# A Physics Canon as a Replayable Algebraic Object

Verifier-Governed Extraction, Representation, and Minimality Classification

David Elliman
Neuro-symbolic Ltd, London, United Kingdom
dave@neusym.ai
ORCID 0009-0001-9792-4036

Draft of 23 August 2026

**Abstract**

Machine-assisted foundational programmes can accumulate locally persuasive claims faster than a human reader can reconstruct their global state. We report a verifier-governed method that turns such a changing corpus into a byte-pinned algebraic object before attempting to represent it. The physical question organizing the demonstration is whether the corpus's talk of outcomes, dephasing, records, export, and redundancy really forces a particular record structure, or merely describes one convenient realization. Source extraction, lifecycle grading, the target process language, equivalence, proof grades, controls, and possible negative outcomes are fixed before a candidate is examined. A provenance-tagged classical-kernel construction then represents the frozen weak source presentation and passes every preregistered signature, preservation, and faithfulness test. The result is deliberately qualified: operational faithfulness holds only for the coarse source quotient justified by the extracted canon, not for an imagined richer theory. Subsequent classification shows that the representation's record and phase skeleton is irreducible and rigid, while other carrier and provenance structure is removable. An exact proper quotient therefore defeats minimality, and a phase invariant defeats initiality. The principal contribution is a replayable instrument that can expose which physical distinctions are load-bearing and select a negative conclusion without retuning its language. No empirical or ontological claim about the demonstration corpus is made.

**Keywords:** AI-assisted science; foundations of physics; formalization; provenance; reproducibility; categorical process theory.

**Artifact status.** The manuscript is a ZERO-AUTHORITY explanation of a banked machine record. The JSON artifacts and their git commits control wherever prose differs. The reproducibility archive is deposited at doi:10.5281/zenodo.22067781; the public bank is pinned at commit `8a980ab`.

## 1 Introduction

The reliability problem of machine-assisted research is usually posed at the scale of an answer: did a model cite the right paper, execute the right code, or derive the stated equation? A long research programme creates a harder problem. Its state is a corpus rather than an answer. Definitions change, claims acquire and lose support, numerical specializations coexist with generic laws, and a correction at one location may leave an older assertion alive elsewhere. A fluent generator can continue producing locally coherent text after the global state has become inconsistent. The damaging failure is therefore not only hallucination; it is an uncontrolled transition in what the programme treats as established.

Blind analysis and preregistration address an analogous danger by fixing a decision rule before an investigator can see which rule is convenient [1–4]. Reproducible computational research adds source, data, and executable workflows [5, 6]. Those practices are necessary here, but not sufficient. In a large theory corpus the objects to be frozen—the source inventory, lifecycle, definitions, admissible models, equivalence relation, proof grades, and the possible negative landings—must themselves be constructed. If that construction reads the desired result too early, a formally preregistered test can still be reverse-engineered.

This paper reports a worked response, called *Algebra Before Bit* (ABB). Its premise is that a framework should not choose the language in which it is represented. A target-neutral process language is fixed first; the framework corpus is then extracted into a typed signature; a candidate representation is evaluated against contracts fixed before the candidate; and only then is minimality classified. Each stage ends at an ordered terminal that can record an exact obstruction or a partial result as readily as a positive one. The workflow extends two earlier pieces of methodology: a rigid ground-truth discipline for speculative research and a truth-maintenance system that relocates trust from a generator to an executable verifier [7, 8].

The figure/ground relation is crucial. The demonstration corpus is a large foundational-physics framework. It supplies difficult source material: processes, records, phases, parameters, equations, open propositions, and a long retraction history. It is the *ground*. The *figure* is the instrument that must ingest those bytes without granting them authority, turn them into a closed presentation, and then accept a negative classification when the evidence requires one. The paper takes no position on whether the corpus describes nature.

## The physical question

The physics-facing question is narrower than the corpus's proposed dynamics but more substantive than a software audit: *when does an operational process amount to a physical record, and which representational structure is actually forced by that claim?* An outcome label, a repeatable instrument, its label-forgotten channel, a locally readable carrier, an exported record, and redundantly broadcast records are not interchangeable notions. In particular, dephasing suppresses selected coherences but does not by itself establish that an outcome was physically invoked, exported to another carrier, reset, erased, or accompanied by a thermodynamic cost. Conflating these steps imports a classical "bit" before the operational algebra has shown where one exists. The order in ABB is therefore algebra before bit: first expose the processes and record capabilities separately, then ask which of them survive every admissible representation.

That question gives the three mathematical units a single physical thread. The signature unit identifies what the source actually says about instruments, dephasing, readout, export, and redundancy. The representation unit asks whether those commitments can coexist in one admitted process model while preserving the full instrument before any lossy decoding. The classification unit then asks which parts of that model are indispensable. Its answer is diagnostic rather than ontological: the two-label record and phase skeleton is load-bearing in the frozen comparison class, whereas generic carrier padding and a target-local provenance distinction can be removed without changing any accepted source claim. Rigidity of the surviving record structure does not imply minimality of the whole representation.

This is the physical content established here: a controlled separation of record formation from its possible realizations, and an exact account of which structure the frozen operational claims do and do not require. The corpus's detailed dynamical mechanisms, proposed physical ontology, and empirical consequences require a companion paper. They are intentionally not inferred from the present representation theorem. This boundary follows the operational distinction between abstract process structure and physical selection familiar from generalized-probabilistic and categorical approaches [9, 10], and it keeps the roles of decoherence and redundant records distinct [11, 12].

**Results, stated with their limits.** The source instrument closes a signature with 2,983 active or conditional entries. A single candidate in the frozen target class passes every one of 32 preregistered rows. The strongest-looking phrase—"operationally faithful"—is also the most restricted: it means injective on a deliberately weak operational quotient having one class per inhabited nonunit source hom, not faithful to a richer source theory. The classification unit then proves that the represented phase is irreducible and that the completed self-equivalence groupoid is terminal, while constructing both a proper acceptance-preserving quotient and a strictly smaller accepted carrier. A separate fixed-point obstruction refutes initiality. In the frozen comparison class the representation is therefore rigid yet non-minimal and non-initial. The primary cascade lands on the earlier exact negative, quotient nonminimality, rather than on a success-labelled terminal.

**Contributions.** The work makes four contributions.

1. It gives a target-blind, byte-covering extraction protocol that distinguishes a source denominator, a lifecycle-filtered signature, and a later representation denominator.
2. It combines domain-separated content identities, role-defined authority chains, exact-set source audits, independent routes, and preregistered terminal cascades into one replayable research instrument.
3. It demonstrates the instrument on a nontrivial categorical representation and classification problem, including exact negative witnesses rather than only positive calibration cases.
4. It treats the repair history as result data. The system's evidential value lies partly in instances where it halted its own author, forced a narrower claim, or routed the programme to a negative terminal.

