What Observation Topology Does Not Determine

Operational counterexamples, constrained invariance, and bounded recovery synthesis

Brian Theory
AI-assisted research draft

13 November 2026

D09 — Classification: explanatory negative-results note. The existing corpus already separates homotopy preservation from recovery. The present constructions sharpen the finite-deadline and encoding qualifications, supply a simple stochastic progress bridge, and delimit a finite search for safe recovery certificates. They do not establish a new general topology/control theorem.

Abstract

An observation order describes relations among observable signatures. It does not, without an additional model, describe the transitions an apparatus can execute. We formalize this separation with a fixed two-state observation order supporting every finite-deadline recovery probability between zero and one as the transition law varies. Even fixing the directed support graph does not provide a positive uniform finite-deadline bound without a rate floor. Conversely, a deterministically redundant observation recoding changes the homotopy type of a positive-observation order complex while preserving every attainable task risk under matched processing costs. We prove an operational-equivalence criterion by translating policies and a restricted recovery theorem from an observable integer rank with a uniform conditional progress probability. The theorem gives a binomial-tail bound and expected recovery time at most rank divided by progress probability, attained by a simple chain. We extend the recoding example to arbitrary finite orders on a fixed signature partition. A finite controller-product construction separates invariant safety from recovery and gives an exhaustive search only for a strong monotone certificate class; a three-state chain recovers almost surely without admitting that certificate. A separate established product-chain test gives complete recovery synthesis within a declared bounded deterministic-controller class. A diagnostic-memory fixture illustrates both a one-bit benefit and its failure under ambiguous diagnostic noise. Exact small-structure calculations and Markov recursions verify the examples. The result is a guide to the extra assumptions required for cognitive or physical interpretations of structural invariants, not a rejection of the pure structural mathematics.

1 The overlap audit changes the thesis

The handoff proposed investigating a bridge from structural invariants to prediction, recovery, and control. That is a reasonable question, but the broad negative answer is already present in B-Theory. The current *Causal Self-Diagnosis with Noisy Tests*, version 0.3 of 5 November 2026, separates formal permission from executable physical safety in Section 10. Proposition 11.2 gives three-state counterexamples in both directions between homotopy preservation and a specified recovery task; Proposition 10.1 states a conditional implementation bridge [10]. Repeating that separation with a new vocabulary would not establish a new contribution.

The structural manuscripts have different conclusions. *Binary Order Thinning*, local v3, proves a rooted-incidence classification under a one-bit, height-at-most-two restriction [13, Theorems 3.3

and 4.2]. *Observation Refinement*, v0.5, treats supplied carrier splitting and already warns that some nerve constructions depend on presentation [12, Section 6 and Remark 6.2]. The 9 November two-bit addendum expressly separates its structural witness from physical safety and predictive performance [14, Section 6]. We do not challenge those structural results or infer a stronger task claim from them.

The narrower purpose here is to make three tests explicit: whether the invariant constrains a fixed-deadline probability, whether the description survives an operationally redundant recoding, and which dynamical premises restore a useful bound. The resulting mathematics is elementary. Its value is in auditing interpretation and formulating the next genuine research problem.

External baseline. Controller abstraction already has operational relation frameworks. Feedback refinement relations connect admissible inputs and successor sets so that a symbolic controller can be implemented through quantized state information [9, Definition V.2 and Theorem VI.3]. For stochastic decision processes, bisimulation-based metrics incorporate transition probabilities and rewards and support value comparisons [6, Sections 4–5]. These antecedents make clear why a bare topological invariant cannot replace the missing transition and task data. The rank theorem below is proved directly and is not offered as a replacement for either framework.

2 Keep the structural and operational objects separate

2.1 Positive observation orders

Let S be a finite nonempty set of apparatus states, and let $\Phi = (f_i)_{i\in I}$ be Boolean observations. Define the signature

$$O_\Phi(s) = (f_i(s))_{i\in I} \in \{0,1\}^I.$$

The positive-observation topology on S is generated by the truth regions $f_i^{-1}(1)$. Its specialization preorder, with the convention used here, is

$$s \preceq_\Phi t \quad \Longleftrightarrow \quad f_i(s) \le f_i(t) \quad \text{for every } i.$$

Identify equal signatures to obtain the finite T_0 poset P_Φ. Its order complex $\Delta(P_\Phi)$ has strict chains as simplices. Ordinary Betti numbers below use coefficients in $\mathbb{F}_2$; for our one- and two-point examples the homology is torsion-free and the same ranks hold over every field.

The partition into full signatures is a different object from this positive topology. Observing a full bit value distinguishes zero from one even when only the positive truth region is a generating open. Closing the observation algebra under complements changes the topology, and this must not be mistaken for obtaining an additional physical measurement.

2.2 Controlled dynamics and targets

A separate finite action alphabet $\mathcal{A}$ and kernels $K_a(s,s')$ determine the apparatus dynamics. At time t, the agent observes $O_\Phi(S_t)$, chooses an admissible action based on its observed history, and then S_{t+1} is drawn from $K_{A_t}(S_t,\cdot)$. In all examples the kernels and observation maps are known, and only the current state is observed through the stated interface. A recovery target is a declared subset $G \subseteq S$, with

$$\tau_G = \inf\{t \ge 0 : S_t \in G\}, \qquad V_H(s) = \sup_\pi \mathbb{P}_s^\pi(\tau_G \le H).$$

The deadline H is a nonnegative integer common to every compared policy. First entry is the objective; maintaining membership after entry is a different task. No resource cost is included unless explicitly stated. There is no inference that an edge in P_Φ is an executable action. A model imposing that relationship must state it as an additional assumption.

For prediction we specify an output function and a loss, separately from Φ. A topological similarity between observations does not determine this loss. Policy code, sensing costs, and memory constraints can matter as well; operational equivalence below is stated only under a matched implementation contract.

