

The MIND: Five Finite Constructions

First research edition, with P1 and P2 mathematical addenda

Research prepared for Brian Theory
AI-assisted mathematical synthesis; not external peer review

5 October 2026

Abstract

We develop five bounded constructions concerning observation, resolution, factorization, distributed consistency, and actuator calibration. Their common question is which distinctions suffice for a specified task. Exact derivations and finite computations establish properties of the models. They do not establish a marker of phenomenal consciousness. The constructions deliberately retain counterexamples to attractive but invalid bridges. The accompanying source register records reading depth, and the local baseline records manuscript versions and hypotheses. No claim of mathematical priority is made. The P1 addendum gives an exact predictive-acquisition criterion for the bounded RF poset class, with minimal and transcript-level obstructions. The P2 addendum characterizes the exact mutual-information interval for binary supports with fixed marginals and a positive cell floor, while preserving causal and recoding distinctions.

Contents

1	Contracts before interpretation	2
2	M1: distinctions, quotients and predictive closure	3
2.1	Model specification	3
2.2	Minimal dynamic completion	3
2.3	Symmetrization is not identification	4
2.4	Test and abandonment rule	4
3	M2: equal group index, unequal information gain	4
3.1	Model specification	4
3.2	A decision-sensitive extension	5
4	M3: three obstructions to identifying structure with causal information	5
4.1	Model specification	5
4.2	Support does not determine dependence	5
4.3	Effective information and intervention redistribution	6
4.4	Information interactions depend on source coordinates	6
4.5	Limits and programme	6
5	M4: local consistency without a global assignment	6
5.1	Model specification	6
5.2	Noise threshold	7
5.3	Workspace interpretation and its failure test	7

6	M5: calibration, memory and the limits of a self-model interpretation	7
6.1	States, observations and control	7
6.2	Exact performance and ablations	8
6.3	What is self-related, and what is not?	8
7	Additional negative results and review targets	9
7.1	A directed-order distance need not be a metric	9
7.2	Plasticity is not automatically observable change	9
8	P1: predictive refinement under a repair-safe library	9
8.1	Fixed-state acquisition contract	9
8.2	Compatibility with deterministic prediction	10
8.3	P1-C1: a minimal supplied-but-unsafe pullback	11
8.4	P1-C2: exact transcripts need not have a least refinement	11
8.5	Checked scope	11
9	P2: attainable information over a binary logical support	12
9.1	Declared constraints and prior construction	12
9.2	Sharp continuum bounds	12
9.3	Finite grids and a quantitative rounding bound	13
9.4	An identical observational law with different interventions	14
9.5	Admissible recoding and verification limits	14
10	Open theorem and experiment targets	15

1 Contracts before interpretation

The local baseline identifiers B01–B09 refer to sections of `00-context/baseline.md`; canonical versions are in `00-context/local-versions.json`. Source and idea identifiers are stable across the library. M1–M5 are constructions, not additional idea counts. The executable companion is `experiments/finite_models.py`; it uses only the Python standard library. Its outputs are in `experiments/results.json`. There is no Monte Carlo seed because all examples use finite enumeration or finite sums. Decimal entropies use floating point; counting and calibration probabilities use exact integers or rational numbers where applicable.

Type	Carrier and composition	Required additional semantics
CM / LM	2×2 truth tables or formula entries; native XOR–AND product	Operand order, polarity and matched frames; no probability law is implicit.
Support relation	Subsets or 0–1 matrices with OR–AND relational composition	This is a different semiring from native CM composition.
Stochastic channel	Nonnegative matrices with row sums one; ordinary real product	Distributions, observation protocol and interventions must be supplied.
Fuzzy relation	Values in $[0, 1]$, selected t-norm and aggregation	Membership is not automatically probability; thresholding need not preserve composition.

Modal chain ring	Nonzero vectors and explicitly allowed effects over a specified ring	Possible/impossible outcomes, primitive effects and zero divisors matter; no Born rule follows.
Hilbert model	Complex vectors or density operators and quantum channels	Physical realization, measurement and decoherence are additional commitments.

An *interpretation* associates a formal variable with a measured capacity. An exact map within mathematics is not that interpretation. We distinguish computational tasks, representational mechanisms, physical implementation, and claims about experience. All five models operate at the first two levels.

2 M1: distinctions, quotients and predictive closure

2.1 Model specification

Let W be a finite latent carrier and $\Phi = (\phi_i)_{i \in I}$ a family of Boolean tests. The observation is $q(x) = (\phi_i(x))_{i \in I}$, with realized carrier $Q = \text{im } q$. An available action may acquire one new test ψ at a declared cost; latent deterministic dynamics is $T : W \rightarrow W$. One transition is one time unit; a test-acquisition step is a separate operation. The boundary is the analyst's chosen W , and the agent only receives requested test outcomes. The present model is noiseless; noisy observations require a channel and do not inherit exact fiber equality as an empirical fact.

The positive topology is generated by $\phi_i^{-1}(1)$. On realized signatures, the upward-open specialization order is coordinatewise comparison. If both positive and negative truth regions are included, the result is the partition topology of signature fibers. Thus observational equality and positive order retain different data (B02).

For $W = \{00, 01, 10, 11\}$, observing the first coordinate gives three positive opens. Observing both coordinates gives six positive opens, whereas observing both signs of both coordinates gives sixteen. The complete signature map is injective in the second case, although $\{00\}$ is not positively open. This example prohibits reading positive topology as a complete yes/no partition.