Section 2 defines the instrument and its threat model. Sections 3–5 report the signature, representation, and classification units. Section 6 analyses the repair ledger. Sections 7 and 8 state the limits and replay protocol.

# 2 The verifier-governed instrument

## 2.1 Threat model: fluent transition proposals

The principal adversary is not a malicious intruder. It is a fluent generator, prompted by an owner who wants a programme to succeed, proposing changes faster than the owner can reconstruct their consequences. Familiar researcher degrees of freedom remain present [13]; machine generation increases their throughput. Four attacks matter especially:

**Target leakage.**
A source concept is included, excluded, or retyped because it helps the representation that will later be proposed.

**Denominator drift.**
A class, equivalence, proof grade, or comparison order is changed after a failed attempt, making the failure disappear.

**Epistemic rounding.**
"Not attempted", "attempted but unresolved", "refuted", and "proved" are compressed into one convenient status.

**Authority laundering.**
A readable note, index, literature source, or later-child result becomes a mathematical premise although it was licensed only as a locator or explanation.

The defended asset is not truth in the world. It is a replayable relation between frozen inputs, typed claims, evidence, and landings. The architecture can show that a result follows from a stated comparison class, not that the comparison class is physically privileged.

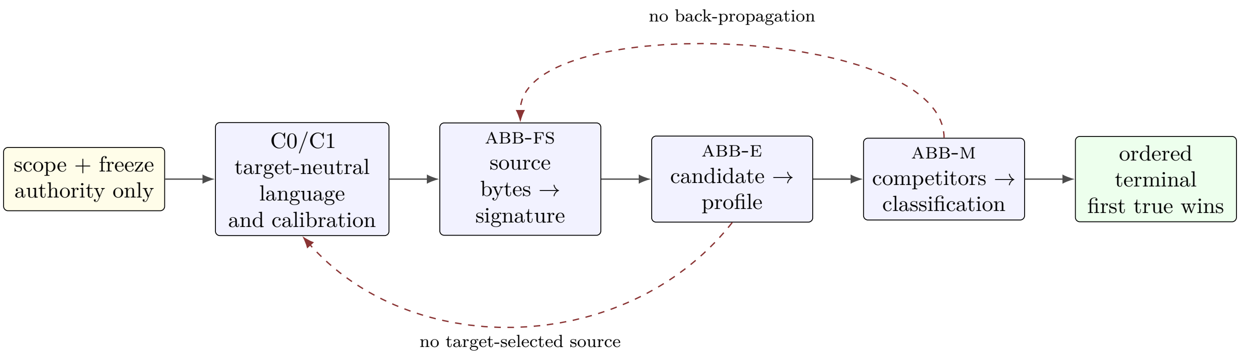


Figure 1: **The governed order.** Each arrow consumes a banked result. Dashed reverse arrows are prohibited. The framework is first read semantically at ABB-FS, after the target-neutral language and the extraction schema have been frozen. Every unit ends in an ordered terminal cascade.

### 2.2 Roles, gates, and the non-self-authorizing chain

Every unit separates construction from countersign. The constructor may produce candidate artifacts but cannot authorize them. An independent review recomputes identities, checks source inventories and proofs, and either countersigns or returns a narrow hold. Only the human-controlled commit path banks a countersigned byte set. A terminal packet cannot countersign or bank itself.

Authority is role-defined rather than a permanently closed filename list. The current instance comprises a chain of scope and freeze pairs. Each freeze record binds the predecessor chain, the reviewed scope hash, frozen denominators, terminal order, and the exact actions licensed next. A later amendment must become another link; it cannot claim authority merely by being placed near the chain. This design is the direct repair of an early version whose closed enumeration could not lawfully include its own successor.

Figure 1 shows the resulting order. A downstream stage may consume the banked result of an upstream stage, but no stage may use a downstream output to tune its own definitions.

### 2.3 Content identities and declared recipes

A hexadecimal digest is useful only if another reader can reproduce it. Each identifier family therefore declares:

1. a UTF-8 domain separator terminated by a null byte;
2. the exact payload, including omitted self-identifier fields;
3. JSON encoding, key-order and array-order rules;
4. offset bases, nested digests, and text transforms where applicable;
5. whether sorting is forbidden; and
6. the serialization contract for the containing artifact.

Self-excluding identifiers hash the complete object with the identifier field removed. Registries usually hash ordered complete row objects rather than a list of row digests. Consequently, changing a statement and recomputing its row identifier still breaks the containing registry. The hash is not a proof of meaning; it is a precise alarm against byte drift.

This is stricter than ordinary data provenance [14]. Provenance records where an object came from. Here the recipe also determines whether a future reviewer can reconstruct the identity without guessing an offset, serializer, or inner hash. Section 6 reports cases in which stored values were correct but undeclared recipes made them unverifiable; the review halted until the recipe was made local and complete.

| surface | fixed before evidence | what later evidence may do |
|---|---|---|
| source | byte inventory, lifecycle grammar, extraction schema | classify every byte span; retain disagreements |
| model | target class, equivalence, optional-structure strata | construct an admitted candidate or exact obstruction |
| proof | statement, premise ledgers, grade ladder, route intersection | prove, refute, remain partial, or disclose a bound |
| decision | terminal names, order, selector, one-hot expectations | set condition values; never reorder the cascade |
| provenance | manifest and grade policy | bind exact objects; never promote a directory or pattern to proof |

Table 1: The principal anti-postselection surfaces. The right column is allowed to change the result, never the denominator in the middle column.

### 2.4 Two routes and common-origin accounting

Two routes are used where a single route could silently bake in its own conclusion. In extraction, a registry-oriented route classifies index rows while a byte-oriented route partitions the canonical files directly. In proof stages, analytic and constructive ledgers fix premises and plans without reading one another. Their first mutual read is an intersection packet.

Textual similarity does not establish a common premise. A pair enters the intersection only through a proved equivalence or one shared replacement. Everything not selected is explicitly classified "not common". Common content is counted once. Agreement that rests on the shared intersection is called *verification depth*, not independent convergence. This prevents two implementations of one argument from being marketed as two independent confirmations.

### 2.5 Ordered terminals as typed scientific outcomes

A scope freezes an ordered list of possible outcomes before construction. The final packet evaluates every condition, retains every established condition, and selects the first. One-hot controls independently make every branch reachable through the same selector used for the live verdict. This turns negative findings into ordinary outputs rather than exceptions.

The ordering is epistemic rather than optimistic. Authority failures and postselection precede mathematical claims; bounded nonexhaustive searches precede global classification; exact counterexamples precede later success conditions. A proof may therefore complete an obligation while still making a more affirmative unit-level landing impossible. The distinction is intentional: obligation completeness says that the question was answered; terminal order says what the answer permits the programme to claim.