3 A fixed observation order permits arbitrary recovery probabilities

Theorem 3.1 (Finite-deadline non-determination)**.** *Fix $H \geq 1$. There is a single two-state observation order, a single target, and a single action such that, as its transition parameter varies, recovery probability by H attains every value in $[0, 1]$. The observation order complex is contractible for every parameter.*

Proof. Take $S = \{0, 1\}$, $f(s) = s$, initial state zero, and target $G = \{1\}$. The positive order is the chain $0 \prec 1$, whose order complex is an interval. There is one action with

$$K_p(0, 1) = p, \quad K_p(0, 0) = 1 - p, \quad K_p(1, 1) = 1, \qquad 0 \leq p \leq 1.$$

Starting at zero, recovery has not occurred by step H iff all H transitions remain at zero. Thus $V_H(0) = 1 - (1 - p)^H$. For any desired $v \in [0, 1]$, choose $p = 1 - (1 - v)^{1/H}$. Neither the observation map nor its order depends on p. □

No function of this fixed observation complex alone can give a positive uniform lower bound or an upper bound below one on $V_H(0)$. This is a model-class impossibility: the class permits arbitrary p. It does not refute a task bound proved with a rate restriction, an energy constraint, or a transition-preserving abstraction.

Proposition 3.2 (A fixed support graph still lacks a deadline guarantee)**.** *For the same family restricted to $0 < p < 1$, every member has the same directed transition support, every member reaches 1 almost surely, and*

$$\inf_{0<p<1} V_H(0) = 0, \qquad \sup_{0<p<1} \mathbb{E}_0 \tau_G = \infty.$$

Proof. The support contains $0 \to 0$, $0 \to 1$, and $1 \to 1$ for every allowed parameter. The geometric tail is $(1 - p)^H$, which tends to zero as $H \to \infty$ for each fixed $p > 0$. Its sum gives $\mathbb{E}\tau_G = 1/p$. Letting $p \downarrow 0$ proves both asserted uniform statements. □

The order of quantifiers matters. For each fixed positive rate there is eventual recovery with probability one; there is no common finite deadline guaranteeing a prescribed positive recovery probability across all rates. Qualitative reachability is not a quantitative performance guarantee.

Prediction can vary independently as well. Retain the same observation order and use a symmetric flip kernel with known flip probability $p \in [0, 1/2]$. After observing the current bit, the best zero-one forecast of the next realized bit has error p. Hence identical observation topology supports every optimal one-step error between zero and $1/2$. If the target is instead the next-step probability itself and the kernel is known, it can be reported exactly. This illustrates why specifying the prediction target and loss is as necessary as specifying the dynamics.

What this adds to the current diagnosis report. The earlier report already shows failure of directed recovery under homotopy-preserving thinning. Here the observation structure is held completely fixed while a parameter changes the finite-deadline risk. The distinction is useful analytically, but it is a continuation of the same warning rather than an independent discovery of a topology/control gap.

4 Operationally redundant recoding can change topology

4.1 An exact two-state example

Keep $S = \{0, 1\}$ and $f(s) = s$. With $\Phi = \{f\}$ the signature poset is a chain and $\beta_0 = 1$. Replace the recorded symbol f by the deterministically computed pair $(f, 1 - f)$, so $\Psi = \{f, 1 - f\}$. The signatures are $(0, 1)$ and $(1, 0)$, which are incomparable; P_Ψ is a two-point antichain, $\Delta(P_\Psi)$ is two isolated points, and $\beta_0 = 2$. The homotopy types differ.

Nevertheless the original observation is recovered by reading the first coordinate, and the new one is computed from the old. The complete observational information is identical. In an interface where postprocessing is available at matched cost, every policy can make exactly the same decisions. This example is about recoding an obtained bit. It does not authorize physically executing the complement of a forbidden query. In the acquisition report, oriented physical tests can have different permission even when their answer bits are complements [11, Example 8].

Theorem 4.1 (Operational equivalence under reversible observation coding)**.** *Consider the same state space, initial law, admissible actions, kernels, terminal task, and loss with observation maps $O : S \to Z$ and $O' : S \to Z'$. Suppose maps $g : O(S) \to O'(S)$ and $h : O'(S) \to O(S)$ satisfy*

$$O' = g \circ O, \qquad O = h \circ O'.$$

Suppose the policy classes are closed under these translations and the declared costs of the corresponding implementations agree. Then every attainable finite-horizon task risk under one observation map is attainable under the other. In particular their optimal risks coincide for every common horizon.

Proof. Given a policy using O, translate every observed symbol $O'(S_t)$ to $h(O'(S_t))$ and feed the translated history to that policy. Use the same random seed when coupling the policies. Initially the states and action distributions agree. If the coupled state histories agree through time t, the translated observations and selected actions agree, so applying the same transition randomness makes the next states agree. Induction proves equality of trajectory laws and terminal decisions. Task loss and corresponding implementation cost therefore agree. The map g gives the reverse translation, proving equality of attainable risk sets. □

This theorem is an elementary simulation argument, not a new comparison-of-experiments theorem. It establishes information equivalence under an explicit closure and cost assumption. If computing the recoding consumes time, additional memory, or a paid operation that the baseline cannot use, that assumption fails. Resource-sensitive comparisons must charge those operations rather than invoke the theorem for free.

A criterion for proposed cognitive invariants. Suppose a numerical descriptor is claimed to determine exactly a capability that is invariant under the matched recodings of Theorem 4.1. Either the descriptor is invariant, or the map from descriptor to capability must identify its varying values. For a merely conservative bound, equality is not necessary: two presentations can give different valid

bounds on the same capability. Soundness must hold throughout the permitted class, but tightness need not agree. The earlier draft incorrectly required equality of such bounds; this revision removes that requirement. Positive-observation Betti numbers fail the first test in the example. They may still describe an intentionally restricted observation language, but that language restriction is then part of the interpretation. The theorem does not say that all topology is representation dependent; it identifies a concrete construction and a concrete equivalence relation.