Proposition 1 (Exact deterministic descent). *For a surjection $q : W \rightarrow Q$ there exists a unique $\bar{T} : Q \rightarrow Q$ satisfying $qT = \bar{T}q$ if and only if*

$$q(x) = q(y) \implies q(Tx) = q(Ty).$$

Proof. Necessity follows by applying the commuting equation to x, y . For sufficiency define $\bar{T}(q(x)) = q(Tx)$. The hypothesis makes this independent of the representative; surjectivity gives existence on every element and uniqueness. $\square$

For $q(a, b) = a$ and $T(a, b) = (b, a)$, states 00, 01 are currently indistinguishable but their successors have different first bits. The coarse state is therefore not a sufficient predictive state. Enumeration of all $4^4 = 256$ transitions finds exactly 64 that descend to this two-cell partition; this verifies only the count for that carrier. The proof supplies the general claim.

2.2 Minimal dynamic completion

For a finite set of actions with transition maps T_a , successively refine the current equivalence relation using the signatures

$$\sigma_{k+1}(x) = (q(x), (\sigma_k(T_a x))_a), \quad \sigma_0 = q.$$

Interpreting signatures by their equality classes, this process stabilizes after at most $|W| - |Q|$ strictly splitting rounds. At stabilization each action descends. Induction shows that two states remain equivalent exactly when they give identical q observations along every finite action word. This is a standard automata/bisimulation construction, proposed as a bridge to B03, not a new consciousness theorem. If the supplied observation library cannot realize the required splits, the abstract refinement is not implementable by the agent. An RF-IDENT-safe implementation additionally requires the precise order and query contract in B06.

2.3 Symmetrization is not identification

Let $R = \{(0, 1), (1, 2)\}$. Adding converse pairs makes R symmetric but not transitive: $(0, 2)$ is absent. Reflexive-transitive closure then identifies all three points in one equivalence class. Symmetrization preserves undirected adjacency; equivalence closure forgets graph distance and all within-component distinctions. This is a candidate reconstruction of one operation used in bi-logic, not a derivation of Matte-Blanco’s psychology (see logic dossier). The separate stages must remain visible.

2.4 Test and abandonment rule

Compare dynamic completion against the original partition, the fully resolved state and a finite-history predictor. Measure predictive error and observation cost under held-out action sequences. A new B-Theory contribution requires a useful constraint or invariant beyond standard partition refinement. Abandon a claimed exact quotient when descent fails; retain it only as an explicitly approximate or history-dependent model.

3 M2: equal group index, unequal information gain

3.1 Model specification

The four states are leaves of a rooted binary tree at depth two. Distances are $d(x, x) = 0$, $1/2$ for distinct siblings and 1 between root subtrees. The old observation is the first bit. Either action resolves the left or right old cell into its two leaves. States are static during acquisition, the boundary is the four-leaf tree, and one action costs one declared probe unit. A probability measure μ describes sampled states; it is extra structure beyond the tree.

The full tree automorphism group has order eight. Its pointwise stabilizer of the two old cells has order four; resolving either cell leaves a stabilizer of order two. Both refinements therefore have index two and log index one. These are stabilizers inside the full combinatorial tree group. If one instead restricts to measure-preserving automorphisms for nonuniform μ , the group changes and the following group counts cannot simply be reused.

Proposition 2 (Local entropy increment). *For a partition refined by splitting a cell C into $C_1, \dots, C_r$,*

$$H_\mu(\mathcal{P}') - H_\mu(\mathcal{P}) = \mu(C) H \left(\frac{\mu(C_1)}{\mu(C)}, \dots, \frac{\mu(C_r)}{\mu(C)} \right),$$

with zero contribution from a null cell.

Proof. Expand $-\sum_A \mu(A) \log_2 \mu(A)$ in both partitions; unchanged cells cancel. Substitute $\mu(C_i) = \mu(C)p_i$ in the remaining sum. $\square$

For leaf masses $(.7, .1, .1, .1)$ the left increment is $.8h_2(.875) = .4348515546$ bits; the right increment is $.2$ bits. For the uniform measure both increments are $.5$ bits, still distinct from the one-bit log index. B05’s group and entropy quantities answer different questions.

3.2 A decision-sensitive extension

For target variable Y and loss L , choose an acquisition action to reduce Bayes risk, not automatically to maximize $H(X)$. If only the right-cell distinction changes the optimal response to Y , the smaller entropy increment can have greater task value. The refinement preserves old outcomes; changing the cost or the target can reverse its desirability. This is compatible with Blackwell comparison of experiments, which assumes a specified decision problem and does not erase acquisition costs.

Exact labels are a noise-free idealization. Empirically estimate a channel $P(z | x, a)$, compare held-out risk to a constant or random probe, and check sensitivity to tree encoding, nonuniform sampling and biased missing data. Abandon the symmetry-index-as-information proposal; retain separate symmetry, entropy and utility coordinates. A resolution hierarchy alone supplies no latent time, unconscious process, or spontaneous symmetry-breaking mechanism.

4 M3: three obstructions to identifying structure with causal information

4.1 Model specification

Here states are binary variable tuples or four microstates. Observations are support, joint frequencies, or interventions, explicitly distinguished below. Actions either change an intervention distribution or select a coarse-graining. Channels describe one-step transitions. There is no learning update. The boundary and variable factorization are analyst choices, and all distributions are population-level specifications, not estimated data.

