Causal Self-Diagnosis with Noisy Tests

Observation Laws, Irreversible Access, and Task Recovery

Research development draft prepared for Brian Theory

AI-assisted derivation, computation, and independent internal review

5 November 2026 Version 0.3

Abstract

A controller can compensate for a malfunction while remaining unable to attribute it to its actuator or environment. We study this distinction in finite diagnostic models with specified experiment permissions. Positive readout likelihoods preserve exact posterior support, so a fixed support-only permission rule restricts all acquisition to the initially available library.

For history-dependent access, initial statistical separation is insufficient. With static known finite-alphabet channels and always-legal idempotent repetition, an oracle revealing each test's complete observation-law class characterizes arbitrarily reliable finite diagnosis. The closure of attainable randomized terminal decision channels is the convex hull of finite legal oracle-tree channels, giving Bayes and minimax limits. Four hidden states are minimal for the exhibited gap between reliable finite-policy families and one ongoing consistent policy. Known persistent nuisance and sampling prices have distinct, explicitly delimited extensions.

We compare order-complex homotopy preservation with a specified recovery task. Neither implies the other; minimal three-state witnesses and a task-dependent availability predicate make this distinction concrete. A conditional implementation bridge states the additional physical obligations. An anchored motor model and independent stateful calculations expose effects of initial motion, persistent preparation error, heat, and depleted resources. The original 48,000 episodes are supplemented by 240,000 independent simulated episodes and exact confirmatory-power calculations.

Five simulated specialist reviews and two cross-reviews sharpen the proofs, source attribution, and experimental scope. A reconstructed historical witness challenges one printed necessity claim under its exact definition. The paper concerns functional diagnosis. It claims neither subjective awareness, hardware validation, external peer review, nor publication priority.

Status. A working paper, not a submitted or externally reviewed manuscript. Author attribution for submission remains undecided. Proved results refer to the arguments and internal checks here, not formal proof-assistant certification.

Contents

1 Useful control without causal self-diagnosis 3

2 Related work and limits of the novelty claim 3

3 A fixed-state diagnostic contract 4

4 What noisy support permits 5

5 Cumulative orders: a finite oracle characterization 6

6 Complete observation laws and attainable decisions 9

7 Self, world, and the decision that needs the distinction 12

8 Cost and risk as parts of the contract 14

9 Exact experiments and independent checks 15

10 From formal permission to physical experiments 17

11 Preserving a specified recovery task 19

12 Observation fidelity and finite-horizon robustness 22

13 An anchored motor experiment: model and test protocol 23

14 Stateful experiments, resource limits, and confirmatory power 26

15 Conclusions and remaining boundaries 30

A Assumption failures and boundary cases 30

B Verification and novelty ledger 32

C Reproduction 33

D A source-audit witness for a printed necessity claim 33

E Historical coverage and remaining originals 35

F Simulated review and revision ledger 35

1 Useful control without causal self-diagnosis

Suppose an action produces the opposite of its intended outcome. The actuator may have reversed, or the environment may have reversed its effect. A controller can learn the resulting input–output relation while remaining unable to attribute the change. That ambiguity matters only for particular downstream tasks: repair, forecasting after an intervention, or transfer may require distinctions that immediate control does not.

We use *causal self-diagnosis* operationally: estimating a variable assigned to the system's own mechanism within a specified causal model and intervention interface. The assignment to self is supplied by the physical model, not recovered from a statistical partition alone. No criterion for consciousness is proposed.

Two B-Theory manuscripts motivate the study. R4 compares actuator-tracking controllers under a known finite model. Its explicit, compiled, and external trackers implement the same mapping; gains over restricted baselines concern evidence processing, not the self-model label [25]. The acquisition report distinguishes a sufficient predictive representation from one that permitted observations can actually acquire [31]. We connect these questions by adding noisy, potentially confounded diagnostic channels and explicit losses.

The central modeling issue is what a test changes. A noisy answer changes belief weights but need not eliminate a state. A test may also change a formal observation order, or an intervention may change the physical system. These are different updates. Treating them as interchangeable can produce an incorrect theorem.

The paper studies diagnostic access under explicit constraints. It gives a support-only reduction, develops an ideal-oracle characterization for cumulative availability, connects identifiable distinctions to regret and interventions, and tests the physical interpretation with counterexamples, an implementation contract, and a proposed motor experiment. Exact computations and internal independent checks accompany the arguments. The support-only theorem and statistical tools are elementary specializations of established ideas. We give a finite decision-tree characterization for static diagnostic models with irreversible query availability and idempotent repeats, combining adapted simulation, finite-state common randomness, and amplification. Section 6 extends this reduction to complete observation laws and terminal channels. The cumulative-order results are derived here as a concrete extension; this does not assert publication priority.

2 Related work and limits of the novelty claim

Model-based diagnosis already treats competing fault explanations and discriminating measurements. Reiter's consistency-based account [27] and de Kleer and Williams' probabilistic measurement selection [8] are direct antecedents. Naming an actuator variable does not create a new diagnosis paradigm.

Controlled sensing and active sequential hypothesis testing select informative actions under error and cost criteria [7, 21, 20]. Pairwise separation by action-conditioned laws is an established identifiability condition. We restrict the executable library and ask for labels grouping hypotheses. This is closely related to equivalence-class determination [11] and decision-region determination [15]. Exact version-space termination eliminates incompatible hypotheses. Our repeated full-support setting permits arbitrarily reliable finite reporting while generally forbidding zero-error finite certification.

Stochastic diagnosis already distinguishes exact fault certainty from asymptotically high confidence [33, 2]. The recovered 2005 original gives a sufficient AA-diagnosability condition and examples showing it is not necessary. Active information acquisition and information-state dy-

namic programming are also established [34]; its erratum makes policies measurable relative to the observations they generate. The recovered unreliable-observation paper [35] motivates the focused source-level check in Appendix D. Later criteria require care about the fault-prefix quantifier [6]; the relevant 2017 full theorem remains unavailable [32]. We claim neither the support/confidence distinction nor active diagnostic sensing as new.

Qualitative POMDP safety uses belief supports rather than positive probability magnitudes [16, 5]. High-probability safe exploration instead uses confidence bounds under additional assumptions [29]. A zero-risk support rule and a small failure-probability rule are different problems. Our topological predicate is not a physical safety theorem.

Identifiability relative to observations and interventions is standard causal reasoning [23, 13]. Our XOR model concerns parameter attribution, not graph discovery or an average causal effect. Predictive state representations explain why useful prediction need not recover a privileged latent description [19]. Robotic self-modeling and damage recovery have substantial experimental precedent [4]; we supply no robot data or architectural superiority claim.

Repetition and majority voting to simulate a noiseless decision tree are established [10]. Comparing a noiseless experiment with its noisy garbling is classical [3]; sequential versions also exist [9]. Renou and Venel explicitly treat restrictions on feasible histories, controlled information, and early actions constraining later access [28]. Their Section 4.3 is a particularly close antecedent. This is a close antecedent to constrained experiment comparison; our claim is a finite specialization with legality-preserving simulation. Recent precedence-constrained decision trees are particularly relevant [30]. Their monotone prerequisites differ from the irreversible incompatibilities in our cumulative inclusion rule. This distinction is not a completed priority assessment.

3 A fixed-state diagnostic contract

Let H be finite and nonempty, with known prior $\pi(h) > 0$ for every state. The hidden diagnostic target is fixed during the experiments in Sections 4–5. Concrete apparatus coordinates may change, as modeled separately in Section 13. A finite library $\mathcal{E}$ contains repeatable tests. Test e has finite alphabet $\mathcal{Y}_e$ and known kernel $K_e(y \mid h)$, with

$$K_e(y \mid h) > 0 \qquad (e \in \mathcal{E},\ y \in \mathcal{Y}_e,\ h \in H). \tag{1}$$

The conditional law of each observation given the entire past, selected test, and state is this kernel. Repetition therefore supplies fresh information; one-time marginal distributions alone are not the assumption.

A policy is one history-dependent selection, stopping, and reporting rule used under every state. Private randomness is independent of the state and noise. All tests and answers are recorded. The target is $\ell : H \to \mathcal{L}$. States with the same label need not share observation laws.

Let b_t be the posterior and $S_t = \{h : b_t(h) > 0\}$ its exact support. A *support-only* admissibility rule is any fixed predicate

$$L : \mathcal{E} \times (2^H \setminus \{\varnothing\}) \to \{0, 1\}.$$

Test e is allowed only if $L(e, S_t) = 1$, evaluated before its answer. No downward monotonicity is assumed. Robust physical safety $L(e, S) = \mathbf{1}\{S \subseteq A_e\}$, for a set of safe states A_e, is one special case; formal structural rules may mean something else.

Definition 3.1. A label is *reliably acquirable* if for every $\delta > 0$ there is a legal policy with a finite deterministic observation bound and $\max_h \mathbb{P}_h(\widehat{\ell} \neq \ell(h)) \leq \delta$. This is a family of finite procedures, not a zero-error stopping certificate.

For finite H and positive prior, vanishing maximum, statewise, and Bayes errors are equivalent for a sequence of policies:

$$\pi_{\min} \max_h \mathrm{err}_h \le \sum_h \pi(h)\mathrm{err}_h \le \max_h \mathrm{err}_h, \qquad \pi_{\min} = \min_h \pi(h) > 0.$$

Guarantees concern the supplied model, not unspecified noise, omitted faults, or a continuum of arbitrarily close alternatives.

4 What noisy support permits

Theorem 4.1 (Support persistence)**.** *After every finite reachable history, $S_t = H$. The executable library is always $\mathcal{E}_0 = \{e : L(e, H) = 1\}$. Conversely, every adaptive policy using only $\mathcal{E}_0$ is legal.*

Proof. Posterior mass is proportional to $\pi(h) \prod_{s=1}^{t} K_{e_s}(y_s \mid h)$, a product of positive terms. Common policy factors cancel because identical histories receive identical selection probabilities under every state, including randomized policies. Substitute $S_t = H$ into L. No monotonicity is needed. □

Define $h \equiv_0 g$ when $K_e(\cdot \mid h) = K_e(\cdot \mid g)$ for every $e \in \mathcal{E}_0$. This equivalence of experimental laws is different from posterior support.

Theorem 4.2 (Labels and residual error)**.** *A label is reliably acquirable exactly when it is constant on every $\equiv_0$ class. If $h \equiv_0 g$ but $\ell(h) \neq \ell(g)$, every legal policy satisfies $\mathrm{err}_h + \mathrm{err}_g \ge 1$. The limiting optimal Bayes error is*

$$R_* = 1 - \sum_{C \in H/\equiv_0} \max_{l \in \mathcal{L}} \sum_{\substack{h \in C \\ \ell(h) = l}} \pi(h). \tag{2}$$

Proof. For an equivalent pair every legal test has the same response law. Induction gives identical adaptive transcript laws, including stopping and reports. The two correct labels are disjoint events under this common law, proving the error bound and necessity.

For sufficiency, cycle through all tests in $\mathcal{E}_0$, use maximum likelihood over H, and report the selected state's label. The bound below tends to zero when all differently labelled pairs are separated. A constant target needs no tests; an empty library can acquire only a constant target.

An oracle revealing the equivalence class cannot worsen performance. Within a class, all transcript likelihoods agree and posterior ratios remain the prior ratios. Its best report is the largest prior label mass, giving (2). Cycling consistently learns the class and approaches this oracle error. □

4.1 Sample bounds and stopping distinctions

For each pair define affinity $\rho_e(h, g) = \sum_y \sqrt{K_e(y \mid h) K_e(y \mid g)}$. With deterministic sample counts n_e, maximum likelihood satisfies

$$\mathbb{P}_h\{\ell(\widehat{h}) \neq \ell(h)\} \le \sum_{g : \ell(g) \neq \ell(h)} \prod_{e \in \mathcal{E}_0} \rho_e(h, g)^{n_e}. \tag{3}$$

A wrong-label maximizer has likelihood at least that of h. Markov's inequality for the square root of its likelihood ratio gives the product of affinities; apply the union bound. Affinity equals one exactly when the two laws agree.

If each test is used m times, define $C_{\min} = \min_{\ell(h)\neq\ell(g)} - \sum_e \log \rho_e(h,g) > 0$ for a nonconstant separated label. Maximum error is at most $(|H|-1)e^{-mC_{\min}}$. Summability shows that reports under one infinite cycling policy eventually remain correct almost surely. The agent need not know when its last error occurred.

In contrast, no nonconstant label has a zero-error procedure terminating almost surely under every state. Finite transcript laws are mutually absolutely continuous, as are finite stopped transcript laws by decomposing an event over the possible stopping times. A label reported with positive probability under one state therefore has positive reporting probability under a differently labelled state. Unbounded expected time does not remove this obstruction when termination remains almost sure.

The results extend to standard Borel observation spaces when each test's laws are mutually absolutely continuous across states. Use a common dominating measure and positive densities almost everywhere. Support persistence then holds almost surely simultaneously over finite times. Equal topological support alone is insufficient (Appendix A).

4.2 A noiseless RF example and an explicit noisy lift

The acquisition manuscript thins a finite order with an oriented Boolean test: retain $u \le v$ only if $q(u) \le q(v)$. The test is admissible when the order-complex inclusion is a homotopy equivalence [31]. For $P = \{a\} + \{x,y\} + \{z\}$, root admissibility is $q(a) \le q(z)$; every test is legal on the middle antichain. This is a structural rule, not a physical safety guarantee.

In coordinate order (a,x,y,z) take

$$q = 0110, \qquad r = 1100, \qquad \ell = (0,1,2,0). \tag{4}$$

Noiselessly, legal query q returns zero on a homogeneous label-zero branch, or one on $\{x,y\}$, where r is legal and separates labels one and two.

Now define a *support-reset noisy lift*: after each noisy answer, evaluate the predicate on the original order induced on exact posterior support. Every query returns its nominal bit through a fresh known BSC with $0 < \varepsilon < 1/2$. Support remains P, only q is available, and its laws at x, y agree. Limiting uniform-prior Bayes error is $1/4$; limiting minimax error is $1/2$. Arbitrarily small positive noise changes feasibility under this exact support convention. No general impossibility of noisy active learning follows.

