(-1)^C \mid \mathcal{F}]$ after observing Z, where $\mathcal{F}$ is the complete available record. Between observations, $J = Z_{t+1} \oplus Z_t = F \oplus G$. Direct Bayes conditioning gives

$$m_{t+1} = m_t \lambda_J, \quad \lambda_0 = \frac{1-p-q}{1-p-q+2pq}, \qquad \lambda_1 = \frac{q-p}{p+q-2pq}. \tag{9}$$

Indeed F, G are independent of the old belief, and the two multipliers are $\mathbb{E}[(-1)^F \mid F \oplus G = J]$. The numerator for $J = 1$ is $(1-p)q - p(1-q) = q - p$. With a noiseless routine reading $V = B \oplus C$ of a fresh fair target, minimum classification error is $(1 - |m|)/2$. With the original independent symmetric sensor pair it is $[1 - (1-2e)|m|]/2$. Negative multipliers reverse the preferred orientation; taking $|m|$ is essential.

Proposition 9.1 (Unanchored hidden-reference obstruction)**.** *With the above fair initialization and no symmetry-breaking operation, all routine and reference records leave the current target's Bayes error equal to* $1/2$*, even with a reference reading at every tick.*

Proof. After the first Z, $m = 0$, and (9) preserves zero. Equivalently, jointly complementing C, H and every fresh target preserves the prior, all increment laws, and all observations. The two target values therefore have equal posterior probability. Adaptive actions whose effects and observations preserve this symmetry do not change the argument. □

Thus a physical service cycle preserving unknown reference polarity does not establish diagnostic capacity. Nor can repeated noisy references generally be reduced to the latest reference age: with fixed C, independent reference error $1/10$, and three readings, majority error is $7/250$, below the single-reading error $1/10$. This is an elementary counterexample to a latest-only sufficient statistic, not a novelty claim.

9.2 A priced symmetry-breaking operation

To obtain an exact resource comparison, couple the hidden bits to the following D05 physical counters. At the start of a tick let $x = (r, s, a, b)$, with integer store $0 \le r \le R$, service lives $1 \le s, a \le 4$, and time since probe $0 \le b < 4$. Fix positive integers k, d, R, g. There are four ordinary actions:

$$\begin{aligned}
\mathsf{Harvest}(x) &= (\min(R, r+g), s-1, a-1, b+1),\\
\mathsf{Sensor}(x) &= (r-k, 4, a-1, b+1),\\
\mathsf{Actuator}(x) &= (r-k, s-1, 4, b+1),\\
\mathsf{Probe}(x) &= (r-k, s-1, a-1, 0).
\end{aligned}$$

Harvest needs $a \geq 1$; sensor service needs $a \geq 1$; actuator service and probe need $s \geq 1$. Charged actions require sufficient resource. Every successor must remain in the displayed safe counter domain. These are the original D05 transitions, copied here to make the coupled contract self-contained; the original D05 files are unchanged.

At every tick, before action choice, the controller freely observes $Z = C \oplus H$. A fresh independent fair B is generated before that choice, but $V = B \oplus C$ is available only on a probe. Service actions do not reveal or reset either hidden bit. A probe may be augmented by a trusted restoration, denoted Probe^+: immediately before its reading it sets $H = 0$ and rereads Z, so C is known. It costs $k + d$ rather than k and has exactly the same counter successor and duration as a probe. No hidden-state transition occurs between this restoration and reporting. It neither changes C nor reveals B except through V. This perfect, instantaneous, polarity-setting intervention is an explicit supplied capability, not inferred reliability of a faulty reference. It is available only as part of Probe^+.

After the action and report, both hidden bits independently flip with the *same* probability $p \in (0, 1/2)$. Record the post-restoration Z when restoration occurs; otherwise retain the ordinary current Z. The next J compares the next reading with this post-action value. Counters are observed exactly; hidden bits are not. Computation, free Z sampling and memory are unpriced.

Require counter safety indefinitely with probability one, and conditional mean target error at each selected probe at most α, $0 \leq \alpha < 1/2$, for every positive-probability observation/action history. Conditioning includes which action was selected. Averaging a reliable and an unreliable action over private coins is not a substitute for this requirement. There is no abstention; $b < 4$ requires a probe in each block of four actions. This is stronger than a long-run average-error objective.

