Machine-Native Reasoning

Evidence, Methods, and Accountable Revision

Reviewed Mathematical Core

Komorebi P.

Core edition 0.2

9 October 2026

Internal proof review

A continuous argument from accessible evidence to probability-valid operations, correction, controlled inquiry, and accountable revision.

This is a newly written selection, not a concatenation of the research archive. Only the mathematical statements developed and reviewed here belong to its core. Historical experiments, speculative architectures, and unresolved integration claims are not imported as conclusions. The review record separates prose-proof checking, finite computational checks, historical formal evidence, and document validation. This edition is not independently peer reviewed or proof-assistant verified.

The shape of the argument

A machine can retain an answer without retaining a method. It can possess a method without having evidence to initialize it. It can execute that method correctly while using a record that no longer permits correction. It can improve one part of the computation while making the complete process more expensive. These are not unrelated qualifications. Each is a failure at a particular connection between evidence, representation, and action.

This book follows those connections in order. It begins with the information actually available to a reasoner. It asks what state that information can initialize, which operations the state supports, what must be remembered when further evidence arrives, and how one bounded investigation can acquire what is missing. Only then does it discuss improving the investigation and maintaining commitments across several participants.

The running example is deliberately small. A sensor reports a physical bit under an unknown convention, with a third nuisance bit that never affects the specified task. The example is enlarged only when a chapter explicitly changes its contract. It is not a model of a whole mind. Its purpose is to make hidden information, permitted actions, and each passage between results inspectable.

Chapter	Question answered	What the next chapter still needs
1	What does the record determine?	Operations that can be performed on the retained state.
2	What survives a change of representation?	A record capable of later correction.
3	What must memory retain?	Decisions and recipients that use that memory.
4	What can the retained information justify?	An investigation that can actually acquire it.
5	What fits one executable budget?	A justified way to improve the discovery procedure.
6	What search work can be removed?	Evidence that remains applicable after a revision.
7	Which guarantees and gains survive change?	Explicit treatment of other participants and obligations.
8	What can be maintained together?	A clear boundary between the proved chain and its applications.

The second admission block joins the signal-model, Bayesian-refinement, and artifact-planning sources at their common probability interface. Chapters 2, 3 and 5 are expanded in place: no whole source manuscript is appended.

There is one notation and one proof sequence. Repeated source arguments appear once. Original source names identify provenance, not new modules assumed to interoperate. The chapter endings explain exactly which output becomes the next chapter's input.

What the review label means

The earlier 1,919-page collection preserved source documents and navigation. Its own assembly report did not claim a proof review of every imported page. It remains a historical source reader. This book replaces breadth as the admission rule with result-level inspection.

A formal statement is admitted here only when its domains and quantifiers are explicit, its proof is given in full, its source relationship is identified, its use in a later chapter is checked, and its applicable finite construction or sharp boundary is exercised by a documented exact check. This is an internal mathematical review. It does not convert a finite check into a proof of an infinite statement.

The review is also narrower than accepting a source wholesale. A theorem selected from *Machine Societies* does not import a general claim about spontaneously formed institutions. A linear reconstruction argument does not certify an emergent workspace. A commuting equation does not establish that a neural network learned its terms. Those stronger interpretations are outside the core.

The preceding core reran the exact suites of four recent source books. Their receipts remain historical regression evidence, not new runs in this revision. The current build reruns the complete selected-core suite and the additional admission checks. The selection's proofs were reviewed separately. Older formal developments retain their own evidence status, but no historical Lean result is attributed to this rewritten text. No new Lean compilation was performed.

The companion review files name every retained result, its assumptions, its source anchor, its computational check, and its remaining limit. They also account for every source document in the former collection. Omission means either that its full evidence was not re-established for this edition or that it is unnecessary to the continuous argument. It is not a claim that the omitted work is false or valueless.

Contents

The shape of the argument i

What the review label means ii

1 From evidence to a reusable state 1

2 Operations must survive their representations 4

3 Memory that remains open to correction 8

4 From corrected records to decisions and explanations 15

5 One investigation, one resource budget 17

6 Improving discovery without changing its meaning 23

7 What remains valid after a revision 26

8 Shared work and maintained obligations 29

9 One continuous chain, with no hidden promotion 32

Result and source map 34

Reading sources and review boundaries 35

Chapter 1

From evidence to a reusable state

The first question is not how much state to store. It is which distinctions are accessible.

Let Ω be a finite nonempty set of possible worlds. A world may include a physical state, a fixed model index, and an origin tag. An observation record is a function $r : \Omega \to H$. A target is a function $t : \Omega \to Y$. All codomains used below are nonempty. We work on attained records $r(\Omega)$, so no arbitrary interpretation of an impossible record is needed.

Proposition 1.1 (A record determines a target exactly on its fibres)**.** *There exists $d : r(\Omega) \to Y$ with $d(r(\omega)) = t(\omega)$ for every world if and only if*

$$r(\omega) = r(\omega') \implies t(\omega) = t(\omega').$$

Proof. Necessity follows because a function assigns one value to one input. For sufficiency, define $d(h)$ to be $t(\omega)$ for any world with $r(\omega) = h$. Such a world exists for attained h, and the hypothesis makes the choice immaterial. The resulting decoder satisfies the required equation. □

Source relationship. Grounded Abstraction [GA], Proposition 3.2 and Theorem 3.4; Machine Understanding [MU], the representation and sender-access criteria. This is the shared elementary factorization argument, stated once.

The proposition is an information statement, not a runtime bound. The table defining d could be large. Nor does the mere existence of d teach a learner how to find it. A proof that a suitable decoder exists and a justified account of its acquisition remain separate.

1.1 The sensor used throughout the book

Write a world as

$$\omega = (b, s, n) \in \{0, 1\}^3.$$

The physical target bit is b. The sensor convention is s. The nuisance bit n never affects the operations or outputs of the initial contract. The display is

$$y = b \oplus s,$$

where $\oplus$ is addition modulo two. A permitted control f flips the physical bit and leaves s, n unchanged. Consequently the next display is $1 - y$.

An agent given only y can predict every future display under repeated flips: after k flips it is $y \oplus (k \bmod 2)$. It cannot determine the current physical bit. Worlds $(0, 0, 0)$ and $(1, 1, 0)$ share display zero and have different targets. The obstruction is the fibre of the available record, not a shortage of computational effort.

Calibration is a separate, nondestructive observation $c(\omega) = s$. When calibration is permitted and trustworthy, the joint record (y, s) determines $b = y \oplus s$. Adding a blank field called "sensor convention" does not provide this information. The observation operation, its interpretation, and its cost must be part of the contract.

Available record	Number of values	Determines display?	Determines physical bit?
y	2	Yes	No
(y, s)	4	Yes	Yes
(b, s, n)	8	Yes	Yes, with unused detail

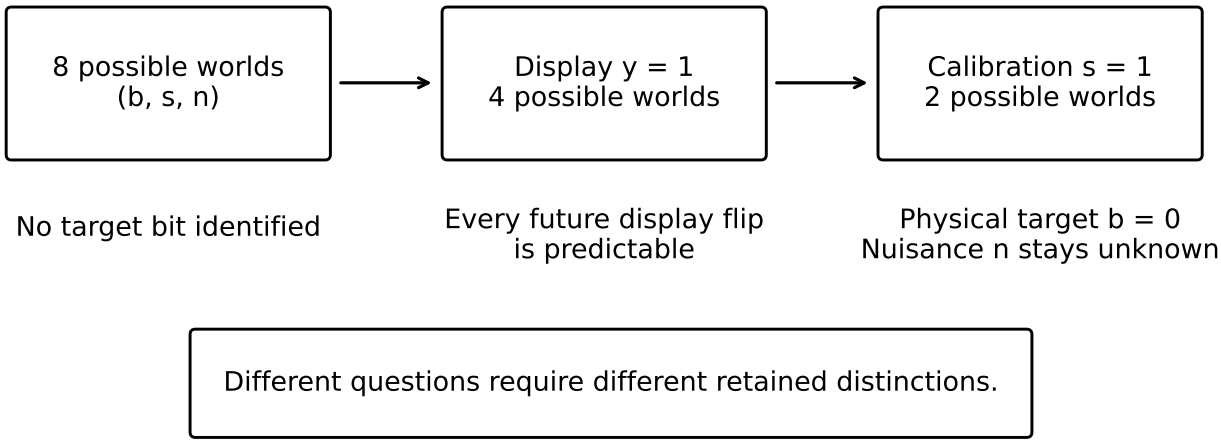


Figure 1.1: One exact observation sequence. The last two worlds differ only in the nuisance bit. Knowing every hidden coordinate is unnecessary for the protected physical target.

1.2 One answer versus every permitted continuation

Let X be a finite nonempty state space. Let $\mathcal{A}$ be a finite alphabet of total deterministic operations $F_a : X \to X$, and let $q : X \to Y$ be the protected output. A word $w = a_1 \cdots a_k$ executes a_1 first. Thus $F_w = F_{a_k} \cdots F_{a_1}$, and the empty word denotes the identity.

Define behavioral equivalence by

$$x \sim x' \quad \Longleftrightarrow \quad q(F_w x) = q(F_w x') \text{ for every finite word } w.$$

The empty word protects the present output. The other words protect its permitted future uses.

Theorem 1.2 (Evidence-initialized behavioral state)**.** *A finite predictor initialized only from $r(x)$, updated only by the operation names, and exact for all protected continuation outputs exists if and only if*

$$r(x) = r(x') \implies x \sim x'.$$

When it exists, $X/\sim$ supplies a predictor with the smallest possible number of initialized states. Its updates are $\bar{F}_a([x]) = [F_a x]$ and its readout is $\bar{q}([x]) = q(x)$.

Proof. If two initial records are equal, a deterministic predictor has the same initialized state. Applying the same subsequent operations produces the same predicted outputs. Exactness therefore requires behavioral equivalence.

For the converse, equivalent states have equivalent successors: an output after F_a followed by w is one of the outputs in the defining family. Hence $\bar{F}_a$ is well defined. The empty word makes $\bar{q}$ well defined. The fibre premise gives a well-defined initializer $d(r(x)) = [x]$. Induction on word length proves exactness. Finally, behaviorally distinct states cannot initialize the same predictor state, since some continuation separates their outputs. Every exact predictor therefore needs at least one initialized value for each class, and the quotient attains that number. □

Source relationship. [GA], Theorem 3.4, and [MU], Theorem 4.1. The two source versions are consolidated, not counted as independent discoveries.

In the sensor, the display task has two behavioral states. Changing the protected output to (y, b) yields four behavioral states. The old record still has only two values and cannot initialize the enlarged quotient. The new task has exposed a real information requirement.

1.3 Why a statistical law is not an observed record

A parameter can determine a distinct probability law while every finite sample remains possible under several parameters. For instance, Bernoulli parameters $1/3$ and $2/3$ assign positive probability to every finite binary string. No estimator based on one such string can be certainly correct for both parameters. The laws differ, but their finite supports overlap.

A correct uncertainty record can be the appropriate output in this setting. Calling the most probable world "identified" would silently replace an error-free claim with a probabilistic one. Chapter 3 gives the corresponding probability update explicitly.

What is now available

We have a criterion for an initialized predictive state, but no automatic guarantee that an arbitrary coordinate manipulation preserves it. The next chapter studies operations on a representation. Its input is a declared state map and operation family, not an assumed learned encoder.

Chapter 2

Operations must survive their representations

A representation is useful only to the extent that its intended operations remain meaningful.

There are two different ways to change representation. An exact split representation retains the whole mathematical state in another space. A compression retains only part of it. The first can preserve arbitrary composition. The second needs additional conditions.

2.1 Exact packets and the hidden-return defect

Let V, Z be vector spaces over a field. Let $A : V \to Z$ be analysis and $S : Z \to V$ synthesis, with $SA = I_V$. Then AS is an idempotent on Z. The map $\mathcal{E}(T) = ATS$ represents a full-state operator $T : V \to V$ on resolved coordinates. Direct multiplication gives

$$\mathcal{E}(T)\mathcal{E}(U) = ATSAUS = ATUS = \mathcal{E}(TU).$$

The resolved identity is AS on the compatible image $A(V)$, not necessarily I_Z on arbitrary coordinates.

For a compression, reverse the splitting. Let $C : V \to Z$ and $R : Z \to V$ satisfy $CR = I_Z$, and put $P = RC$. The compressed operator is $\Phi(T) = CTR$.