An additional supplied test $s = 0100$ is root legal. The q, s signatures separate the labels. Repeating each n times and using majority decoding gives error at most $2\exp[-2n(1/2-\varepsilon)^2]$. At $\varepsilon = 1/5, n = 21$ the bound is 0.0456454; using the exact binomial tail gives the sharper union bound 0.00193939.

This lift is a choice: a noisy answer does not make prior tests constant on its support, so the noiseless branch-cancellation reasoning does not apply. Retaining previous structural changes gives a different model.

5 Cumulative orders: a finite oracle characterization

Let P be a finite poset with carrier H, and $\mathcal{E}$ a finite set of nominal Boolean tests $q_e : H \to \{0,1\}$. The physical hidden state is fixed. Let $J \subseteq \mathcal{E}$ record the distinct executed tests. Define an order on the full carrier by

$$u \le_J v \quad \iff \quad u \le_P v \text{ and } q_e(u) \le q_e(v) \text{ for every } e \in J.$$

Test e is legal at J when

$$A(J,e) = \mathbf{1}\{\Delta(P_{J\cup\{e\}}) \hookrightarrow \Delta(P_J) \text{ is a homotopy equivalence}\} = 1. \tag{5}$$

Repeating an executed test is structurally idempotent and legal. Each repetition returns a fresh independent BSC observation with common known $0 < \varepsilon < 1/2$.

The structural carrier is never reduced by a noisy answer. Initial-state support remains full, but the order can change. Theorem 4.1 therefore does not freeze $A(J,e)$.

Example 5.1 (Root separation is insufficient)**.** On the diamond $a < x, y < z$, supply $q = 0110, s = 1101$ and target $\ell = (0,1,2,0)$. Both tests are legal initially, with endpoint words 00 and 11, and their joint signatures separate all different labels.

After q, the strict edges are ax, ay, az, a tree. After s, the edges are ax, az, xz, yz; the filled triangle axz and attached edge make a contractible complex. After both, only ax, az remain and y is isolated. The second distinct test increases the component count from one to two, so is illegal in either order. Repetitions cannot create the missing distinction; neither first test alone resolves every label.

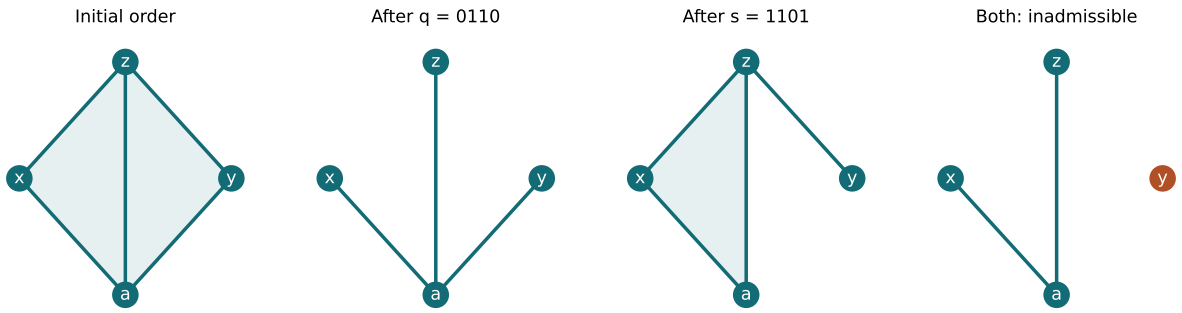


Figure 1: Cumulative-order counterexample. Lines show all strict comparabilities, not only Hasse edges. Chain triangles are filled in the order complex. The final inclusion is inadmissible because y becomes a separate component.

5.1 Keep structural history separate from candidate states

An ideal-oracle problem has state (J, C): J records structural history and C records states consistent with exact nominal answers. Legality remains $A(J,e)$ on the full carrier; it is not recomputed on C. A new test returns $q_e(h)$ and updates J and $C_b = \{h \in C : q_e(h) = b\}$.

Let $\mathcal{T}(C)$ mean that ℓ is constant on C. Define

$$W(J,C) = \mathcal{T}(C) \vee \exists e \in \mathcal{E} \setminus J : A(J,e) = 1 \ \wedge \bigwedge_{\substack{b\in\{0,1\} \\ C_b \neq \varnothing}} W(J\cup\{e\}, C_b). \tag{6}$$

Each recursive step enlarges J, so the recursion terminates. A test constant on C is allowed: it can change structure and unlock a later test without adding nominal information on that branch.

Theorem 5.2 (Arbitrarily reliable finite diagnosis)**.** *The noisy cumulative-order model reliably acquires ℓ if and only if $W(\varnothing, H)$ holds. The equivalence also holds if noisy policies may have almost surely finite stopping times under every state.*

Proof. If W holds, select its finite ideal tree. Every path uses at most $m = |\mathcal{E}|$ distinct tests. At a node, execute its prescribed test an odd number k of times and branch on the majority. First execution is legal by the tree; repetitions are idempotent. Even after an erroneous earlier majority, future structural choices remain legal because legality depends on J, not the true state or inferred C. If a majority selects a nominally empty child, stop and return an arbitrary label. This makes a total policy on every possible noisy history.

Conditioned on the past and true state, a majority error has probability at most $\exp[-2k(1/2-\varepsilon)^2]$. If no majority is wrong, the ideal tree is followed and returns the correct label. Thus maximum error is at most $m\exp[-2k(1/2-\varepsilon)^2]$, using at most mk samples.

Conversely, take a noisy policy with maximum error $\delta < 1/|H|$. Simulate it using the ideal oracle: on the first request for e, obtain and cache $q_e(h)$; every requested noisy response is then generated from the cached bit and a fresh BSC coin. Repetitions require no additional oracle query. Structural histories and legality match the noisy execution.

Run this simulation separately for all finitely many states using one common random tape. Each run has its correct marginal law; runs need not be independent. The union bound gives positive probability that all state runs are correct. Almost sure termination under every state gives simultaneous termination with probability one. Fix a tape on which every run terminates correctly. Prune cached repetitions and deterministic computation. This leaves a deterministic ideal tree with at most $|\mathcal{E}|$ new tests per path. Each nonempty branch corresponds to an actual state and terminates. The remaining legal choices witness (6). □

This is a characterization given a legality predicate, not an efficient general homotopy-equivalence algorithm. Its (J, C) state space is exponential. The theorem concerns a family of finite policies with arbitrarily small error. It does not automatically imply a single indefinitely running consistent policy: a mistaken irreversible branch can permanently remove useful access.

Corollary 5.3. *If $W(\varnothing, H)$ fails, every almost surely terminating noisy policy satisfies $\sum_h \mathrm{err}_h \geq 1$, so maximum error is at least $1/|H|$ and Bayes error at least $\pi_{\min}$.*

Proof. A sum below one would give a positive-probability common tape correct for all states, and hence a successful ideal tree. □

5.2 Limiting Bayes decision value

Let $\mathcal{D}$ be finite and $g(d, h)$ bounded. In this subsection samples have no cost, and the allowed sample budget grows. Define the unnormalized value

$$V(J,C) = \min\Bigg\{ \min_{d\in\mathcal{D}} \sum_{h\in C} \pi(h)g(d,h), \\ \min_{\substack{e\notin J \\ A(J,e)=1}} \sum_{\substack{b\in\{0,1\} \\ C_b\neq\varnothing}} V(J\cup\{e\}, C_b)\Bigg\}. \tag{7}$$

An empty inner minimum is omitted; priors are not renormalized at branches. Let B_N be the optimal noisy Bayes terminal loss with at most N samples.

Theorem 5.4 (Exact limiting value)**.** $\lim_{N\to\infty} B_N = V(\varnothing, H)$. *With $G = \max_{d,h} g(d,h) - \min_{d,h} g(d,h)$, an amplified ideal policy using at most mk samples has excess Bayes loss at most $Gm\exp[-2k(1/2-\varepsilon)^2]$.*

Proof. An ideal oracle can simulate any noisy policy using cached bits and private noise. Randomization cannot improve a finite Bayes decision-tree optimum; conditioning on the random tape and pruning cached queries gives the recursion (7). Thus V lower-bounds every noisy risk.

Take an optimal deterministic ideal tree and amplify as above. Couple it to the exact ideal run. Until a majority error their decisions coincide. The loss difference is at most G on the failure event. The same union bound gives the upper bound. The sequence B_N is nonincreasing, completing the limit argument. □

In Example 5.1, uniform-prior ideal Bayes error is $1/4$. The limiting randomized minimax error is instead $1/3$. Every deterministic ideal policy errs on at least one representative of $\{a, x, y\}$, so the sum of those three error probabilities is at least one under any mixture. For attainment, mix equally three ideal policies: query q and choose label one on its middle branch; query q and choose label two there; or query s and choose label one on its $\{a, x, z\}$ branch. Unique-label branches are correct. These policies err respectively only at y, only at x, or at a, z. Repetition approximates the mixture. A fully deterministic ideal policy has worst-state zero–one error one; randomized terminal reporting and randomized test selection are substantive in a minimax comparison.

5.3 A general availability rule and a strict quantifier gap

The proofs of Theorems 5.2 and 5.4 use only a finite predicate $A(J, e)$, the update $J \mapsto J \cup \{e\}$, static nominal bits, and always-legal idempotent repetition. They therefore apply to any such availability rule. Homotopy inclusion is one specialization. A resource-consuming or arbitrary history-dependent rule is not automatically covered.

For this general class, the finite-policy-family distinction can be strict. Let $H = \{(b, c) : b, c \in \{0, 1\}\}$ and let the target be c. A test g returns nominal bit b. Test e_0 returns c when $b = 0$ and zero otherwise; e_1 returns c when $b = 1$ and zero otherwise. Selecting either specialized test permanently blocks the other; g and repetitions of the selected test remain available. All responses have fresh common-support BSC noise with known error below one half. An ideal policy reads g, then e_b, so finite majority-amplified policies achieve arbitrarily small error.

No single ongoing policy can have statewise error tending to zero. If it never chooses a specialized test, its observations are independent of c. Otherwise some first choice e_i occurs by a finite time n with positive probability. Before that choice only g is observed. Finite-history absolute continuity gives positive probability to the same event under the wrong value $b = 1 - i$. On that event the two states with $c = 0, 1$ have identical available future observation laws: the correct specialized test is blocked forever. The sum of their error probabilities is bounded below by this positive event probability for every $t \geq n$. Thus they cannot both have vanishing error. The lower constant depends on the policy; there is no contradiction with arbitrarily accurate finite families.

This last availability example has not been shown here to arise from order-complex inclusions. It substantiates the general quantifier warning, not a separate topological realization claim.

6 Complete observation laws and attainable decisions

The nominal bit may omit distinctions carried by the readout. For known channels, a more general oracle returns each test's complete observation law.

6.1 The actual-law oracle

Let H be finite and nonempty, $\mathcal{E}$ finite, and $\mathcal{D}$ finite and nonempty. The hidden state is fixed. Test e has finite alphabet $\mathcal{Y}_e$ and a known probability vector $K_e(\cdot \mid h)$; zero entries are allowed. Conditional on h, the selected test, and the controller's entire available past, including its private randomization, the next response has this same vector. Private coins are independent of the hidden state and response innovations. A fixed $A(J, e)$ governs access, and execution updates J to $J \cup \{e\}$. Every executed test remains legal to repeat without further structural change. Stopping and reporting are always allowed.

Define $h \sim_e g$ when their complete test-e vectors agree, and write $r_e(h) = [h]_{\sim_e}$. The law-oracle answer is $r_e(h)$. For candidate set C, its children are $C_{e,r} = \{h \in C : r_e(h) = r\}$. Replace the binary children in (6) and (7) by these nonempty children, retaining the same full-carrier legality predicate. Empty disjunctions mean failure; a minimum over no legal new tests is $+\infty$. Empty children are homogeneous by convention and have unnormalized loss zero. Their noisy implementations stop with an arbitrary decision.

Theorem 6.1 (Actual-law characterization)**.** *A label is arbitrarily reliably acquirable by bounded finite policies exactly when the law-oracle recursion succeeds. Allowing statewise almost surely finite stopping does not change feasibility. For a full-support prior and bounded terminal loss, the limiting finite-horizon Bayes value equals the law-oracle value.*

Proof. A law answer determines its response vector, so it permits state-independent simulation of each response from a fresh uniform number. Cache answers on first use. The common-tape argument of Theorem 5.2 preserves J and legality; pruning cached repetitions gives at most $m = |\mathcal{E}|$ new queries on any path. This proves necessity and the terminal-loss lower bound, including statewise almost sure stopping.

For sufficiency, decode law classes by maximum likelihood from k fresh readings. For distinct class vectors $p_{e,r}$ and $p_{e,s}$,

$$\mathbb{P}_r(\widehat{r} \neq r) \leq \sum_{s \neq r} \left(\sum_y \sqrt{p_{e,r}(y) p_{e,s}(y)} \right)^k .$$

On a sequence where a competing likelihood exceeds the true one, the true likelihood is bounded by their geometric mean. Summing proves the bound without division by zero. Distinct vectors have affinity below one by the equality case of Cauchy–Schwarz. Thus maximum local error η_k tends to zero. Amplify each node of a legal ideal tree; total report error is at most $m\eta_k$ in at most mk samples. Wrong class decisions preserve structural legality, since it depends on J, not the inferred candidate set. Empty children are totalized by immediate reports. Loss range width G gives excess terminal Bayes loss at most $Gm\eta_k$. □

If any test has distinct vectors, let $\rho_* < 1$ be the largest affinity over those pairs. Then $\eta_k \leq (|H| - 1)\rho_*^k$ for $k \geq 1$. If no such pair exists, take $\eta_k = 0$; if $\rho_* = 0$, one reading distinguishes each class exactly. The earlier support-persistence and finite-certification impossibility require positive likelihoods or mutual absolute continuity. This theorem does not. Its permission rule is stipulated independently of posterior support: changing that rule, or the carrier of an RF poset, is a new model.