4.2 Every finite order can be presented on the same information

Proposition 4.2 (Arbitrary-order recoding)**.** *Fix a finite observation map O with $m \geq 1$ nonempty signature classes, labeled by L. For every partial order $\leq_L$ on L, there is a Boolean recoding of O whose signature poset is isomorphic to $(L, \leq_L)$ and whose full signatures recover O exactly. Under the implementation assumptions of Theorem 4.1, all these presentations have the same attainable risks.*

Proof. For each $i \in L$ record $f_i(s) = \mathbf{1}\{i \leq_L O(s)\}$, identifying the original signature with its label. If $x \leq_L y$, transitivity gives $f_i(x) \leq f_i(y)$ for every i. Conversely, coordinatewise domination implies $f_x(y) = 1$ since $f_x(x) = 1$, hence $x \leq_L y$. Antisymmetry makes these signatures distinct, so a lookup on the finite image recovers the original label. Thus the signature order is exactly the prescribed order, and reversible policy translation proves the risk claim. □

This is the elementary representation of a finite poset by principal upper-set indicators (equivalently, principal downsets as signature supports; see 4, Section 1), used here to audit operational interpretation; no new order-representation theorem is claimed. In particular, the same m observable classes admit a chain presentation with contractible order complex and an antichain presentation with m components. Any function of the order alone that is invariant under *all* these recodings is constant over all m-element orders. This does not say that every topological space has an m-vertex order-complex realization. Nor does it say that a physical observation language permits every lookup at equal cost: that is the explicit matched-postprocessing contract. Restrictions to positive or otherwise executable formulas can prevent these recodings, but must then be part of the claimed invariant's scope.

5 A conditional bridge from observable progress to recovery

A useful positive theorem must include what an agent can control and observe. Let $q : S \to M$ be the actual observation map, $\bar{r} : M \to \{0, 1, \ldots, d\}$ a known rank, and $r = \bar{r} \circ q$. Let $G = \{s : r(s) = 0\}$. A known feedback action $a(q(s))$ is assumed legal for every state compatible with the observation. Fix $p \in (0, 1]$. Suppose, for every $s \notin G$, it has

$$K_{a(q(s))}\{s' : r(s') \leq r(s) \mid s\} = 1, \tag{1}$$

$$K_{a(q(s))}\{s' : r(s') \leq r(s) - 1 \mid s\} \geq p. \tag{2}$$

These conditions are uniform over hidden states in an observation class. A controller that needs the unobserved true state to choose the promised action is not covered. The theorem also applies to a history-dependent process if the two conditions hold conditionally after every positive-probability history before hitting G.

Theorem 5.1 (Uniform rank-progress bound)**.** *Under* (1)–(2)*, starting at rank* $r_0 \le d$ *and following the specified feedback policy, for every integer* $H \ge 0$*,*

$$\mathbb{P}(\tau_G > H) \le \mathbb{P}(\text{Binomial}(H, p) < r_0), \tag{3}$$

$$\mathbb{E}\tau_G \le r_0/p. \tag{4}$$

For $r_0 = 0$ *both quantities are zero. For* $p = 1$*, recovery occurs within* r_0 *steps. A chain that independently decreases its rank by exactly one with probability* p *and otherwise stays put attains equality in both bounds.*

Proof. At each step before recovery, the conditional probability q_t of a strict rank decrease is at least p. Because the state sets are finite, represent each conditional transition using a fresh independent uniform random variable U_t, allocating the interval $[0, q_t]$ to strict-decrease outcomes. Define $B_t = \mathbf{1}\{U_t \le p\}$. Then the B_t are iid Bernoulli(p) and every $B_t = 1$ before recovery forces a rank decrease by at least one. After recovery, define the auxiliary decrease indicator to be one; it is no longer used to update rank. Nonincrease in (1) means that r_0 such successes are enough to hit zero. Hence the actual hitting time is at most the time T_{r_0} of the r_0th success in the coupled Bernoulli sequence. The event $T_{r_0} > H$ is exactly fewer than r_0 successes among the first H trials, proving (3). The waiting times between Bernoulli successes are independent geometric variables of mean $1/p$, so $\mathbb{E}T_{r_0} = r_0/p$. In the stated chain, rank decreases occur exactly at those successes, giving equality. □

The proof uses conditional probability bounds, not independence of the apparatus's successive progress events. Independence belongs to the comparison sequence produced by the coupling. If strict progress probabilities decline without a positive common lower bound, the theorem no longer supplies a finite expected time. If backward moves are possible, counting successes alone is insufficient.

5.1 What role can an order play?

Suppose an observed finite order supplies an integer height function decreasing along every executable progress transition toward G. If the chosen feedback action also satisfies the no-backtracking and rate-floor conditions, the maximum relevant height supplies d. This is a legitimate use of the order as a progress certificate. The force of the result comes from the relation between the order and the kernels. The order complex's homology alone does not establish that relation.

The equality example has states $0, 1, \ldots, d$, absorbing zero, and transitions $j \to j - 1$ with probability p and $j \to j$ otherwise. Threshold observations $\mathbf{1}\{j \le k\}$ reveal the rank and form a chain of signatures. Its order complex is a simplex, hence contractible, for every d. Yet the mean hitting time is d/p. Even among contractible examples, the height and rates determine a quantitative timescale that homotopy type discards.

A weaker connectivity statement is possible under different assumptions: if every executable transition follows an edge of a given complex, a state cannot reach a target in another connected component. To turn connectivity into a sufficient recovery condition, one must additionally provide executable directions, legal actions, observability of the needed choices, and suitable rates or fairness. A connected picture by itself provides none of those properties.