Proposition 2.1 (Exact accounting of the discarded excursion)**.** *For any endomorphisms T, U of V,*

$$\Phi(TU) - \Phi(T)\Phi(U) = CT(I - P)UR.$$

Thus compression preserves this composition precisely when the discarded excursion on the right vanishes. Full split encoding $T \mapsto ATS$ preserves composition without that extra condition.

Proof. Expand the product $\Phi(T)\Phi(U) = CTRCUR = CTPUR$ and subtract it from $CTUR$. The identity for full split encoding follows from $SA = I_V$ as computed above. These computations make no approximation and require no orthogonality. □

Source relationship. Sufficient Resolution [SR], Theorems 2.1 and 2.3. The present proof reviews only these algebraic identities, not the complete forty-seven-page source or its full formal implementation.

Counterexample 2.2 (A discarded coordinate returns)**.** Take $V = \mathbb{R}^2$, let $C(x_1, x_2) = x_1$, let $R(z) = (z, 0)$, and let $T(x_1, x_2) = (x_2, x_1)$. The compressed one-step map is zero. Yet $T^2 = I_V$, so $\Phi(T^2) = I_Z$, while $\Phi(T)^2 = 0$. The discarded second coordinate was not irrelevant to the two-step question.

This is the mathematical reason that a short successful readout is not enough. A representation may preserve what is visible now and lose a route by which hidden information becomes visible later. Packets identify the retained and residual components, but the decomposition alone does not prove the residual harmless.

2.2 A method can be transported without reconstructing everything

Suppose source states $z \in Z$ have operations F_a, target states $u \in U$ have operations $G_{\tau(a)}$, and $T : Z \to U$ is a translation. The operation-name map τ is fixed. Let q and $\bar{q}$ be matched readouts.

Theorem 2.3 (Forward transport of a reusable method)**.** *If the source domain is invariant, and*

$$TF_a = G_{\tau(a)}T, \qquad q = \bar{q}T,$$

then for every finite word w,

$$TF_w = G_{\tau(w)}T, \qquad qF_w = \bar{q}G_{\tau(w)}T.$$

Proof. The empty word gives $TI_Z = I_U T$. Assume the first equality for a prefix. Invariance allows the next local equality to be applied on its successor, yielding $TF_aF_w = G_{\tau(a)}TF_w = G_{\tau(a)}G_{\tau(w)}T$. Induction proves the first claim. Composition with the matched readout proves the second. □

Source relationship. [MU], Theorem 6.1; the same local-to-continuation argument appears in OSL and Neuralese. No claim about a learned latent translator is imported.

Boundary. The theorem is forward. It does not imply an inverse, matching of new target operations, cheap computation of T, or that the source solver causally uses the represented state.

The sensor representation (y, s) translates to (b, s) through $T(y, s) = (y \oplus s, s)$. Flipping the physical bit complements the first coordinate in either representation. The square commutes, and so does every repeated flip. The nuisance bit is never reconstructed because it has no role in this contract.

2.3 One quotient for several possible dynamics

A stronger situation arises when the operator itself is uncertain. Let V be finite dimensional, let T_a^θ be finitely many candidate linear operators, and let A_a be one complete candidate system chosen as reference. A model index θ is fixed during the initial comparison. Let $Q : V \to Y$ be the same linear readout for every candidate.

Define $D_a^\theta = T_a^\theta - A_a$ and the propagated discrepancy space

$$E = \operatorname{span}\{A_w D_a^\theta v : v \in V,\ a, \theta,\ w \text{ finite}\}.$$

It can be constructed by starting with the sum of the discrepancy ranges and repeatedly adding their images under the reference actions. In finite dimension this increasing sequence stops. Immediate invisibility $QD_a^\theta = 0$ alone does not guarantee $QE = 0$.

Theorem 2.4 (A common observable quotient)**.** *The following conditions are equivalent:*

1. *$QT_w^\theta = QA_w$ on all of V, for every fixed candidate θ and every word w.*
2. *$E \subseteq \ker Q$.*
3. *There is a surjective linear $r : V \to Z$, common updates G_a, and a readout $\bar{Q}$ with $rT_a^\theta = G_a r$ and $Q = \bar{Q}r$ for every a, θ.*

Under these conditions, V/E is a common quotient and also predicts the output when the candidate index changes between operations.

Proof. Assume the first condition. For every a, θ, w, equality as operators on all of V gives

$$QA_wD_a^\theta = QT_w^\theta T_a^\theta - QA_wA_a = 0.$$

The last equality applies the premise to the word executing a and then w. Spanning proves $E \subseteq \ker Q$.

Conversely, E is invariant under every A_a by construction, and contains the entire range of each difference D_a^θ. Hence $T_a^\theta E \subseteq E$, and every candidate induces the same map $G_a([v]) = [A_a v]$ on V/E. Since $QE = 0$, the readout descends. This proves the third condition. Composing its equalities yields the first condition. The same induction permits a different θ at each step because the quotient map and update remain shared. □

Source relationship. [GA], Theorem 13.1, "Exact common quotient for structured model uncertainty." The all-initial-vectors hypothesis is retained.

For example, let $Q(x_1, x_2, x_3) = x_1$ and let every candidate act as

$$T^\theta(x_1, x_2, x_3) = (x_1 + x_2, x_2, \theta x_3).$$

Then E is contained in the third coordinate and remains invisible. The two-coordinate method $(x_1, x_2) \mapsto (x_1 + x_2, x_2)$ is exact, although the hidden dynamics are not identified.

The conclusion would be different if a later action moved the third coordinate into the first. That change enlarges E beyond $\ker Q$. It invalidates the use of this quotient, not the theorem that justified the former one.

2.4 An algebraic packet is not automatically a probability

The previous sections concern linear representations. Bayesian correction will require more: a nonnegative state with a well-defined total mass. Those requirements belong to a declared domain, not to the spelling of its coordinates.

Let $\mathcal{K}$ be a convex cone in a real vector space, and let λ be linear and strictly positive on every nonzero member of $\mathcal{K}$. Its normalized states satisfy $v \in \mathcal{K}$ and $\lambda v = 1$. A coordinate change can send $\mathcal{K}$ to a cone containing negative coordinates without changing those states' probability meaning. Conversely, an arbitrary signed component of a valid state need not itself be valid.

Proposition 2.5 (When complete packets are genuine probability components)**.** *Let $P_1, \dots, P_m$ be linear maps with $\sum_i P_i = I$. If every P_i preserves $\mathcal{K}$, each normalized $v \in \mathcal{K}$ has the mixture decomposition*

$$v = \sum_{i:w_i>0} w_i v_i, \qquad w_i = \lambda P_i v \geq 0, \quad \sum_i w_i = 1, \quad v_i = P_i v / w_i.$$

Every zero-weight component is zero. If $\mathcal{K} = \mathbb{R}^d_{\geq 0}$ and the P_i are also idempotent, they must be disjoint diagonal coordinate selections. Orthogonal completeness alone does not imply a probability mixture.

Proof. Cone preservation places $P_i v$ in $\mathcal{K}$, so its mass is nonnegative. Strict positivity says that a member of this cone with zero mass is zero. Applying λ to the complete decomposition gives $\sum_i w_i = \lambda v = 1$. Dividing every nonzero component by its mass proves the mixture identity.

On the nonnegative coordinate cone, preservation is equivalent to nonnegative matrix entries, as is seen by applying a matrix to every standard basis vector. Because the matrices sum to I, their nonnegative off-diagonal

entries must all vanish. An idempotent diagonal entry satisfies $x^2 = x$, hence is zero or one. The diagonal sums are one, so each coordinate is selected by exactly one map. Finally, completeness and orthogonality without positivity do not give the first premise, as the following example shows. □

Source relationship. Large Signal Model Foundations [LSM], Proposition 30, physical page 40. Only this finite packet argument and the joint-closure result below are admitted from that source at this stage.

Counterexample 2.6 (An exact detail component has negative probability mass)**.** For $p = (3/4, 1/4)^\top$, the orthogonal average and detail projectors are

$$P_{\text{av}} = \tfrac{1}{2}\begin{pmatrix} 1 & 1 \\ 1 & 1 \end{pmatrix}, \qquad P_{\text{det}} = \tfrac{1}{2}\begin{pmatrix} 1 & -1 \\ -1 & 1 \end{pmatrix}.$$

They sum to the identity and produce $(1/2, 1/2)^\top$ and $(1/4, -1/4)^\top$. The components reconstruct p exactly, but the detail has zero total mass without being zero. Normalizing it as a belief is undefined. Clipping its negative entry and normalizing the resulting sum would instead produce $(3/5, 2/5)^\top$, a different belief.

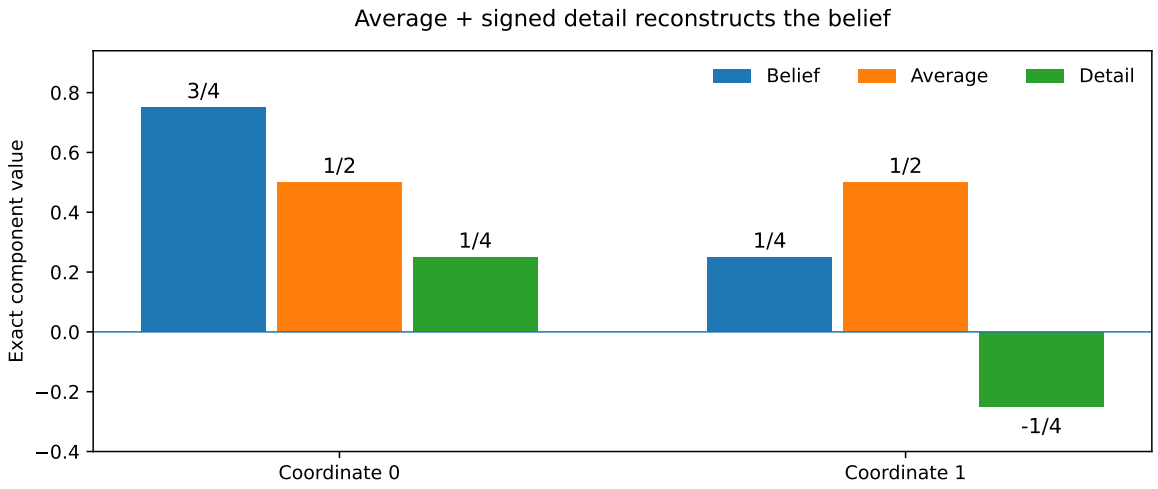


Figure 2.1: Exact reconstruction and probability interpretation are different requirements. The signed detail is useful algebraically but is not a separate probability distribution. Values are from Counterexample 2.6.

This distinction supplies the interface to the next chapter. A signal packet may describe a discrepancy, a coordinate, or a component of an operator. A belief represents nonnegative mass over alternatives. Moving between them requires a domain-preserving map. No learned signal architecture is inferred from this algebraic admission.

What is now available

We have a declared method that preserves the selected outputs. But a predictive quotient need not retain the relationships needed for new evidence. The next chapter changes the retained object from a current physical coordinate to a set or distribution of candidate worlds.

Chapter 3

Memory that remains open to correction

A useful memory preserves the distinctions by which its own conclusions can be revised.

Let $t : \Omega \to Y$ be the current target. A candidate set $V \subseteq \Omega$ records worlds still possible. The empty set is allowed and means that the represented candidates contradict the admitted evidence. It is not a successful identification.

Let $\mathcal{E} \subseteq \mathcal{P}(\Omega)$ contain Ω and be closed under intersection. In this chapter the world does not change when an evidence event is received. The update is $V \mapsto V \cap E$. Physical actions and their effects will return in Chapter 5.

3.1 A state sufficient for every later filter

Define the correction signature

$$\Sigma(V) = \big(t(V \cap E)\big)_{E \in \mathcal{E}}.$$

It records what each admitted future filter would leave of the possible target values.

Theorem 3.1 (Minimal exact memory for a correction contract)**.** *Equality of Σ is a congruence for every update $V \mapsto V \cap E_0$. The signature supports the current target set and its exact update after every finite evidence sequence. Any deterministic retained state supporting the same readouts and updates on all candidate subsets must distinguish every two distinct signatures. Thus the image of Σ has the least possible cardinality for this contract.*

Proof. If $\Sigma(V) = \Sigma(W)$, then for each E,

$$t((V \cap E_0) \cap E) = t(V \cap (E_0 \cap E)) = t(W \cap (E_0 \cap E)).$$

Intersection closure supplies the coordinate $E_0 \cap E$, so the updated signatures agree. The Ω coordinate is the present target set. Induction gives every finite sequence.