6.2 A finite polytope of terminal channels

A deterministic legal law-oracle tree realizes a function $f : H \to \mathcal{D}$. There are finitely many such functions, since no branch repeats a test. Let their set be $\mathcal{F}$, including immediate reports, and

define

$$Q_f(d \mid h) = \mathbf{1}\{f(h) = d\}, \qquad \mathcal{P} = \operatorname{conv}\{Q_f : f \in \mathcal{F}\}.$$

For a noisy policy Π, put $Q_\Pi(d \mid h) = \mathbb{P}_h(\Pi \text{ reports } d)$. Use the metric $\|Q - Q'\| = \max_h \|Q(\cdot \mid h) - Q'(\cdot \mid h)\|_{\mathrm{TV}}$.

Theorem 6.2 (Terminal-channel closure)**.** *Every legal statewise almost surely terminating randomized noisy policy has $Q_\Pi \in \mathcal{P}$. Moreover,*

$$\overline{\bigcup_{N<\infty} \{Q_\Pi : \Pi \textit{ legal, randomized, and stops by } N\}} = \mathcal{P}.$$

Proof. Simulate the policy with cached law answers and a state-independent tape containing private coins and response uniforms. Almost surely, all finitely many hidden-state runs terminate. Partition this set by their terminal report vector $f \in \mathcal{D}^H$:

$$E_f = \bigcap_{h \in H} \{\text{the } h\text{-run terminates and reports } f(h)\}.$$

These events are measurable, since each finite report event is a countable union over stopping times. Every positive-mass E_f has a representative tape whose pruned legal tree realizes f. Consequently $Q_\Pi = \sum_f \mathbb{P}(E_f) Q_f \in \mathcal{P}$. This avoids any measurable selection of entire trees.

Conversely, amplification approximates each tree channel within $m\eta_k$. A state-independent private choice of tree implements any finite convex mixture with the same bound. The finite polytope is closed, proving equality. □

Bayes loss is linear in Q; randomized minimax loss $\max_h \sum_d Q(d \mid h) g(d, h)$ is continuous. Their finite-budget limits are minima over $\mathcal{P}$; minimax is a finite linear program over tree mixtures. Private randomization is essential for the full convex hull. Closure does not imply finite attainment of a zero-error vertex, consistency of one ongoing policy, or convergence of arbitrary constrained optima on unattainable boundaries.

Adapted simulation and controlled experiment comparison are established [9, 28]. The claim is a finite reduction for the stated idempotent-access class. Recent sequential stopping/budget comparisons [36, 18] use different decision classes; they further caution against a broad priority claim.

6.3 Four states are necessary for the consistency gap

Proposition 6.3. *For $|H| \le 3$, a successful law-oracle tree implies a legal deterministic sequence of distinct tests whose joint law signatures separate all differently labelled states. One ongoing legal policy then has reports eventually permanently correct almost surely. Four states are therefore minimal for the latch's finite-family versus ongoing-consistency gap.*

Proof. Stop a successful tree at homogeneous nodes. With at most three candidates, a nonconstant answer partition has at most one child containing two or more states; all other children are singleton leaves. Follow the sole potentially nonterminal child, with constant-answer queries having one nonempty child. This gives a legal root-to-leaf test sequence. Every differently labelled pair either separates on the sequence or remains together down its sole nonterminal branch; it cannot persist to a homogeneous leaf. Legality of the entire sequence is independent of the true state. Execute it and cycle its now-legal tests. Affinity bounds give summable errors at completed cycles; Borel–Cantelli gives eventual permanent correctness. Hold the latest completed-cycle estimate between cycle ends. Constant targets need no tests. □

The four-state latch supplies sufficiency for the gap under common BSC noise. This is a cardinality result for general availability, not a topological realization of that latch.

6.4 Persistent nuisance, small separation, and prices

Known finite persistent noise can be included in the static state. If a vector $Z = (Z_e)$ is drawn once jointly with h and repeated queries return Z_e, use state (h, Z), target $\ell(h)$, and deterministic kernels. Retain the declared $A(J, e)$; do not silently rebuild the RF order. The law oracle is exactly the cached-response experiment. For fair θ and independent $N \sim \text{Bernoulli}(0.2)$, repeated $\theta \oplus N$ has Bayes error 0.2, attained after one read. The augmented prior on (θ, N) is $(0.4, 0.1, 0.4, 0.1)$; each observed class has mass 0.5 and maximum target mass 0.4. This represents finite known static nuisance, not arbitrary temporal dependence. Persistent-response active diagnosis is established [1].

Known state-dependent noise is different: rates 0.1 and 0.2 split actual-law classes even when the nominal bits agree. The law-oracle value is zero, matching the observed limiting risk. Tiny known gain or noise differences can remove a qualitative obstruction while demanding impractical budgets. For fixed $0 < p < 1$,

$$\rho(p, p + d) = 1 - \frac{d^2}{8p(1 - p)} + O(d^3).$$

Sample requirements diverge as laws approach equality. Estimated kernels need quantitative uncertainty bounds.

For nonnegative reading prices λc_e, $c_e \le c_{\max}$, the optimal total objective satisfies

$$0 \le C(\lambda) - V \le Gm\eta_k + \lambda m c_{\max} k \quad (k \ge 1).$$

The lower bound is oracle domination; amplifying an optimal ideal tree proves the upper bound. Hence $C(\lambda) \to V$ as $\lambda \downarrow 0$. Exponential decoding gives an $O(\lambda \log(1/\lambda))$ upper rate, not an optimality claim. Fixed positive prices do not vanish as a cap grows.

First-use resource registers can be included if repeated tests remain legal stutters and their kernels stay unchanged. Per-read depletion does not qualify. A fair binary target, one BSC test of error 0.2, and a battery allowing one read has error 0.2 at every positive horizon, although the noiseless query identifies the target.

Independent exact checks cover 8,192 oracle/report-map comparisons, 5,120 noisy Bellman values, 640 general-channel values, and 8,704 small-state path checks. These are bounded corroborations, not proofs of priority or unrestricted correctness.

7 Self, world, and the decision that needs the distinction

Let $h = (\theta, w) \in \{0, 1\}^2$, where θ is actuator reversal and w environmental reversal. Command a produces physical output $y = a \oplus \theta \oplus w$. A routine measurement is

$$o = y \oplus E, \qquad r = o \oplus a = \theta \oplus w \oplus E, \qquad E \sim \text{Bernoulli}(\eta).$$

Here E is measurement error, not physical execution error. The numerical experiments score classification of the noiseless effective mechanism.

Proposition 7.1 (Routine symmetry)**.** *Every adaptive command policy has identical routine transcript laws at (θ, w) and $(1 - \theta, 1 - w)$. For known $\eta \ne 1/2$, the two parity classes $\phi = \theta \oplus w$ are distinct in distribution.*

Proof. The response kernel depends only on ϕ. Equal histories give identical action distributions and then identical response distributions. Induction includes randomized commands. Different parity values give different Bernoulli laws when $\eta \neq 1/2$. □

Once ϕ is learned, $a = g \oplus \phi$ attains target g in the original environment. After transfer to a known environment $w_{\text{new}} = 0$, the correct command is $a = g \oplus \theta$. With independent fair θ, w, routine histories leave θ fair, so routine-only transfer error is $1/2$ even with unlimited data.

Assign the RF witness (4) to $a = 00, x = 01, y = 10, z = 11$. Its label distinguishes single-reversal cases while merging both parity-zero states. It is conditional discrepancy attribution, not full pair identification or actuator identification everywhere. This interpretation does not validate RF as physical safety.

7.1 Task sufficiency and regret

For a finite decision set and loss g, put $g_*(h) = \min_d g(d, h)$. A class C admits a decision with zero regret at every state exactly when

$$\bigcap_{h \in C} \arg\min_d g(d, h) \neq \varnothing. \tag{8}$$

A common minimizer proves sufficiency. Conversely, regrets are nonnegative; zero expected regret requires every positive-probability decision to minimize loss at every state. Randomization does not relax the criterion. Pairwise intersections are insufficient with ties: $\{u, v\}, \{v, z\}, \{u, z\}$ intersect pairwise but not jointly.

For a positive distribution p on C, residual Bayes regret is

$$\mathcal{R}(p) = \min_d \sum_h p(h) g(d, h) - \sum_h p(h) g_*(h).$$

It is zero exactly under (8). A nondestructive observation Z has information value $\mathcal{R}(p) - \mathbb{E}_Z \mathcal{R}(p_Z) \geq 0$ before its cost. This is finite decision theory, not a consciousness score. If the experiment changes the physical state, this formula alone does not isolate its value.

Repair need not require attribution. Suppose leaving an actuator incurs 10θ and replacement costs 3. Repair is optimal when $p(\theta = 1) > 0.3$. At prior $1/2$, always repairing costs 3. A globally admissible nondestructive diagnostic $Z = \theta \oplus N$, with fresh $N \sim \text{Bernoulli}(0.1)$ and cost 0.2, yields posterior 0.9 or 0.1 with equal probability. Repair on the first outcome and leave on the second: terminal loss $(3+1)/2 = 2$, or 2.2 including testing. Clairvoyant expected loss is 1.5. These numbers describe one optional test, not unrestricted sequential optimality. Free repair can remove all regret without diagnosis; cost-free compensating commands can also make physical repair irrelevant.

7.2 Interventions that preserve or break ambiguity

Known toggles of θ or w change parity by a known offset but preserve the symmetry. More generally, if a symmetry preserves observation kernels and commutes with every allowed transition, adaptive experiments preserve its transcript equivalence. The name *intervention* does not imply informative attribution.

A trusted reset $w := 0$ preserves θ and provides noisy direct measurements of it through the routine channel. Resetting $\theta := 0$ instead preserves w. Combined with pre-reset parity information, that can reveal the original actuator state; without such a record it may erase the initial coordinate. Knowing the current reset coordinate is different from learning its origin.

For a one-use known environmental shift, preserve θ and let $q_w = \mathbb{P}(w_{\text{new}} = 1 \mid w_{\text{old}} = w)$. Within original parity class ϕ, the two candidates are $(0, \phi)$ and $(1, 1 - \phi)$. Under $a = 0$ their observed probabilities of one differ by

$$(1 - 2\eta)(q_0 + q_1 - 1). \tag{9}$$

Thus the pair is separated in distribution exactly when $q_0 + q_1 \neq 1$ and $\eta \neq 1/2$. Identity, toggling, and independent flipping with known probability have sum one and fail. Reset to zero has sum zero. A fresh independent environment with known bias ρ has sum 2ρ, so helps unless $\rho = 1/2$. Repetition requires a specified replenishment protocol; one separated distribution is not unlimited independent data.

For a fixed known finite schedule of affine deterministic interventions over $\mathbb{F}_2$, write $h_t = A_t h_0 \oplus b_t$. A residual has coefficient row $[1\ 1]A_t$ on the initial state. Stack these rows into M. Noiseless observations identify a label precisely when it is constant on fibers of M; both initial coordinates are identified precisely when $\operatorname{rank} M = 2$. Toggles change offsets, not rank. A baseline row $[1\ 1]$ plus reference row $[1\ 0]$ has rank two. With fresh informative BSC noise the same criterion concerns separation of joint laws, not exact finite recovery. Adaptive schedules require the corresponding branchwise analysis.

8 Cost and risk as parts of the contract

8.1 Diagnostic value and information lower bounds

For nondestructive tests with costs $c_e > 0$ and finite terminal decisions, the Bayes recursion is

$$U_0(b) = \min_d \sum_h b(h) g(d, h), \tag{10}$$

$$U_n(b) = \min \left\{ U_0(b),\ \min_{e \text{ legal}} \left[c_e + \sum_y p(y \mid b, e) U_{n-1}(b^{e,y}) \right] \right\}. \tag{11}$$

Cumulative structure requires J as an extra argument; physical actions require transitions and possibly a joint origin–current belief. This is standard dynamic programming.

In the support-only model, uniform label error at most $\delta < 1/2$ implies for each differently labelled pair, under finite expected stopping time,

$$\sum_{e \in \mathcal{E}_0} \mathbb{E}_h N_e \, D_{\mathrm{KL}}(K_e^h \| K_e^g) \geq \mathrm{kl}(1 - \delta, \delta). \tag{12}$$

Here kl is binary relative entropy. The stopped likelihood ratio accumulates observation information; policy factors cancel. Data processing to the event of reporting $\ell(h)$ gives the bound. For unbounded stopping, use the stopped change-of-measure theorem [17, Lemma 1], not optional stopping without hypotheses. Finite alphabets and positive kernels make increments bounded; positive costs and finite expected cost give the needed integrability. Minimizing $\sum_e c_e x_e$ over $x_e \geq 0$ subject to these pairwise information constraints gives an expected-cost lower bound under h. No attaining algorithm is claimed.

8.2 A probabilistic gate changes the rule

If $z_e(h)$ is hazard probability, a one-use Bayesian chance rule permits the experiment when $\sum_h b(h) z_e(h) \leq \alpha$. It cannot override a failed exact predicate without changing the contract. Its guarantee depends on the specified prior, likelihood, and hazard model.

Take states $0, 1, 2$, prior $(1/2, 1/4, 1/4)$, and a harmless screen whose zero probabilities are $(0.9, 0.1, 0.1)$. A diagnostic has hazard zero at state zero and one at states one and two. After n zero screens,

$$p_{\text{bad}} = \frac{1}{1 + 9^n}.$$

Two zeros give $1/82 \approx 0.012195$, below an explicitly allowed 0.02 threshold, while all three states remain supported. Under an actually bad state the diagnostic remains hazardous with probability one.

A policy screening exactly twice and unlocking only after two zeros has gate probability 0.41 and joint bad-and-gate probability 0.005. Gate probabilities conditional on safe and bad states are respectively 0.81 and 0.01. The conditional risk is $0.005/0.41 = 1/82$. These are different quantities.