6 Constrained invariance and a finite certificate search

6.1 A precise controller and uncertainty contract

Let $C \subseteq S$ be a supplied safe set and $G \subseteq C$ the recovery target. Fix a finite nonempty family of known rational kernels K_a^θ, $\theta \in \Theta$, a finite observation map $q : S \to Y$, and legal action sets $A(s)$. A deterministic controller has memory M, action table $\alpha : M \times Y \to \mathcal{A}$, and update table $u : M \times Y \times Y \to M$. Its timing is

$$A_t = \alpha(M_t, q(S_t)), \qquad M_{t+1} = u(M_t, q(S_t), q(S_{t+1})).$$

Thus updating uses the next observation, never the hidden successor state. For an integer $b \geq 0$, the memory limit $|M| \leq 2^b$ charges persistent controller memory; it does not charge the current input symbol or the fixed program tables. It is not a total implementation-memory bound.

On $X = S \times M$ define the induced kernel

$$P^\theta((s,m),(s',m')) = K^\theta_{\alpha(m,q(s))}(s,s')\mathbf{1}\{m' = u(m,q(s),q(s'))\}.$$

A supplied nonempty product domain $D \subseteq C \times M$ contains the allowed initial product states. The controller must be legal at every $(s,m) \in D$ and satisfy

$$P^\theta(x, D) = 1 \quad (x \in D,\ \theta \in \Theta). \tag{5}$$

This is a closed safe domain for this controller, not a claim to have found the maximal viability kernel. It includes conditions at target states so that safety can continue after first entry. An integer rank $R : Z_D \to \mathbb{Z}_{\geq 0}$ is defined on the nonempty observable cells $Z_D = \{(q(s), m) : (s,m) \in D\}$; write $\rho(s,m) = R(q(s),m)$ and $T = D \cap (G \times M)$. Require $\rho^{-1}(0) = T$. Consequently target membership must be observable relative to D; mixed target/non-target cells are excluded.

Fix $p \in (0,1]$. For every $x \in D \setminus T$ and every θ, require

$$P^\theta(x, \{x' : \rho(x') \leq \rho(x)\}) = 1, \qquad P^\theta(x, \{x' : \rho(x') < \rho(x)\}) \geq p > 0. \tag{6}$$

All successor sets here are subsets of D. Checking each model row is stronger than checking only trajectories of one unknown fixed model: it also covers an adversary choosing θ afresh, using the past and even the hidden current state, before each transition. This stronger switching contract is intentional. The certificate may reject problems that are solvable when a single unknown parameter remains fixed.

Proposition 6.1 (Safe recovery and cost)**.** *Under* (5)–(6)*, the controller remains in C almost surely at every time, including after recovery. From initial rank r_0, the bounds of Theorem 5.1 hold uniformly under the stated switching uncertainty. If the nonnegative charge per transition, including sensing and updating, is at most $c_{\max}$, then the charge C_τ accumulated before first entry satisfies*

$$\mathbb{E}C_\tau \leq c_{\max} r_0 / p.$$

If additionally $P^\theta(x,T) = 1$ for all $x \in T$ and all θ, the target itself is invariant after entry.

Proof. Induction using (5) keeps the product in D at every fixed time; the countable union of zero-probability exit events has probability zero. Taking conditional mixtures of the row inequalities preserves nonincrease and the progress floor, even when the model choice depends on hidden state. The rank coupling applies up to first entry. Pathwise $C_\tau \leq c_{\max}\tau$ gives the expectation bound. The additional closure condition on T proves target invariance by the same induction. □

An expected-cost bound is not a hard budget or a replenishment guarantee. If an exhausted store disables an action, include the store in S and verify closure there. Otherwise a small tail of long waits can exhaust any finite reserve. D05 concerns recurrent resource viability in its own explicit model; the present result neither reproduces its threshold nor solves maintenance with hidden health.

6.2 Exactly what finite enumeration decides

Proposition 6.2 (Completeness within the monotone certificate class)**.** *For finite rational data as above, a fixed finite memory alphabet and supplied nonempty D, exhaustive enumeration decides whether a deterministic controller and an integer rank satisfy legality,* (5)*, target identification, and* (6) *with some $p > 0$. It suffices to enumerate ranks in $\{0, \dots, |Z_D| - 1\}$. The decision is only about this certificate class.*

Proof. There are finitely many action and update tables. For each, legality and closure are finite checks. If a nontrivial certificate exists, it must use zero: otherwise a cell of minimum positive rank would have no lower successor in D, contradicting positive progress. List its distinct rank values as $0 = r_0 < r_1 < \cdots < r_k$. Replacing r_i by i preserves all strict and weak comparisons and the zero set; $k \le |Z_D| - 1$. For each candidate compressed rank compute the minimum, over all non-target product states and model rows, of the probability of a strictly lower rank. It is an exactly computable rational number; accept precisely when it is positive and every nonincrease check holds. If $D = T$, take rank zero and $p = 1$ by convention. These steps enumerate every certificate after compression and accept only valid ones. □

The same argument permits enumerating a finite collection of domains D and all memory sizes up to 2^b, with an explicitly supplied initialization rule or an enumeration of maps from initial observations to memory. This is a finite, generally expensive decision procedure, not a complexity improvement. Arbitrary randomized controllers, continuous uncertainty sets without an exact row oracle, and general almost-sure recovery are outside its completeness statement. The earlier program `audit_extensions.py` implements an exact certificate verifier and small fixed-controller checks. Section 7 adds a separate cycle-admitting bounded-controller baseline; neither implementation is a scalable general POMDP solver.