If signatures differ, one admitted event produces distinct target sets. A memory that merged those candidates would update identically and return the same answer on that event, contradicting exactness. Hence each signature needs its own state. The signature construction attains the lower bound. □

Source relationship. [GA], Theorem 5.2. This is a behavioral quotient of the correction process, not a claim that all candidate sets must be stored literally.

3.2 An exact count for partition evidence

Suppose $C_1, \ldots, C_m$ partition Ω into nonempty cells and the admitted events are *exactly* their unions. Let

$$D_t = \{(i, y) : C_i \cap t^{-1}(y) \neq \varnothing\}, \qquad r = |D_t|.$$

The occupancy memory retains whether a candidate exists in each attainable cell–target pair.

Corollary 3.2 (The size of partition correction memory)**.** *On all subsets of* Ω*, correction signatures are exactly distinguished by these occupancy records. There are* 2^r *signatures, and* r *fixed binary bits are necessary and sufficient to encode them.*

Proof. An individual-cell filter reveals the target values occupied in that cell, so the signature determines occupancy. Conversely, occupancy determines the target values left by any union of cells. Every subset of D_t occurs: select one world from each included nonempty cell–target intersection and none from excluded intersections. Those intersections are disjoint. Thus there are precisely 2^r possible records. The preceding theorem and binary counting give necessity, while the occupancy vector supplies sufficiency. □

Source relationship. [GA], Corollary 5.3. The word "exactly" in the evidence premise is essential.

For the sensor, the four cells are fixed pairs (b, s) with either nuisance value. Calibration and display evidence generate their unions. The target is b, so there are four attainable cell–target pairs. Candidate membership uses eight bits for eight worlds. Correction memory uses four bits for sixteen possible occupancies. The current target set alone has four possibilities and needs only two bits, but is not sufficient for later calibration.

For example,

$$V = \{(0,0,0),(1,1,0)\}, \qquad W = \{(0,1,0),(1,0,0)\}$$

have the same possible target set $\{0,1\}$. After the event $s = 0$, one leaves target zero and the other target one. Their different associations are operationally relevant.

Counterexample 3.3 (A new evidence event changes the required memory)**.** Take one cell $\Omega = \{u, v\}$ and a constant target. The singleton candidate sets have the same old occupancy. If the event $\{u\}$ is added, filtering one singleton gives a target value while filtering the other gives contradiction. A record sufficient for the old event algebra is insufficient for the larger one.

3.3 Noisy correction is not elimination

A likelihood can be small without being zero. Deleting a candidate because an observation was surprising can invalidate a correct noise model. We therefore specify a separate probability contract.

Let $M \in \mathcal{M}$ be a fixed model with prior $p(m)$. Let an experiment have likelihood $L_m(o)$ on a finite nonempty observation alphabet O. A reporting channel $B(o, z)$, independent of the model conditional on o, creates a retained report Z. The joint law is

$$\Pr(m, o, z) = p(m)L_m(o)B(o, z), \qquad Q_m(z) = \sum_o L_m(o)B(o, z).$$

Posterior distributions are defined only on positive-probability events.

Theorem 3.4 (When a report preserves every Bayesian correction)**.** *The following are equivalent:*

1. *For every prior, the posterior after o equals that after z on every positive-probability pair (o, z).*
2. *The equality holds for one prior positive on every model.*
3. *There is a stochastic kernel $\rho(o \mid z)$ such that*

$$L_m(o)B(o, z) = Q_m(z)\rho(o \mid z) \quad \textit{for all } m, o, z.$$

The reconstruction in the last condition refers to the actual joint experiment, not an unrelated simulation with matching marginals.

Proof. The first condition implies the second. Under a positive prior, define $\rho(o \mid z) = \Pr(o \mid z)$ whenever z has positive probability. For a positive pair, posterior equality yields

$$p(m)L_m(o)B(o, z) = \Pr(o, z)p(m \mid o) = \rho(o \mid z)\Pr(z)p(m \mid z).$$

Substituting Bayes' formula and cancelling positive $p(m)$ gives the factorization. A zero-probability pair under a positive prior has $L_m(o)B(o, z) = 0$ for all m; its conditional weight is zero. If z is impossible for every model, assign any probability row to $\rho(\cdot \mid z)$ because both sides vanish.

Conversely, the factorization makes the likelihood of (o, z) a model-independent multiple of the likelihood of z. Normalization under any prior therefore gives $p(m \mid o, z) = p(m \mid z)$. Since the common channel only reads o, the same positive-pair posterior equals $p(m \mid o)$. This proves the first condition. □

Source relationship. [MU], Theorem 12.1. The full-support and zero-probability cases are included in the proof.

The theorem permits removal of an independent nuisance flag. It does not permit removal of every association. Consider two priors on (b, s), with rows indexed by b and columns by s:

$$P = \begin{pmatrix} 1/2 & 0 \\ 0 & 1/2 \end{pmatrix}, \qquad P' = \begin{pmatrix} 0 & 1/2 \\ 1/2 & 0 \end{pmatrix}.$$

Both individual marginals are uniform. A noisy calibration reports s correctly with probability $3/4$, independently conditional on the fixed world. The probability of report one is $1/2$ under both priors. After that report, however, the target posteriors are $3/4$ and $1/4$. Separate target and sensor marginals cannot support this correction. Their association matters even though neither marginal changed.

These examples do not identify a universally optimal memory. They identify the least information required by a specified correction contract. A new likelihood, a newly admitted model, or a different target can change that requirement.

3.4 A record must close under prediction as well as correction

The static Bayesian theorem preserves one update of a fixed candidate model. A continuing reasoner also predicts what happens between observations. The same compressed record must support both operations. Otherwise a valid posterior can become an inadequate input to the very next prediction.

First consider a known finite state set X with row-stochastic transition matrices $K_a(x, x')$. We retain column-vector probability distributions. A partition π has aggregation $C_\pi(z, x) = 1_{\{\pi(x)=z\}}$, so prediction is $p \mapsto K_a^\top p$.

For each permitted observation symbol o, let $\ell_o(x) > 0$ be its specified likelihood. When these are a complete report alphabet, assume also $\sum_o \ell_o(x) = 1$.

A mandatory partition π_0 with at least two blocks protects the target distinctions. We seek a refinement that supports every prediction and each separate Bayesian update for every prior. The observation is an operation available in its own right. This is stronger than inspecting only a chosen combined transition-and-report schedule.

Theorem 3.5 (Least refinement closed under prediction and correction)**.** *First split π_0 by the complete vectors $(\ell_o(x))_o$, obtaining π_{init}. Then repeatedly split current blocks by*

$$\left(\sum_{x' \in B} K_a(x, x') \right)_{a,B},$$

where B runs over the current blocks. The process stops after at most $|X| - |\pi_{\text{init}}|$ strict refinement rounds. Its final partition is the unique coarsest refinement of π_0 for which all specified predictions admit stochastic coarse transitions and all separate Bayesian updates admit fixed strictly positive coarse likelihoods for every prior.

Proof. At a stable partition the probability of reaching each coarse block depends only on the current block. These sums define a stochastic coarse transition and yield $C_\pi K_a^\top = \bar{K}_a^\top C_\pi$, first on point masses and then by linearity. Likelihoods are constant within final blocks. Grouping the numerator and normalizer by those blocks therefore commutes with every Bayesian update.

For minimality, take any admissible refinement η of π_0. It has at least two blocks. Fix a strictly positive fine likelihood ℓ and a strictly positive coarse likelihood $\bar{\ell}$ inducing its assumed update. For states x, y in one block and z in a different block, compare priors assigning half their mass to x, z and to y, z. Their coarse priors agree. Equality of the posterior odds gives

$$\ell(x)/\ell(z) = \bar{\ell}(\eta(x))/\bar{\ell}(\eta(z)) = \ell(y)/\ell(z).$$

Thus $\ell(x) = \ell(y)$, and η refines the first likelihood split. Coarse likelihoods may differ by a positive scale for each observation without changing this conclusion.

If η refines a current partition, each current block is a union of η blocks. Exact prediction under η makes the probability of reaching each such union constant inside every η block. Hence η refines the next split too. Induction proves that every admissible refinement refines the terminal one. Each strict round increases the block count, proving the bound. Coarseness determines the partition uniquely, up to naming its blocks. □

Source relationship. [LSM], Proposition 9 and Proposition 12, physical pages 11–14; Bayesian necessity also appears in PRIQ [PRIQ], Theorem 5.6. The proof includes the two-block and strict-positivity premises rather than importing them by a cross-reference.

For the sensor, the mandatory display partition has two blocks. A noisy calibration that reports s correctly with probability $3/4$ splits it into the four (b, s) blocks. The nuisance bit remains merged. The physical flip maps these four blocks among themselves, so the transition stage introduces no further split. The representation is now closed for both prediction and correction, not merely for today's display.

Refinement can also propagate backward through dynamics. In a separate four-state example, the mandatory partition is $\{0, 1, 2\} \mid \{3\}$ and the deterministic successors are $(0, 2, 3, 3)$. A constant likelihood adds no information. Transition signatures first isolate state 2, then separate 0 from 1. The final partition is discrete. The reason for the second split is a future observable difference, not a newly sampled measurement.

The construction uses complete tables. It does not acquire those tables, certify their physical accuracy, or guarantee a finite encoding of every posterior coordinate. Its least-state statement concerns the specified separate operations and all priors. Different admissible schedules or reachable-prior restrictions can justify another quotient.

3.5 Retain the joint law of the report and the next state

A general experiment can change the state and report correlated information about the change. Describing its transition and report marginals separately is then insufficient. Let

$$K_{a,o}(x', x) = \Pr(X_{t+1} = x', O_{t+1} = o \mid X_t = x, A_t = a)$$

be a nonnegative finite instrument with $\sum_{o,x'} K_{a,o}(x', x) = 1$ for every a, x. This notation is column-oriented: $K_{a,o}p$ is an unnormalized post-experiment distribution. It specifies an experiment's physical change and its observation in one object.

For a surjection $\pi : X \to Z$, let C sum probabilities over its fibres. Write

$$z_{a,o}(p) = \mathbf{1}^\top K_{a,o}p, \qquad \Psi_{a,o}(p) = K_{a,o}p / z_{a,o}(p)$$

when the denominator is positive. Zero mass means an impossible report under that prior, not an arbitrarily assigned posterior.

Theorem 3.6 (One quotient for reports, posteriors, and feedback)**.** *A nonnegative normalized coarse instrument $\bar{K}_{a,o}$ satisfies*

$$CK_{a,o} = \bar{K}_{a,o}C \quad \textit{for every } a, o$$

if and only if, for all a, o, z',

$$\sum_{x' : \pi(x') = z'} K_{a,o}(x', x)$$

is constant as x varies within any one input fibre. Such a coarse instrument is unique. Equivalently, it preserves for every prior both the report probabilities and the aggregated positive-probability posteriors:

$$z_{a,o}(p) = \bar{z}_{a,o}(Cp), \qquad C\Psi_{a,o}(p) = \bar{\Psi}_{a,o}(Cp).$$

Starting from Cp_0, these identities preserve every finite joint action/report law and every corresponding coarse posterior under one common policy that uses only the previously observed action/report history. This includes randomized policies with the same history-conditional action probabilities.

Proof. Applying the matrix identity to a point mass proves necessity of the fibre condition. Conversely, use the common sums to define each column of $\bar{K}_{a,o}$. They are nonnegative. Summing across z', o partitions the fine successor set and gives column mass one. The definition proves the matrix equality on point masses, hence on every vector. Surjectivity gives uniqueness.

Because $\mathbf{1}_Z^\top C = \mathbf{1}_X^\top$, the matrix equality gives equal report masses. Dividing it by a common positive mass gives equal aggregated posteriors. Conversely, equality of masses and positive-mass posteriors gives the unnormalized matrix equality on every prior. If the common mass is zero, both unnormalized vectors are zero by nonnegativity. Point masses finish that direction.