Theorem 9.2 (Exact reliability surcharge in the equal-drift family)**.** *The coupled model admits some indefinitely viable initial counter state if and only if*

$$\min\{g, R\} \geq 3k + d. \tag{10}$$

If only counter safety and recurring probes are required, with no diagnostic reliability constraint, nonemptiness instead holds exactly when $\min\{g, R\} \geq 3k$*. Hence the two feasibility regions differ strictly.*

Proof. For sufficiency, start at $(0, 2, 3, 0)$ and repeat $\mathsf{Harvest}, \mathsf{Sensor}, \mathsf{Actuator}, \mathsf{Probe}^+$. The service lives and monitor follow $(2, 3, 0), (1, 2, 1), (4, 1, 2), (3, 4, 3), (2, 3, 0)$. After every harvest the store contains at least $3k + d$ whenever (10) holds. The three other actions consume exactly $3k + d$. Every probe is exact by restoration, so any allowed α is met.

For necessity put $h = \min\{g, R\}$ and suppose $h < 3k + d$. When $p = q$, equation (9) gives $\lambda_1 = 0$, regardless of the earlier belief. Thus on any tick with $J = 1$, an ordinary probe has conditional error $1/2$ and is forbidden by the reliability requirement. A probe on such a tick must be Probe^+.

Consider a run of n consecutive such ticks after an initial transition. Each event has conditional probability $\beta = 2p(1-p) > 0$, independent of the earlier action history, using the post-action reference value as specified above. Consequently the finite run has probability $\beta^n > 0$. On that run, every probe incurs the surcharge. For any safe n-action segment, counts obey $N_S, N_A, N_P \geq n/4 - 1$ because each service and probe is required in every four consecutive actions. Also $N_H = n - N_S - N_A - N_P \leq n/4 + 3$. Resource accounting then implies

$$(3k + d)(n/4 - 1) \leq kN_S + kN_A + (k + d)N_P \leq R + hN_H \leq R + h(n/4 + 3).$$

Therefore $(3k + d - h)n/4 \leq R + 3h + 3k + d$. Choose a finite integer n violating this inequality. No policy can both remain safe and satisfy the reliability constraint on the positive-probability run,

contradicting almost-sure indefinite viability. The proof does not assume that an infinite all-change run has positive probability. Without the error constraint the same accounting uses ordinary probes and gives threshold $3k$; the corresponding four-action cycle attains it, as in D05. □

For example, $k = d = 1$, $R = g = 3$ permits the original recurring-probe cycle but no reliable coupled policy; $R = g = 4$ permits reliable indefinite operation. The stronger threshold is not a universal cost of knowledge. It depends on equal flip rates, hard historywise error control, bounded storage, the four-tick obligations, and the particular trusted restoration action. Unequal rates leave $\lambda_1 \neq 0$ in general. Average-error constraints, imperfect restoration, or paid reference sampling require new analysis.

9.3 Prior art and evidence status

Jointly maintaining information channels is not a new research theme. Luo and colleagues study transmission, waiting and renewal of a wearing channel in remote estimation [5]; their linear-Gaussian estimation objective and restoration downtime differ from the binary hard-reliability contract here. Van Oosterom and colleagues already study a system and its deteriorating imperfect sensor using inspection and maintenance [7]; their full accepted manuscript was read in this second audit. Section 2 uses age-indexed observation matrices ordered by Blackwell degradation, a Bayesian system-state belief, and perfect instantaneous inspections. Its conclusion explicitly proposes randomly changing observation relations and dependence on earlier readings. Thus hidden reference health lies within an already announced direction; our distinction is the explicit hard-reliability resource contract, not the general research programme. The contribution claimed here is the supplied finite-model proof and counterexample, not historical priority for coupling sensing and maintenance. The equal-drift reset of posterior contrast is what makes the resource necessity proof exact. It avoids any uncertified belief-grid approximation. The accompanying exact-arithmetic and finite-state checks corroborate the argument; they do not replace its all-policy proof. Reviews are internal AI reviews, not external human peer review.