6.3 Certificate failure is not recovery failure

Consider fully observed states a, b, g, one action, safe set $C = \{a, b, g\}$, target g, and transitions

$$a \longrightarrow b \quad \text{surely}, \qquad b \longrightarrow \begin{cases} a & \text{with probability } 1/2, \\ g & \text{with probability } 1/2, \end{cases} \qquad g \longrightarrow g.$$

Starting at a, $\mathbb{P}(\tau_g > H) = 2^{-\lfloor H/2 \rfloor}$ and $\mathbb{E}_a \tau_g = 4$; starting at b the mean is 3. Thus safe recovery is almost sure. A positive-floor monotone state rank cannot exist: the deterministic $a \to b$ transition must decrease rank, whereas the positive-probability $b \to a$ transition forbids that decrease. Even finite deterministic memory cannot fix this particular one-step certificate. Along the infinite supported path $a, b, a, b, \dots$, ranks never increase and every a-step must strictly decrease, contradicting boundedness. The infinite path has probability zero; each of its finite prefixes nevertheless has positive probability, so every tested transition is relevant.

The example instead satisfies the standard block bound $\mathbb{P}(\tau_g > H) \le (1/2)^{\lfloor H/2 \rfloor}$ from any non-target state: conditionally on not having hit, the next two steps succeed with probability at

least $1/2$. Iterating conditional probabilities proves the bound without assuming independent blocks. More generally a conditional success probability at least $\varepsilon \in (0, 1]$ in each block of integer length $L \geq 1$ gives tail $(1-\varepsilon)^{\lfloor H/L \rfloor}$ and mean at most L/ε.

It also has a nonmonotone expected-progress certificate: $V(a) = 4$, $V(b) = 3$, $V(g) = 0$ has one unit of expected decrease at each non-target state. For a bounded nonnegative process with conditional expected decrease at least δ before hitting zero, telescoping the stopped process gives $\delta \sum_{t<n} \mathbb{P}(\tau > t) \leq V_0$; monotone convergence gives $\mathbb{E}\tau \leq V_0/\delta$. This is standard additive drift, available under weaker hypotheses than our monotone-rank premise [8, Theorems 3.1–3.3]. It does not, by itself, give the binomial tail.

Prior art and actual increment. Safe almost-sure reachability under partial observation and support-based winning regions are already studied by Junges et al. [7, Sections 2–4 and 6]; their productivity ranks are not our stronger every-transition nonincrease condition. Finite-memory products and substantially stronger small-policy synthesis results also predate this revision [5]; for example Asadi et al. [3, Theorem 1 and Corollary 2] address limit-sure reachability with small memory. Robust multi-environment models are treated by van der Vegt et al. [15]; we do not equate fixed-environment and switching semantics. Our increment is an explicit, reproducible certificate contract, its quantitative consequence, and a counterexample to completeness for recovery, integrated with the observation-topology warning. None is asserted historically novel.

7 Established recovery synthesis beyond monotone ranks

7.1 A complete bounded-controller baseline

The preceding certificate is sufficient but unnecessarily restrictive for many recoverable systems. The proper comparison is to established finite-memory reachability synthesis, not to an assumption that general recovery has no available methods. Chatterjee et al. [5, Section 4, Proposition 4] already characterize winning finite-memory strategies by paths to the goal from every reachable state-memory pair; their SAT method and deterministic-strategy variant substantially predate this note. Andriushchenko et al. [2, Sections 1–3] explicitly discuss finite controller design spaces, product-chain evaluation, and alternatives to naive enumeration. Controller timing conventions differ across these works, so we do not transfer their memory-size complexity claims to our syntax without a translation.

Here is a self-contained specialization with the timing of Section 6. Fix finite nonempty $S, \mathcal{A}, M, \Theta$, a nonempty initial set $I \subseteq S$, safe set C and target $G \subseteq C$. The family of rational kernels is known, but a single unknown $\theta \in \Theta$ remains fixed throughout each episode. The same deterministic tables α, u, and initializer $\iota : q(I) \to M$ must work for every model and every $s_0 \in I$. Write $I_M = \{(s, \iota(q(s))) : s \in I\}$, $T = G \times M$, and P^θ for the induced product chain.

Two reachable sets are needed. Let U_θ be all states reachable from I_M in the *actual, unstopped* product chain. For analysis only, replace every row at T by a self-loop; let V_θ be the reachable set of this stopped chain. This alteration does not modify the executing controller. In particular, it supplies no ability to recognize a hidden target or halt upon entry.

Proposition 7.1 (Standard finite-product recovery test)**.** *A candidate controller is legal, remains safe almost surely for all time, and reaches G almost surely from every allowed initial state in every fixed model if and only if, for each θ:*

1. *every product state in U_θ projects into C and its selected action is legal;*

2. every product state in V_θ has a directed positive-probability path to T in the stopped chain.

Consequently exhaustive enumeration of all finite deterministic tables and initializers decides this objective for $|M| \le 2^b$, where $b \ge 0$ is an integer. Unlike Proposition 6.2, no observable monotone rank or observable target membership is required.

Proof. Every finite positive-support path has positive probability. Thus an unsafe or illegal reachable state prevents the first property, while absence of such a state ensures it by induction. This test uses actual post-target dynamics because safety is perpetual.

If a stopped-chain state has no path to T, reaching it before the first hit has positive probability, and thereafter hitting is impossible. Conversely, if every reachable state has a path to T, each has a simple such path with at most $|V_\theta \setminus T|$ transitions. The finite chain has a positive minimum on its positive probabilities. Following such a path therefore gives a uniform positive probability of a hit in a bounded block. Repeating the conditional block argument proves almost-sure hitting. Equivalently, every reachable bottom strongly connected component in the stopped chain must be a target singleton. Finite table enumeration is exhaustive by construction; each test is finite. Padding a smaller controller with unreachable memory states, or explicitly enumerating all sizes, handles the memory upper bound. □

An illegal selected action may be temporarily completed by a self-loop when constructing the graph, then rejected if reachable; its arbitrary completion cannot rescue a candidate. A model-specific reachable set must be computed separately for each fixed model. Combining all edges first changes the semantics and can reject valid controllers. Stopping for the second check also matters: a controller starting in G and subsequently moving to a safe trap has already met first-entry recovery, even if it never returns. It must still pass the unstopped safety check.