For a history h, keep unnormalized states $u_h, \bar{u}_h$. Initially $\bar{u}_\varnothing = Cu_\varnothing$. After choosing action a with common probability $\gamma(a \mid h)$ and receiving o, the next states are $\gamma(a \mid h)K_{a,o}u_h$ and $\gamma(a \mid h)\bar{K}_{a,o}\bar{u}_h$. The intertwining identity preserves their aggregation. Induction gives equal mass for every history, including zero-mass histories. Normalization on positive histories gives the posterior statement. Stopping decisions based on the same history can be included as matching terminal choices. □

Source relationship. Selected synthesis of the positive-transfer and filtering arguments in [PRIQ], Theorems 5.2, 6.1 and 6.3, and the projected-law construction in AVOC [AVOC], Theorem 4.5. The joint instrument and feedback formulation is proved here. No entire nonlinear source theorem or Lean implementation is imported.

Counterexample 3.7 (Correct marginals with an incorrect correction)**.** Two fine current states have the same retained label. From the first, the next retained state and report are $(0, 0)$ or $(1, 1)$, each with probability $1/2$. From the second they are $(0, 1)$ or $(1, 0)$, again equally likely. In both cases the next-state marginal is uniform and the report marginal is uniform. After report one, however, the next retained state is certainly one in the first case and certainly zero in the second. No common filter based only on the old label can realize both. The joint fibre condition detects exactly the lost association.

Preserving only normalized posteriors is weaker still. A report coin independent of the hidden state leaves every prior unchanged whether its success probability is $1/4$ or $3/4$. Those two instruments have identical posterior maps but different report laws. A planner that prices later branches needs those probabilities. In the exact instrument theorem the evidence scale is fixed, rather than silently discarded during normalization.

3.6 Small law error and a rare conditional history

Approximate models must be compared before declaring their posteriors close. Let $\widetilde{K}$ be a normalized instrument on the same coarse state and report alphabets. Suppose, at step t, every relevant action and fine state satisfy

$$d_{\mathrm{TV}}\left(\left(\sum_{x':\pi(x')=z'} K_{a,o}(x', x)\right)_{o,z'}, (\widetilde{K}_{a,o}(z', \pi(x)))_{o,z'}\right) \le \epsilon_t.$$

This is a bound on a complete report/next-state distribution, not on an isolated matrix entry. Assume $d_{\mathrm{TV}}(Cp_0, \widetilde{p}_0) \le \delta_0$ with nonnegative δ_0, ϵ_t. Both models run the same history-based policy. Coarse states recorded in the comparison below are analytical coordinates, not extra observations supplied to that policy.

Proposition 3.8 (A whole-history bound with an explicit conditioning cost)**.** *At horizon H, the distributions $P_H, \widetilde{P}_H$ of the full coarse-state and action/report histories satisfy*

$$d_{\mathrm{TV}}(P_H, \widetilde{P}_H) \le d_H := \min\left\{1, \delta_0 + \sum_{t=0}^{H-1} \epsilon_t\right\}.$$

The same bound holds after forgetting any of those history coordinates. For an observed history event E with probabilities $p_E, \widetilde{p}_E > 0$, the conditional final coarse-state distributions obey

$$d_{\mathrm{TV}}\big(P_H(Z_H \in \cdot \mid E), \widetilde{P}_H(Z_H \in \cdot \mid E)\big) \le \min\left\{1, \frac{d_H}{\max(p_E, \widetilde{p}_E)}\right\}.$$

If either event probability exceeds d_H, the other is necessarily positive.

Proof. Condition on a complete coarse history with positive fine probability. Its hidden current-state distribution is a mixture of states in the last coarse fibre. Convexity of total variation makes the true next report/coarse-state row differ from the declared row by at most ϵ_t. Mixing over a common action kernel preserves this bound. Extend the two old history laws first through the same declared row. Total variation cannot increase under a common stochastic kernel, by the triangle inequality and row sums. Replacing those rows with the true conditional rows adds at most ϵ_t, weighted by a total history mass of one. Hence the history distance at the next step is at most the old distance plus ϵ_t. Induction and the cap one prove the first claim. Summing over forgotten coordinates is also a stochastic map, proving the second.

For conditioning, let u, v be the restrictions of the two history laws to E, with masses $p_E, \widetilde{p}_E$. The difference on the complement has ℓ^1 norm at least $|p_E - \widetilde{p}_E|$. Thus

$$\|u - v\|_1 + |p_E - \widetilde{p}_E| \le 2d_H.$$

Adding and subtracting v/p_E bounds the conditional distance by d_H/p_E. Interchanging the laws gives $d_H/\widetilde{p}_E$, so the smaller of these bounds has the displayed maximum denominator. Projection to the final state cannot increase the distance. Finally, the probability of any event changes by at most total variation, which proves positivity when one mass exceeds d_H. □

Source relationship. Finite probability-instrument specialization of [LSM], Propositions 43–44, with the joint retained-state record made explicit. The symmetric conditioning form follows by applying the same elementary bound in both directions. [PRIQ], Section 9 supplies the complementary stepwise posterior-error perspective.

Counterexample 3.9 (A rare report carries the entire discrepancy)**.** Two joint laws place mass $1 - \delta$ on the same ordinary outcome. On a report event E of mass $\delta > 0$, one puts all mass on retained state zero and the other on retained state one. Their whole-law distance is δ, while their conditional distance after E is one. The factor involving event probability is therefore essential. This is a finite-law statement, not an empirical estimate of a sensor error.

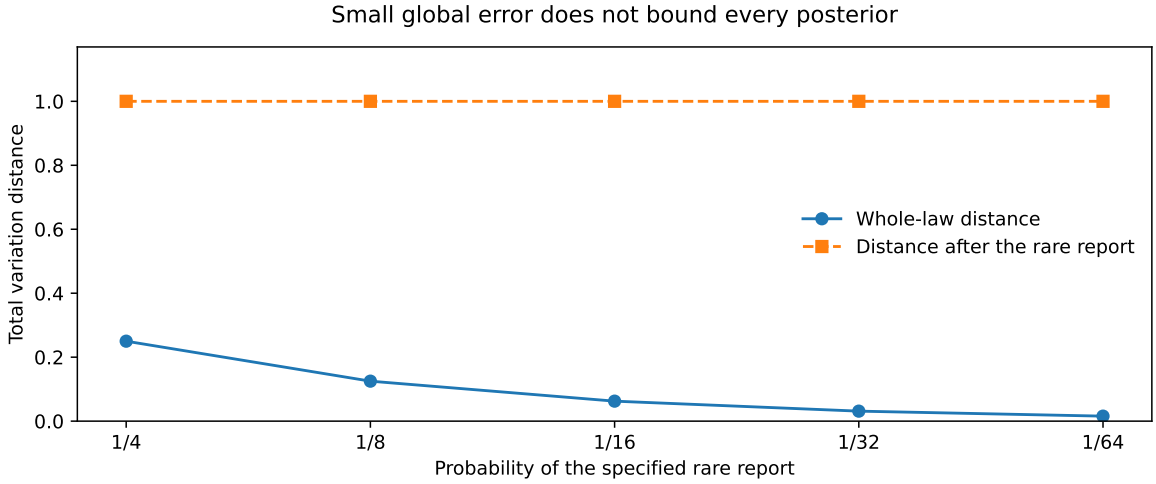


Figure 3.1: In Counterexample 3.9, shrinking the rare-event mass reduces whole-law error but leaves conditional error maximal. The points are exact specified probability laws, not sampled observations.

The bound compares one policy across two models, not independently optimized controllers. Chapter 5 supplies the uniform comparison needed for re-optimization. First, the next chapter asks what decisions and explanations the corrected record can justify.

Chapter 4

From corrected records to decisions and explanations

The same record, model, and continuation must support the decision being certified.

Correctness of a probability calculation is not a guarantee that a real process follows its likelihood model. Within the declared model, however, a small error in the retained distribution can support a precise decision bound. We state that bridge before introducing communication, so that an explanation has a definite object to preserve.

4.1 One quantitative decision contract

Let p and $\widehat{p}$ be distributions on the same finite world set. For each world let v_ω be a response and $\widehat{v}_\omega$ its approximation in one metric space. Assume

$$d_{\mathrm{TV}}(p, \widehat{p}) \le e, \qquad \sup_\omega d(v_\omega, \widehat{v}_\omega) \le d_0.$$

For each action a in a finite nonempty set, let the loss $\ell(a, v)$ lie in $[0, 1]$ and be K-Lipschitz in v. Define the model risks

$$R(a) = \sum_\omega p(\omega)\ell(a, v_\omega), \qquad \widehat{R}(a) = \sum_\omega \widehat{p}(\omega)\ell(a, \widehat{v}_\omega).$$

Theorem 4.1 (A decision bound for one actual approximation chain)**.** *Put $\xi = \min\{1, e + Kd_0\}$. For every action, $|R(a) - \widehat{R}(a)| \le \xi$. If $\widehat{a}$ minimizes the approximate risk, then*

$$R(\widehat{a}) - \min_a R(a) \le \min\{1, 2\xi\}.$$

If the approximate gap between its unique minimizer and every alternative exceeds 2ξ, that action is also the unique true model-risk minimizer.

Proof. Add and subtract $\sum_\omega p(\omega)\ell(a, \widehat{v}_\omega)$. The response term has absolute value at most Kd_0. The expectation of a function in $[0, 1]$ changes by at most total variation, giving the other term at most e. Risks lie in $[0, 1]$, so their difference is also at most one.

For a true minimizer a_*,

$$R(\widehat{a}) \le \widehat{R}(\widehat{a}) + \xi \le \widehat{R}(a_*) + \xi \le R(a_*) + 2\xi.$$

The loss range supplies the cap one. Finally, for any other a, subtract the two error bounds to obtain $R(a) - R(\widehat{a}) \ge \widehat{R}(a) - \widehat{R}(\widehat{a}) - 2\xi > 0$. □

Source relationship. [GA], "A common quantitative grounding certificate"; [MU], "A joint posterior and recipient decision certificate." This is a specialization to a common finite candidate set and bounded loss.

The phrase "model risk" matters. The theorem cannot repair an omitted candidate, a false sensor convention, or an action outside the stated loss function. It also cannot combine a posterior bound from one history with a response bound from another. The record must bind the actual calculation and its premises.

For the noisy calibration example, use zero-one loss for selecting the target bit. Under P and report one, choosing one has risk $1/4$, while choosing zero has risk $3/4$. An approximation with total variation at most $1/10$ and no response error preserves the unique decision because the gap $1/2$ exceeds $2/10$. Under P' the correct decision reverses. A valid error bound around the wrong prior is not a substitute for retaining the correct association.

4.2 An explanation belongs to a recipient

An artifact is not an explanation merely because its producer can interpret it. Fix a finite case set X, a sender-access map $\alpha : X \to H$, and a finite artifact catalogue C. Recipient i has a background record $\beta_i(x)$, a demanded output $t_i(x)$, and a fixed bounded procedure L_i. A timeout, malformed artifact, wrong output, or budget overrun counts as failure.

Let $C(x)$ be the artifacts that succeed for every recipient in case x. For an attained sender value h, put

$$C_h = \bigcap_{x:\alpha(x)=h} C(x).$$

Theorem 4.2 (A common artifact for fixed recipients)**.** *A sender restricted to $\alpha(x)$ can provide one exact artifact for all recipients if and only if every C_h is nonempty. When feasible, the smallest possible number of distinct transmitted artifacts is*

$$\min\{|C| : C \subseteq C,\ C \cap C_h \neq \varnothing \text{ for every attained } h\}.$$

Proof. A sender seeing h must choose the same artifact for all cases in that fibre. It must therefore select a member of C_h, proving necessity. Choosing one member of each nonempty finite intersection gives a valid encoder. The image of any encoder intersects every C_h. Conversely, any catalogue meeting all those sets supplies such an encoder by choosing from the intersections. Finite minimization gives the formula. □

Source relationship. [MU], "The fixed-recipient covering criterion." The earlier designed-decoder coloring criterion is not substituted for the installed-recipient problem.

Counterexample 4.3 (Individual acceptability is not a common message)**.** Two cases have the same sender record. In one, the sole acceptable artifact is c_0; in the other it is $c_1 \neq c_0$. Each case has a valid explanation if its hidden identity is known. Their intersection is empty, so the sender cannot select a valid common explanation from its actual information.