### 2.6 A target-neutral algebra before the framework read

The extraction and representation stages rely on an upstream language that is fixed without reading the framework canon. Its base is a typed symmetric monoidal process grammar with causal and convex enrichment. Objects represent system types; morphisms represent transformations; states and effects are the special homs $I \to A$ and $A \to I$; and scalars are $I \to I$. Operational equivalence is contextual rather than syntactic: two processes are identified only when every admitted single-hole closing context gives the same scalar. The single-hole restriction is material. It lets a later theorem reduce a closed context around one state hole to an admitted effect without assuming local tomography; a grammar allowed to copy the hole could test nonlinear expressions such as $\rho \otimes \rho$ and invalidate that reduction.

Two independently drafted presentations—an ordered-vector-space route and a free symmetric-monoidal route—are reconciled through a common core rather than chosen by convenience. Records are not introduced as a primitive word. They are separated by operational capability: outcome families, record-valued instruments, repeatability, label-forgotten channels, dephasing, export to another carrier, and redundant broadcast. The hierarchy does not grant any adjacent implication for free. In particular, a classical record may copy without exhibiting a distinct export carrier, and dephasing may remove coherence without creating a durable record.

The model category is fixed before any physical axiom is selected. Optional axioms are maintained as discriminators with countermodel obligations, not declared truths of nature. This allows an exact exclusion theorem to be useful without converting its antecedent into a physical commitment. A free theory with no mixed record–system generator, for example, forces every $X \otimes A \to X \otimes B$ diagram to factor through record-only and system-only components. Convex closure then yields only separable controlled families, and extremality of the identity stochastic channel obstructs a nontrivial label-preserving controlled extension. Adding the appropriate mixed generator removes the obstruction. The result is a structural discriminator: it says what the bare language cannot express, not which extension nature chooses.

### 2.7 Calibration controls that have both pass and fail paths

Before the framework signature is read, three familiar theories calibrate the representation vocabulary. They are controls, not candidate answers. Each must realize its declared structure and fail a discriminator that prevents a label-only pass.

For the classical control, the complete sixteen deterministic strategies in the $(2, 2, 2)$ Bell scenario satisfy

$$S = E_{00} + E_{01} + E_{10} - E_{11}, \qquad |S| \leq 2.$$

Eight attain $+2$ and eight attain $-2$. An independent factorization proof writes $S = A_0(B_0 + B_1) + A_1(B_0 - B_1)$; exactly one Bob bracket vanishes and the other has magnitude two. Convex closure lifts the vertex result to the whole local polytope. The pass side contains the simplex, its bipartite product, and a repeatable classical record; the fail side is inability to cross the CHSH bound.

The quantum control uses exact arithmetic in $\mathbb{Q}(\sqrt{2})$. A frozen two-qubit behaviour attains $S = 2\sqrt{2}$ and hence $S^2 = 8$, with every probability equal to $1/4 \pm \sqrt{2}/8$ and all remote-setting marginals exactly $1/2$. Its fail path is universal broadcasting of the nonorthogonal pair $|0\rangle, |+\rangle$. Pure marginals force broadcasting to cloning; a Kraus decomposition then requires a coefficient-overlap square of two, while Cauchy–Schwarz permits at most one. The contradiction is algebraic and covers every finite Kraus rank, not a finite search. It reproduces the operational separation between classical broadcasting and quantum no-broadcasting citebarnum2007.

The box-world control fixes the Popescu–Rohrlich behaviour in the same scenario. Its conditional table has remote-setting-independent marginals and $S = 4$, so $S^2 = 16 > 8$. It passes no-signalling and fails the quantum Tsirelson discriminator. The separation is therefore computed from two bound certificates, not inferred from the words "quantum" and "post-quantum". These controls exercise the general-probabilistic and process-theoretic distinctions discussed in [9, 15–19].

Calibration also demonstrates the epistemic grammar. A universal theorem is not promoted from one successful control, a bounded search is not a proof, and absence of a countermodel is not entailment. The upstream unit deliberately closed at a typed partial terminal because one phrase in its frozen closure denominator remained undefined. Later parent-level authorization carried that partial and two exact refutations forward as design constraints; it did not erase them. Thus ABB-FS begins with a calibrated language but not an all-clear assertion about that language.

## 3 From canonical bytes to a framework signature (ABB-FS)

### 3.1 The source denominator is not the signature

The first semantic read occurs only after a blind extraction schema fixes block kinds, lifecycle statuses, parameter outcomes, provenance grades, signature fields, and record distinctions. The source perimeter contains a canonical assertion file and a correction/retraction ledger. A derived

| stage | count | interpretation |
|---|---|---|
| nonblank canonical blocks | 4,820 | complete classified byte denominator |
| included blocks | 4,030 | source content retained after reconciliation |
| excluded blocks | 790 | retained with a named exclusion reason |
| active or conditional entries | 2,983 | representation signature |
| open-annex entries | 1,042 | obligations or claims not assumed true |
| inactive retained entries | 5 | retracted inclusions kept for audit only |

Table 2: The ABB-FS funnel. The three signature partitions close exactly: $2{,}983 + 1{,}042 + 5 = 4{,}030$.

claim index is available only as a locator; an index mismatch cannot be repaired by trusting the index.

The byte-oriented route partitions both canonical files into contiguous half-open spans. Blank separators remain in the partition, so the coverage claim has an exact denominator. Every nonblank block is classified once. The registry-oriented route maps every derived-index row to source bytes. The reconciliation retains disagreements and records source blocks absent from the index rather than silently treating them as nonexistent.

The final funnel is shown in Table 2. The distinction between 4,820 classified source blocks and 2,983 active signature entries is load-bearing. Exclusion requires a named reason; openness prevents a source statement from becoming an axiom; retraction keeps evidence for history while removing eligibility.

The active entries cover objects, processes, records and instruments, actions and phases, relations, parameters, and theorem obligations. Every entry contains exact source bindings, typed or explicitly non-applicable endpoints, dependencies, relation and lifecycle grades, provenance, parameter links, record-boundary tags, and eligibility. Repeated terminology cannot identify opaque source types across blocks. Coherent isomorphism is not silently replaced by equality.

### 3.2 Parameterization without erasing specialization

The source contains numerical constants, finite alphabets, item numbers, dates, bibliographic metadata, path fragments, and genuine parameter families. A byte-level tokenizer fixed before classification produces 64,381 canonical candidates. The classifier retains 55,525 framework parameter or concrete specialization components and excludes the rest as metadata. Those retained occurrences link to 8,336 parameter entries.