10 Imperfect restoration: a certified range and a failure case

The previous theorem supplies perfect symmetry breaking. To test that assumption, replace only the restoration operation: it sets H to an unobserved fresh $E \sim \mathrm{Bernoulli}(\eta)$, $0 < \eta < 1/2$, independent of the old state and all other randomness. It then rereads $Z = C \oplus E$. The controller knows η, but not whether restoration succeeded. A single such operation is available per augmented probe; it has the same surcharge d and timing as before. Physical service lives remain observed. This models a calibrated stochastic reset, not arbitrary imperfect repair or an empirically certified mechanism.

Here reliability is required conditional on the entire report-time record, including the selected action, the restored reference readout, and the current routine readout. This strengthens any interpretation that averages over the restored reference result. The fresh fair target makes the routine readout independent of orientation, so it does not alter the orientation posterior. Put $a_\eta = 1 - 2\eta$. For a pre-restoration orientation contrast m and restored sign $z = (-1)^Z$, Bayes' rule gives

$$m^+ = \frac{m + z a_\eta}{1 + z a_\eta m}, \qquad \mathbb{P}(z \mid m) = \frac{1 + z a_\eta m}{2}. \tag{11}$$

Both denominators are positive. The exact conditional diagnosis error is $(1 - |m^+|)/2$. If one ignores the old record, the reference has average error η; that is not a bound on each conditional error after combining the record with a conflicting readout.

Proposition 10.1 (Reachable failure of nominal-error reasoning)**.** *Take $p = q = \eta = 1/10$, initially fair orientation, and the four-action restored-probe cycle. After a first restored sign $+1$, four successive $J = 0$ transitions, and an opposite restored sign, the report-time Bayes error is*

$$\frac{541529}{1319290} > \frac{2}{5},$$

despite nominal reset error $\eta = 1/10$.

Proof. The first posterior contrast is $4/5$. Each $J = 0$ multiplies it by $40/41$. Before the next reset it is $m = (4/5)(40/41)^4 = 2048000/2825761$. Substitution of $z = -1$ in (11) gives the stated error. All events have strictly positive probability: the initial sign has probability $1/2$, each no-change event has probability $41/50$, and the final conflicting sign has probability $(1 - (4/5)m)/2 > 0$. Physical feasibility can be ensured with $k = d = 1, R = g = 4$. Thus this is a reachable failure within the priced cyclic interface, not merely an arbitrary prior. □

Theorem 10.2 (A certified imperfect-restoration regime)**.** *Retain equal drift $p = q \in (0, 1/2)$ and define*

$$u = \frac{1 - 2p}{1 - 2p + 2p^2}, \qquad v = u^4, \qquad A_4 = \frac{\eta(1 + v)}{1 - (1 - 2\eta)v}.$$

If $v < 1 - 2\eta$ and $A_4 \le \alpha < 1/2$, the imperfect-restoration model has a nonempty reliable recurrent kernel exactly when $\min(g, R) \ge 3k + d$. If $0 \le \alpha < \eta$, that kernel is empty at every resource level. No classification of the intermediate range is asserted.

Proof. In the four-action cycle, four hidden transitions separate consecutive restored probes. Their contrast multipliers are either u or zero, so the pre-restoration contrast satisfies $|m| \le v$, independently of prior reset outcomes. Before the first probe the contrast is zero under the supplied fair initialization. When $v < a_\eta$, formula (11) gives

$$|m^+| \ge \frac{a_\eta - v}{1 - a_\eta v}, \qquad \frac{1 - |m^+|}{2} \le \frac{\eta(1 + v)}{1 - a_\eta v} = A_4.$$

The lower bound follows by choosing the conflicting sign and maximizing $|m|$ up to v; differentiation of $(a_\eta - x)/(1 - a_\eta x)$ gives $(a_\eta^2 - 1)/(1 - a_\eta x)^2 < 0$. Thus every possible report in the affordable restored-probe cycle meets the requirement. Resource sufficiency is exactly the previous cycle argument.

For necessity, a tick with $J = 1$ still erases the old contrast. An ordinary probe then has error $1/2$; an imperfectly restored probe has error exactly η for either restored outcome. For any $\alpha < 1/2$, every allowed probe on an all-change run must therefore pay restoration. The finite-run resource argument of Theorem 9.2 gives the same necessary bound $\min(g, R) \ge 3k + d$. If $\alpha < \eta$, no probe on such a run is permitted at all. Four consecutive all-change ticks have positive probability and force a probe by the monitor, contradicting infinite feasibility regardless of resources. □

For $p = 1/4$ and $\eta = 1/50$, $u = 4/5$, $v = 256/625$ and $A_4 = 881/18962 \approx 0.04646$. Thus $\alpha = 1/20$ admits the sharp resource threshold even though restoration is not perfect. The value A_4 is a conservative sufficient bound; its attainability or optimality over all belief policies is not claimed. The theorem does not prove feasibility or infeasibility throughout $\eta \le \alpha < A_4$, nor when $v \ge 1 - 2\eta$.