If, for every allowed screening test and outcome, every good-state to bad-state likelihood ratio is at most $\Lambda > 1$, prior bad mass p_0 satisfies after any n screen observations

$$p_{\text{bad}}(H_n) \geq \frac{p_0}{p_0 + (1 - p_0)\Lambda^n}.$$

Bound each state-pair transcript ratio and then mix over states. Reaching $\alpha < p_0$ therefore requires at least $\lceil \log[p_0(1-\alpha)/(\alpha(1-p_0))]/\log \Lambda \rceil$ screens. This is necessary, not a guarantee that favorable data occur.

For a statewise gate guarantee, suppose a finite screening policy unlocks with probability at least $1 - \beta$ under safe h and at most γ under unsafe g, with $1 - \beta > \gamma$. The unlock event gives

$$\sum_e \mathbb{E}_h N_e \, D_{\text{KL}}(K_e^h \| K_e^g) \geq \text{kl}(1 - \beta, \gamma).$$

Identical screen laws force $\beta + \gamma \geq 1$. Zero false unlocking cannot coexist with positive finite-time unlocking power under mutually absolutely continuous finite transcript laws.

A per-use threshold does not bound episode risk by the same threshold. If conditional budgets α_t sum to at most Δ on every path, conditional expectation and a union bound give any-hazard probability at most Δ under the Bayesian prior-predictive law. A time-uniform confidence set with statewise simultaneous coverage $1 - \Delta$ also bounds unsafe selection by Δ if actions are safe throughout that set [14]. It is not exact posterior support; applying it to an arbitrary topological rule creates no physical guarantee.

9 Exact experiments and independent checks

All experiments are synthetic consequences of declared finite models. Exact fractions are used for posterior updates, decision values, and enumeration. Simulations are separate implementation checks, not empirical evidence about robots or consciousness.

9.1 Task-specific optimal policies

Use independent fair θ, w and two globally admissible nondestructive ports. Routine sensing returns $\phi = \theta \oplus w$ through error 0.1; body sensing returns θ through error 0.2. The body port is a trusted measurement assumption. Each sample costs c, at most six are allowed, and early stopping is permitted. Terminal zero–one losses concern ϕ (present control), θ (reference-environment transfer), or (θ, w) (full causal diagnosis). These objectives exclude execution noise and additional intervention costs.

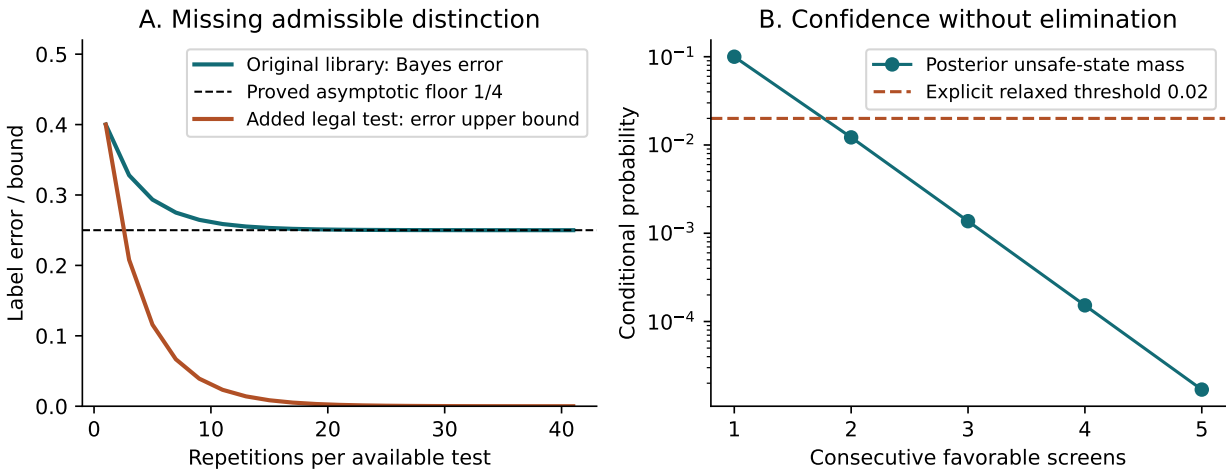


Figure 2: Left: exact uniform-prior Bayes error for the original RF library and an exact binomial union bound for the augmented library, at noise 0.2. This compares an error with an upper bound, not two measured policy errors. Right: risk along favorable screen histories; no original state reaches zero likelihood.

The Bellman solver evaluates both responses to every test. Ties favor stopping, then port order; terminal label ties favor the lowest label. At $c = 0.03$:

Task	First test	Label error	Mean tests	Total cost
Present control	routine	0.015040	2.3924	0.086812
Transfer	body	0.073280	2.7424	0.155552
Full causal state	routine	0.119872	4.5448	0.256216

Total cost is error plus 0.03 times expected tests. These are optima only within the stated horizon, ports, known model, and losses. With routine sensing alone, transfer optimally stops immediately: error 1/2 and zero sampling cost. Routine symmetry proves the stronger unlimited-data attribution obstruction.

Under the fair prior, θ and ϕ are independent. Their posterior odds are 4^{d_θ} and 9^{d_ϕ}, where each d is observed ones minus zeros for its port. This supplies a distinct sufficient-statistic simulation and shows that a matched compiled or external controller can reproduce the same policy.

9.2 Bounded exhaustive verification

For the diamond, all 16 Boolean tables, all sublibraries of size zero through three, and all 15 target partitions give 697 libraries and 10,455 library–target cases. A noiseless tree recursion agrees with the acquisition report's specialized classification in every case: 3,916 are noiselessly reportable. Under the support-reset noisy lift, 3,744 satisfy initial-test separation; 172 noiselessly feasible cases are lost. Under cumulative admissibility, the oracle recursion reports 3,632; 112 initially separated cases are blocked. These counts describe one bounded family, not statistical estimates or proofs for larger posets.

Additional checks cover 240 support/query legality cases by directly constructed comparability graphs; 1,023 binary histories through length nine for positive support in the RF witness; and 4,096 paired paths across all 128 deterministic adaptive command policies of horizon three for

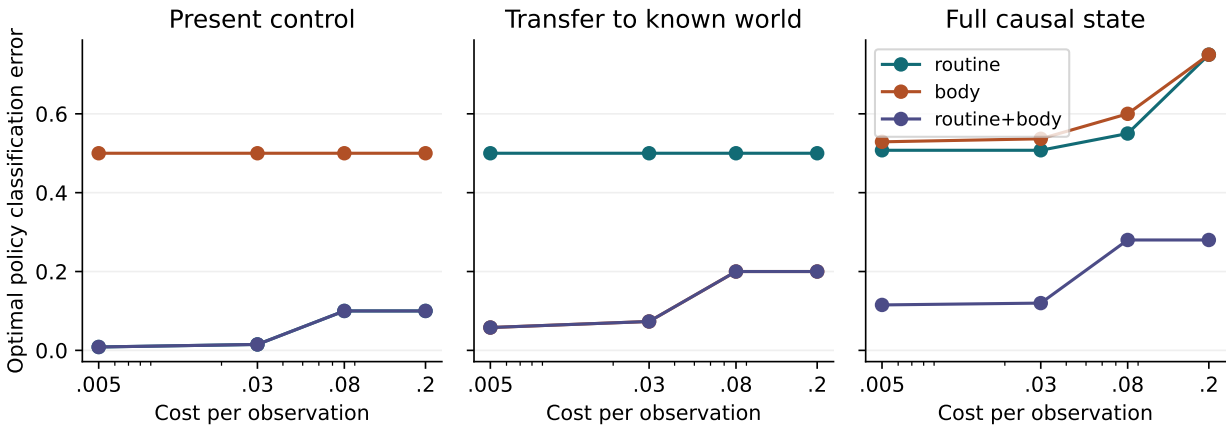


Figure 3: Label error of the optimal cost-minimizing policy with six tests available. Only the four displayed costs were computed; lines guide the eye. High costs can favor immediate stopping. Useful diagnostic access depends on the terminal task.

XOR symmetry. The graph test is justified here because components of these diamond suborder complexes are contractible. It is not a general homotopy-equivalence test.

An independent evaluator enumerates all 723 deterministic sensing/stopping trees of depth at most three for two ports. Across three tasks and four costs, its unnormalized path calculations match all 12 Bellman values. A Monte Carlo implementation instead reconstructs beliefs from integer sufficient statistics and runs 100,000 episodes per task with seed 20261103. Estimated total costs are 0.086650, 0.1555168, and 0.2557106, with approximate 95% normal half-widths 0.0007791, 0.0016509, and 0.0020388. Exact values lie within the intervals. These intervals describe simulation variability, not uncertainty in exact model values.

The previous report and supporting calculations were recovered and byte-hashed. Its R4 checker is preserved separately and rerun for continuity. That is not an audit of the historical R4 implementation or validation of physical assumptions.

10 From formal permission to physical experiments

The general availability model describes diagnostic access under irreversible constraints. A finite register can implement it, but that observation alone does not make the constraint useful for a physical device. We separate three questions: whether a test is formally permitted, whether it is executable, and whether its execution preserves the physical conditions needed by the task.

10.1 Homotopy does not preserve directed recovery

Section 11 proves a minimal three-state example: a homotopy-preserving query deletes the only directed recovery from one start to its target. The following limits explain why an implementation contract is necessary.

If every directed start–target reachability pair must be preserved, a suborder $R \subseteq P$ of a transitively closed order must equal P: every pair in P must remain in the reachability relation, which is R. For query thinning this forces q to be monotone on P. Universal directed recovery therefore forbids any deletion. A specified target or task may admit a weaker, useful condition.

A viability example sharpens the difference. Add an absorbing damage state d. Before refinement, x can repair to z or wait and enter d; afterward only the latter action remains. From a, moves to x and z are available; z can hold. The controlled invariant subset of the safe set $\{a, x, z\}$ changes from $\{a, x, z\}$ to $\{a, z\}$. The finite-state calculation accompanies the paper. The added dynamics are declared separately from the order.

There is a limited positive interpretation. Suppose each retained undirected edge is an executable, reliable, safe motion in both directions, all vertices are safe, and time and resources impose no constraints. A homotopy-equivalent inclusion preserves connected components, hence existence of a route between a start and goal. Connected-component preservation already suffices for that task. It does not preserve cost: if the removed edge xz costs one and xa, az each cost 100, the shortest route costs 200 after refinement.

10.2 A sufficient implementation contract

Identical abstract labels, kernels, and orders can describe two devices when one test also causes unmodeled damage in one device. No statement using only that tuple can distinguish them. A sound physical interpretation must include the relevant transitions.

Let X be concrete apparatus states, $\zeta : X \to H$ the diagnostic-state map, $G \subseteq X$ a safe set, and $\mathcal{R} \subseteq X \times 2^{\mathcal{E}}$ a relation to abstract test histories. Let $F_e(x)$ contain every possible concrete successor when test e executes. Require related safe initial states and, whenever $(x, J) \in \mathcal{R}$ and $A(J, e) = 1$, require executability, nonempty successors, and

$$\varnothing \neq F_e(x) \subseteq G \cap \{x' : (x', J \cup \{e\}) \in \mathcal{R},\ \zeta(x') = \zeta(x)\}. \tag{13}$$

The hidden diagnostic target must be preserved, and the full conditional observation law must match the model.

Proposition 10.1 (Conditional implementation bridge)**.** *These assumptions transfer finite-step safety to every legal abstract policy. With the declared observation law and target preservation, they also transfer the corresponding finite diagnostic guarantee.*

Proof. Initially the concrete state is safe and related. Given a related safe state, the selected action executes, and every possible successor is safe and related by (13). Induction proves the safety statement. The observation assumption gives the same adaptive transcript law, by induction on histories, and hence the same terminal diagnostic error. □

Recovery requires an additional progress argument. An unchanged J during repetition does not prevent heating, wear, or depletion. Equation (13) must be justified from independently specified dynamics, not by defining safety to mean formal permission. Feedback refinement relations provide an established framework for action and successor conditions, behavioral inclusion, and specification transfer [26]. That framework is distinct from B-Theory's RF condition.

10.3 A latching diagnostic fixture

The general availability example in Section 5 has a direct finite implementation. Keep the hidden target $h = (b, c)$ fixed and use a known apparatus mode $m \in \{\text{free}, 0, 1\}$. Probe g reads b and preserves the mode. Probe e_0 reads $(1 - b)c$, is allowed in modes free or 0, and sets mode 0. Probe e_1 reads bc, is allowed in free or 1, and sets mode 1. The selected specialized probe and g remain repeatable; switching specialized modes within an episode is blocked. Independent BSC readouts leave the hidden target unchanged.

A digital guard implementing this table realizes the availability predicate. More generally, a register storing J can encode any finite $A(J, e)$ with always-legal idempotent repeats. This is a synthetic realization, not a proof that such constraints naturally describe a motor or arise from RF homotopy. A disposable reagent or a probe that consumes its resource on every read does not satisfy this repetition premise.

For flip rate 0.1, repeat g an odd number k of times, use its majority to choose the specialized mode, and repeat that probe k times. Write

$$a_k = \sum_{j=(k+1)/2}^{k} \binom{k}{j} (0.1)^j (0.9)^{k-j}.$$

For target $c = 0$, the error is a_k. For $c = 1$, the correct and incorrect mode branches give error $(1 - a_k)a_k + a_k(1 - a_k) = 2a_k(1 - a_k)$. Thus maximum error is $2a_k(1 - a_k)$ and uniform-prior Bayes error is $3a_k/2 - a_k^2$. At $k = 5$ these are 0.0169734528 and 0.0127667264. They are policy values, not finite-budget optima.

After a fixed k-probe mode decision, infinitely many subsequent readings can identify whether the wrong mode was selected, but cannot recover c on that branch. Its limiting Bayes floor is $a_k/2$, or 0.00428 at $k = 5$. This makes the finite-family versus single-policy distinction operational. The implementation was checked on 72 reachable history/request/hidden-state combinations, and the error formula was independently enumerated for $k = 1, 3, 5, 9, 15$. Neither is a hardware experiment.