Parameterization is conservative. A concrete value remains attached to its family; it is not replaced by a free symbol. Equal lexemes do not identify denotations, and a family declaration does not assert that every family member is physically realized. The relevant invariant is therefore

$$\text{specialization} \longrightarrow (\text{family}, \text{domain}, \text{constraints}, \text{source value}),$$

not the weaker transformation “number $\mapsto$ parameter”.

### 3.3 Lifecycle and proof provenance

The provenance denominator contains 1,730 imported theorem, proof-obligation, and equation/relation rows. It assigns exactly one of four grades: `PREVIOUSLY_PROVED`, `PHYSICALLY_OPEN`, `TO_BE_REDERIVED`, or `ASSUMED`. The census is respectively 19, 33, 1,678, and zero.

Promotion to PREVIOUSLY_PROVED requires a committed machine-readable blob and a checkable positive result. A path mention, directory, pattern, missing object, or confident sentence is ineligible. This policy is deliberately conservative: it classifies evidence available to the source corpus, not the truth of the theorem. A paper, bibliography entry, or source path never becomes a proof premise merely by being recoverable.

### 3.4 The record firewall

Source language often compresses several operational and physical notions into one word. The signature keeps thirteen notions distinct, including effects, complete record-valued instruments, conditioned outcomes, label-forgotten channels, dephasing, local readability, export, broadcast, physical invocation, reset, and thermodynamic bookkeeping. These are source distinctions, not thirteen rungs in the target theory.

Six non-implications remain explicit:

1. dephasing does not imply export, reset, erasure, or heat;
2. a pointer basis does not select a physical instrument;
3. defining an instrument does not invoke it;
4. local readability does not imply durable export;
5. export to one carrier does not imply redundant broadcast; and
6. record formation alone supplies neither cadence nor thermodynamic cost.

The point is not terminological hygiene alone. These separations prevent a representation from receiving free physical structure because the source uses one overloaded expression.

### 3.5 Transition to representation

Only after the signature identity is fixed does ABB-FS attach the 32-row representation interface: 16 signature clauses, 11 preservation obligations, and five faithfulness grades. No map, functor, target image, or faithfulness claim exists at this stage. The source-to-target seven-rung wording also remains unresolved until ABB-E either chooses a mapping before candidate work or proves that the acceptance predicate is invariant under the candidate mappings. The latter route is what the next unit establishes.

## 4 An exact, qualified representation (ABB-E)

### 4.1 Freezing the source presentation and target class

The signature is intentionally richer than the formal relations justified by its source. In particular, many blocks name endpoints but do not supply a global typed dependency graph. The strongest presentation justified without inventing cross-block equations is therefore frozen as

```
FREE_TYPED_DECLARATIVE_PRESENTATION__
BLOCK_SCOPED_OPAQUE_TYPES__
NO_CROSS_ENTRY_NORMALIZATION.
```

This weak branch is selected before any target object or candidate is read. It contains 1,328 process-generator declarations, 467 relation-bearing declarations, eight coherence declarations, two parameter families, 890 concrete specializations, 286 theorem obligations, and two proved non-generating separators, totalling 2,983 active entries. A further 2,206 opaque endpoints supplement 225 explicit object declarations, producing 2,431 object symbols and 1,103 process-generator images.

The target class is the complete finite-dimensional causal convex process class fixed in the target-blind calibration units, not a finite sample from that class. The candidate uses finite classical

| candidate component | count |
|---|---|
| object symbols | 2,431 |
| process-generator images | 1,103 |
| relation dispositions | 467 |
| coherence certificates | 8 |
| record rungs / incidence maps | 7 / 12 |
| action and phase intertwiners | 323 |
| parameter entries / concrete specializations | 8,336 / 55,525 |
| active / open / inactive source resolutions | 2,983 / 1,042 / 5 |

Table 3: Principal registries of the accepted ABB-E candidate. Counts are denominators, not evidence of minimality.

carriers and exact substochastic kernels, but no dimension bound or carrier census is used as a surrogate for the target class. Raw provenance-tagged terms and their operational quotient are kept as two different levels. The former remembers the complete source word; the latter identifies only what the frozen weak presentation can operationally test. Optional physical axioms remain unselected.

The seven-rung wording inherited from the parent programme admits two plausible readings at its first rung: one starts from subcausal transformations, the other from effects and scalar outcome probabilities. Rather than select the historically preferred list after seeing a candidate, ABB-E proves invariance. If $G$ denotes the global preservation clauses and $C$ the common six later rungs and incidence data, then $G$ implies both first-rung descriptions, say $A_1$ and $B_1$. Hence

$$G \wedge C \wedge A_1 \quad \Longleftrightarrow \quad G \wedge C \wedge B_1.$$

The proof uses the already-frozen identities $\mathrm{Effect}(A) = \mathrm{Proc}(A, I)$, $\mathrm{Scalar} = \mathrm{Proc}(I, I)$, and $p = e \circ \rho$; it does not identify effects with arbitrary transformations. Thus the complete acceptance predicate is candidate-list-invariant although the historical wording is not silently repaired.

### 4.2 Independent plans before construction

Analytic and constructive ledgers each contain twenty premises and 32 row plans. Neither reads the other. Their first mutual-read packet classifies all $20 \times 20 = 400$ premise pairs: fourteen receive a proved shared replacement and are counted once, while 386 are explicitly not common. Each route retains six route-only premises. The intersection then fixes a proof contract for every row: exact requirement, attachment denominator, admitted grades, negative control, route-specific load-bearing premises, and maximum credit.

This order matters. A candidate cannot motivate an easier interpretation of a relation, a weaker equivalence, or a different proof grade. Conversely, a single-route proof is not discarded: it can earn the declared exact grade, but two routes sharing the same load-bearing denominator earn verification depth rather than independent-convergence credit.

### 4.3 The candidate

The admitted candidate maps the weak source presentation into provenance-tagged classical substochastic kernels. Table 3 shows its complete main registries. Every source entry receives exactly one resolution, all 8,336 parameter entries and 55,525 retained specializations are linked, and the open and inactive annexes remain outside the assumed source theory.

The image of a relation-bearing declaration is grade-sensitive. A target equation can be proved conditionally or its truth withheld; placing an equation string in a registry does not make the source proposition true. The record image similarly preserves a complete instrument before constructing a decoded family. No injectivity of the decoded weights is required. A single supplied

$C_2$ phase decoration is represented with its stabilizer action; the candidate neither asserts that every completion is equivalent nor that the completion fibre is contractible.

### 4.4 The 32-row profile

The first evaluation packet tests the sixteen signature and eleven preservation rows; the second tests five faithfulness rows. All 32 land at `PASS_EXACT` using the preregistered grade `EXACT_CONSTRUCTION_WITH_MACHINE_CHECKABLE_MAP`. The evidence partition closes at 24,886 active and 10,050 open attachments: the first packet evaluates 18,056 active and 8,042 open attachments, leaving 6,830 and 2,008 respectively for the faithfulness packet.