Validation and what remains assumed. The supplied script enumerates hidden-state probabilities independently of the posterior formula, checks the counterexample exactly, and expands all observation branches of a finite prefix of the restored-probe cycle. This is synthetic model validation, not device data. The infinite-horizon assertion rests on the uniform bound and resource proof. The stochastic reset remains an externally calibrated, independent channel; correlated repair faults, uncertain η, unequal drift and time-consuming restoration are unresolved. D05's separate paid exact-reference action consumes a full tick and obeys a different theorem. Neither its $4k$ threshold nor the present instantaneous-surcharge threshold transfers across those timing contracts.

11 Being, Becoming, The MIND, and publication status

The operational capability studied here is recognizing the limited lifetime of evidence used for self-attribution. "Being" refers to the presently available calibrated diagnostic interface; "Becoming" includes changes in that interface and its restoration; "The MIND" refers to the controller's actionable record of reference age and uncertainty. Those interpretations are not additional mathematical results, and no inference to phenomenal consciousness, personal identity, or a uniquely natural self/world boundary follows.

D08 concerns computation quality; this note concerns physical evidence validity. D05 can incorporate reference age while charging the resources needed for recalibration. The preferred publication role is a diagnostic addendum or a section of an integrated D03/D05 study. General imperfect-reference control, device validation, and external priority review remain open; Section 9 solves only its stated equal-drift special case. The arguments have not undergone formal verification or external peer review.

References

[1] B-Theory research programme. Causal self-diagnosis with noisy tests: Observation laws, irreversible access, and task recovery, 2026. URL `https://www.b-theory.com/read-online/causal-self-diagnosis/`. Version 0.3, 5 November 2026. Byline: Research development draft prepared for Brian Theory; AI-assisted derivation, computation, and independent internal review. Submission authorship undecided.

[2] B-Theory research programme. Sensor decisions under unknown dependence: Interval frontiers and paid calibration, 2026. URL `https://www.b-theory.com/manuscripts-html/sensor-paid-calibration/`. Technical addendum, version 0.2, 9 November 2026; not externally peer reviewed.

[3] Laura Balzano and Robert Nowak. Blind calibration of sensor networks. In *Proceedings of the 6th International Conference on Information Processing in Sensor Networks*, pages 79–88, 2007. doi: 10.1145/1236360.1236372. URL `https://nowak.ece.wisc.edu/BalzanoNowakIPSN07.pdf`.

[4] Wassily Hoeffding. Probability inequalities for sums of bounded random variables. *Journal of the American Statistical Association*, 58(301):13–30, 1963. doi: 10.1080/01621459.1963.10500830. URL `https://www.cs.rpi.edu/academics/courses/spring06/random/hoefding.pdf`.

[5] Jiping Luo, George Stamatakis, Osvaldo Simeone, and Nikolaos Pappas. Remote state estimation over a wearing channel: Information freshness vs. channel aging, 2026. URL `https://arxiv.org/html/2501.17473v2`. arXiv:2501.17473v2, 14 May 2026.

[6] Yin Sun and Benjamin Cyr. Sampling for data freshness optimization: Non-linear age functions. *Journal of Communications and Networks*, 21(3):204–219, 2019. doi: 10.1109/JCN.2019.000035. URL `https://arxiv.org/abs/1812.07241`. Relevant technical statements read in arXiv:1812.07241v4, 3 May 2019; journal metadata checked separately.

[7] Chiel van Oosterom, Lisa M. Maillart, and Jeffrey P. Kharoufeh. Optimal maintenance policies for a safety-critical system and its deteriorating sensor. *Naval Research Logistics*, 64(5):399–417, 2017. doi: 10.1002/nav.21763. Accepted manuscript dated 3 September 2017, obtained from Eindhoven University repository; full text including appendix read in second audit.