4.2 Support does not determine dependence

Both distributions

$$P_0 = \begin{pmatrix} 1/4 & 1/4 \\ 1/4 & 1/4 \end{pmatrix}, \quad P_1 = \begin{pmatrix} .4 & .1 \\ .1 & .4 \end{pmatrix}$$

have the all-one Boolean support, which factors as a Cartesian product and has OR–AND Boolean rank one. Yet $I_{P_0}(X; Y) = 0$ and $I_{P_1}(X; Y) = 1 - h_2(.2) = .2780719051$ bits. Consequently no support-only function can recover mutual information for all compatible positive distributions. In particular a Boolean factorization defect is not a probability dependence measure. B07 may expose useful support constraints, but a stochastic lift must be separately specified.

Even a complete observational law is insufficient for causal direction. With $U \sim \text{Bernoulli}(1/2)$, models $X = U, Y = X$ and $X = U, Y = U$ have the same distribution $X = Y$. Under $\text{do}(X = 1)$ they give respectively $\mathbb{P}(Y = 1) = 1$ and $1/2$. The distinction is between structural equations and conditioning, not between two matrix notations.

4.3 Effective information and intervention redistribution

Use the four-state channel

$$K = \begin{pmatrix} 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 1 \\ 1/3 & 1/3 & 1/3 & 0 \end{pmatrix}.$$

The macro partition $\{0, 1, 2\} \mid \{3\}$ yields deterministic alternation and is strongly lumpable. Under a uniform intervention over microstates, $I(X; X') = h_2(1/4) = .8112781245$ bits. Under a *new uniform macro* intervention, macro mutual information is one bit. Under the macro input distribution induced by uniform micro interventions, it is again .8112781245 bits. The apparent gain compares different intervention ensembles; it violates no data-processing inequality. This finite example illustrates a known causal-emergence distinction [3]. It is not a calculation of IIT's Φ .

4.4 Information interactions depend on source coordinates

Let A, B be independent fair bits and $Y = A \oplus B$. Then $I(Y; A) = I(Y; B) = 0$ but $I(Y; A, B) = 1$. Under the invertible recoding $(C, D) = (A \oplus B, B)$, $I(Y; C) = 1$. A variable-wise synergy assertion therefore requires a justified factorization into sources; global invertibility does not preserve it. This does not select a unique partial-information decomposition. Physical locality or independently intervenable variables could justify a factorization, but a convenient Boolean encoding alone cannot.

4.5 Limits and programme

A support-to-stochastic lift can be studied as a family of all compatible distributions, reporting a range of information measures instead of a single score. With small samples, an unobserved event is not a proven structural zero. Matched tests should vary sampling priors, causal models, coarse-graining and invertible recoding. The theorem target is a characterization of what is identifiable under the stated data; the abandonment rule rejects causal or phenomenal claims when equivalent observational models disagree.

5 M4: local consistency without a global assignment

5.1 Model specification

Variables A, B, C take values in $\mathbb{F}_2$. The three observation contexts are AB, BC, CA . At one time slice each context reports a relation

$$A + B = b_1, \quad B + C = b_2, \quad C + A = b_3.$$

The event sheaf contains all assignments, with ordinary coordinate restriction. The allowed-support subpresheaf imposes these relations. Each pair's support projects surjectively onto the shared single variable. Equipping each pair with the uniform distribution on its two allowed assignments also gives matching overlap marginals. Actions may flip a constraint or transmit a selected assignment; these are different interventions. No neural timing or workspace access mechanism is specified by this static consistency test.

Proposition 3 (Parity obstruction). *A global assignment exists if and only if $b_1 + b_2 + b_3 = 0$ in $\mathbb{F}_2$. When it exists there are exactly two assignments.*

Proof. Summing all equations cancels each variable twice, proving necessity. Given A , set $B = A + b_1$ and $C = B + b_2$; the third equation holds exactly when the parity condition holds. The two choices of A give the two solutions. $\square$

For $(b_1, b_2, b_3) = (1, 1, 1)$ no global section exists, although pairwise projected supports and marginals agree. This does not contradict sheaf gluing: we have *not* chosen mutually compatible individual sections on the contexts. The full event sheaf glues such sections; the issue is extending these allowed sets or marginal distributions [4].

The cover nerve is a triangle with no filled face. It has $\beta_0 = \beta_1 = 1$ for both even and odd parity. Its ordinary homology cannot distinguish the two cases. One needs the restriction/constraint data, not just nerve shape. Order-complex homology in B03–B04 is another construction again. A bridge must provide a functor or an explicit comparison, not rename the invariant.

5.2 Noise threshold

Suppose each edge is anti-correlated with probability $1 - \eta$, has equal single-variable marginals, and the two satisfying and two violating outcomes are each equiprobable within their class. Every deterministic assignment has at least one equal edge, so a joint distribution requires $3\eta \geq 1$. Conversely at $\eta = 1/3$, the uniform distribution on the six nonconstant assignments has these marginals. At $\eta = 1$, uniform 000, 111 does too. Their convex mixture with weight $(3\eta - 1)/2$ on the latter realizes every $\eta \in [1/3, 1]$. Thus the threshold is exact for this symmetric family. Finite samples require confidence bounds on all three edge errors before claiming violation; measured context changes may instead break the common variable identification. This toy is not evidence of physical quantum contextuality in a brain.