The five grades must not be compressed into the single word "faithful". They are:

**FAITH-01: typed.** Object and hom typing is preserved on the frozen presentation.

**FAITH-02: equation.** Every source equation admitted at the appropriate grade is preserved; unproved propositions are not promoted.

**FAITH-03: operational.** A source-only completion certificate, fixed before target images, is injective on the weak operational quotient.

**FAITH-04: raw.** The complete provenance word has an explicit left inverse on the raw source grammar.

**FAITH-05: reflection.** Target predicates are reflected only when an exact source witness is present.

FAITH-03 is the most easily overstated result. Its tested source domain has one operational class per inhabited nonunit hom of the frozen weak presentation. Distinct tags separate those classes; decoded-weight injectivity plays no role. This does not prove faithfulness for a richer quotient carrying additional source equations. FAITH-04 is stronger at the raw provenance level but says nothing about target-local structure outside the source image. FAITH-05 reflects 975 witnessed predicates and withholds 437 unsupported ones, of which 35 are attached to active entries and 402 to open entries.

### 4.5 The result includes its qualifications

Seven qualifications are inseparable from the accepted profile:

1. every pass is relative to the frozen weak declarative presentation;
2. FAITH-03 is limited to its coarse source quotient;
3. the 467 relation rows retain their proposition-obligation grades;
4. FAITH-05 is limited to exact source provenance and withholds the 437 unsupported target predicates;
5. FAITH-04 concerns the raw provenance grammar, not all target-local morphisms;
6. one accepted candidate provides no minimality result; and
7. no physical axiom, empirical truth, or novelty claim is selected.

Removing any one would change the theorem. With them bound, terminal 17 of the frozen eighteen-terminal cascade is earned: a conditional representation of the frozen weak presentation exists. The cascade does not authorize a strong functor from a richer source theory, an absolute consistency claim, or minimality. Those are separate questions.

# 5 Minimality and classification (ABB-M)

## 5.1 The comparison problem

Existence cannot receive free credit in a minimality test. The comparison unit therefore reconstructs the *representation-generated core*: the least target subtheory closed under the accepted images, named record data, composition, tensor product, causal and convex structure, and the witnesses needed by the complete ABB-E profile. Operational equivalence is recomputed using contexts internal to that core; ambient target material is not inherited.

The comparison denominator is fixed before competitors exist. It contains 2,431 object coordinates (4,862 state/effect slots), 3,534 hom bases, 413 record carriers, and 8,336 parameter families. Hom coordinates range over an unbounded generative shortlex grammar, not a term-depth cutoff. Carrier profiles are compared componentwise and pointwise in every parameter:

$$P(R') \preceq P(R) \quad \Longleftrightarrow \quad P_i(R';\lambda) \leq P_i(R;\lambda) \quad \text{for every coordinate } i \text{ and admissible } \lambda.$$

Strictness requires at least one strict coordinate. Weighted sums, lexicographic preferences, and finite parameter sampling are forbidden. Structured equivalence must preserve the profile; cardinality alone never identifies representations.

Six notions are kept separate: phase irreducibility, quotient minimality, carrier minimality, uniqueness of a minimal equivalence class, rigidity of a self-equivalence groupoid, and initiality in the comparison category. The scope includes explicit counterfixtures to every tempting implication between them.

As before, analytic and constructive ledgers are fixed blind. Their $25 \times 25 = 625$ pair intersection proves seventeen shared replacements (twelve locks, two method locks, and three bars), classifies 608 pairs as not common, and leaves eight route-only premises on each side. Eleven proof contracts are frozen before a competitor, automorphism, mediator, or solver output is read.

## 5.2 Core admission and phase irreducibility

The representation-generated core passes all admission conditions and re-proves the complete 32-row ABB-E result internally, with the seven qualifications retained. Its comparison into the ambient candidate sends each internal witness to the corresponding ambient witness, so core admission does not assume the conclusion it is meant to test.

Phase irreducibility admits a short universal proof. Let $i : B \hookrightarrow P_2$ be an arbitrary operational monomorphism supporting the frozen two-label record and phase data. Record support forces normalized states $b_0, b_1$ and read effects $\epsilon_0, \epsilon_1$ to factor through $B$. Define

$$r = b_0 \circ \epsilon_0 + b_1 \circ \epsilon_1.$$

For each basis state $\delta_y$ of the classical carrier, $(i \circ r)\delta_y = \delta_y$. Equality on the classical basis is equality of the channel, so $i \circ r = \mathrm{id}_{P_2}$. Then $i \circ (r \circ i) = i \circ \mathrm{id}_B$; operational monicity cancels $i$ and gives $r \circ i = \mathrm{id}_B$. Thus $i$ is an isomorphism. Every phase-supporting subobject is phase-equivalent to $P_2$, establishing irreducibility without a cardinality argument, an optional physical axiom, or a uniqueness claim about the phase-completion fibre.

## 5.3 An exact proper quotient

Irreducibility does not imply quotient minimality. The generated core contains a target-local dephasing realization $D$ supporting one record incidence. Its classical kernel is idempotent,

$$D^2 = D,$$

but its raw provenance terms distinguish the one-token word $D$ from the two-token word $D^2$. None of the frozen 32-row contracts requires that target-local provenance distinction. The congruence

| object-carrier census | one-point | two-point | total |
|---|---|---|---|
| accepted ABB-E candidate | 1,375 | 1,056 | 2,431 |
| strict competitor | 2,021 | 410 | 2,431 |

Table 4: Carrier counterexample. Exactly 646 generic endpoints shrink. With state and effect coordinates frozen separately, this creates 1,292 strict profile coordinates; every other coordinate is equal and none is larger.

generated by $D^2 \sim D$, extended as the protected target-touched two-sided tensor ideal, is stable under composition, tensor product, and affine structure while leaving the pure source image byte-for-byte intact.

The quotient is proper: it is not raw-hom faithful and hence cannot be structured-adjoint equivalent to the original candidate. Nevertheless the post-quotient representation re-passes all 32 rows, preserves the seven qualifications, and keeps all seven collision policies non-live. FAITH-04 survives because source raw words retain their complete provenance; it does not protect unrelated target-local words. This is an exact witness that the accepted representation is not quotient-minimal in the frozen comparison class.

### 5.4 A strictly smaller carrier

Carrier minimality fails independently. The accepted candidate assigned tagged two-point carriers to 646 generic opaque action endpoints for which the source demanded no two-point distinction. A complete competitor replaces those carriers by tagged singletons while retaining every one of the 410 two-label record outputs, all incidence data, the $C_2$ phase decoration, parameters, and raw provenance.