In the sensor, a sender observing y can transmit one bit to a recipient holding s. The recipient's fixed XOR procedure returns $b = y \oplus s$. The sender need not know b: its message combines with the recipient's background.

A decoder that ignores this message fails where background alone cannot distinguish the target. More strings do not fix that procedure. Installed information, interpreter work, and message size are separate resources.

What is now available

We know which artifacts work given sufficient sender information. The next chapter asks whether that information can be acquired within one budget, counting every stored branch of the controller.

Chapter 5

One investigation, one resource budget

A possible answer, an affordable observation, and a small program must belong to the same workflow.

Let X be a finite nonempty state space. It includes every origin tag, permission, background record, and history variable that affects the stated behavior. An experiment e is allowed on a domain $D_e \subseteq X$, has a deterministic successor $f_e : D_e \to X$, a report $o_e : D_e \to O_e$, and a fixed nonnegative integer resource vector c_e.

A current support $V \neq \varnothing$ records possible states. The experiment is permitted only if $V \subseteq D_e$. After report o, the possible *post-experiment* states are

$$V_{e,o} = \{f_e(x) : x \in V,\ o_e(x) = o\}.$$

Only nonempty report branches are retained. If the target concerns a past value, that value must survive as an immutable tag. Resetting the current state is not recovering its history.

A terminal artifact a is acceptable only if every fixed recipient succeeds in every state of V. Its worst-case terminal cost is $g_a(V)$, computed componentwise. A controller is a literal finite tree with one record per experiment or terminal node. Stored memory counts all branches, including those not visited in one execution.

5.1 The complete finite frontier

Let $\mathcal{P}_h(V)$ be the nondominated pairs (v, m) of worst-case online resource vectors and stored-record counts attainable by successful trees with at most h experiment steps. At depth zero the available pairs are

$$\mathrm{Pareto}\{(g_a(V), 1) : a \text{ succeeds throughout } V\}.$$

At greater depth retain those stops and, for each uniformly permitted experiment and one feasible child pair (v_o, m_o) for every nonempty branch, form

$$\left(c_e + \max_o v_o,\quad 1 + \sum_o m_o\right).$$

A profile dominates another when it is no larger in every resource coordinate and strictly smaller in at least one. The maximum is componentwise. An impossible child contributes no tree; it is not a zero-cost branch.

Theorem 5.1 (Exact joint feasibility for explicit controllers)**.** *This recurrence gives exactly the nondominated resource profiles of successful depth-bounded controllers. Every retained pair has a concrete tree witness. A common online and storage budget is feasible if and only if it bounds one retained pair. Removing dominated child profiles does not change feasibility under any componentwise budget.*

Proof. At zero depth the only possibilities are acceptable terminals, whose profiles are exactly the stated stops. At positive depth a successful tree either stops or begins with one uniformly allowed experiment. Every possible report needs a successful child on precisely $V_{e,o}$. By induction those children have profiles in the lower-depth attainable set. Worst-case online cost is the root cost plus their componentwise maximum. Complete tree storage is one plus the sum of all child records.

Conversely, combining any allowed root with a successful child for every report constructs a successful tree with that profile. This proves exact enumeration before pruning. Replacing a dominated child by a dominating one cannot increase either a coordinatewise maximum or a sum of record counts, and does not change its supported initial set. Thus Pareto pruning preserves every feasible budget and every retained profile has a witness. □

Source relationship. [MU], "A joint resource certificate for fixed recipients," and [SII], Theorem 6.2. The scalar-cost repair case is a specialization, not a different composition theorem.

The recurrence can be expensive. It is a finite mathematical characterization, not an efficient general learner. It assumes its transition and acceptance tables. Their discovery and checking are preparation costs, not free consequences of the existence theorem.

Counterexample 5.2 (Two separate optima do not form one controller)**.** There are four immutable cases, each requiring its own terminal answer. A direct observation of the case costs four units and gives a five-record tree including four terminals. Unit-cost equality tests can identify the case in at most three tests with seven records. Terminal answers have zero online cost in this particular example. The frontier is $(4, 5)$ and $(3, 7)$. The mixed budget $(3, 5)$ is infeasible: an all-equality successful tree needs four leaves and at least seven records; any tree using the direct observation on a reachable path costs at least four.

5.2 When a checked candidate can end the search

Consider $n \geq 1$ immutable cases with total nondestructive binary unit-cost queries. Every unit-cost repair succeeds in exactly one case, and every case has a repair. A controller must succeed certainly in every case.

Theorem 5.3 (A candidate attaining both lower bounds is optimal)**.** *Every successful controller has worst-case work at least* $1 + \lceil \log_2 n \rceil$ *and literal storage at least* $2n - 1$*. A checked controller attaining both bounds and the requested horizon realizes the entire nondominated work/storage frontier.*

Proof. Distinct cases cannot end at the same successful repair leaf, so at least n leaves are reachable. A binary tree of query depth d has at most 2^d leaves, giving $d \geq \lceil \log_2 n \rceil$. The terminal repair adds one unit.

Let l be the number of leaves and i_1, i_2 the counts of unary and binary internal nodes. Counting tree edges gives $i_1 + 2i_2 = i_1 + i_2 + l - 1$, hence $i_2 = l - 1$. Total records are $2l - 1 + i_1 \geq 2n - 1$. A successful candidate at both coordinatewise lower bounds dominates every other attainable cost profile. Its validity and horizon check make it an actual member of the frontier, leaving that pair as the whole nondominated frontier. □

Source relationship. [SII], Theorem 11.1. The assumptions exclude destructive resets and repairs that succeed on several cases.

For eight possible defective cells, equality-first diagnosis requires seven queries in the worst case, while three coordinate-bit queries identify the cell. Including repair gives costs eight and four, respectively. Both literal trees have fifteen records. Faster execution has not eliminated the cost of storing the procedure.

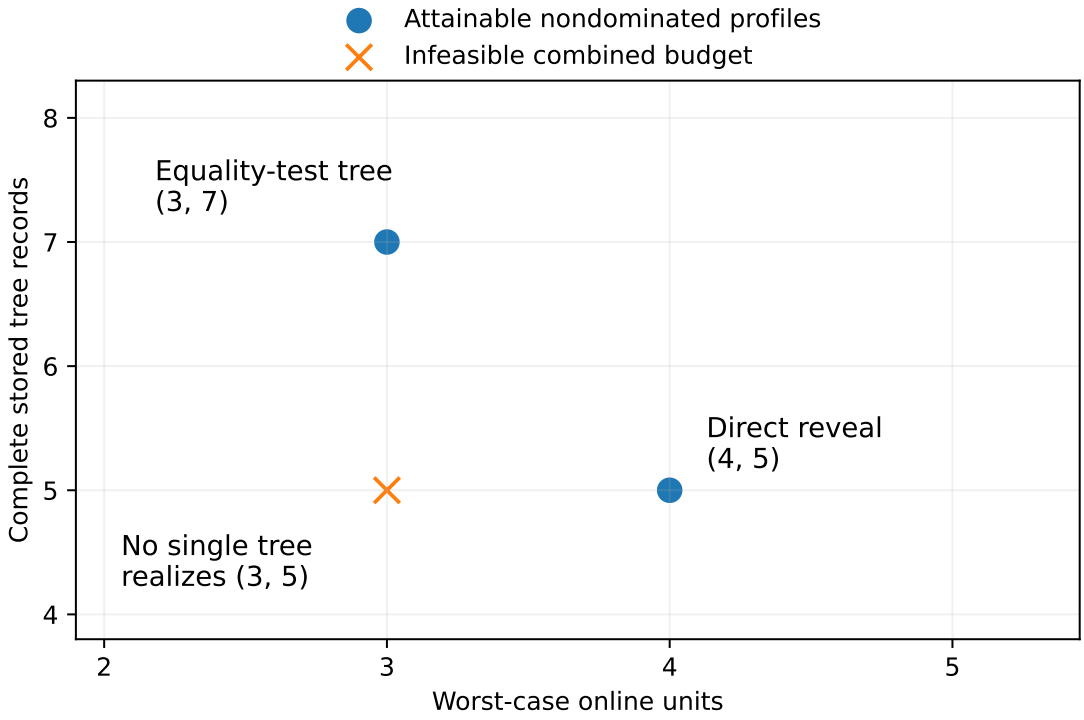


Figure 5.1: The four-origin example has two nondominated profiles. Combining the lowest coordinate from each does not construct a feasible controller. Terminal answers have zero online cost in this example.

This eight-cell example is separate from the eight-world sensor. The sensor has only four action-relevant (b, s) pairs and an irrelevant nuisance coordinate. Confusing hidden-world cardinality with the number of required repair distinctions would import a false lower bound.

5.3 Expected-cost inquiry uses the report probabilities

The preceding frontier certifies worst-case execution and complete tree storage for deterministic observations. A noisy report requires a different optimization. Expected cost weights each future branch by its probability. A low expectation is not a hard execution budget, and neither replaces the cost of storing the whole controller.

Use the finite instrument of Theorem 3.6, with fine belief p and coarse belief $q = Cp$. At q, a finite set of experiments $A(q)$ is permitted, and a nonempty finite set of terminal decisions $D(q)$ is available. These catalogues must be identical for every fine belief with the same q. Stopping with decision d has loss $\ell_d(\pi(x))$ in fine state x, so its conditional risk is $\langle q, \ell_d \rangle$. An experiment has a finite nonnegative charge $c_a(q)$, identical on the fine and coarse side. All history variables affecting permissions, charges, or targets must be included in the declared state.

Let $J_h(p)$ be the least expected charge plus terminal loss attainable using at most h further experiments, with terminal decisions permitted at any step. It is defined by finite observation trees, not by a hidden-state controller.

Its coarse recurrence is

$$\bar{J}_0(q) = \min_{d \in D(q)} \langle q, \ell_d \rangle,$$
$$Q_h(q,a) = c_a(q) + \sum_{o:\bar{z}_{a,o}(q)>0} \bar{z}_{a,o}(q) \bar{J}_{h-1}\big(\bar{\Psi}_{a,o}(q)\big),$$
$$\bar{J}_h(q) = \min\left\{ \bar{J}_0(q), \min_{a \in A(q)} Q_h(q,a) \right\}.$$

An empty experiment catalogue leaves only stopping. Impossible reports do not receive invented posteriors.

Theorem 5.4 (The same compressed belief supports the same inquiry)**.** *Under the exact instrument quotient and the common loss, charge, and catalogue assumptions above,*

$$J_h(p) = \bar{J}_h(Cp) \quad \textit{for every prior and every finite } h.$$

Every first-action value and optimal first-action set agrees. A minimizing coarse-belief policy produces a fine-model optimal history-based policy, including its stopping decision. Finite minimizing trees exist. The result supplies neither an efficient synthesis algorithm nor a worst-case storage or resource guarantee.

Proof. At zero horizon the terminal risks agree by fibrewise summation. Suppose the equality holds at $h-1$. For any permitted first experiment, Theorem 3.6 gives the same probability of each report and the same coarse posterior on every positive branch. The induction hypothesis gives equal optimal continuation costs on each such branch. The first charge agrees, so the entire action value agrees. Stopping has the same value as well. Taking the minimum over the identical catalogues proves both value equality and equality of minimizing first choices.

Choose one minimizing first choice and attach minimizing lower-horizon trees to its finitely many positive branches. If the choice is terminal, no child is required. Induction constructs a finite tree attaining the value. The same induction bounds every other admissible tree below by that minimum. An observed history initializes the same coarse belief in the fine execution, so selecting from that belief realizes the coarse policy without observing the fine hidden state. □

Source relationship. The finite backup and projected-law proof mechanisms of [AVOC], Theorems 4.5 and 9.2, applied here to the belief transition supplied by Theorem 3.6. This is a prose composition, not a newly kernel-checked instance of AVOC. The established partially observed planning context is [KLC].

Counterexample 5.5 (A correct probability quotient can hide a forbidden action)**.** Two fine states share all report and transition probabilities and are given the same coarse label. A terminal repair is permitted in one state but forbidden in the other. The instrument quotient is valid, but a coarse interface that offers the repair in both states is not. Exact belief and value formulas do not repair a missing permission distinction. The catalogues must factor through the retained state, or the state must be refined.