5.3 Workspace interpretation and its failure test

A proposed workspace may broadcast an assignment and ask modules to accept or reject it. Added communication can reveal disagreement; it cannot make the odd parity constraints jointly true without revising a constraint or changing the task. Compare a centralized constraint solver, distributed message passing and naive majority voting at matched bandwidth and latency. Test solution accuracy and contradiction detection. Neither global consistency nor a nonzero obstruction is claimed sufficient or necessary for experience.

6 M5: calibration, memory and the limits of a self-model interpretation

6.1 States, observations and control

An actuator has an unknown fixed polarity $\theta \in \{0, 1\}$, initially fair. Calibration actions a_j are independent fair bits known to the full agent. The measured result is

$$y_j = a_j \oplus \theta \oplus \epsilon_j, \quad \epsilon_j \stackrel{\text{iid}}{\sim} \text{Bernoulli}(\eta).$$

The agent stores $z_j = y_j \oplus a_j$, and estimates θ by majority vote with random fair tie-breaking. This is a maximum-likelihood/MAP rule for known $\eta < 1/2$ and a fair prior. Calibration lasts n fast observation steps; θ is constant over calibration and the following execution block. The chosen boundary includes the actuator and its action/sensor loop; an external object with the same input-output law would admit the same inference.

During execution, a target g_t is supplied, the agent chooses $a_t = g_t \oplus \hat{\theta}$, and the output is $o_t = a_t \oplus \theta \oplus \nu_t$, with independent execution noise $\nu_t \sim \text{Bernoulli}(\mu)$. The agent is not allowed to learn from execution feedback in this experiment. Relaxing that restriction a different controller. Forward-model motor control motivates this functional construction [5], but does not establish its conscious interpretation.

6.2 Exact performance and ablations

The probability of correct calibration is

$$A_n(\eta) = \sum_{k=0}^n \binom{n}{k} \eta^k (1-\eta)^{n-k} (\mathbf{1}_{2k < n} + \frac{1}{2} \mathbf{1}_{2k=n}), \quad A_0 = 1/2.$$

At $\eta = .2$, the values for $n = 0, 1, 3, 5$ are $.5, .8, .896, .94208$. With utility $U_n = A_n - .03n$, the best member of this *four-option menu* is $n = 3$. Utility is accuracy minus probe cost, not survival or thermodynamic free energy. No global optimal-stopping claim is made.

Probes	A_n	U_n	True survival	IID shortcut
0	.50000	.50000	.49271000	.34375000
1	.80000	.77000	.78757400	.81443540
3	.89600	.80600	.88193048	.92066609
5	.94208	.79208	.92722159	.95564458

Survival here is an *operationally defined* event: at most two output failures during six executions, with $\mu = .1$. It is not demonstrated autopoiesis. Conditional on a correct estimate, failures are $\text{Bin}(6, \mu)$; conditional on an incorrect estimate they are $\text{Bin}(6, 1 - \mu)$. Hence

$$\mathbb{P}(\text{survive}) = A_n B_{6, \mu}(2) + (1 - A_n) B_{6, 1-\mu}(2),$$

where $B_{T,p}(b) = \sum_{k=0}^b \binom{T}{k} p^k (1-p)^{T-k}$. Replacing failures by independent trials at the averaged error rate loses their common dependence on $\hat{\theta}$ and gives the erroneous last column.

A no-memory decoder retaining only the final calibration sample has accuracy A_1 , even if it paid for three probes. An action-blind decoder sees only $y_1, \dots, y_n$. Conditional on either θ , that vector is uniform, so its best accuracy remains $1/2$. Replacing its action trace by independent fair bits has the same consequence; arbitrary within-episode permutation is *not* asserted to remove every information trace. A trivial fixed prior estimate has accuracy $1/2$; an oracle has calibration accuracy one and still suffers execution noise. Comparisons require identical measurement and cost budgets; a one-probe policy and a three-probe memory ablation have different costs despite equal final calibration accuracy.

6.3 What is self-related, and what is not?

The estimated parameter controls the consequences of the agent’s own actions, which gives an operational body-parameter relation. There is no second-order confidence report, introspection, narrative identity, or phenomenal variable. The sufficient memory can be an integer vote count. A compiled finite-state controller that stores that count and maps it directly to actions has exactly the same performance without a separately named self-model module. Thus the ablations establish contributions of memory and action-conditioned data, not the necessity of an explicit self-representation architecture.

An informative next test changes θ after calibration and compares recalibration, change-point detection, a history-matched lookup controller, and an externally observing estimator. Report target success, proper scoring rule loss, calibration cost and recovery time separately. The model should be abandoned as a claim of privileged self-awareness if those matched alternatives explain its entire advantage. It remains useful as an exact test of adaptive control and observational access.

7 Additional negative results and review targets

7.1 A directed-order distance need not be a metric

R003 [6] proposes, in its arXiv v1 printed p. 15, footnote 15, zero on the diagonal, shortest directed order-path length between comparable points, and infinity otherwise. For the strict order $x < z, y < z, x \parallel y$, this gives

$$d(x, y) = \infty > d(x, z) + d(z, y) = 2.$$

Thus the formula fails the triangle inequality on general strict partial orders, and is not natural-number-valued there. This is a counterexample to that formula under its stated generality, not a rejection of the paper’s phenomenological programme. An undirected graph path metric on each connected component repairs the triangle inequality but changes the semantics; it can connect incomparable points. A total-order restriction is another, narrower route and requires explicit assumptions about allowed edges.