11 Preserving a specified recovery task

Ordinary homotopy is not a directed control specification. A more direct alternative is to declare which recovery outcomes must remain possible. This subsection first states a finite graph property. Its interpretation as apparatus behavior requires the implementation assumptions below.

Let P be a finite poset on H, $T \subseteq H$ a nonempty set of acceptable recovery configurations, and $D \subseteq H$ a fixed set of protected starts. For a suborder $R \subseteq P$, put

$$V_T(R) = \{x \in H : \exists t \in T,\ x \le_R t\}.$$

This is a reachability set, not an infinite-horizon controlled invariant set. The order is reflexive, so starts already in T satisfy the objective. Assume $D \subseteq V_T(P)$, and define

$$\mathcal{F}_T = \{J \subseteq \mathcal{E} : D \subseteq V_T(P_J)\}, \qquad A_T(J, e) = \mathbf{1}\{J \cup \{e\} \in \mathcal{F}_T\} \quad (J \in \mathcal{F}_T).$$

Thus a diagnostic test may remove recovery options while retaining at least one acceptable outcome from every protected start. Unlike preservation of every ordered pair, this requirement permits nonmonotone tests. The protected set D is fixed independently of noisy answers; substituting an inferred candidate set would need a different safety argument.

Proposition 11.1 (Task feasibility and the oracle)**.** *For $x \le_P t$, define $B(x, t) = \{e \in \mathcal{E} : q_e(x) = 1,\ q_e(t) = 0\}$. Then*

$$J \in \mathcal{F}_T \quad \iff \quad \forall x \in D\ \exists t \in T : \quad x \le_P t \text{ and } J \cap B(x, t) = \varnothing.$$

In particular, $\mathcal{F}_T$ is closed under taking subsets, and an executed test is always legal to repeat. With static nominal bits and fresh BSC readings, Theorems 5.2 and 5.4 hold with A_T in place of homotopy admissibility.

Proof. The displayed condition is exactly the definition of $x \leq_J t$, with the quantifiers written out. Removing tests cannot destroy it. Repeating a test leaves J unchanged. The availability predicate depends on J, not the answer, so both the cached-oracle simulation and majority amplification preserve legality, including after a mistaken majority. These are precisely the premises used in the two earlier proofs. □

Here later execution cannot unlock a previously unavailable test: if $J \subseteq K \in \mathcal{F}_T$ and $K \cup \{e\} \in \mathcal{F}_T$, then $J \cup \{e\} \in \mathcal{F}_T$. A query constant on a candidate set can therefore be pruned from a successful ideal tree for this particular rule; deleting it enlarges future availability and loses no branch information. This pruning need not be valid for arbitrary availability predicates.

Proposition 11.2 (Smallest separation from RF)**.** *Even with $D = H$, $\varnothing \neq T \subsetneq H$, and a nonconstant Boolean test, RF and preservation of V_T imply neither one another. Three states are necessary and sufficient for each failed implication.*

Proof. For RF without task preservation, use $a < x < z$, $T = \{z\}$, and $q = 010$. The triangle axz collapses onto the retained tree $ax \cup az$ by removing the triangle and its free edge xz, but x no longer reaches z.

For task preservation without RF, use $a < x, a < y$, $T = \{x, y\}$, and $q = 101$ in order (a, x, y). The edge ax is removed and ay remains. Every state still reaches T, but the number of connected components increases from one to two. Both examples use exactly the original maximal elements as acceptable terminal configurations.

There is no nonempty proper T on one state. On two states, universal initial reachability of a proper T forces a two-element chain with its maximum as the sole target. Retaining its edge preserves both properties; deleting it destroys both. This proves minimality without relying on the finite computation. □

There is a useful boundary. If $T = \{t\}$ and $V_T(P) = H$, then t is a greatest element. Task preservation keeps it greatest in P_J, so both order complexes are cones and their inclusion is a homotopy equivalence. For one Boolean test, preservation is equivalent to $q(t) = 1$ or $q \equiv 0$. Consequently all individually legal tests are jointly legal in this special case. Pure reachability to a single universally reachable target does not produce cumulative exclusions.

It still permits informative strict thinning. On $a < x < z$ with target z, the test $q = 101$ removes $a < x$ but preserves recovery. Together with $r = 011$ it produces the three distinct signatures $10, 01, 11$, while the joint order retains $a < z$ and $x < z$. Thus a task constraint can permit full diagnosis without preserving every directed pair.

A three-state cumulative obstruction. On $a < x, a < y$, take $T = \{x, y\}$ and tests $p = 110$, $s = 101$. Each test alone keeps recovery from all three states, through x or y respectively. Together they remove both outgoing edges from a. Their joint signatures $11, 10, 01$ separate the three identity labels, but either first test blocks the other, and its nonsingleton answer cell is unresolved. Hence the ideal oracle predicate is false. Under the fresh-noise assumptions, the limiting uniform-prior Bayes error and randomized minimax error are both $1/3$. Every deterministic ideal policy makes at least one error. Policies erring only at a, only at x, or only at y are attainable, and their equal mixture attains the minimax bound. Repeated readings approximate these ideal policies. This is a failure of initial separation to imply cumulative feasibility; it is not a separation between arbitrarily accurate finite families and a single ongoing consistent policy.

Cost and robust control refinements. Reachability does not preserve the cost of recovery. Give the chain's separately executable edges $a \to x$, $x \to z$, $a \to z$ costs $1, 1, 10$. The test 101 preserves RF and reachability to z, but increases the best cost from a from 2 to 10. The direct edge here is an independent primitive action; it is not a shorthand for the now unavailable cheap path. Given independently specified edge costs, one can instead require

$$d_{J\cup\{e\}}(x, T) \le B_x \qquad (x \in D),$$

where d_J is the optimal recovery cost and the budgets B_x are fixed before diagnosis. Exact cost preservation uses $B_x = d_\varnothing(x, T)$. Per-step allowances should not be confused with a fixed total allowance. These budgets concern subsequent recovery; they do not make repeated diagnostic samples free of physical cost.

For a finite fully observed recovery system with safe set G, target $T \subseteq G$, enabled recovery actions $U_J(x)$, and successor sets $F_J(x, u)$, robust finite reachability is computed by

$$Z_0 = T, \qquad Z_{k+1} = Z_k \cup \{x \in G : \exists u \in U_J(x), \quad \varnothing \ne F_J(x, u) \subseteq Z_k\}.$$

The stabilized set W_J is exactly the set from which a controller can force safe arrival in T in finitely many steps: the first entry stage supplies a strictly decreasing rank. Preserving $D \subseteq W_J$ gives another finite availability predicate when the supplied dynamics depend only on J and repeated diagnosis leaves that dynamics unchanged. Infinite-horizon safety uses instead the greatest fixed point of $Z \mapsto G \cap \{x : \exists u,\ \varnothing \ne F_J(x, u) \subseteq Z\}$. These are standard finite reachability and invariant-set constructions, not a new control-synthesis method. If the physical state is not observed, the controller must use a common action on each information set; pointwise existence of a recovery action is insufficient.

What a physical use would require. The symbols in the preceding poset are counterfactual recovery configurations. During diagnosis the unknown origin or mechanism label h can remain fixed while the apparatus mode J removes options for a later recovery phase. The actual recovery configuration must be represented separately when it changes. A transitive comparability is not by itself an executable action, and the thinning rule does not establish that probing a real device disables exactly those actions.

A physical implementation must independently supply the diagnostic transition law, recovery actions and their successors, safe and target sets, costs, and the observation interface. Each legal test must execute without changing the diagnostic label, leave the apparatus in the protected winning region for its new mode, and produce the declared conditional observation law. Repetitions must satisfy the same obligations despite wear, temperature, and resource use. Recovery must have an implementable policy at the controller's actual information level. A relation from the concrete state to the abstract state must preserve enabled inputs and contain all concrete successors in the abstract successors, with compatible safety and target labels. These are the familiar local obligations behind controller refinement; see Definition V.2 and Theorem VI.3 of [26]. Establishing them for a device would supply the missing physical content. Defining safety as "the chosen abstract test is legal" would not do so.

12 Observation fidelity and finite-horizon robustness

12.1 Extra information hidden in the noise

The oracle lower bound requires simulation of the *complete* readout law from the nominal answer and public history $\mathcal{H}_t$:

$$K_e(y \mid h, \mathcal{H}_t) = \widetilde{K}_e(y \mid q_e(h), \mathcal{H}_t). \tag{14}$$

There may be no additional hidden-state information within a nominal-answer class. The original common BSC satisfies this condition. Having every error rate below one half is insufficient.

For two equiprobable states let $q = 0$ in both, but let the probability of readout 1 be 0.1 and 0.2, respectively. Every finite history has positive probability under both states. The nominal oracle always returns zero and has label error 1/2. Independent actual readings identify the state by their frequencies. The equal-prior MAP error after n readings is

$$R_n = \frac{1}{2} \sum_{k=0}^{n} \binom{n}{k} \min\{0.1^k 0.9^{n-k},\, 0.2^k 0.8^{n-k}\} \longrightarrow 0. \tag{15}$$

The convergence follows, for example, by thresholding the empirical mean at 0.15 and applying Hoeffding's inequality under both states. Any distinct fixed rates, however close, produce the same asymptotic obstruction to nominal-oracle domination.

Samples n	1	21	101	501	1001
R_n	0.450000	0.261171	0.0757204	0.000721426	$3.33741\,10^{-6}$

At $n = 2001$, $R_n \approx 9.66883\,10^{-11}$. These are exact integer binomial sums converted to decimals. Independent binary-path enumeration matches the count formula in 39 cases: lengths 0–12 for rate pairs $(0.1, 0.2)$, $(0.10, 0.11)$, and $(0.2, 0.2)$. They do not contradict the fixed-library separation theorem, which uses actual laws rather than nominal bits.

12.2 What a small model discrepancy does guarantee

Proposition 12.1 (Finite-horizon perturbation bound)**.** *Use the same fixed hidden-state space, prior, policy, independent private randomization, terminal loss, and admissible policy class in two models. The policy stops by N. Conditional on every fixed h, matching relevant controller history, and selected action, suppose the complete next-step laws differ by at most η in total variation. For padded transcript T_N,*

$$\max_h \|\mathcal{L}_P(T_N \mid h) - \mathcal{L}_Q(T_N \mid h)\|_{\mathrm{TV}} \le d_N = 1 - (1-\eta)^N \le \min\{1, N\eta\}.$$

The joint laws of (H, T_N) also differ by at most d_N under the common prior. Statewise and Bayes classification errors differ by at most d_N. Expected terminal losses in a common interval of width G differ by at most Gd_N. The value bound persists when taking infima over the same admissible policy class.

Proof. Fix h and couple private coins identically. While histories agree, actions agree. A maximal coupling keeps each next step equal with conditional probability at least $1 - \eta$; multiply through N. The coupling inequality gives the statewise bound. Couple H identically and integrate for the joint bound. Couple randomized reporting too, so classification error is a common event of target and report. Rescale losses for the bounded-loss result, then take infima. □

Stopping can be represented by a common absorbing symbol through time N. An observation-only discrepancy bound does not control hidden damage or changed interlock transitions: include those outcomes in the one-step law or justify them separately. Likewise, the proposition does not compare optima under different admissibility predicates. For $\eta = 0.01$, $d_{10} \approx 0.095618$, $d_{100} \approx 0.633968$, and $d_{1000} \approx 0.999957$. A fixed small discrepancy supplies no useful infinite-budget transfer by itself.

12.3 A weaker premise for amplification, but not domination

Let E_j indicate an erroneous binary readout. Independence can be weakened for the majority upper bound to

$$\mathbb{E}[E_j \mid \mathcal{F}_{j-1}, h] \le \varepsilon < 1/2$$

at every step, where $\mathcal{F}_{j-1}$ includes the entire past. For $\lambda > 0$, the conditional exponential moment is at most $1 - \varepsilon + \varepsilon e^{\lambda}$. Iterating this bound and applying the usual Bernoulli Chernoff bound gives majority error at most $\exp[-2k(1/2 - \varepsilon)^2]$. This permits some dependent readings. It does not imply (14): the state-dependent-rate example satisfies the conditional error bound while defeating oracle necessity. A single reused noise bit also fails the conditional error premise after informative observations.

Observation marginals alone are insufficient. For fair H, compare $Y = H \oplus E$ with $Y = 1 - H \oplus E$, where $E \sim \text{Bernoulli}(0.1)$. Both Y marginals are fair, yet reporting Y has errors 0.1 and 0.9. The relevant target-conditional or joint laws differ. The product bound is sharp: an all-zero record versus independent Bernoulli contamination of rate η has transcript TV exactly $1 - (1 - \eta)^N$.

13 An anchored motor experiment: model and test protocol

13.1 Measurement model and scope

A proposed bench study uses an enclosed rotary motor with configurable command polarity, an external direct or reversing transmission, upstream shaft and downstream endpoint angle sensors, and an independently established reference transmission. The motor/transmission interface defines the self/world boundary. Mechanical reversal may change inertia, gain, friction, or backlash; those changes should be measured rather than assumed away. A camera-coordinate flip is a measurement-interface control, not a physical transmission reversal.

Let θ, w, γ, β denote actuator reversal, transmission reversal, body-sensor orientation error, and endpoint-sensor orientation error. For a signed command $u \in \{-1, +1\}$, an initially resting first-order model is

$$\dot{\omega} = \big((-1)^{\theta} K u - \omega\big)/\tau, \qquad \dot{\alpha} = \omega, \tag{16}$$

$$\Delta\alpha = (-1)^{\theta} K u [T - \tau(1 - e^{-T/\tau})], \qquad \Delta\psi = (-1)^{w} r\, \Delta\alpha, \tag{17}$$

$$B(t) = c_B + (-1)^{\gamma}\alpha(t) + \nu_B(t), \qquad Y(t) = c_Y + (-1)^{\beta}\psi(t) + \nu_Y(t). \tag{18}$$

Here K is gain, τ the time constant, r the transmission ratio, and c_B, c_Y fixed measurement offsets. Define signed differences $D_B = u[B(T) - B(0)]$ and $D_Y = u[Y(T) - Y(0)]$. Offsets cancel; orientation errors do not.