5.4 The sensor chooses when another calibration is worth its cost

Return to the sensor after display $y = 0$. Conditional on that display, assign equal probability to $(b, s) = (0, 0)$ and $(1, 1)$, with any nuisance distribution inside each pair. Only noisy calibration and terminal bit selection are offered in this example. A calibration reports s correctly with probability $3/4$, with fresh conditionally

independent noise and no change to the hidden world. It costs $1/20$ in the same stated loss units as a wrong terminal choice, whose loss is one. A correct terminal choice has loss zero.

The retained belief on the four (b, s) pairs is an exact quotient of the eight-world belief. Its two diagonal masses initially equal $1/2$. After one calibration the favored target has probability $3/4$. A second matching report raises that probability to $9/10$, while conflicting reports restore $1/2$.

Report limit	Optimal expected cost	An optimal initial choice
0	$1/2$	Stop and select either bit.
1	$3/10$	Calibrate once, then select.
2	$3/10$	Calibrate once, then stop.
3	$11/40$	Calibrate, then use an adaptive continuation.

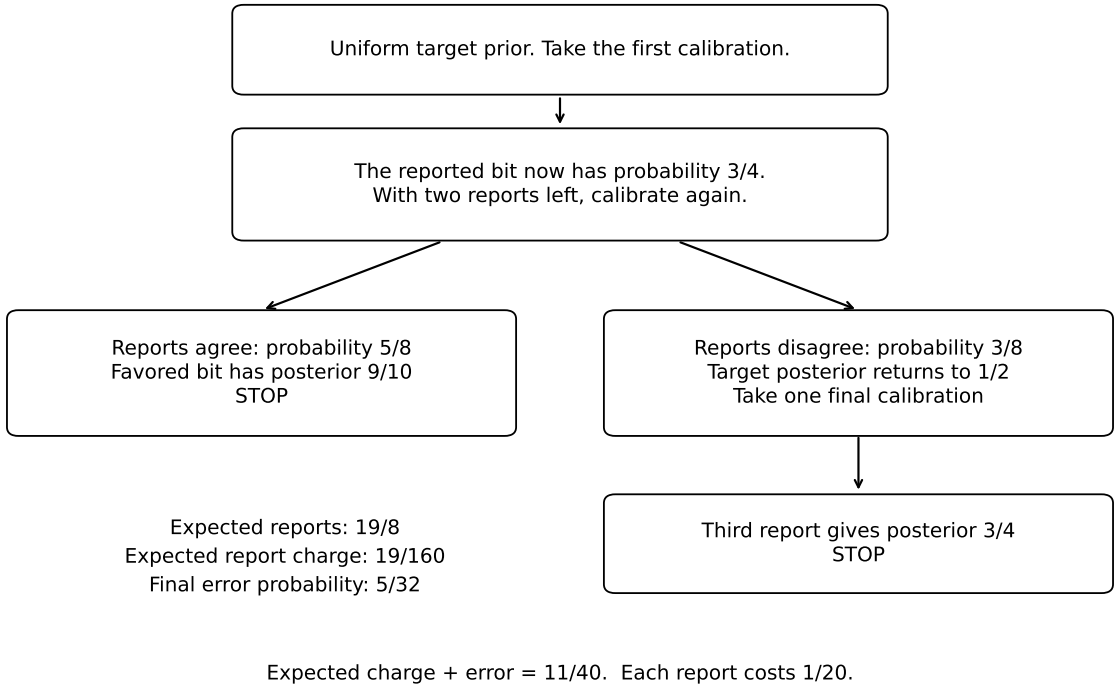


Figure 5.2: An optimal three-report policy in the unchanged eight-world sensor with the explicitly added noisy calibration channel. All masses and costs are exact. The diagram groups sign-symmetric report histories, not distinct hidden states supplied to the agent.

With three reports available, the optimal policy takes a second report after the first. If the first two agree, it stops. If they disagree, it takes the third. The disagreement probability is $3/8$. Thus the expected report count is $2 + 3/8 = 19/8$, the expected charge is $19/160$, and the error probability is $5/32$. The total is $19/160 + 5/32 = 11/40$.

A horizon of two does not improve on one in this model: after the first result, paying for one additional report does not lower expected classification error enough to recover its charge. A horizon of three creates the option

to investigate only the unresolved branch. The value comes from one contingent policy, not from separately optimal answers selected after the hidden bit is disclosed.

The two-valued predictive display state from Chapter 1 is too coarse for this inquiry. The four-pair probability record supports every calibration update and decision while still discarding the nuisance bit. This is the promised chain from evidence to representation to a useful experiment choice.

5.5 An approximate model supports a different, quantitative claim

Exact policy transport and approximate planning must not share one unqualified guarantee. Fix a finite horizon H and one finite class $\mathfrak{P}$ of admitted action/report trees, shared by the true and approximate models. The complete loss of a tree is the same function of its coarse-state and action/report path in both models, taking values in $[0, B]$. For example, a terminal loss in $[0, 1]$ and at most H charges bounded by $c_{\max}$ give $B = 1 + Hc_{\max}$.

Assume Theorem 3.8's bound d_H holds uniformly over every tree in that class. An early stop is padded by a common absorbing terminal step for the history comparison. State-dependent charges are permitted when included in the same coarse-path loss. A model-dependent posterior score used as an uncharged surrogate loss would be a different problem.

Proposition 5.6 (Uniform model error controls a changed policy)**.** *Let $J(\pi)$ and $\widetilde{J}(\pi)$ be a tree's expected losses in the true and approximate models. Then*

$$|J(\pi) - \widetilde{J}(\pi)| \le Bd_H \quad (\pi \in \mathfrak{P}).$$

If $\widetilde{\pi}$ minimizes the approximate loss and π_* minimizes the true loss over this same class, then

$$J(\widetilde{\pi}) - J(\pi_*) \le \min\{B, 2Bd_H\}.$$

These conclusions do not certify a zero-risk constraint or preservation of admissibility outside the common class.

Proof. The expectation of a function in $[0, B]$ changes by at most B times total variation. Apply the common-policy history bound separately to each fixed tree, proving the first assertion uniformly. Now

$$J(\widetilde{\pi}) \le \widetilde{J}(\widetilde{\pi}) + Bd_H \le \widetilde{J}(\pi_*) + Bd_H \le J(\pi_*) + 2Bd_H.$$

The range of loss gives the cap B. The proof compares each tree with itself across the models before using the optimality inequality. It never applies a same-policy bound directly to two different controllers. A bad event assigned probability zero in one model can have small positive probability in the other, so an average bound supplies no sure-feasibility conclusion. □

Source relationship. Direct combination of the newly admitted whole-history bound with the selection argument already used in Theorem 4.1; [AVOC], Sections 8–10, provide the source context for guarded planning and model-error value bounds. This finite fixed-policy-class formulation is stated and proved in full here.

With true calibration accuracy $37/50$ and planning accuracy $3/4$, the one-step defect is $1/100$. For $H = 3$ and $B = 23/20$, the uniform regret bound is $69/1000$. Exact observed performance is reported separately, not substituted for this guarantee. Expected-risk and hard-budget checks remain distinct. Next we study changes to the compiler producing the inquiry.

Chapter 6

Improving discovery without changing its meaning

Removing redundant work is an improvement only when the removed work could not change the justified result.

Let $\mathcal{G} = (p_0, \ldots, p_{N-1})$ be a finite ordered grammar of pure, total candidate programs. Let D be a finite context set and $t : D \to Y$ a fixed labeling. A teacher either accepts a candidate that agrees everywhere or returns a context where it disagrees. The teacher's work is part of the development account.

6.1 Finite counterexample-guided acquisition

The learner retains every returned counterexample. It proposes the earliest candidate consistent with all retained examples, or reports exhaustion of the grammar when none remains.

Proposition 6.1 (What finite counterexamples establish)**.** *The learner never submits the same rejected candidate twice. Within at most N teacher calls, it returns the earliest globally correct candidate in the declared grammar or exhausts that grammar. A timeout before this process finishes is neither acceptance nor exhaustion.*

Proof. Each counterexample refutes the submitted program and remains a constraint forever. That program cannot be submitted again. A globally correct candidate never violates a truthful counterexample. With only N candidates, either an accepted one is found or all are ruled out. If the accepted candidate were preceded by a globally correct candidate, that earlier candidate would still be consistent and would have been selected first. This proves the ordering claim. Interruption changes none of these logical implications and cannot supply a completed search result. □

Source relationship. [SII], Theorem 15.1. This is finite exact synthesis in a supplied language, not a statistical generalization guarantee.

Agreement on D is not validity outside D. That distinction is especially important when a candidate is used to prune future searches. A wrong proposal can be rejected by a final checker. A wrong pruning rule may remove the only candidate the checker ought to see.

6.2 A learned search representation can preserve the exact investigation

Suppose a deterministic normalizing map n is sound in the following direction:

$$n(p) = n(q) \implies p(z) = q(z) \quad \text{for every admitted context } z.$$

The map need not detect every semantic equivalence or produce a mathematically canonical normal form. Each successful contextual rewrite must be justified under the actual pure, total semantics. Keep the earliest original AST for every attained normal form, not a replacement AST with an uncharged new interpreter.

Theorem 6.2 (Ordered semantic compression preserves the teacher trace)**.** *Assume all counterexamples are retained, candidate and example orders are fixed, and the teacher is deterministic and extensional, returning the first disagreement in a fixed context order. Searching only the earliest representatives yields exactly the same submitted ASTs, counterexamples, and final acceptance or grammar exhaustion as the original learner. With the same AST interpreter it uses no more retained-example evaluations. Preparation, dispatch, and stored-index costs are separate.*

Proof. A removed expression has an earlier representative with the same values on every admitted context. If that representative was accepted, the original search ended before reaching the removed expression. Otherwise an already retained example refuted the representative, either during a consistency test or from the teacher. Persistence of examples means that the same example refutes the later equivalent expression. The original learner therefore never submits that removed expression to the teacher.

Deleting only those unsuccessful consistency tests preserves the next submitted AST. A deterministic teacher then returns the same counterexample or acceptance. Induction over submissions proves equality of the complete trace, including exhaustion. Every deleted evaluation has nonnegative cost in the named evaluation count. The argument makes no claim that constructing or consulting the representative index is free. □

Source relationship. [SII], Theorems 21.1 and 21.2. Sound equational replacement is the premise; empirical agreement on a few probe points does not supply it.

Counterexample 6.3 (A probe-equivalent equation is not a general identity)**.** The expressions $\max(x, 1)$ and x agree for probes $x \in \{1, 2, 3\}$. At zero they differ. Grouping them by those probe values alone can remove a needed candidate. The identity $\max(x, 0) = x$ is valid over nonnegative integers, but not over all integers. The semantic domain is part of the equation's certificate.

The retained object may therefore improve later search without containing its answers. This is the precise process-level role of a checked equation library. The grammar, primitive semantics, teacher, and proof rules can still be supplied. Their presence must remain visible instead of being redescribed as discoveries.

6.3 A reusable rule for stopping an unpromising read

A diagnostic compiler inspects a binary query on m immutable cases. In the prefix already read, a cases give zero and b give one. Each child is allowed at most C cases because the remaining binary depth has that leaf capacity. Assume integers

$$C \geq 1, \quad 2 \leq m \leq 2C, \quad a, b \geq 0, \quad a + b \leq m.$$

Proposition 6.4 (A sound prefix exclusion)**.** *The unread reports can be completed with both final branch counts at most C if and only if* $\max(a, b) \leq C$*. A scan may therefore stop when one count first exceeds C. In a fixed ordered depth-first compiler with identical scope, terminal, and fallback rules, replacing full root scans by this early rejection preserves the first returned tree or the completed depth-bounded failure result.*

Proof. Counts only increase, so an overfull child cannot be rescued. Otherwise the two remaining capacities are $C - a$ and $C - b$, both nonnegative. Their sum $2C - a - b$ is at least the number $m - a - b$ of unread rows. Distributing the unread cases between those capacities gives a feasible completion.

Every root rejected early would also be rejected after its full scan. A root not rejected early is read completely and produces the same ordered child supports. Induction on the finite remaining depth therefore yields identical recursive searches, backtracking, and first completion. The same failure condition invokes the same fallback. The skipped rows are only suffixes whose contents could not change rejection. A resource timeout is excluded from the completed-failure claim. □

Source relationship. [SII], Proposition 16.1 and Theorem 16.3. The conclusion concerns capacity feasibility, not the existence of a completing query in the installed language.