7.2 Plasticity is not automatically observable change

R001 [7] motivates an objection to static substitutes. The main-text inference from nontrivial weight change to necessarily changed input-output function requires qualifications: $f_t(x) = w_t x - w_t x = 0$ has changing parameters and a constant function. Likewise, finite-horizon adaptive dynamics can be written on an augmented state (h_t, W_t) and composed for a fixed number of steps. This does not imply cheap simulation, equal physical intervention responses, or a fixed-window solution for every horizon. Compare input histories, memory, parameter access, resource budgets and admissible perturbations explicitly. The article’s supplementary proofs were not successfully retrieved in this edition, so no whole-paper refutation is claimed.

8 P1: predictive refinement under a repair-safe library

This addendum develops IL002/IL005 and M1. Its authoritative report, full proofs, provenance, independent reviews and bounded checks are in `p1/README.md`. It extends the local RF target from complete identification to reported classifications. The inherited RF and RFM sources are the canonical records listed under B06, not a theorem for arbitrary height-two posets. Neither novelty nor empirical relevance is established here.

8.1 Fixed-state acquisition contract

Let $P = \{a\} + M + \{z\}$, with distinct least and greatest elements and an antichain M . Fix an explicitly supplied finite library $\mathcal{Q}$ of oriented, noiseless Boolean queries. The unknown $x \in P$ remains fixed during acquisition. Background signatures define the post- T_0 order but are not freely observed (A1). Complements are separate supplied queries. There is no cross-branch consumption, changing library, extra history-dependent cost or state-changing acquisition.

At support S , a query q retains $v \leq w$ exactly when $q(v) \leq q(w)$. Legality requires the resulting order-complex inclusion to be a homotopy equivalence, before conditioning. Branch cancellation leaves the induced initial order. Set $\mathcal{Q}_s = \{q : q(a) \leq q(z)\}$. The inherited local rule is

$$\text{safe}(S, q) \iff \begin{cases} q(a) \leq q(z), & a, z \in S, \\ q(a) \leq q(v) \text{ for all } v \in S, & a \in S, z \notin S, \\ q(v) \leq q(z) \text{ for all } v \in S, & z \in S, a \notin S, \\ \text{true}, & a, z \notin S. \end{cases}$$

With both endpoints, 00 and 11 permit elementary edge–triangle collapses, 01 deletes nothing, and 10 has nonzero relative H_1 . On an endpoint star a descending edge disconnects the complex. A useful legal query on an endpoint-containing support therefore belongs to $\mathcal{Q}_s$.

Proposition 4 (P1-T1: target classification). *For a finite-valued target $c : P \rightarrow C$, a legal tree reports $c(x)$ exactly if and only if $\mathcal{Q}$ separates differently labelled middle pairs and $\mathcal{Q}_s$ separates every differently labelled pair involving an endpoint.*

Proof. For necessity, follow each pair’s shared history to its first separator. An endpoint pair requires a useful legal separator in $\mathcal{Q}_s$. For sufficiency, split $\mathcal{Q}_s$ into equal-endpoint tests $\mathcal{E}$ and ascending-endpoint tests $\mathcal{A}$. Peel informative $\mathcal{E}$ tests on the branch retaining both endpoints. Detached branches are antichains, classifiable using $\mathcal{Q}$. The remaining core agrees with the endpoints on every $\mathcal{E}$ test. If $\mathcal{A}$ is empty, the endpoint separation condition makes the core homogeneous. Otherwise stop if homogeneous, or ask any $\mathcal{A}$ test to produce a lower and an upper star. A remaining leaf with a different label from its endpoint has an $\mathcal{A}$ separator: an $\mathcal{E}$ separator is excluded by core membership. This is legal throughout the respective star. Peel such leaves into antichain branches, classify those, and stop when each endpoint branch is homogeneous. $\square$

Direct scans give $O(n^2 + mn^2)$ decision and construction time, including order validation, and at most $n - 1$ internal nodes. Depth optimality does not follow.

Proposition 5 (P1-T2: finest attainable transcript). *Partition P by $\mathcal{Q}_s$ signatures. Keep each class containing an endpoint intact and refine every other class by full $\mathcal{Q}$ signatures. Let R be the resulting equivalence relation. Then c is legally reportable exactly when $R \subseteq \ker c$. Some legal tree realizes exactly R as its leaf partition.*

Proof. R is the equivalence closure of the unavoidable pairs in P1-T1. Apply that proposition to the R -class label. Correct classification separates distinct R classes; no legal tree separates states within one. Its leaves are therefore precisely those classes. $\square$

8.2 Compatibility with deterministic prediction

Separately supply an output $b : P \rightarrow O$ and known deterministic maps $T_u : P \rightarrow P$. These define counterfactual predictions; they are not executed or freely observed during acquisition. M1 gives

$$E_0 = \ker b, \quad E_{k+1} = E_k \cap \bigcap_u (T_u \times T_u)^{-1}(E_k).$$

The stable E_* is the largest transition congruence inside $\ker b$, with at most $n - |b(P)|$ strict refinements.