The competitor instantiates all fifteen candidate registries, is admitted to the core, and re-passes the complete representation profile. It is pointwise no larger over all parameter domains and strictly smaller on 1,292 object-coordinate slots. Profile invariance under structured equivalence then proves inequivalence. No weight, finite sample, or appeal to object cardinality as identity is used.

### 5.5 The minimal world, rigidity, and initiality

The carrier-minimal census is a two-component orbit groupoid. Its phase component is either the trivial or swap action of $C_2$, distinguished by the conjugation-invariant fixed-point counts two and zero. Within a component, classes are indexed by admissible source-fixing target-local raw congruences modulo structured automorphisms. A family

$$\Gamma_n = \langle D^{n+1} \sim D^n \rangle, \qquad n \geq 1,$$

witnesses countably many inequivalent classes in each component. The family proves infinitude; it is not itself asserted to exhaust all congruence parameters. Exhaustion at the structural level instead comes from the normal-form theorem: each carrier-minimal object is structured-equivalent to a least-profile phase representative quotiented by an admissible $\Gamma$.

The accepted candidate is nevertheless rigid. Any structured self-equivalence must preserve the separately named record states $b_0, b_1$ and read effects, eliminating the only candidate nontrivial swap. Every self-equivalence is connected to the identity by a unique comparison-normalized 2-isomorphism; the self-equivalence groupoid is therefore equivalent to the terminal groupoid. This says nothing about minimality and does not make the external completion fibre contractible.

Initiality fails by a different invariant. Consider the carrier-minimal representative with trivial $C_2$ action. An admissible mediator with an invertible comparison would have to conjugate the accepted swap action to the trivial action. Conjugation preserves fixed-point count, but the swap

| row | question | status | result |
|---|---|---|---|
| M-01 | core admission | proved exact | generated core carries the complete qualified ABB-E profile |
| M-02 | irreducibility | proved exact | every phase-supporting subobject is an equivalence |
| M-03 | quotient minimality | refuted exact | $D^2 \sim D$ gives a proper acceptance-preserving quotient |
| M-04 | carrier minimality | refuted exact | 646 endpoints admit singleton carriers |
| M-05 | hierarchy vector | classified exact | irreducible; not quotient- or carrier-minimal |
| M-06 | minimal census | classified exact | two phase components with congruence-orbit moduli |
| M-07 | unique E minimum | not applicable | the E class is not among the carrier minima |
| M-08 | rigidity | classified exact | self-equivalence groupoid is terminal |
| M-09 | initiality | refuted exact | fixed-point invariant obstructs a mediator |
| M-10 | implication lattice | classified exact | witnessed non-collapses among the six notions |
| M-11 | literature disposition | proved exact | twelve-work manifest reconciled; no proof premise borrowed |

Table 5: The ABB-M result ledger. "Exact" is relative to the frozen source presentation, target class, equivalence, and comparison denominator.

fixes zero labels and the identity fixes two. The hom-groupoid to that minimum is empty. This refutes initiality universally over possible mediators; it is not inferred from nonminimality.

Table 5 collects the eleven exact result rows. The status census is three proved, three refuted, four classified, and one not applicable; there are no partial, open, bounded-nonexhaustive, or barred rows.

The terminal cascade has conditions true at slots nine and ten. Frozen order selects slot nine: the ABB-E representation is not quotient-minimal and an exact proper quotient is supplied. The carrier counterexample remains ledgered at slot ten. Rigidity is recorded orthogonally; it cannot rescue a minimality claim. The selected negative is therefore the headline of the classification, not an embarrassment hidden behind the existence theorem.

## 6 What the instrument caught

Methodological rules are easy to endorse abstractly. Their evidential value is clearer when they inconvenience the project that introduced them. The banked repair ledger records several such episodes. Table 6 lists the most instructive ones; each changed a governed artifact before countersign and each left a machine-checkable trace.

Table 6: Observed refusals and repairs. These are historical events in the banked chain, not hypothetical adversaries.

| surface | proposed or discovered defect | enforced repair | scientific consequence |
|---|---|---|---|
| authority perimeter | an early child scope enumerated a closed predecessor file list that could not lawfully absorb a successor amendment | replace the enumeration by a role-defined chain and list the current files only as its instance | later authority cannot be gained by filename proximity |

| surface | proposed or discovered defect | enforced repair | scientific consequence |
|---|---|---|---|
| epistemic status | optional work and attempted-but-unresolved work shared an "open" route, making a later terminal unreachable | introduce a typed unattempted-by-election status, preregister elections, and forbid retroactive demotion after work begins | not attempting an optional theorem no longer masquerades as failed work |
| digest recipes | stored candidate and context hashes replayed only after the reviewer guessed hidden field lists, offset bases, and a second context radius | declare the seven-field identifier payload, inner hash, global offset rule, and block-clipped ±112-byte context recipe | correct values became independently verifiable rather than merely plausible |
| provenance type | two recovered directory paths were labelled and hashed as committed files although the stored objects were git trees | relabel them as tree-existence-only, bind stable tree OIDs, and make them ineligible for proof grade | the proof-grade census remains supported only by file blobs |
| selector rule | the declared PRES-08 predicate reproduced 383 record-tagged attachments but not 33 additional rows | state the actual selector as the union of three clauses with all record-tagged entries and replay both active and open sets | a conservative superset is now reproducible rather than asserted |
| minimality inference | a draft inferred quotient minimality from FAITH-04 raw source injectivity | split the protected pure-source image from target-local raw structure and test a generated tensor ideal | the accepted $D^2 \sim D$ quotient refuted the intended conclusion |

The final row is especially probative. The pipeline had already proved the representation and could easily have treated raw faithfulness as protection against quotienting. Instead, its frozen denominator distinguished source raw terms from target-local raw terms. That distinction produced the primary negative terminal. The methodology did not merely correct metadata; it changed the mathematical headline against the programme's local incentive.

Not every adversarial control records an observed mistake. For example, route-output echo, terminal hardwiring, automatic thermodynamic authority, and control-mode branching are planted mutations exercised through the live acceptance predicates. They show that a detector can fail the relevant state; they are not evidence that the historical constructor attempted the mutation. Keeping observed repairs separate from planted refusals prevents the control battery from becoming an embellished narrative of discovery.

The countersign role is similarly bounded. It does not certify a paper after sampling it. It reconstructs declared digest families, exact sets, source spans, controls, and load-bearing mathematics, and it can issue a narrow hold. Several holds above were declaration-only: the mathematical bytes remained fixed while the recipe or predicate was made reproducible. Others forced a substantive retreat. Both are useful. A verifier that only catches false equations but permits unverifiable identities still leaves the research state under generator control.