This is completeness for bounded *deterministic* controllers with the printed access and timing contract. It is not a theorem that deterministic controllers suffice at a fixed memory budget, and it does not solve arbitrary unbounded-memory or randomized policies. It improves the preceding certificate's coverage by using standard control mathematics, not a new topological invariant.

7.2 Quantitative evaluation and the uncertainty boundary

For an accepted controller let

$$n = \max_{\theta \in \Theta} |V_\theta \setminus T|.$$

If $n = 0$, the first hit is at time zero. Otherwise let $\delta \in (0, 1]$ be the least positive probability among transitions out of non-target states in the V_θ, over all models. A simple path to T has at most n edges and probability at least δ^n. Hence, uniformly over the finite fixed-model family and allowed initial states,

$$\mathbb{P}(\tau_G > H) \le (1 - \delta^n)^{\lfloor H/n \rfloor}, \qquad \mathbb{E}\tau_G \le n/\delta^n, \qquad H \in \mathbb{Z}_{\ge 0}. \tag{7}$$

When $\delta = 1$, the zeroth power is one; the bound says recovery occurs within n steps. The conditional block proof does not assume independent actual blocks. This is a standard, often loose bound: the same path-probability and block argument appears explicitly in Akshay et al. [1, Section 4, Theorem 5], using maximum shortest-path distance in place of our possibly larger n. Their sampling-based certification problem differs from our supplied finite family, but the mathematics of the bound is a direct antecedent. The finite positive rate floor here does not contradict Proposition 3.2, whose model class allows rates arbitrarily close to zero.

Exact evaluation is preferable when feasible. Let Q_θ be the submatrix on $V_\theta \setminus T$. Almost-sure hitting implies that its powers tend to zero, and

$$(I - Q_\theta)h_\theta = \mathbf{1}$$

has a unique finite solution giving expected hitting times. For rational inputs, exact Gaussian elimination avoids rounding a small positive escape probability to zero. A bounded rational nonnegative per-transition expected charge replaces $\mathbf{1}$ by its rowwise expectation. This is stopped cost, not a continuing maintenance bill. The resulting values can be compared across enumerated controllers within this fixed class; no global optimality over larger policy classes is implied.

Fixed uncertainty is weaker than adversarial switching. With one action and states a, b, g, consider two deterministic models with absorbing g:

$$K^A : a \to g,\ b \to a; \qquad K^B : a \to b,\ b \to g.$$

Both fixed models reach g within two steps from a or b. A switching adversary selecting B at a and A at b instead produces the cycle $a, b, a, b, \ldots$. Thus modelwise acceptance does not provide the switching guarantee of Section 6. This distinction belongs alongside the existing fixed-environment literature [15]; it is not a new robust-control separation claim.

7.3 A diagnostic-memory fixture and its noise failure

The new computation includes a small synthetic fixture, not a hardware experiment. A hidden fixed fault is $h \in \{L, R\}$. Initial states h_0 share observation `start`; the sole legal test moves to a diagnostic state $d_{h,y}$ whose observed report is $y \in \{L, R\}$. The report equals h with probability $1 - \eta$. The sole legal advance action then moves to h_w, where both faults give the same observation `work`. Two repairs are legal there: the matching repair reaches an absorbing goal with probability p and otherwise remains at h_w; the other repair causes absorbing damage. Goal and damage permit only a hold action. The safe set excludes damage. The two fixed candidate rates are $p = 1/4$ and $p = 3/4$, and both faults are permitted initial states.

For $\eta = 0$, one persistent bit stores the diagnostic report across the change to the ambiguous work observation. Always using the corresponding repair and preserving memory on failed retries is safe and recovers almost surely. For either fault its exact tail and mean are

$$\mathbb{P}(\tau_G > H) = \begin{cases} 1 & H < 2, \\ (1-p)^{H-2} & H \geq 2, \end{cases} \qquad \mathbb{E}\tau_G = 2 + 1/p.$$

The two compulsory setup transitions are charged, giving means 6 and 10/3. With no persistent bit, the work action is the same for both faults and damages one, so no deterministic memoryless controller meets the robust safety objective. The benefit concerns retained information under this interface; it is not an advantage of self-model terminology.

For two memory states we exhaustively enumerate a normalized set of 1,024 controllers: two work-action choices at each memory state, and eight binary memory-update choices. Initial memory is fixed to zero without loss by relabeling. The relevant updates are two `start`→report entries, four report→`work` entries, and two `work`→`work` entries. All other updates are unreachable before first entry or enter absorbing states where the sole action is hold, so they cannot affect safety or recovery and may be set to zero. These reductions cover every controller's behavior in this fixture; the 1,024 tables are not distinct behavioral equivalence classes. Exactly 32 pass. There are only

two normalized one-memory-state controllers, and neither passes. The model includes ten declared states; with perfect reports the two misleading diagnostic states are unreachable.

Adding $\eta = 1/10$ removes every winner in both enumerations. More strongly, no policy under this interface can guarantee both safety and recovery when $0 < \eta < 1$: each report, and thus each pre-repair observation history, has positive probability under both faults. Any chosen first repair damages one compatible fault with positive probability. Randomization cannot make both wrong-action probabilities zero, and withholding repair cannot reach the goal. Memory does not create a missing reliable diagnostic distinction. This failure is as central as the one-bit success case; an imperfect-sensor extension requires a different risk tolerance, a safe diagnostic action, or another declared interface.