Proposition 6 (P1-T3: predictive acquisition). *A legal tree reports the E_* class if and only if $R \subseteq E_*$. This is equivalent to the existence of any legally reportable transition congruence $F \subseteq \ker b$.*

Proof. P1-T2 proves the first equivalence. Any such F satisfies $R \subseteq F \subseteq E_*$ by reportability and maximality. Conversely E_* itself supplies F when the inclusion holds. $\square$

If the coarsest predictive partition is blocked, every finer predictive partition is blocked. A witness is an R -equivalent pair separated by E_* , with a finite action word giving unequal b outputs. If compatible, acquisition may retain extra transcript information and discard it in the reported class. No quotient-order or topological property of this compression is asserted.

There is a dual closure:

$$G_0 = R, \quad G_{k+1} = \text{Eq} \left(G_k \cup \bigcup_u (T_u \times T_u)(G_k) \right).$$

Its limit G is the least transition congruence containing R , since every containing congruence contains every iterate. Compatibility is equivalently $G \subseteq \ker b$. Reportable predictive congruences are then exactly the transition congruences F with $G \subseteq F \subseteq E_*$: G is the finest and E_* the coarsest reported partition.

8.3 P1-C1: a minimal supplied-but-unsafe pullback

On $a < x < z$, take $b = 010$, $\mathcal{Q} = \{010, 110\}$ and $T(a) = x$, $T(x) = x$, $T(z) = z$. This T is order preserving. The base partition is $\{a, z\} \mid \{x\}$, but $bT = 110$ separates a, z , so E_* is discrete after one round. The pullback is already supplied. Nevertheless $R = \{a, z\} \mid \{x\}$: 110 is unsafe initially and remains the descending test 10 on the zero branch $\{a, z\}$ of 010. Ordinary acquisition has depth two; RF acquisition is impossible. Replacing 110 with the actually supplied complement 001 repairs it.

Three states are minimal for strict deterministic predictive refinement: on at most two states a nonconstant base is injective, while a constant base stays constant after every word. Two queries are minimal when the nonconstant base is safely obtainable and ordinary acquisition obtains a strict refinement: one query cannot split its own fibres. These are conditional minimality claims.

Passive histories $(b(x), b(Tx))$ are 01, 11 and 00 and identify the initial state after one transition. This is not RF acquisition: the zero branch becomes $\{x, z\}$, where the second query is unsafe. For the timing control $T = (a, z, x)$ and $\mathcal{Q} = \{010, 100\}$, locally legal two-time sampling is possible under an explicitly added image-support update rule. Such state-changing protocols lie outside the fixed-state theorem and do not forecast before the later sample.

8.4 P1-C2: exact transcripts need not have a least refinement

Let $P = \{a\} + \{x, y\} + \{z\}$, with identity dynamics and target $B = \{a, x, z\} \mid \{y\}$, and supply 0001, 0011, 0111, 0101. A legal tree using 0001 then 0011 has leaves $P_1 = \{a, x\} \mid \{y\} \mid \{z\}$. One using 0111 then 0101 has leaves $P_2 = \{a\} \mid \{y\} \mid \{x, z\}$. Both report B . Every root query splits $\{a, x, z\}$, so no exact leaf partition is B . The only common coarsening of P_1, P_2 that still refines B is B , which is unavailable. Thus no coarsest attainable exact-leaf refinement exists, although the coarsest predictive reported partition does. This distinguishes two meanings of *implementable refinement*.

8.5 Checked scope

The coordinator checks 964 cases across nine named libraries: 75 target-classification cases, 888 predictive cases, and one exact-leaf audit. A separate verifier adds 188 checks, including simplicial

collapse and relative-homology certificates. All assertions pass. Caps are eight states, eight queries, twelve libraries and 4096 combined cases; exhaustive transition enumeration stops at four states. No all-library search, general-poset classification, transition learning, noise model or empirical test is claimed. The general statements depend on proofs, not extrapolation from these checks.

9 P2: attainable information over a binary logical support

9.1 Declared constraints and prior construction

This addendum develops IF001, IR002 and M3. The variables $X, Y \in \{0, 1\}$ have named intervention ports and a declared two-variable partition. Fix $a = \mathbb{P}(X = 1)$, $b = \mathbb{P}(Y = 1)$ in $[0, 1]$, exact support $S \subseteq \{0, 1\}^2$, and $\epsilon > 0$. Feasible laws have zero mass outside S and mass at least ϵ on each allowed cell. The floor is an additional weight constraint, not a consequence of logical support. All information is in bits, at the population level; no estimation or temporal mechanism is claimed. Real probability arithmetic remains distinct from the local CM/LM algebra over $\mathbb{F}_2$.

The Bernoulli coupling representation and Fréchet bounds are established results [2]. At fixed marginals, maximizing mutual information is minimizing joint entropy on a transportation polytope [1]. The floor-clipped binary characterization below is an elementary derived consequence, not a priority claim. The narrow search and inspected passages are recorded under P2S001–P2S002 in the source register and in `p2/prior-art.md`.

Every joint law with these marginals has the form

$$P_t = \begin{pmatrix} 1 - a - b + t & b - t \\ a - t & t \end{pmatrix}, \quad J(t) = I_{P_t}(X; Y) = h_2(a) + h_2(b) - H(P_t).$$

Here $t = \mathbb{P}(X = 1, Y = 1)$, rows index X , columns index Y , and $0 \log_2 0 = 0$. The entropy formula also handles degenerate marginals.