With anchored orientations $\gamma = \beta = 0$, $r = 1$, matched gain across hidden states, and independent Gaussian frame errors of standard deviations σ_B, σ_Y,

$$D_B \mid \theta \sim N((-1)^\theta \mu, 2\sigma_B^2), \tag{19}$$

$$D_Y \mid \theta, w \sim N((-1)^{\theta \oplus w} \mu, 2\sigma_Y^2), \qquad \mu = K[T - \tau(1 - e^{-T/\tau})]. \tag{20}$$

Thresholding at zero gives a BSC with error $\Phi(-\mu/(\sqrt{2}\sigma))$, where Φ is the standard-normal CDF. Thus the binary model follows from an explicit measurement interface. For k independent differences of a port measuring θ with equal prior,

$$\mathbb{P}(\theta = 1 \mid D_1, \ldots, D_k) = \text{logistic}\left(-\frac{\mu}{\sigma^2} \sum_j D_j\right).$$

This decoder uses signed differences; it is not claimed optimal among all processing of every raw frame.

The fault parameters are fixed during an episode, while angle and velocity move. Repeated probes need verified settling/reset or an initial-motion model. Wear, heat, and changing polarity must be modeled when relevant. The originating R4 polarity-shock task [25] is a different contract from this fixed-target episode. The present equations and parameters are not hardware operating limits.

13.2 Orientation anchors and a remaining gauge ambiguity

Without orientation calibration, the nominal measurements identify

$$b = \theta \oplus \gamma, \quad y = \theta \oplus w \oplus \beta, \quad y_{\text{ref}} = \theta \oplus \beta \quad \text{after setting } w := 0.$$

Flipping θ, γ, β together preserves all three observations. A reference transmission alone therefore cannot identify θ. However $y \oplus y_{\text{ref}} = w$ identifies the original transmission bit.

Measurements and independent anchors	θ	Original w
Body and endpoint; neither sensor anchored	No	No
Add known reference transmission; neither anchored	No	Yes
Anchor body sensor; no reference	Yes	No
Reference transmission and anchored endpoint sensor	Yes	Yes
Anchor both sensors; no reference	Yes	Yes

Enumeration of all 16 orientation configurations verifies the table. Identification means distinction in observation law, not finite noisy certainty. Calibrating one unanchored sensor against another does not remove the ambiguity. An independent displacement or equivalent physical reference is required. If the original environment must be restored after reference testing, restoration reliability and cost belong in the task model.

13.3 Executed simulation and negative controls

Six conditions each contain 8,000 episodes, balanced at 2,000 per (θ, w) cell, with five probes per port. Parameters are $\tau = T = 0.08$ seconds, normalized $K = 1$, $\sigma_B = 0.025$ and $\sigma_Y = 0.016$ radians. Thus $\mu \approx 0.0294303553$ radians. Condition $i = 0, \ldots, 5$ uses seed $20261105 + 1009i$. A

self-contained JavaScript implementation supplies raw-frame examples, statewise results, and all seeds.

Under nominal independent noise, analytical transfer errors are approximately 0.059926 for five body signs, 0.031348 for five body differences, and 0.001817 for five reference differences. Simulated errors are respectively 0.061250, 0.032625, and 0.001500. Routine sensing predicts present effective polarity with error 0.001625, but its parity-only actuator transfer decoder has error 0.5.

Condition	Body	Reference	Gain-aware routine
Nominal	0.032625	0.001500	0.492500
Body sensor reversed	0.967125	0.002125	0.503875
Endpoint sensor reversed	0.028875	0.997750	0.498750
Common body noise	0.198875	0.001875	0.499750
Reference wrong in 15% of episodes	0.029500	0.155000	0.505875
Unequal actuator gain	0.017625	0.000625	0.185875

Entries are transfer errors; body and reference columns use continuous differences. The gain-aware routine decoder assumes $K = 1$ for $\theta = 0$ and $K = 1.6$ for $\theta = 1$; it is correctly specified only in the final row. That row demonstrates information about θ through magnitude. The parity-only routine decoder remains at 0.5 in every row. Even quantizing signs may retain leakage, because gain changes sign-error probabilities.

The common-noise condition reuses a command-aligned Gaussian difference, so repetition adds no body information; error remains near the one-probe value 0.202587. Sensor-reversal controls deliberately leave the decoder unaware of the new orientation. They test reliance on calibration. They do not show unavoidable failure for a decoder told the reversal. The imperfect-reference condition selects a reversed reference in 15% of episodes and retains that selection throughout the episode.

For an illustrative repair decision, leaving a reversed actuator costs 10, successful replacement costs 3, each measurement costs 0.05, and reference setup costs 0.4. Replace when posterior reversal probability exceeds 0.3. The nominal body and reference scores are 1.893875 and 2.156750, compared with 3 for always replacing and 5 for never replacing. Body reversal raises its score to 6.516250. These are counterfactual terminal scores: physical repair and replacement success were not simulated.

All 48,000 episodes were reproduced from the saved code after reconnecting the local executor. Eighty-eight paired noiseless trajectory points also match under $(\theta, w) \mapsto (\theta \oplus 1, w \oplus 1)$; adding the same observation noise couples the paired routine laws. The finite trajectory check corroborates the algebra, not an unrestricted simulation theorem. No hardware measurements are included.

13.4 A proposed confirmatory protocol

The following protocol is a research design, pending apparatus selection and engineering review.

1. Commission independently: establish sensor orientations, gain, time constant, backlash, settling, and noise using separate calibration data. Actual apparatus design supplies mechanical and electrical limits.

2. Fix the observation interface. Compare sign-only readings with continuous differences; audit current, timestamps, fixture identity, settling duration, and auxiliary channels. Expose those only in explicitly defined interface conditions.

3. Freeze target labels, likelihoods, policies, exclusions, loss functions, setup/restoration costs, and stopping rules before confirmatory outcomes.

4. Blind the configuration schedule, balance all four hidden cells, randomize trial order and command direction, and lock predictions before revealing configurations. Balance handling and timing to limit setup cues.

5. Score present control, reference-environment transfer, and repair separately. Include the reversal, common-noise, imperfect-reference, and unequal-gain controls above.

6. Implement actual enclosure, travel, current, settling, fixture, and resource checks. Posterior confidence cannot override these interlocks. An adapter latched for an episode can instantiate the separate availability model; reset occurs outside the episode.

Section 14 quantifies the proposed confirmation, its decoder-dependent power, cost criteria, resource limits, and calibration precision. The 4,800 nominal episodes do not include perturbation controls. The acquisition arms and independent sampling unit must be frozen after a separate pilot.

14 Stateful experiments, resource limits, and confirmatory power

The motor calculation and the irreversible latch address different claims. The motor supplies an explicit observation model for attribution and decision value; the latch supplies a stipulated restriction on future measurements. The original motor simulator does not implement that restriction. It generates each port from rest and does not execute reference installation, restoration, or repair. Imposing a latch on a motor adapter would define a deliberately constrained benchmark, whose physical implementation would need its own transition contract.

An independent simulation study examines several assumptions omitted by that static calculation. It uses standard-library Python, distinct random-number generation, and explicit motion and resource state. Its 240,000 episodes supplement the original JavaScript study's 48,000 episodes; they are not a recount of those results. Neither study contains hardware observations.

The diagnostic target must also be distinguished from current apparatus state. Write $h = (\theta_{\text{origin}}, w_{\text{original}})$ for the original configuration, while tracking current actuator condition and transmission setting separately. Reference installation changes w_{current}; repair may change θ_{current}. These interventions do not rewrite the original label. Attribution of the original fault, present control, and successful repair consequently require separate scores.

14.1 Initial motion changes the experiment

For constant target speed $v = (-1)^{\theta} K u$, the first-order model has the exact transition

$$\omega' = v + (\omega - v)e^{-T/\tau}, \qquad \alpha' = \alpha + vT + (\omega - v)\tau(1 - e^{-T/\tau}).$$

A signed displacement therefore has mean

$$u(\alpha' - \alpha) = (-1)^{\theta} K[T - \tau(1 - e^{-T/\tau})] + u\omega\tau(1 - e^{-T/\tau}).$$

The second term disappears only when the initial velocity is zero or its effect is explicitly incorporated into the likelihood. Zero-drive settling for duration S multiplies velocity by $e^{-S/\tau}$, but also changes position.

Using the paper's parameters, independent enumeration of all 32 five-command sequences gives body-decoder error 0.070783 when independent random command signs are applied without settling. Alternating signs give 0.273952, compared with 0.031348 for pulses independently starting at rest. The second alternating pulse even has negative signed displacement. Repeating the same command instead gives error 0.000025117, because accumulated velocity increases the signal. These are different physical interventions, despite having the same nominal number of pulses.

A separate Runge–Kutta integration checks the analytic transition in 108 cases. Raw-frame Monte Carlo results provide a further implementation check: each condition contains 40,000 episodes, balanced between the two actuator labels.

Condition	Analytical error	Simulated error
Independent starts at rest	0.031348	0.032375
No settling, random commands	0.070783	0.070975
No settling, alternating commands	0.273952	0.274125
Settling for 0.4 seconds	0.031350	0.029800
Persistent difference offset	0.124178	0.124625
Temperature-dependent gain	—	0.065250

Condition $i = 0, \ldots, 5$ uses seed $20261105031 + 104729i$. The final condition stipulates gain $K = 1 - 0.5H$ before each pulse, held constant during that pulse, with the thermal state defined below. This extension isolates body sensing; it does not simulate a reference installation or an additional transmission factor.

Given the conservative pulse-end bound $|\omega| \leq K$, a sufficient settling duration for subsequent mean contamination below 0.01μ is

$$S \geq \tau \log\left(\frac{K\tau(1 - e^{-T/\tau})}{0.01\mu}\right) \approx 0.41172 \text{ s}.$$

This is a model calculation, not a hardware operating prescription. Settling does not reset cumulative travel, backlash, temperature, or wear. Repositioning must therefore be represented and charged separately.

14.2 Preparation error and correlated readings

Fresh observation noise and persistent preparation error have different consequences. Suppose the selected latch route stores $q_e(h) \oplus Z$, where $Z \sim$ Bernoulli(0.02) is sampled once. Subsequent readings independently flip that stored bit with probability 0.1. The complete observation process still factors through $q_e(h)$, so the nominal oracle can simulate it. Repetition, however, learns the corrupted stored value.

For the two-block majority policy, let a_k denote the majority error of k fresh 0.1 flips and set

$$b_k = 0.02 + 0.96a_k.$$

Its errors for $c = 0$ and $c = 1$ are respectively b_k and $(1 - a_k)b_k + a_k(1 - b_k)$. Exact rational calculation gives Bayes error 0.032256057344 at $k = 5$, tending to 0.02. Thus this modification defeats amplification while preserving the oracle's simulation premise.

A continuous analogue is

$$D_j = (-1)^\theta \mu + S + \epsilon_j, \qquad S \sim N(0, s^2), \quad \epsilon_j \sim N(0, 2\sigma_B^2),$$

with S fixed throughout an episode. The sign-of-sum decoder has error

$$\Phi\left(-\frac{\mu}{\sqrt{s^2+2\sigma_B^2/k}}\right) \longrightarrow \Phi(-\mu/s).$$

For $s = 0.02$ radians, five-read error is 0.124178 and the limiting error is 0.070576. Calibration or preparation uncertainty cannot be made fresh merely by acquiring another frame.

For exchangeable Gaussian differences with unchanged marginal variance and correlation ρ, five-read error instead becomes

$$\Phi\left(-\frac{\sqrt{5}\mu}{\sqrt{2}\sigma_B\sqrt{1+4\rho}}\right).$$

It rises to 0.057846 at $\rho = 0.10$ and 0.094059 at $\rho = 0.25$. Even reusing a single noisy baseline frame changes the dependence structure. Acquisition specifications must identify shared frames, calibration variables, and reference-setting errors, rather than reporting only marginal noise.

14.3 Finite resources and the implementation bridge

Consider the illustrative, dimensionless thermal recurrence

$$H_{\text{peak}} = e^{-T/2}H_{\text{pre}} + 0.12, \qquad H_{\text{next}} = e^{-S/2}H_{\text{peak}}, \qquad T = 0.08.$$

Its stationary peak is

$$\frac{0.12}{1-e^{-(T+S)/2}}.$$

Starting cold, repetition without waiting crosses the stipulated limit one on pulse ten. Waiting 0.16 seconds crosses it on pulse twenty-four. The cold-start sequence remains below the limit when $S \geq 0.175667$ seconds. A separate thirty-use resource still blocks pulse thirty-one, irrespective of cooling.

These checks expose an abstraction issue. A guard correctly refusing an exhausted or overheated test may preserve physical safety while violating the abstract requirement that every permitted repeat executes. One remedy is to fix a finite horizon N and establish adequate initial reserves and transition bounds through N. Another is to include resources and motion in the abstract state and establish an appropriate new result. The always-repeatable theorem does not automatically cover that enlargement.

Actual engineering must supply the safe set, disturbance bounds, travel and current limits, shutdown behavior, and reliable setup transitions. Rejection, timeout, and abort are outcomes with information and cost; silently discarding them changes the experiment. None of the dimensionless thresholds above validates an apparatus.

14.4 Operating characteristics of the proposed confirmation

Specify three acquisition arms—routine, body, and reference—and freeze one primary decoder per arm. With twelve cell-by-arm error rates, one-sided exact binomial upper bounds at $\alpha = 0.05/12$ have simultaneous coverage at least 0.95. At 400 episodes per cell, the requirement $U < 0.10$ is equivalent to observing at most 24 errors. The bounds for 24 and 25 errors are 0.098392 and 0.101447. Integer arithmetic verifies the corresponding binomial tail inequalities.