What the validation does and does not establish. The cycle-admitting checker accepts all 32 recovering chains in the previous 49-chain study, including the eight rejected by monotone ranks. A further exhaustive grid uses two actions at each non-target state, chooses each of its four successor supports from the seven nonempty subsets of three states, and assigns uniform positive probabilities on each support. Across 2,401 models and 9,604 deterministic memoryless policy evaluations, 2,240 models admit safe almost-sure recovery from both non-target states. Exact arithmetic also verifies 1,344 block-tail inequalities and 164 fixture-tail identities. This is a small exhaustive model study and a functional validation of the implementation. We do not claim comparative performance against published SAT/inductive solvers, realistic sensor calibration, or external experimental validation.

8 Verification and editable figures

The inherited source `code/reproduce.py` performs two bounded cross-checks. The new audit reran it and obtained byte-identical data tables and check summary. This is reproduction of the author implementation, not a new independent algorithm. First, it enumerates all 16 subfamilies of the four Boolean maps on a two-state set. It records their signature partitions and order-complex Betti numbers. For the family $\{f\}$ and the family $\{f, 1-f\}$, the code verifies two distinct signature classes in both, with β_0 respectively one and two. These are exact finite calculations over the stated construction, not an exhaustive result for arbitrary observation spaces.

Second, it propagates the full rank-chain distribution by the Markov recursion and independently evaluates the binomial tail for $d \in \{1, 3, 5\}$, $p \in \{0.05, 0.2, 0.7, 1\}$, and $H = 0, \ldots, 60$. All 732 comparisons agree to within 2.45×10^{-15} in ordinary floating arithmetic. The general equality is proved in Theorem 5.1; numerical agreement is an implementation check. No Monte Carlo trajectories are used and no random seed is needed.

The new `code/audit_extensions.py` uses exact rational arithmetic independently of the inherited floating implementation. It reproduces all 732 rank-chain equalities, reconstructs all 219 labeled four-element partial orders from the recoding proposition, and examines all 49 one-action three-state chains formed by choosing a nonempty successor set for each of two non-target states, assigning uniform positive mass, and making the third state absorbing. Of these 49 chains, 32 recover almost surely from every state; only 24 admit a monotone state-rank certificate. The remaining eight exhibit the certificate's incompleteness. All 2,352 tested tail inequalities for accepted ranks hold exactly. The program checks controller-product construction, a robust floor of 1/4, rank compression, unsafe successors, illegal actions, unobservable targets, target exit despite safety, and the all-target case. These bounded checks are not a general synthesis benchmark or a formal verification of the printed proofs.

The figure intentionally plots dynamical quantities. No Betti-number curve is labeled as

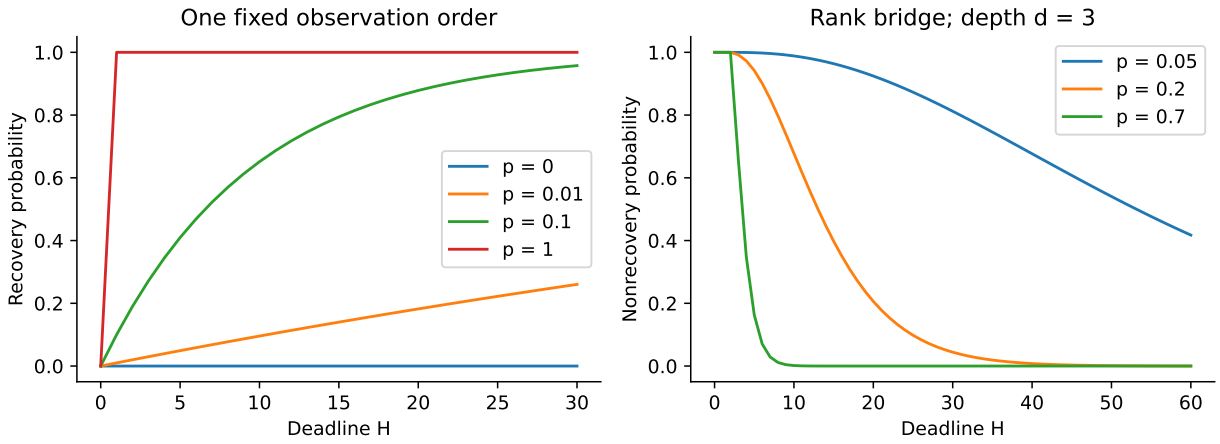


Figure 1: Left: different recovery probabilities with exactly the same two-state observation order. Right: exact nonrecovery probabilities in a depth-three rank chain. The right-hand curves attain the bridge bound. Rates and deadlines are dimensionless discrete-time model parameters, not physical measurements. PDF, editable SVG, tabulated data, and generating Python code are included.

information gain, safety, or awareness. A new topology-based proposal should be judged against the directly available kernel/rank calculation: if its invariant adds neither a valid implication nor a computational advantage, it may be only a visualization.

9 Failure boundaries and a focused research prospectus

When the negative conclusion does not apply. A topological property of a carefully defined state-transition system can constrain motion, obstruction, or coordination. Our no-go construction does not address every such theory. It says that the particular observation structure cannot determine unrestricted, separately specified dynamics. A geometry derived from physical collision constraints, for example, would already contain additional operational information. The specification must identify what is included in the construction and how policy execution relates to it.

When the positive conclusion fails. A single observation class may contain two states requiring incompatible progress actions. A rank defined on true states then fails the observation-factorization premise. Unknown or drifting transition probabilities can also violate a claimed floor p. A certified finite-horizon bound would need simultaneous lower confidence bounds and an account of failure probability; simply substituting an empirical rate into (3) is unjustified. Sensor refinement can change action legality or physically damage the apparatus even when it adds no topological defect. Those changes must be represented in the transition contract, as the existing diagnosis report already emphasizes.