9.2 Sharp continuum bounds

Proposition 7 (Full-support characterization). *For $S = \{0, 1\}^2$, put*

$$L = \max\{\epsilon, a + b - 1 + \epsilon\}, \quad U = \min\{a - \epsilon, b - \epsilon\}.$$

The family is empty if and only if $L > U$. Otherwise, with $t_0 = \min\{U, \max\{L, ab\}\}$, its exact attainable-information set is

$$[J(t_0), \max\{J(L), J(U)\}].$$

Both extrema and every value between them are attained. The minimum law is unique. If $L < U$, all maximum laws are endpoints; both endpoints can tie.

Proof. The four cell inequalities give exactly $L \leq t \leq U$. On a feasible interval all cells are positive and

$$J'(t) = \log_2 \frac{t(1 - a - b + t)}{(a - t)(b - t)},$$

$$J''(t) = \frac{1}{\ln 2} \left(\frac{1}{t} + \frac{1}{1 - a - b + t} + \frac{1}{a - t} + \frac{1}{b - t} \right) > 0.$$

The numerator minus denominator in the odds ratio is $t - ab$. Thus J decreases below ab and increases above it on the positive-cell domain, giving the clamped minimizer. Strict convexity gives the endpoint maximum. Continuity and the intermediate-value theorem give the entire closed image interval, including the case $L = U$. $\square$

Independence is feasible exactly when the four product cells $ab, a(1-b), (1-a)b, (1-a)(1-b)$ are all at least ϵ . For a nonempty feasible family, otherwise the minimum is strictly positive, despite rectangular support. For $a = b = 1/2$ and $0 < \epsilon \leq 1/4$, the information interval is $[0, 1 - h_2(2\epsilon)]$.

Proposition 8 (Proper-support characterization). *For a proper support, every forbidden cell fixes t : the cells $(1, 1), (1, 0), (0, 1), (0, 0)$ respectively require $t = 0, a, b, a + b - 1$. The family is feasible if and only if all these requirements agree and the resulting table satisfies every support/floor constraint. When feasible, its attainable-information set is the singleton $\{J(t)\}$.*

Proof. There is at least one forbidden cell; its affine coefficient of t is 1 or -1 . It fixes the only free parameter. Checking the remaining constraints is necessary and sufficient. Normalization is built into P_t . $\square$

This includes degenerate marginals; empty support is infeasible. It does not identify weights from support alone when marginals are unknown. For $S = \{(0, 0), (0, 1), (1, 0)\}$, $a = b = 1/3$ and $0 < \epsilon \leq 1/3$, the unique table has three equal masses and $I = 2h_2(1/3) - \log_2 3 \approx 0.2516291674$ bits.

If full support is constrained only by strict positivity and fixed $0 < a, b < 1$, the parameter interval is open, with endpoints $\max\{0, a + b - 1\}$ and $\min\{a, b\}$. Its information image is $[0, M)$, where M is the larger endpoint information on the closed nonnegative-cell completion. Strict convexity excludes an interior maximizer. Thus a supremum must not be reported as attained within exact full support. The positive floor changes this family and makes it compact.

9.3 Finite grids and a quantitative rounding bound

Proposition 9 (Aligned rational grids). *For full support, let N be a positive integer, suppose $Na, Nb, N\epsilon$ are integers, and assume $L \leq U$. The feasible grid points are $t = k/N$, $k = NL, \dots, NU$. The grid maximum equals the continuous maximum. Its minimum is the smaller of $J(t_-), J(t_+)$, where*

$$t_- = \min\{U, \max\{L, \lfloor Nab \rfloor / N\}\}, \quad t_+ = \min\{U, \max\{L, \lceil Nab \rceil / N\}\}.$$

Writing $J_{N,\min}$ for that minimum,

$$0 \leq J_{N,\min} - J(t_0) \leq \frac{1}{2\epsilon \ln 2 N^2}.$$

If t_0 is an endpoint, the gap is zero.

Proof. Alignment places both endpoints on the grid. Monotonicity on either side of ab gives the two-candidate minimum rule. When $t_0 = ab$ is interior, choose a nearest grid point, at distance at most $1/(2N)$. The second derivative is at most $4/(\epsilon \ln 2)$ throughout the segment, and $J'(ab) = 0$, so Taylor's theorem gives the bound. The minimum over the grid can only be smaller than this candidate value. Endpoint minimizers are already on the grid. $\square$

The nearest parameter need not minimize asymmetric J ; evaluate both candidates. The finite set of grid values is not the whole interval between its extrema. If Na or Nb is nonintegral the grid family is empty. If only the floor is unaligned, the feasible grid endpoints round inward and the continuous maximum need not be attained on the grid.

(a, b, ϵ)	Feasible t	Continuous MI, bits	Grid $N = 20$ min
$(1/2, 1/2, 1/10)$	$[1/10, 2/5]$	$[0, 0.2780719051]$	0
$(3/10, 2/5, 1/20)$	$[1/20, 1/4]$	$[0, 0.2512271879]$	0.0058021490
$(1/4, 1/4, 1/10)$	$[1/10, 3/20]$	$[0.0270944047, 0.1436583459]$	0.0270944047
$(1/2, 1/2, 1/4)$	$\{1/4\}$	$\{0\}$	0

Displayed decimals approximate exact entropy expressions. In the second row independence has $t = 3/25$, absent from the $N = 20$ grid; the $N = 40$ minimum is about 0.0003569247. In the third row independence $t = 1/16$ violates the floor, and the continuous minimum is positive.