For $C = 4$, a fifth zero already closes a root. Reading more rows cannot reduce that branch. By contrast, a four-versus-four split merely survives this necessary test. If no remaining query distinguishes either child, the root still fails.

Counterexample 6.5 (A changed repair defeats the old premise)**.** Four cases previously required four different repairs. Admit one new repair successful in all four. No identifying query is now needed. The old singleton-leaf lower bound and its pruning interpretation are not applicable. The correct response is to recheck the repair scope, not to preserve an obsolete optimization because it once passed.

What is now available

The resulting method is semantically justified only under its premises. We must still determine whether it saves total resources, and which proof records remain usable when a later revision changes those premises. Both questions belong to the next chapter rather than being inferred from a successful synthesis.

Chapter 7

What remains valid after a revision

A mathematical proof can remain true while no longer applying to the installed system.

A certificate for ordinary integer addition is not a certificate for an operation with the same printed name and a different implementation. A grammar quotient depends on the grammar order. A comparison depends on the actual controller and the prices used to value its events. These are distinct dependencies.

7.1 Reuse follows semantic dependencies

Let an artifact state be $a = (a_1, \dots, a_d)$. A finite directed acyclic graph represents claimed guarantees. For vertex v, let D_v be its direct artifact dependencies and let $\mathrm{pa}(v)$ be its parent guarantees. The declaration is semantically complete if the meaning of its claim and local verification rule is determined by those components and parent claims. External assumptions are included explicitly as premises.

After changing a to a', mark a vertex stale when a direct input changes and propagate staleness to all descendants. Call the resulting set Δ.

Theorem 7.1 (Selective reuse with complete dependencies)**.** *Assume the graph is finite and acyclic, direct dependencies are semantically complete, local rules are sound, and the applicable external premises hold. Assume also that every old witness established its old claim before the revision. Old witnesses still establish every claim outside Δ. Re-establishing affected claims in topological order, using revised values and established parents, restores the corresponding revised claims. Without an affected parent's re-establishment, this reuse rule supplies no justification for its dependent claim.*

Proof. A vertex outside Δ has unchanged direct values and no stale parent. In topological order, induction shows that every parent retains its meaning and proof. Completeness of the dependency declaration leaves the local argument applicable. For a stale vertex, once its revised parents are established, soundness of the newly checked local rule gives the revised conclusion. A second topological induction proves restoration. An unchecked parent leaves a missing premise, so the rule alone cannot justify a dependent conclusion. This is not a proof that the conclusion is false. □

Source relationship. [SII], Theorem 24.1. It is a theorem about a declared complete dependency graph, not a method that infers complete dependencies from arbitrary programs.

Counterexample 7.2 (A missing dependency defeats a perfect replay)**.** A stored claim is $u = v$, initially with both values zero. Its registry records only u. Change v to one. A flawless comparison with the recorded u still passes, while the claim is false. A hash or replay cannot recover the premise the registry omitted.

Adding the universal repair from the preceding chapter changes a successful-repair set. It invalidates a cardinality certificate and any dependent performance statement. It need not invalidate an arithmetic equation used

elsewhere. A new valid repair also does not inherit the old controller's speed comparison. Rechecking can legitimately conclude that the new route is slower, cheaper, or incomparable under the stated resources.

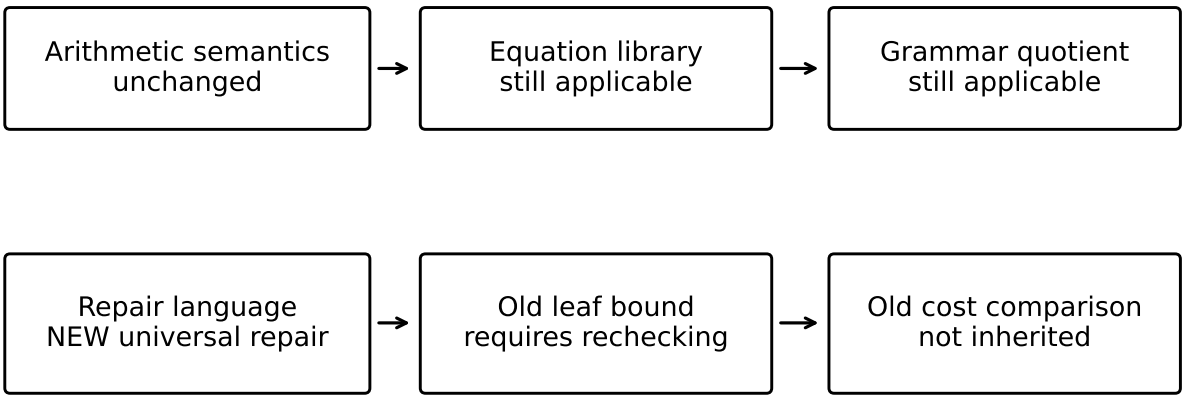


Figure 7.1: A change to the repair language affects the diagnostic justification but need not change arithmetic semantics. The depicted graph is a declared example, not an automatically discovered dependency registry.

7.2 Semantic improvement and resource improvement are separate

Fix named additive event types. Let $K \geq 0$ be an extra preparation vector, let s_i be the old-minus-new per-use vector for workload type i, and let n_i count its uses. The net event difference is

$$\Delta = \sum_i n_i s_i - K.$$

A nonnegative price vector w converts these events into one declared additive measure. Storage capacity, peak memory, and elapsed time need their own models unless a justified conversion is specified.

Proposition 7.3 (Which gains survive every admitted price?)**.** *The saving $w \cdot \Delta$ is nonnegative for every nonnegative price vector if and only if every coordinate of Δ is nonnegative. If prices lie in the convex hull of a finite nonempty set W, their minimum saving equals* $\min_{w \in W} w \cdot \Delta$.

Proof. A componentwise nonnegative vector has nonnegative inner product with every nonnegative vector. If coordinate j is negative, assigning price one only to that coordinate gives a negative result. This proves the first assertion. A point in the stated convex hull is a convex combination of its listed prices, so its saving is the same convex combination of their savings. It cannot be smaller than the minimum listed value, and a listed minimizer attains that value. □

Source relationship. [SII], Proposition 25.1. No hardware performance law follows from this additive accounting identity.

For the source's checked weighted-search comparison, the event difference is $(1552, 1552, -1868)$ in reads, rational additions, and comparisons. Equal prices give 1236 saved units. Prices $(1, 0, 1)$ give a loss of 316. Both conclusions concern the same valid controller. The main text uses this as a transparently rechecked source calculation, not as a new measured timing result.

7.3 Evidence across an adaptive sequence

A procedure that searches many candidates can select a favorable error in its evaluator. The right statement must therefore condition on the information that produced the candidate, not merely repeat a marginal error bound for a candidate fixed in advance.

Proposition 7.4 (A conditional error budget for a lineage)**.** *Let $\mathcal{H}_{j-1}$ contain the information available before attempt j. Suppose the event A_j of accepting a false claim satisfies*

$$\Pr(A_j \mid \mathcal{H}_{j-1}) \le \delta_j \quad \textit{almost surely}, \qquad \sum_{j \ge 1} \delta_j \le \delta,$$

for deterministic nonnegative δ_j. Then the probability of ever accepting any false claim is at most δ, even when candidates and the stopping time depend on earlier history.

Proof. Taking expectations gives $\Pr(A_j) \le \delta_j$. A union bound gives the result for every finite set of attempts. The finite unions increase to the event of any false acceptance, so continuity of probability gives the infinite statement. Stopping removes attempted events and cannot enlarge their union. □

Source relationship. [SII], Proposition 26.1. The mathematical input is conditional validity of each actual attempt. A numerical error schedule alone does not establish that premise.

This proposition is not an empirical result about a population of agents. It explains what evidence would need to remain valid when an evaluator is used adaptively. Deterministic exact checks in this book are not treated as statistical samples and are not assigned confidence intervals through this formula.

What is now available

A revision can preserve its meaning, identify stale obligations, and state its resource advantage under explicit prices. None of this yet determines what several participants can know together or which authority they should exercise. The final technical chapter keeps information sharing and institutional permissions separate while applying the same preservation discipline.

Chapter 8

Shared work and maintained obligations

Joint information can support a task without creating a single decision maker.

One agent may have a display record and another a calibration record. Neither need know the physical target bit alone. Together their records can determine it. This is an informational relation before it is any claim about collective agency.

8.1 What several accessible records can reconstruct

Let V be a finite-dimensional vector space over a field. Participant i has a linear record $A_i v$. Stack the records into $Av = (A_1 v, \dots, A_m v)$, and let Qv be the protected linear target.

Proposition 8.1 (Target-relative joint reconstruction)**.** *There is a linear decoder D with $Q = DA$ if and only if* $\ker A \subseteq \ker Q$. *Full-state reconstruction is possible exactly when A is injective. The intersection identity* $\ker A = \bigcap_i \ker A_i$ *determines which ambiguities sharing removes.*

Proof. If $Q = DA$, a vector annihilated by A is also annihilated by Q. Conversely, define $D(Av) = Qv$ on $\operatorname{im} A$. If $Av = Aw$, then $v - w \in \ker A \subseteq \ker Q$, so this definition is independent of the representative. It is linear and extends to the full record space by extending a basis. For $Q = I_V$ the kernel condition is injectivity. A stacked record vanishes exactly when each component record vanishes, giving the intersection identity. □

Source relationship. Target-relative specialization of the reconstruction and factorization machinery in Sufficient Resolution [SR] and Stratified Relational Workspace Theory [SRWT]. The proof here does not import their complete noise, dynamics, or identity theories.

Over $\mathbb{F}_2$, Agent A observes $A_A(b, s) = b + s$ and Agent B observes $A_B(b, s) = s$. The joint matrix is invertible. Either agent alone lacks a direction needed to recover b. The nuisance bit may remain unknown because the target map ignores it.

The theorem concerns accessible records. It does not send them. A communication protocol, truthful interpretation, permitted disclosure, and the recipient's executable decoder remain separate requirements. Chapter 4 states the corresponding fixed-recipient criterion.

8.2 Permissions are part of an operational state

Now let a finite institutional state x have a protected label $\ell(x)$, an enabled-action set $D(x)$ within a fixed finite alphabet, and a deterministic successor $F(x, a)$ for enabled actions. Labels may include resource balances, outstanding obligations, and offices, but need not retain every historical byte.

Theorem 8.2 (A quotient that preserves legal operations)**.** *For a surjection $\pi : X \to Z$, quotient labels, enabled sets, and successors satisfying*

$$\ell = \bar{\ell}\pi, \qquad D = \bar{D}\pi, \qquad \pi F(x, a) = \bar{F}(\pi x, a)$$

exist exactly when equal fibres have equal labels, equal enabled sets, and equal projected successors for every enabled action. The quotient data are unique and preserve every legal finite action and label sequence.

Proof. Factoring each quantity proves necessity. For sufficiency choose a representative of each fibre to define its label, enabled set, and successor. The three equalities make those definitions independent of the representative. Surjectivity gives uniqueness. Equal enabled sets admit the same first action, and successor equality preserves membership in one fibre. Induction on the sequence length preserves legality and all labels. □

Source relationship. [MS], "Exact governance quotient." This extends the deterministic behavioral construction by explicitly protecting action availability. It does not derive moral or external legal authority from the state description.

A summary that preserves only the next sensor response can merge states in which a repair is authorized with states in which it is forbidden. It may remain a good predictor while becoming an invalid action interface. The necessary correction is to preserve the enabled-set distinction, not merely to attach a reassuring name to the old code.

8.3 Continuing should not erase the ability to finish

Use a finite *observed* configuration space Z. Under partial observation, the configurations would instead have to be justified public possibility states. At z there are finite catalogues of running actions and ending actions, each with a nonempty successor set. A declared safe set is S, and an ending target is $G \subseteq S$. A safe hold action is available at G.

A stop command is assumed authentic and observable before the next action is selected. Its delivery and authentication are premises, not consequences of the following fixed-point calculation. Define

$$\begin{aligned} T_0 &= G, \\ T_{k+1} &= T_k \cup \{z \in S : \exists a \in A_{\text{end}}(z),\ \text{Post}(z, a) \subseteq T_k\}. \end{aligned}$$

Let T_* be the stable set. Then define

$$L_0 = T_*, \qquad L_{k+1} = \{z \in L_k : \exists a \in A_{\text{run}}(z),\ \text{Post}(z, a) \subseteq L_k\},$$

and let L_* be its stable value. Both sequences stabilize because Z is finite.