A narrower remaining research problem. Section 6 settles finite enumeration only for its explicit deterministic monotone certificate class. Section 7 now supplies the established cycle-admitting solution for bounded deterministic controllers under a finite fixed-model family. The remaining task is to justify useful computational or interpretive advantages over such established methods under an executable observation and uncertainty contract. Compare block and expected-drift certificates against established POMDP reach-avoid methods, charging their memory, sensing,

and verification costs. Any claimed structural advantage must improve a valid implication or a measured search cost over those baselines. Randomized policies at a fixed memory budget, general uncertainty representations, learned support/rate guarantees, and physical execution remain separate obligations; no new general control theorem or computational advantage is established. Sampling-only certification is already a distinct research problem in current primary work [1], not solved by substituting empirical rows into our checker.

Being, Becoming, and The MIND. The operational lesson is that a system's organization and its ability to recover depend on different specified data. A representation may change while capabilities remain constant, or remain fixed while the apparatus loses capability. A model of an agent's own limitations should track the relationship between observations, actions, and transitions rather than infer capacity from a structural picture alone. These are functional and methodological conclusions; they supply no criterion for subjective awareness.

Merge and publication recommendation. The revision retains the explanatory classification after separate AI-agent specialist and hostile reviews of each revision; these are internal reviews, not external human peer review. This draft should be integrated as an explanatory quantitative appendix to the existing diagnosis/recovery line or retained as a short methods note. Its headline separation is already established in B07, and its rank bound is an elementary stochastic comparison. It remains distinct from D04 because it audits the interpretation of a structural descriptor rather than choosing a representation from data. If a later D04 manuscript proposes topology as a selection score, these counterexamples should be part of that manuscript's validation rather than counted as a separate discovery. No independent novel research-paper status is recommended on the present evidence.

References

[1] S. Akshay, Suguman Bansal, and Gaurav Mahla. Certifying almost-sure reachability in unknown MDPs, 2026. URL `https://arxiv.org/abs/2610.18119v1`. arXiv:2610.18119v1, 13 October 2026; Section 4, Theorems 4–5 inspected; preprint.

[2] Roman Andriushchenko, Milan Češka, Sebastian Junges, and Joost-Pieter Katoen. Inductive synthesis of finite-state controllers for POMDPs, 2022. URL `https://arxiv.org/abs/2203.10803v2`. arXiv:2203.10803v2; Sections 1–3 inspected.

[3] Ali Asadi, Krishnendu Chatterjee, Raimundo Saona, and Ali Shafiee. Limit-sure reachability for small memory policies in POMDPs is NP-complete. In *Proceedings of the Forty-First Conference on Uncertainty in Artificial Intelligence*, volume 286 of *Proceedings of Machine Learning Research*, pages 238–256, 2025. URL `https://proceedings.mlr.press/v286/asadi25b.html`.

[4] George M. Bergman. On lattices and their ideal lattices, and posets and their ideal posets, 2008. URL `https://arxiv.org/abs/0801.0751v2`. arXiv:0801.0751v2; Section 1, principal downsets.

[5] Krishnendu Chatterjee, Martin Chmelík, and Jessica Davies. A symbolic SAT-based algorithm for almost-sure reachability with small strategies in POMDPs. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 2016. URL `https://arxiv.org/abs/1511.08456`. Full text arXiv:1511.08456v1.

[6] Norm Ferns, Prakash Panangaden, and Doina Precup. Metrics for finite Markov decision processes. In *Proceedings of the Twentieth Conference on Uncertainty in Artificial Intelligence*, pages 162–169, 2004. URL https://www.cs.mcgill.ca/~prakash/Pubs/Ferns_MetricsForMDPs.pdf.

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

[8] Timo Kötzing and Martin S. Krejca. First-hitting times under additive drift, 2018. URL https://arxiv.org/abs/1805.09415. arXiv:1805.09415v1; Theorems 3.1–3.3 inspected.

[9] 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. doi: 10.1109/TAC.2016.2593947. URL https://arxiv.org/abs/1503.03715. Full text read in arXiv:1503.03715v3, 2 January 2017.

[10] Research development draft prepared for Brian Theory. Causal self-diagnosis with noisy tests: Observation laws, irreversible access, and task recovery, 2026. URL https://www.b-theory.com/manuscripts/causal-self-diagnosis-r1-v0.3-2026-11-05.pdf. Version 0.3, 5 November 2026; AI-assisted derivation, computation, and independent internal review. Submission authorship unresolved. Accessed 13 November 2026.

[11] Brian Theory. From observable states to acquirable predictive models, 2026. URL https://www.b-theory.com/manuscripts-html/acquirable-predictive-models/. Revised technical report, 7 October 2026; site reading copy 1. Accessed 13 November 2026.

[12] Brian Theory. Successive concept differentiation as observation refinement: Carrier splitting, interaction-core homology, and descent, 2026. URL https://www.b-theory.com/manuscripts-html/observation-refinement/. Referee-repair draft v0.5, 29 September 2026; site reading copy 3. Accessed 13 November 2026.

[13] Brian Theory. Binary order thinning of finite posets: Homology, collapses, and sharp limits, 2026. URL https://www.b-theory.com/manuscripts-html/binary-order-thinning/. Reviewed local v3, 27 September 2026; site reading copy 3 with October 5 editorial polish. Accessed 13 November 2026.

[14] Brian Theory. Signed incidence and a nine-state torsion witness in simultaneous two-bit order thinning, 2026. URL https://www.b-theory.com/manuscripts-html/signed-incidence-two-bit/. Version 1.0 working release candidate, 9 November 2026. Accessed 13 November 2026.

[15] Marck van der Vegt, Nils Jansen, and Sebastian Junges. Robust almost-sure reachability in multi-environment MDPs, 2023. URL https://arxiv.org/abs/2301.11296. arXiv:2301.11296v2, 6 December 2024; introduction and model contract inspected.