9.4 An identical observational law with different interventions

Let $U \sim \text{Bernoulli}(1/2)$ and $E \sim \text{Bernoulli}(1/5)$ be mutually independent. Compare acyclic structural models

$$D : X = U, \quad Y = X \oplus E; \quad C : X = U, \quad Y = U \oplus E.$$

Both have the full-support observational table $P = \begin{pmatrix} 2/5 & 1/10 \\ 1/10 & 2/5 \end{pmatrix}$, the same floor $1/10$, and $I(X; Y) = 1 - h_2(1/5)$. A surgical intervention replaces only the equation for X and leaves the exogenous law and the equation for Y unchanged. It gives

	$\mathbb{P}(Y = 1 \mid \text{do}(X = 0))$	$\mathbb{P}(Y = 1 \mid \text{do}(X = 1))$
D	$1/5$	$4/5$
C	$1/2$	$1/2$

For an independent fair intervention label A that sets $X = A$, $I(A; Y)$ is respectively $1 - h_2(1/5)$ and 0. The common cause U may be hidden; causal sufficiency and a known DAG are not assumed. The observational interval therefore identifies no unique causal channel. This is neither an IIT Φ nor a PID calculation.

9.5 Admissible recoding and verification limits

Take the admissible group to be the four local bit flips $(X, Y) \mapsto (X \oplus r, Y \oplus s)$, transporting support, marginals, law and intervention labels. They preserve individual and joint entropies, hence each MI value and the entire interval. An asymmetric marginal constraint must be transported: these maps need not preserve a single numerical marginal slice.

For the preceding law the global invertible map $(C, D) = (X, X \oplus Y)$ gives $\begin{pmatrix} 2/5 & 1/10 \\ 2/5 & 1/10 \end{pmatrix}$, whose coordinates are independent. Whole-state entropy is preserved, but MI drops to zero. This source-mixing map is outside the declared port-preserving group; it is an intentional negative control, not a counterexample to the stated invariance.

The primary suite checks 6320 integer tables and 3930 support/marginal groups for denominators 4, 8, 12, 16, 20, 24, plus 30 named grid, boundary, recoding and causal cases: 10280 total, one run. A separate reviewer checks 3809 cases in one run and reconstructs the proofs. All pass. The respective per-run caps are 25000 and 5000 cases, 60 seconds each, with at most three runs and denominators at most 240. Feasibility uses exact integers/rationals; logarithms use binary64 with absolute tolerance $2 \cdot 10^{-12}$ in the primary suite. No grid enumeration proves the continuum result or certifies the displayed logarithms by interval arithmetic. Search used two primary anchors; exact retrieval counts and locators are in the P2 report. This is a bounded statistical characterization, not a larger-alphabet algorithm, novelty certificate, causal identification or evidence of phenomenal consciousness.

10 Open theorem and experiment targets

1. Section 8 resolves reported predictive acquisition for B06’s fixed bounded-poset contract and supplies an exact-transcript obstruction. A state-changing acquisition model remains separate; do not extrapolate to arbitrary height two.
2. Optimize task risk plus observation cost over a hierarchy, keeping group index, entropy and action value as separate outputs. Compare to ordinary decision trees before introducing additional topology.
3. Section 9 resolves fixed-marginal binary support lifts with a positive cell floor. Larger alphabets and uncertain marginals require a separately motivated extension and ordinary coupling baselines.
4. Relate constraint-presheaf obstruction data to the particular order-complex relative homology used in B03, with an explicit map and counterexamples.
5. Study drifting actuator parameters with calibrated confidence and cost-matched compiled controllers. A cognitive interpretation requires independent behavioral evidence; phenomenal claims need more than this task.

These targets are feasible research programmes, not promises of positive results. Experiments here inspect bounded cases and do not certify scaling. RF-IDENT search, all partitions, and rich sheaf calculations can grow rapidly; the initial next runs should keep declared finite limits.

References

- [1] M. Kovačević, I. Stanojević, and V. Šenk (2015). On the entropy of couplings. *Information and Computation* 242, 369–382. doi:10.1016/j.ic.2015.04.003. Read copy: [arXiv:1303.3235v2](#), Sections 2.1–3.1.
- [2] A. Erdely (2026). A subcopula characterization of dependence for the multivariate Bernoulli distribution. *Journal of Statistical Theory and Applications* 25, article 4. doi:10.1007/s44199-025-00157-4, Section 3, equations (7)–(8).
- [3] E. P. Hoel, L. Albantakis, and G. Tononi (2013). Quantifying causal emergence shows that macro can beat micro. *PNAS* 110, 19790–19795. doi:10.1073/pnas.1314922110.
- [4] S. Abramsky and A. Brandenburger (2011). The sheaf-theoretic structure of non-locality and contextuality. *New Journal of Physics* 13, 113036. [arXiv:1102.0264](#).
- [5] D. M. Wolpert, Z. Ghahramani, and M. I. Jordan (1995). An internal model for sensorimotor integration. *Science* 269, 1880–1882. doi:10.1126/science.7569931.
- [6] J. Kleiner and T. Ludwig (2023). What is a Mathematical Structure of Conscious Experience? [arXiv:2301.11812v1](#), preprint.
- [7] V. N. O’Reilly-Shah, A. M. Selvitella, and A. Schurger (2026). A caveat regarding the unfolding argument: implications of plasticity. *Neuroscience of Consciousness* 2026(1), niag027. doi:10.1093/nc/niag027.