## 6.1 Why this changes the foundations question

The slogan "algebra before bit" names an order of explanation, not a claim that physics is merely syntax. A framework corpus may describe bits, qubits, records, phases, fields, or thermodynamic events. Before asking which of those objects is fundamental, the audit asks a prior question: what typed operations and distinctions does the framework actually commit itself to? That order separates three propositions that are often allowed to travel together.

First, *representability* asks whether a source presentation has an admitted image in a target process theory. It is a mathematical existence claim relative to a denominator. Second, *physical selection* asks whether that target or one optional axiom describes nature. Neither a faithful map nor an exclusion theorem answers it. Third, *minimality* asks whether accepted structure is forced by the source, or is padding chosen by the constructor. A representation can pass every preservation and faithfulness test while failing this third question. The present candidate is an explicit example.

That separation bears on a familiar tension in reconstructions of quantum theory and general probabilistic theories [9, 15–17, 20, 21]. A set of operational principles may characterize a formal theory, but the characterization is meaningful only relative to the allowed models, composites, equivalences, and auxiliary structures. Changing any of those after seeing the desired theory risks converting reconstruction into redescription. The governed denominator makes the relativity explicit. It does not remove judgement; it makes the judgement temporally and cryptographically inspectable.

The record hierarchy provides a concrete example. Decoherence can select a stable pointer structure [11]; redundant records motivate quantum Darwinist objectivity [12]. But a source sentence containing "record" may refer only to an outcome label, a repeatable instrument, a dephasing channel, an exported carrier, or redundant broadcast. Treating these as one primitive imports a physical narrative before the map is tested. The thirteen source distinctions, seven target rungs, and six explicit non-implications are intentionally different inventories. The mapping must say which source content realizes which operational capability, while thermodynamic and invocation claims remain outside unless separately supplied.

Minimality then detects surplus narrative structure. The two-label record skeleton survives at the carrier floor: its read states, effects, and phase action are load-bearing. The generic opaque action endpoints do not. Their two-point carriers can be replaced by singletons without losing any frozen claim. Raw target-local provenance also contains a distinction between one and two applications of an idempotent dephasing kernel that the source image never tests. Quotienting that distinction preserves the theorem. The result is more informative than a global declaration that the candidate is "too large": it locates which structure is forced and which is conventional.

## 6.2 Negative landings as scientific products

An ordered terminal is not a traffic light. A green success label would be actively misleading here. The representation exists, so an existence-only cascade could end triumphantly. The classification contains stronger information: within the declared class, an exact quotient shows that the accepted presentation is not minimal. Selecting that negative while retaining the carrier witness and rigidity theorem gives the reader a structural diagnosis rather than a scalar score.

This is also why terminal order cannot be chosen after construction. If "representation found" preceded "proper acceptance-preserving quotient found", the same evidence would support a more positive headline. The frozen order says in advance that exact defects, collisions, and boundedness conditions govern what the unit may announce. One-hot controls show that the selector can reach every branch, but the live evidence alone determines which condition vector is supplied. Reproducibility therefore includes the decision rule, not merely the witness files.

The distinction between an obligation and a terminal matters as well. An obligation can be discharged by an exact counterexample. At the unit level, that same counterexample may permanently block an all-clear landing. This is not inconsistency in the governance scheme. It records that the mathematical question was answered and that the answer changed what later work is allowed to assume. Parent-level authorization can carry the finding forward as a design constraint, as happened before ABB-FS, but cannot relabel it as a positive theorem.

### 6.3 What transfers to other corpus-scale projects

The categorical details are domain-specific; the governance pattern is not. Any corpus-scale formalization can separate:

1. an exact source inventory from a selected signature;
2. lifecycle and evidential status from semantic content;
3. denominator construction from candidate evaluation;
4. independent proposal routes from their first mutual-read intersection;
5. result grades from readable summaries; and
6. a complete outcome vector from the first selected terminal.

The appropriate mathematical objects may be graphs, programs, differential models, ontologies, or statistical estimands rather than process categories. The invariant requirement is that every later claim resolve to frozen source units and an acceptance rule that existed before the evidence it judges.

A practical implementation should also preserve the distinction between *coverage controls* and *semantic controls*. Byte-covering partitions, exact-set identities, and content hashes prove that no declared input silently vanished. They do not prove a correct classification. Counterexamples, algebraic calibration models, and mutation plants test meaningful consequences. A robust instrument needs both. Coverage without semantics certifies a perfectly indexed error; semantics without coverage permits the inconvenient part of the corpus to disappear.

Finally, the write-up itself must stay downstream. A polished narrative is optimized for comprehension and will necessarily compress machine state. If later mathematics may cite that compression as authority, every omitted qualification becomes a new premise. The zero-authority firewall avoids that feedback loop: prose can be corrected, reformatted, or tailored to a journal without changing the object whose results it explains.

## 7 Scope and limitations

### 7.1 Internal validity, not physical truth

The result is conditional on pinned source bytes, a weak source presentation, the target class, and the frozen comparison category. It does not show that the source framework is empirically adequate, that its open propositions are true, or that classical kernels are nature's ontology. Optional physical axioms remain unselected. The record firewall expressly denies automatic inferences from dephasing or record formation to physical invocation, durable export, reset, erasure, heat, or a Landauer cost.

Nor does exactness imply absolute consistency. The calibration and representation packets provide explicit admitted models and exact proofs at their declared grades. They do not prove that every part of the surrounding informal corpus is consistent. Open and retracted source material remains in the annexes precisely so it cannot disappear from view or enter as an axiom.

### 7.2 The weak-presentation limitation

The central representation theorem is intentionally about the weakest cross-entry structure justified by the extracted source. It is not a strong functor from a fully normalized formal theory. In particular, FAITH-03's one-class-per-inhabited-hom quotient can be much coarser than a future presentation that formalizes all 467 relation-bearing declarations. The candidate's success therefore establishes a baseline: the canon can be represented without contradiction at this grade. It does not settle whether a richer quotient is representable or faithful.

This limitation is not repaired by FAITH-04. The provenance left inverse protects raw source words, while the minimality counterexample acts in a target-local dephasing sector. Raw bookkeeping can coexist with operational coarseness and with quotientable ambient structure. Treating the three as one faithfulness notion would erase the main negative result.

### 7.3 Relative minimality

"Not minimal" means not minimal in the comparison class frozen before the competitors were read. A broader target class could add smaller objects or new equivalences; a narrower class could exclude the two witnesses. The paper makes no absolute resource-minimality claim. Similarly, the $\Gamma_n$ family proves infinitely many inequivalent minimal classes but is not a complete enumeration of all congruence parameters. The normal-form classification describes the complete class abstractly as admissible congruence orbits.

Rigidity, irreducibility, quotient minimality, carrier minimality, uniqueness, and initiality remain distinct. The accepted candidate itself witnesses why: it is rigid and irreducible but fails both minimality tests and initiality. The paper therefore supplies no incomplete implication lattice disguised as a total hierarchy.