Theorem 8.3 (Safe continuation that retains an ending strategy)**.** *A starting configuration belongs to L_* exactly when one policy can run safely indefinitely without a stop command and, after a stop command at any decision epoch, reach G on every possible successor path using only the admitted ending actions. Before the command a policy preserving L_* suffices. After the command it can reach G within the first-entry rank of the current state in the T_k sequence.*

Proof. Induction shows that membership in T_k is equivalent to an ending-only policy reaching G within k transitions while staying safe. The chosen action moves every successor to a smaller rank. Outside T_*, every

ending action has a successor outside T_* or outside safety. Choosing such a successor at every step defeats any purported surely safe ending. Thus T_* is exactly the sure-ending region.

At L_* a running action keeps every successor inside L_*. Repeatedly using such actions is safe and retains membership in T_*. A command can therefore switch to the decreasing-rank ending policy at any epoch.

Conversely, if a policy has the stated guarantee, every pre-command reachable configuration must admit an ending immediately and hence lie in T_*. For any finite number of further running steps, it must keep all paths in that set. Induction in the decreasing construction places the initial state in every L_k, hence in L_*. No action is chosen after secretly observing its successor. □

Source relationship. Fully observed finite specialization of [MS], "Continued operation with an available authorized ending." The more general public-possibility construction is not silently assumed in this version.

Counterexample 8.4 (Safe inactivity without an ending)**.** A state z has a safe idle self-loop and no ending action reaching G. It is safe forever but outside T_*. In a second state u, an ending is available, while the only running action spends its reserved resource and moves to z. Then u belongs to T_* but not L_*. Being able to stop now does not establish that ordinary operation preserves that ability.

These results do not establish society formation, strategic compliance, or a just constitution. They establish information, action, and continuation properties of a specified finite institution. Their place at the end of the book is earned by the preceding account of accessible evidence, common artifacts, feasible controllers, and dependency-preserving revision.

Chapter 9

One continuous chain, with no hidden promotion

The sensor story can now be followed without changing the meaning of a record midway through the argument.

Begin with eight candidate worlds (b, s, n). The agent's initial display y determines its two-state predictive class under flips, but not the physical target bit. The first chapter proves both claims. The second gives the operator on the retained class and the coordinate map that becomes usable after calibration. Neither requires the nuisance bit.

When display and calibration observations arrive, candidate sets are filtered. The third chapter proves why the joint association must be retained, and why sixteen occupancy states suffice for the exact partition-evidence contract. Under a separately declared noisy calibration, four probability masses replace the occupancy bits. The same numerical posterior cannot be inferred from separate marginals. The enlarged chapter now requires a joint report/next-state law, not only correct separate marginals. Its finite refinement makes the required distinctions explicit, and its history bound states the separate price of conditioning on a rare report.

The fourth chapter places that posterior and any response approximation under one loss function. It also specifies what a recipient can decode from a message and its prior record. The fifth asks whether the sender can acquire the missing observation and transmit a successful artifact within a single online and stored-controller budget. The complete tree is counted even if an execution visits only one path. A separate noisy-calibration calculation then uses the exact same retained belief to select a finite expected-cost inquiry. Uniform whole-history error supports a quantitative policy comparison, but not an automatic hard-budget or zero-risk guarantee.

The sixth chapter then changes the discovery procedure rather than the hidden world. Checked equations can remove equivalent candidates while preserving every teacher proposal and counterexample. A separately justified capacity rule can stop reading an impossible diagnostic root. Neither step invents the primitive grammar, the teacher, or the theorem used to authorize pruning.

The seventh chapter binds those uses to their premises. Adding a universal repair changes the required diagnostic distinctions. The old lower bound becomes inapplicable; arithmetic identities need not. A fresh cost comparison may find a saving or a loss. The eighth chapter finally combines different participants' records and asks what actions they are entitled to perform under their declared institutional state. Safe indefinite operation and preservation of an authorized ending are checked separately.

9.1 A worked end-to-end record

Take the actual illustrative world $(b, s, n) = (0, 1, 1)$. The complete protocol begins with all eight worlds possible:

Event	Remaining worlds	What is justified
Before observation	All eight	No physical target identified.
Display $y = 1$	Four worlds	Predict every later display flip.
Calibration $s = 1$	Two worlds differing only in n	Physical target $b = 0$.
Send y to a recipient holding s	Same two worlds	Recipient computes $b = 0$ through its fixed XOR decoder.
Commit the authorized target action	As declared by its effect	One action justified across the current support.
Record the fulfilled obligation	Includes the action's receipt	Continuation or ending can refer to the actual completed act.

The witness does not claim to have discovered its sensor semantics. Its exact checking confirms the stated scenario. Another physical process requires its own evidence. A failed receipt, an untrusted calibration, or an altered action effect changes the guarantee rather than being averaged away as a small numerical error.

9.2 Why the core is narrower than the archive

This book does not append the old Neuralese experimental history, proposed PRISM architecture, full SRIC lineage, or every formal guide. Their presence would make the reader repeatedly switch between a proved conditional statement, a proposed mechanism, a historical execution, and an implementation plan. Those are useful source objects, but they are not one continuous proof.

Selected source results are restated and proved where they carry the argument. Selected stochastic-value arguments from AVOC now enter through complete prose proofs and explicitly matched finite models. Its complete formal development, and the specialist Certified Bayesian Inference construction, remain candidates for a later formal source volume. No pinned kernel replay was performed for these additions, so their source pedigree is not used as a blanket assurance label.

Nor is a source excluded merely because it contains open questions. Its proved results can be considered individually. A claim enters when its exact form, dependencies, and evidence have been reviewed at the level promised for this edition. This is a stricter and more useful rule than treating every page of a final-named document as equally checked.

9.3 The result of the compilation

The core establishes a connected conditional account of accessible state, meaningful operations, corrective memory, usable explanations, bounded investigations, justified search changes, and maintained obligations. The arguments share several elementary constructions, but each use preserves its distinct type and hypotheses.

No theorem here turns a semantic interface into a trained neural model, a finite investigation into general intelligence, a shared record into a collective subject, or an invariant into legitimate external authority. The mathematical contribution is the explicit chain and the failures that break it. Those failures can be investigated without changing the standard of proof or requiring a large implementation campaign.

Result and source map

The following map gives the role of each admitted formal result. Exact source anchors, excerpts, hashes, and detailed hypothesis review are in the accompanying machine-readable ledger and review files. Source relationships do not imply that every source page is admitted.

Result	Preserved object	Source relationship
1.1	A target determined by an available record	GA; shared fibre criterion.
1.2	Every declared finite continuation	GA and MU, consolidated.
2.1	Exact split encoding and named composition loss	SR, split and hidden-return identities.
2.3	Source operations and matched readouts	MU; OSL and Neuralese use the same conditional principle.
2.4	One output method across candidate dynamics	GA, common observable quotient.
2.5	Probability interpretation of a complete packet family	LSM, positive packet criterion.
3.1	Every admitted future set correction	GA, correction-signature theorem.
3.2	Exact partition-evidence memory cardinality	GA, occupancy corollary.
3.4	All posterior updates in the actual reporting chain	MU, finite channel factorization.
3.5	Coarsest partition closed for prediction and update	LSM and PRIQ, finite closure.
3.6	Joint next-state/report law, posteriors and feedback	PRIQ and AVOC, selected composition.
3.8	Whole-history error and conditioning cost	LSM, law and normalization bounds.
4.1	One bounded-loss decision calculation	GA and MU, specialized common-risk bound.
4.2	One artifact under fixed recipient procedures	MU, intersection and covering criterion.
5.1	One executable tree under all declared budgets	MU and SII, finite Pareto synthesis.
5.3	A complete optimality claim for a checked tree	SII, binary singleton-repair scope.
5.4	One optimal finite-horizon expected-cost inquiry	AVOC, projected values and finite plans.
5.6	Uniform policy evaluation before selection	LSM/AVOC, common-law specialization.
6.1	Exact finite candidate rejection and acceptance	SII, finite counterexample refinement.
6.2	The same proposed ASTs and teacher transcript	SII, sound equations and ordered quotient.
6.4	The same completed ordered diagnostic search	SII, capacity and streaming arguments.
7.1	Applicability of dependency-bound witnesses	SII, complete declared dependency graph.
7.3	Gains under one complete additive ledger	SII, price-relative dominance.
7.4	Conditional false-acceptance control	SII, adaptive lineage error budget.
8.1	A target from jointly accessible linear records	SR and SRWT, target-relative specialization.
8.2	Legal action availability and label histories	MS, exact governance quotient.
8.3	Continued operation with an ending policy	MS, fully observed finite specialization.

Reading sources and review boundaries

GA

Komorebi P. *Grounded Abstraction: Identifiability, Representation Change, and the Limits of Learning*. Version 1.0.0, 9 October 2026. Exact source archive identified in core 0.1, whose source-suite replay receipts are retained historically.

MU

Komorebi P. *Machine Understanding: Representation, Explanation, and Transfer*. Version 1.0.0, 8 October 2026. Exact source archive identified in core 0.1, whose source-suite replay receipts are retained historically.

SII

Komorebi P. *Self-Improving Intelligence: Representation, Discovery, and Recursive Revision*. Version 1.0.0, 8 October 2026. Exact source archive identified in core 0.1, whose source-suite replay receipts are retained historically.

MS

Komorebi P. *Machine Societies: A Theory of Persistent Plural Computational Organization*. Version 1.0.0, 8 October 2026. Exact source archive identified in core 0.1, whose source-suite replay receipts are retained historically.

SR

Komorebi P. *Sufficient Resolution*. Standardized traditional paper, 8 October 2026. Selected split, hidden-return, and factorization arguments reviewed from the source text and reproved here. No new Lean replay is claimed.

SRWT

Komorebi P. *Stratified Relational Workspace Theory*. Collected source manuscript as identified in the original compendium catalogue. Only the elementary reconstruction relationship used here is admitted. Full dynamics, emergence, persistence, and identity claims are not re-audited by this selection.

LSM

Large Signal Model Foundations, edition 0.4.0. Selected packet-positivity, joint finite closure, and whole-law/conditioning arguments: Propositions 9, 12, 30, 43 and 44. Exact source PDF and selected-page hashes are in the admission record. This book does not import the full observable-learning, cone-discovery, or software claims.

PRIQ

Hypergraph Lumpability and Bayesian Refinement, version 0.1.4. Selected positive normalization and finite filtering composition: Theorems 5.2, 5.5–5.6, 6.1 and 6.3, plus the posterior-error context in Section 9. The nonlinear analytic and hyperedge results are outside this admission.

AVOC

AVOC: Artifact Value, Guarded Operations, and Calibrated Planning. Traditional publication edition based on formal v0.41.0. Selected conditioning, projected stochastic value, and finite planning proof mechanisms: Theorems 3.6, 4.5 and 9.2. Their use here is a reviewed prose specialization. No new Lean compilation or theorem-level refinement of the complete host implementation is claimed.

KLC

Classical partially observed planning. Leslie Pack Kaelbling, Michael L. Littman and Anthony R. Cassandra. *Planning and acting in partially observable stochastic domains*. Artificial Intelligence 101 (1998), 99–134.

Author-hosted manuscript and publisher record checked for the belief-state and finite-policy context. This reference locates the established planning framework, not a novelty claim for the recursions here.

These sources identify the project's provenance. They are not a comprehensive literature review or a priority claim for factorization, automata quotients, sufficient statistics, dynamic programming, program synthesis, graph induction, or controlled invariant sets. Those mathematical traditions precede this exposition. The present book makes its selected arguments self-contained rather than relying on unstated external theorems.

What was newly checked

Every numbered theorem, proposition, and corollary in this core received a statement-level review of its domains, quantifiers, proof, and subsequent use. The separate selected-core program checks explicit finite instances and negative controls through alternative elementary calculations where practical. The complete selected-core checks and six new admission groups were run in this revision. The four source-suite regressions were performed for core 0.1 and are retained as historical receipts, not counted as new source-wide replays.

No source-wide proof review follows from those regressions. No earlier test count is called a new mathematical theorem. The rewritten text has no proof-kernel certificate. Same-environment rebuilds establish reproducibility of the document and the recorded computations, not independence of the author or universal absence of error.