For four independent cells with common error p, the probability that all pass is

$$\left[\Pr\{\text{Binomial}(400,p) \leq 24\}\right]^4.$$

It equals 0.995928 for the continuous body model, but only 0.095909 for the five-sign body decoder, whose stipulated error is 0.059926. Exhaustive sample-size searches using finite binomial sums give minimum sizes of 266 and 963 episodes per cell, respectively, for at least 90% four-cell power at the same confidence level. The reference model requires 73 per cell under its stipulated error 0.001817. These planning results depend on the model and should not replace an independent pilot.

Bonferroni coverage does not require independence between arms. Each binomial bound nevertheless requires its marginal sampling model, and the power product additionally assumes independent cells. Persistent apparatus or session effects may invalidate episode-level independence. Fixed-sample intervals also do not automatically permit optional stopping.

The stated 4,800 episodes cover the three nominal arms only; negative controls require additional episodes. Two decoders evaluated on shared recordings form a paired comparison and expand the confirmatory family if both are primary. Routine attribution is a negative control, rather than an arm expected to satisfy the attribution threshold.

14.5 Joint error and decision-cost criteria

A possible joint 95% confidence family allocates 0.04 across the twelve error bounds and 0.01 across two anchored-arm cost comparisons. At 400 episodes the error cutoff remains 24. For 1,600 independent episodes per arm and the stipulated loss range of width ten, Hoeffding's inequality supplies the upper cost bound

$$\overline{L} + 10\sqrt{\frac{\log(1/0.005)}{2(1600)}} = \overline{L} + 0.406906.$$

Predeclare success when this bound is below the uninformed baseline three.

Independent Gaussian integration gives nominal expected costs 1.890983 for body and 2.158251 for reference. Under these stipulated losses, sufficient lower bounds on passing the cost criterion are 0.999999859 and 0.997644, respectively. Combining cost and error success probabilities requires a union bound or a joint calculation, not an assumption that the statistics are independent.

Reference sensing improves terminal decision loss over body sensing by only 0.132732; its additional setup charge 0.4 outweighs that improvement. Actual settling, restoration, failed trials, and repair expenses must be included. If replacement succeeds with probability s, its expected loss is $3 + 10(1 - s)p$, giving threshold $p > 3/(10s)$. The baseline and confidence calculation must also change; unsuccessful replacement may enlarge the loss range. These results concern counterfactual decisions under declared prices, not demonstrated repair performance.

14.6 Calibration precision and finite claims

The earlier 1,293-per-stratum Hoeffding count estimates sixteen Bernoulli rates to within 0.05 simultaneously. It bounds a pairwise difference only by its observed difference plus 0.10; it does not certify a continuous conditional kernel. To transfer a ten-reading error guarantee within 0.01 using Proposition 12.1, require

$$\eta \le 1 - 0.99^{1/10} = 0.0010045287$$

at every relevant conditional step. Even with observed Bernoulli difference zero, the same conservative method with individual radii $\eta/2$ needs 12,806,679 samples per stratum. More efficient parametric calibration requires justified distribution, variance, drift, and dependence assumptions. Finite data cannot establish exact infinite-horizon kernel equality against arbitrary alternatives. Tie practical equivalence margins to the intended finite-horizon risk claim rather than treating failed significance tests as evidence of equality.

15 Conclusions and remaining boundaries

Functional self-diagnosis is task-relative diagnostic access: which distinctions about a system's mechanisms can permitted experiments acquire, and which decisions benefit? Present control, original-cause attribution, transfer, and repair need different labels and losses.

The complete-law oracle characterizes limiting information under static known finite channels and idempotent access. Its terminal-channel polytope captures Bayes and randomized minimax decisions. Four states are minimal for the displayed gap between reliable finite families and one ongoing consistent policy. Tiny known law differences can remove qualitative ambiguity while making finite budgets decisive; estimated kernels and changing faults require further models.

A specified recovery task gives a direct availability criterion. It and ordinary homotopy imply neither one another. Pointwise recovery opportunity still needs an implementable policy at the controller's actual information level. No naturally occurring apparatus satisfying the exact repeatable-read latch or RF implementation contract has been established.

Stateful motor calculations quantify initial motion, preparation error, heat, depleted resources, and confirmation power. They improve a falsifiable protocol but provide no hardware observations or engineering approval. Actual calibration, transition envelopes, resets, resource reserves, reference restoration, and repair success remain empirical obligations.

Five simulated reviews and two cross-reviews completed the requested internal assessment. Four original historical full texts remain unavailable (Appendix E). The paper is a finite-model synthesis with proofs and counterexamples; publication priority remains unassessed.

Video exposition can show parity compensation, failed attribution, anchored sensing, and an irreversible mistaken choice, followed by the three-state recovery examples and stateful pulse results. Use checked values and identify simulation. No video publication or website changes are included.

A Assumption failures and boundary cases

Topological support. Enumerate the rationals $(r_k)_{k\geq 1}$ and assign mass 2^{-k} to r_k under state zero; use a Gaussian under state one. Both have topological support $\mathbb{R}$, but rational versus irrational outcomes identify the state almost surely. Equality of null sets, not support closure, matters.

Reused noise. Draw one $N \sim \text{Bernoulli}(0.2)$ and return $Y_t = \theta \oplus N$ forever. One-time marginals are positive and separate θ, yet the entire record contains one noisy bit; equal-prior error stays 0.2. Conversely, correlation can identify a state even when one-time marginals agree.

Noise endpoints and unknown noise. For a known BSC, $\varepsilon = 1/2$ is uninformative. $\varepsilon = 0$ or 1 is deterministic and can shrink support. Known $\varepsilon > 1/2$ is informative after reversing decoding. If unknown error rates span both sides of $1/2$, $(\theta = 0, \varepsilon)$ and $(\theta = 1, 1-\varepsilon)$ have the same law. Finite unknown parameters may be included in H; a continuum without a separation gap need not admit a uniform sample bound.

Prior exclusions and numerical pruning. Guarantees concern initial prior support. Assigning a true state probability zero does not make it physically impossible. Replacing small positive weights with zero changes admissible policies. Log-domain arithmetic can prevent underflow; pruning requires a declared approximation or risk contract.

State changes. A trusted reset can put every current state into a safe state while leaving every origin possible. The static theorem does not forbid this. Track origin and current state jointly when both matter. Resource depletion, time-dependent permission, and structural changes require additional state.

Finite oracle arguments. The common-tape proof uses finite H and $\delta < 1/|H|$. It establishes no infinite-hypothesis analogue and no automatic single-policy consistency under irreversible choices. With nonzero sample prices, the terminal-risk limit does not describe the cost-optimal value: amplification may become arbitrarily expensive. If repetitions consume permission, it can fail.

B Verification and novelty ledger

Claim	Evidence	Status and limitation
Support persistence and separation	Proofs, independent audit, finite examples	Elementary synthesis; no priority claim.
Cumulative counterexample	Explicit complexes and edges; two mathematical reviews	Correct for the declared contract; physical use needs the bridge in Section 10.
Oracle feasibility and limiting value	Simulation/amplification proofs; reviewed quantifiers and empty branches	Derived specialization; priority and external review open.
XOR and intervention criteria	Symmetry and rank proofs; 121 rational shift checks	Known identifiability principles, not causal graph discovery.
Optimal-policy table	Exact fractions, 12 exhaustive tree comparisons, 300,000 episodes	Finite supplied model only; no physical data.
Risk gate	Exact posterior/event probabilities and information bounds	Relaxed guarantees; no zero-hazard claim.
Self-model architecture	Matched-information equivalence retained	No superiority or consciousness result.
Physical interpretation	Directed counterexample; conditional bridge; 72 interlock checks	Homotopy alone insufficient; no hardware validation.
Kernel robustness	39 exact leakage comparisons; coupling proof	Finite-horizon bound; factorization remains an assumption.
Motor proposal	48,000 reproduced episodes; 16-state orientation audit	Synthetic study and proposed protocol only.
Historical necessity check	Three-state model and explicit probability bound	Source-defined witness; two simulated source reviews; no official correction.

The reviews are AI-assisted internal assessments, not external peer review. Their objections changed the draft: general support predicates replaced an unnecessarily narrow safety-set model; mutual absolute continuity replaced ambiguous continuous full support; cumulative structure was separated from belief support; empty majority branches were totalized; and Bayes, minimax, zero-error, and finite-family guarantees were distinguished. Code and checks corroborate proofs; they do not prove unrestricted claims.

C Reproduction

Current v0.3 build and verification commands are in `README-v0.3.txt`. Its manuscript is under `review_v03/paper`; reviewer scripts and results are in the corresponding review folders. The original v0.1 files are preserved. The v0.2 source is under `continuation/paper`; recovered Page snapshots retain the earlier disconnection notices as historical records. Current delivery status is in `README-v0.2.txt`. From the package root:

```
python code/experiments.py
python code/independent_checks.py
python code/cumulative_oracle.py
python code/figures.py
python reviews/causal/verify_model.py
python reviews/causal/verify_oracle.py
node continuation/code/robustness_checks.js
node continuation/code/interlock_checks.js
node continuation/code/motor_simulation.js
node continuation/code/historical_counterexample.js
python review_v03/oracle/verify_oracle.py
python review_v03/diagnosis/verify_source_witness.py
python review_v03/topology/enumerate_task_orders.py
python review_v03/experiment/independent_checks.py
python review_v03/code/build.py
```

The first three commands use the Python standard library. Figures also require NumPy and Matplotlib. The simulation is separate from exact calculation. The predecessor checker writes into its current directory; the README specifies a dedicated reproduction directory. The LaTeX source, figure sources, exact JSON outputs, bibliography, source inventory, and independent reviews accompany the PDF.

D A source-audit witness for a printed necessity claim

Under Definition 3.5 and the operational diagnoser construction (17)/(22) of Thorsley, Yoo, and Garcia [35], the necessity direction printed in Theorem 6 fails for this three-state example. Two simulated source reviews independently checked its premises and reconstruction. This is an internal mathematical audit, not an official correction or external peer review. The sufficient assertion is separate; none of our theorems uses its converse. Frequency separation is already explicit in Example 5 of [33]; this is a noisy-mask adaptation.

Start in state 0. Invisible normal event u enters persistent mode 1, or invisible fault f enters mode 2, each with probability 1/2. Mode 1 self-loops on a, b with probabilities $1/5, 4/5$; mode 2 reverses them. The independent sensor reports α from a with probability 9/10, from b with probability 1/10, and β otherwise. There are no later u, f events. All states are live, enabled state/event successors are unique, and there is no invisible cycle. Then

$$p = \mathbb{P}(\alpha \mid N) = 13/50, \qquad q = \mathbb{P}(\alpha \mid F) = 37/50.$$

D.1 Source-defined matrices and recurrence

The sensor-output automaton has a normal invisible transition from 0 to 1 and a fault-suffixed invisible transition from 0 to 2, each of probability 1/2. Subsequent visible events are ordinary

events in both modes; the fault label persists. In augmented order $(0N, 1N, 2N, 0F, 1F, 2F)$, the null-output matrix U has $U^2 = 0$. The absorption equation gives

$$R_y = \widehat{Q}_y + UR_y, \qquad R_y = (I + U)\widehat{Q}_y.$$

The first-observation rows restricted to reachable $(1N, 2F)$ are $[13/100, 37/100]$ for α and $[37/100, 13/100]$ for β. The later cell is $\{(1, N), (2, F)\}$, with

$$T_\alpha = \mathrm{diag}(p, q), \qquad T_\beta = \mathrm{diag}(q, p).$$

In embedded order $(q_0, 0N), (q_1, 1N), (q_1, 2F)$, the component chain is

$$\Pi = \begin{pmatrix} 0 & 1/2 & 1/2 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix}.$$

Both conflicting components are recurrent in the source-defined chain, so the recurrent-certainty criterion fails. Posterior vanishing on records from the other mode does not make a component transient in this chain.

The exact checker compiles all six output matrices, augmented matrices, and the 18-state one-observation chain; solves the absorption equation rationally; extracts reachable cells; and classifies recurrence. It independently sums all 254 visible histories of lengths one through seven from the original event/mask model, matching their likelihoods. All checks pass.

D.2 Posterior and uniform quantifiers

For n visible symbols containing k occurrences of α, every compatible underlying word has exactly $n + 1$ events, with no invisible suffix or ambiguity about event length. Posterior odds are

$$\frac{\mathbb{P}(F \mid Y_{1:n})}{\mathbb{P}(N \mid Y_{1:n})} = \left(\frac{37}{13}\right)^{2k-n}.$$

Definition 3.5 quantifies over prefixes *ending* in a fault. That set is $\{f\}$ here, and each subsequent event emits one observation. For $0 < \rho < 1$, the confidence-failure probability is

$$\mathbb{P}_F\left\{\frac{K_n}{n} \le \frac{1}{2} + \frac{\log(\rho/(1-\rho))}{2n\log(37/13)}\right\}, \qquad K_n \sim \mathrm{Binomial}(n, 37/50).$$

Choose N so the threshold is at most $31/50$ for every $n \ge N$. Hoeffding bounds the event by $\exp(-18n/625)$. Increasing N makes it smaller than any positive tolerance for every later horizon and every required prefix. For $\rho \le 0$ the failure event is empty. This checks the uniform quantifiers, not only almost sure convergence. At $\rho = 0.99$, exact failure at $n = 160$ is $\mathbb{P}\{\mathrm{Binomial}(160, 37/50) \le 82\} \approx 6.0400819\,10^{-10}$. The earlier uniform envelope below 0.01 for $n \ge 160$ remains valid.

Both inspected copies retain the printed if-and-only-if. Equation (21) appears to reverse absorption indices; this reconstruction uses the consistent operational definitions (17)/(22). Referenced appendices were absent from the eight-page copies. Neither review found an excluding premise. A bounded search located no official correction; that is not proof that none exists.

D.3 Later definitions and source gaps

Thorsley [32] is a serious later AA-characterization lead, but only its abstract and first-page preview were recovered. Chen and coauthors [6] was read at Definition 1 and Section IV. Its displayed definition uses all faulty prefixes rather than fault-ending prefixes. The difference is substantive: after $f\beta^m$, posterior odds after another n observations with k occurrences of α are $(37/13)^{2k-n-m}$. For any fixed delay choose $m = n+1$; even $k = n$ leaves odds below one. The later criterion cannot silently substitute for the 2008 definition.

E Historical coverage and remaining originals

The 2005 author copy [33] was checked at Definitions 2–3, Theorems 3–4, and Examples 4–5; the 2007 combined article/erratum [34] at selected information-state/DP results and corrected policy measurability. Both new reviews read the full inspected 2008 paper and rendered its key definitions and theorem pages. Its absent appendices were not read.

Direct comparison now includes de Oliveira [9], pp. 1–8; Renou–Venel [28], Section 4.3, pp. 17–22, discussion pp. 34–35, and Appendix B pp. 44–45. Controlled information and early actions constraining later access are directly checked antecedents. The full proof of their Theorem 3 was not audited. Selected related originals were recovered: Whitmeyer–Williams [36], pp. 1–6; Li–Liu [18], selected pp. 1–10 of pinned v2; and Qiu–Kumar [24], pp. 1–4. Public web reading of the Paoli–Lafortune CDC 2003 precursor does not replace the 2005 journal article or its omitted proof.

Four original-full-text gaps remain: Blackwell 1953 [3], Greenshtein 1996 [12], Paoli–Lafortune 2005 [22], and Thorsley 2017 [32]. Publisher, institutional, author, archive, and metadata routes did not supply them. The apparent Paoli Figshare open-access record has no files and is metadata-only in both inspected versions. These gaps limit priority/subsumption assessment, not the independently stated finite proofs. Retrieval logs, hashes, precise scopes, and failures accompany the paper.

F Simulated review and revision ledger

Five simulated specialist roles and two adversarial cross-reviews completed the requested assessment. They are internal AI reviews, not reports by the real researchers previously shortlisted. Agreement is not a statistical confidence measure or external endorsement. No outreach or hardware action occurred.

Review	Evidence and change	Scope
Oracle mathematics	Complete-law characterization, terminal-channel polytope, and four-state minimality	Finite static known channels; idempotent access.
Oracle cross-review	Explicit private randomization, conditional freshness, and measurable report-vector partition	Independent simulated proof pass.
Stochastic diagnosis	Source-defined matrices and 254 history checks	Exact inspected 2008 conventions.
Diagnosis cross-review	Challenged initialization, invisible events, fault labels, conditioning, and prefix quantifiers	Supported the narrow witness.

Review	Evidence and change	Scope
Topology/control	Minimal witnesses; 3,186 topology cases and 12,848 task triples	Collapse certificates or homology obstructions.
Experimental identification	108 ODE checks, exact cutoff/power analysis, 240,000 independent simulated episodes	No apparatus validation.
Historical decision theory	Stronger controlled-information overlap; retrieval failures documented	Priority remains unassessed.

Exact checks additionally cover 8,192 oracle/report-map cases, 5,120 noisy Bellman values, 640 known-channel values, and 8,704 small-state path cases. The root agent independently verified the 400-trial cutoff and powers with rational integer sums. Earlier versions and outputs were byte-hashed before integration and remain unchanged.

The bridge now specifies an immutable diagnostic-state map and nonempty successors. The robustness statement bounds state-conditional and joint target/transcript laws. Motor sensing is explicitly separated from irreversible availability. Confirmation names arms, decoder-dependent power, sampling units, multiplicity, and costs.

Review reports, scripts, seeded results, inventories, and native-note snapshots accompany this revision. The snapshots retain earlier pending delivery notices; `README-v0.3.txt` gives the completed status. Bounded computations corroborate proofs and do not provide proof-assistant certification.

References

[1] Gowtham Bellala, Suresh Bhavnani, and Clayton Scott. Active diagnosis under persistent noise with unknown noise distribution: A rank-based approach. In *Proceedings of AISTATS*, volume 15 of *PMLR*, pages 155–163, 2011. URL: https://proceedings.mlr.press/v15/bellala11a.html.

[2] Nathalie Bertrand, Éric Fabre, Stefan Haar, Serge Haddad, and Loïc Hélouët. Active diagnosis for probabilistic systems. In *Foundations of Software Science and Computation Structures*, volume 8412 of *Lecture Notes in Computer Science*, pages 29–42, 2014. doi:10.1007/978-3-642-54830-7_2.

[3] David Blackwell. Equivalent comparisons of experiments. *The Annals of Mathematical Statistics*, 24(2):265–272, 1953. doi:10.1214/aoms/1177729032.

[4] Josh Bongard, Victor Zykov, and Hod Lipson. Resilient machines through continuous self-modeling. *Science*, 314(5802):1118–1121, 2006. URL: https://meclab.w3.uvm.edu/papers/2006_Science_Bongard.pdf, doi:10.1126/science.1133687.

[5] Steven Carr, Nils Jansen, Sebastian Junges, and Ufuk Topcu. Safe reinforcement learning via shielding under partial observability. *Proceedings of the AAAI Conference on Artificial Intelligence*, 37(12):14748–14756, 2023. URL: https://ojs.aaai.org/index.php/AAAI/article/view/26723, doi:10.1609/aaai.v37i12.26723.

[6] Jun Chen, Christoforos Keroglou, Christoforos N. Hadjicostis, and Ratnesh Kumar. Revised test for stochastic diagnosability of discrete-event systems. *IEEE Transactions on Automation Science and Engineering*, 15(1):404–408, 2018. URL: https://jchen2020.net/pubs/j_tase2018-r.pdf, doi:10.1109/TASE.2016.2551746.

[7] Herman Chernoff. Sequential design of experiments. *The Annals of Mathematical Statistics*, 30(3):755–770, 1959. doi:10.1214/aoms/1177706205.

[8] Johan de Kleer and Brian C. Williams. Diagnosing multiple faults. *Artificial Intelligence*, 32(1):97–130, 1987. URL: https://dekleer.org/Publications/Diagnosing%20Multiple%20Faults%20AIJ%20reprint.pdf, doi:10.1016/0004-3702(87)90063-4.

[9] Henrique de Oliveira. Blackwell's informativeness theorem using diagrams. *Games and Economic Behavior*, 109:126–131, 2018. URL: https://henriquedeoliveira.com/wp-content/uploads/2020/12/blackwells-informativeness-theorem-using-diagrams.pdf, doi:10.1016/j.geb.2017.12.008.

[10] Uriel Feige, Prabhakar Raghavan, David Peleg, and Eli Upfal. Computing with noisy information. *SIAM Journal on Computing*, 23(5):1001–1018, 1994. URL: https://cadmo.ethz.ch/education/lectures/HS18/SAADS/papers/computing_noisy_information.pdf, doi:10.1137/S0097539791195877.

[11] Daniel Golovin, Andreas Krause, and Debajyoti Ray. Near-optimal bayesian active learning with noisy observations. In *Advances in Neural Information Processing Systems*, volume 23, pages 766–774, 2010. URL: https://proceedings.neurips.cc/paper/2010/hash/1e6e0a04d20f50967c64dac2d639a577-Abstract.html.

[12] Eitan Greenshtein. Comparison of sequential experiments. *The Annals of Statistics*, 24(1):436–448, 1996. doi:10.1214/aos/1033066220.

[13] Alain Hauser and Peter Bühlmann. Characterization and greedy learning of interventional markov equivalence classes of directed acyclic graphs. *Journal of Machine Learning Research*, 13:2409–2464, 2012. URL: https://jmlr.org/papers/v13/hauser12a.html.

[14] Steven R. Howard, Aaditya Ramdas, Jon McAuliffe, and Jasjeet Sekhon. Time-uniform, nonparametric, nonasymptotic confidence sequences. *The Annals of Statistics*, 49(2):1055–1080, 2021. URL: https://arxiv.org/abs/1810.08240, doi:10.1214/20-AOS1991.

[15] Shervin Javdani, Yuxin Chen, Amin Karbasi, Andreas Krause, Drew Bagnell, and Siddhartha Srinivasa. Near optimal bayesian active learning for decision making. In *Proceedings of the Seventeenth International Conference on Artificial Intelligence and Statistics*, volume 33 of *Proceedings of Machine Learning Research*, pages 430–438, 2014. URL: https://proceedings.mlr.press/v33/javdani14.html.

[16] Sebastian Junges, Nils Jansen, and Sanjit A. Seshia. Enforcing almost-sure reachability in POMDPs. In *Computer Aided Verification*, volume 12760 of *Lecture Notes in Computer Science*, pages 602–625, 2021. URL: https://arxiv.org/abs/2007.00085, doi:10.1007/978-3-030-81688-9_28.

[17] Emilie Kaufmann, Olivier Cappé, and Aurélien Garivier. On the complexity of best-arm identification in multi-armed bandit models. *Journal of Machine Learning Research*, 17(1):1–42, 2016. URL: https://jmlr.org/papers/v17/kaufman16a.html.

[18] Zihao Li and Tianhao Liu. Ranking experiments under sequential sampling, 2026. Pinned arXiv:2608.19897v2; selected pages checked. URL: https://arxiv.org/pdf/2608.19897v2.

[19] Michael L. Littman, Richard S. Sutton, and Satinder Singh. Predictive representations of state. In *Advances in Neural Information Processing Systems*, volume 14, 2001. 2001 conference; printed proceedings also cited as 2002. URL: https://web.eecs.umich.edu/~baveja/Papers/psr.pdf.

[20] Mohammad Naghshvar and Tara Javidi. Active sequential hypothesis testing. *The Annals of Statistics*, 41(6):2703–2738, 2013. URL: https://arxiv.org/abs/1203.4626, doi:10.1214/13-AOS1144.

[21] Sirin Nitinawarat, George K. Atia, and Venugopal V. Veeravalli. Controlled sensing for multihypothesis testing. *IEEE Transactions on Automatic Control*, 58(10):2451–2464, 2013. URL: https://arxiv.org/abs/1205.0858, doi:10.1109/TAC.2013.2261188.

[22] Andrea Paoli and Stéphane Lafortune. Safe diagnosability for fault-tolerant supervision of discrete-event systems. *Automatica*, 41(8):1335–1347, 2005. doi:10.1016/j.automatica.2005.03.017.

[23] Judea Pearl. Causal inference in statistics: An overview. *Statistics Surveys*, 3:96–146, 2009. URL: https://ftp.cs.ucla.edu/pub/stat_ser/r350.pdf, doi:10.1214/09-SS057.

[24] Wenbin Qiu and Ratnesh Kumar. Decentralized diagnosis of event-driven systems for safely reacting to failures. In *IFAC World Congress*, 2005. URL: https://skoge.folk.ntnu.no/prost/proceedings/ifac2005/Fullpapers/04835.pdf.

[25] When does an explicit self-model help? calibration and control under internal change in a finite-state pilot. B-Theory research manuscript, 2026. Source 14 October 2026; reading copy 1, 16 October 2026. No named author byline in supplied PDF; not externally peer reviewed. URL: https://www.b-theory.com/manuscripts/r4-self-model-calibration-r1-2026-10-14.pdf.

[26] Gunther Reissig, Alexander Weber, and Matthias Rungger. Feedback refinement relations for the synthesis of symbolic controllers. *IEEE Transactions on Automatic Control*, 62(4):1781–1796, 2017. URL: https://arxiv.org/abs/1503.03715v3, doi:10.1109/TAC.2016.2593947.

[27] Raymond Reiter. A theory of diagnosis from first principles. *Artificial Intelligence*, 32(1):57–95, 1987. URL: https://cse.sc.edu/~mgv/csce580f11/gradPres/reiter-diagnosis.pdf, doi:10.1016/0004-3702(87)90062-2.

[28] Ludovic Renou and Xavier Venel. Comparing experiments in discounted problems, 2025. Pinned arXiv:2405.16458v3; header 5 April, body 8 April 2025. URL: https://arxiv.org/pdf/2405.16458v3.

[29] Yanan Sui, Alkis Gotovos, Joel Burdick, and Andreas Krause. Safe exploration for optimization with gaussian processes. In *Proceedings of the 32nd International Conference on Machine Learning*, volume 37 of *Proceedings of Machine Learning Research*, pages 997–1005, 2015. URL: https://proceedings.mlr.press/v37/sui15.html.

[30] Michał Szyfelbein and Dariusz Dereniowski. Precedence-constrained decision trees and coverings, 2026. Version 1, 24 February 2026; preprint. URL: https://arxiv.org/abs/2602.21312, arXiv:2602.21312.

[31] Brian Theory. From observable states to acquirable predictive models. Technical report, B-Theory, 2026. Revised technical report, 7 October 2026; reading copy 1; not externally peer reviewed. URL: https://www.b-theory.com/manuscripts-html/acquirable-predictive-models/.

[32] David Thorsley. A necessary and sufficient condition for diagnosability of stochastic discrete event systems. *Discrete Event Dynamic Systems*, 27:481–500, 2017. Abstract and first-page preview checked; original full text unavailable. URL: https://link.springer.com/article/10.1007/s10626-017-0236-y, doi:10.1007/s10626-017-0236-y.

[33] David Thorsley and Demosthenis Teneketzis. Diagnosability of stochastic discrete-event systems. *IEEE Transactions on Automatic Control*, 50(4):476–492, 2005. doi:10.1109/TAC.2005.844722.

[34] David Thorsley and Demosthenis Teneketzis. Active acquisition of information for diagnosis and supervisory control of discrete event systems. *Discrete Event Dynamic Systems*, 17:531–583, 2007. Erratum: doi:10.1007/s10626-007-0030-3, pages 585–586. doi:10.1007/s10626-007-0027-y.

[35] David Thorsley, Tae-Sic Yoo, and Humberto E. Garcia. Diagnosability of stochastic discrete-event systems under unreliable observations. In *2008 American Control Conference*, pages 1158–1165, 2008. doi:10.1109/ACC.2008.4586649.

[36] Mark Whitmeyer and Cole Williams. Comparisons of sequential experiments for additively separable problems, 2024. Author manuscript, June 2024. URL: https://whitmeyer.home.blog/wp-content/uploads/2024/06/comparisons_of_sequential_experiments_for_additively_separable_problems_65.pdf.