### 7.4 Automation, review, and cryptographic limits

Hash identities detect byte changes under a declared recipe; they do not interpret physics. JSON pointers can resolve to the intended field while the field's mathematical argument is wrong. Mutation controls can demonstrate discrimination only against their declared plants. Human review remains load-bearing for source semantics, proof scope, model admission, and whether a new adversary class has been omitted.

The demonstration is also a single-corpus study. Its scale is sufficient to exercise omissions, lifecycle changes, parameter ambiguity, provenance, and categorical comparisons, but it is not a statistical estimate of error rates in AI-assisted science. The repair ledger is a trace of this programme, not evidence that the same defects occur with the same frequency elsewhere.

The literature reconciliation is deliberately modest. It uses a fixed twelve-work manifest inherited from calibration, records exact locations, and never imports literature as a proof premise. Candidate-constraint and new-synthesis labels are not global novelty claims. Terminology-independent search and complete citation chaining are not claimed.

### 7.5 Cost and usability

The approach exchanges throughput for reconstructibility. Tens of thousands of row identities, explicit source spans, route ledgers, negative controls, and terminal evidence pointers are expensive to produce and review. Some of that cost is intrinsic to auditing a corpus of this size; some reflects the current JSON-first implementation. A future system should generate readable views from the machine state, provide recipe-aware replay libraries, and make wrong-key or wrong-payload checks harder to perform vacuously. It should not remove the denominators that make those views auditable.

## 8 Reproduction and archive

### 8.1 What is deposited

The reproducibility deposit contains the governed JSON and Markdown artifacts, their scope and freeze records, the exact git commit graph, the three pinned canonical sources needed by extraction, the present LaTeX source and bibliography, and the write-up verifier. The manuscript itself remains ZERO-AUTHORITY: it summarizes the banked state and cannot serve as a premise in the algebraic programme.

The one-command numerical replay is

```
python recent_papers/algebra_before_bit_writeup/verify_writeup_numbers.py
```

from the repository root. A successful run ends with

```
ALL WRITE-UP NUMBERS VERIFIED -- exit 0
```

and a zero process status. The script reads the banked artifacts rather than this manuscript, rebuilds the governed bank inventories, and asserts every paper-specific headline census, result status, terminal identity, and control battery reported here. A changed prose number cannot make the machine record agree; a changed artifact must pass its own earlier identity chain before it can affect the check.

### 8.2 Three levels of replay

Readers can choose three increasing levels of cost.

1. *Write-up replay* runs the numerical verifier and checks that the paper's tables match banked artifacts.
2. *Artifact replay* recomputes domain-separated row, registry, packet, terminal, and freeze identifiers and resolves the cited JSON pointers.
3. *Construction replay* rebuilds the governed artifacts from the pinned source bytes in stage order, exercising adversaries and controls and comparing every output byte with the bank.

These levels distinguish a convenient audit from a full rerun. The first checks reporting fidelity; the second checks state integrity; the third checks the deterministic implementation. None substitutes for reading the load-bearing proofs.

### 8.3 Why the terminal matters for reproducibility

Most reproducibility packages explain how to regenerate a preferred result. Here the terminal order is itself deposited. A replay must recover not only the exact quotient and carrier witnesses but also the vector with conditions true at slots nine and ten and the frozen selection of slot nine. Reordering the terminal names changes the order-sensitive digest. The negative landing is therefore part of the reproducible output, not an interpretive gloss added by this paper.

The archive is deposited at doi:10.5281/zenodo.22067781. Together with the immutable git commit and exact replay instructions, it fixes the explanatory record used for journal submission.

## 9 Conclusion

A large machine-assisted theory project needs an object more durable than a persuasive answer. In this case the object is a chain from canonical bytes to a typed signature, from a frozen signature to a qualified representation, and from that representation to a comparison-class theorem. The chain succeeded in the limited existence question and then refused the stronger minimality story: the accepted representation is irreducible and rigid, yet a proper raw-provenance quotient and a strictly smaller carrier both preserve the complete qualified profile, and a phase invariant obstructs initiality.

That combination is the methodological result. The same instrument that admits an exact construction must be able to select an exact negative without changing its language, denominator, or terminal order. A framework can remain the ground of a demanding demonstration while the replayable, refusal-capable instrument becomes the figure. Whether the source framework is physically right remains outside the claim.

## A Machine-state summary

Table 7 gives a compact route through the deposited machine record. File identities, not this table, are authoritative.

Table 7: Governed unit summary.

| unit | input fixed first | principal machine output | landing |
|---|---|---|---|
| C0/C1 | target-neutral process grammar, models, grades, calibration classes | definitions and calibrated core theorems with typed partials and refutations | definitions fixed; core theorem family partial or unresolved |
| ABB-FS | extraction schema, lifecycle, provenance and record firewall | 2,983-entry fixed signature with 1,042 open and 5 inactive entries | signature and provenance fixed; ABB-E scope required |
| ABB-E | weak source presentation, target class, equivalence, 32 proof contracts | one admitted candidate, 32 exact row passes, seven qualifications | conditional representation of frozen weak presentation |
| ABB-M | representation-generated core, profile order, six separate notions, eleven contracts | exact core, quotient and carrier witnesses, minimal-class census, rigidity and initiality results | not quotient-minimal; carrier counterexample retained as later-established condition |

## B Statements and declarations

**Funding.** No funding was received for this work.

**Competing interests.** The author declares no competing interests.

**Author contributions.** D.E. conceived the study, set the governance and countersign protocol, reviewed the mathematical and machine artifacts, and is accountable for the manuscript and deposited record.

**Use of generative AI.** OpenAI Codex was used under human direction to construct machine artifacts, verification code, and drafts of explanatory text. It is not an author. Every governed mathematical artifact was subjected to the separate countersign process described in the paper; the human author reviewed the claims and remains responsible for accuracy, scope, citations, and the final manuscript.

**Data and code availability.** The governed artifact chain and replay code are in the public git repository pinned in the artifact-status paragraph. The exact paper, source, governed artifact bytes, manifest, checksums, and replay code are archived at doi:10.5281/zenodo.22067781.

**Ethics approval and consent.** Not applicable.

**Zero-authority boundary.** This manuscript and its outline are explanatory outputs. They are barred as mathematical premises in the ABB artifact chain and cannot modify its definitions, evidence grades, or terminals.

## References


[1] Joshua R. Klein and Aaron Roodman. Blind analysis in nuclear and particle physics. *Annual Review of Nuclear and Particle Science*, 55:141–163, 2005. doi: 10.1146/annurev.nucl.55.090704.151521.

[2] Robert MacCoun and Saul Perlmutter. Hide results to seek the truth. *Nature*, 526:187–189, 2015. doi: 10.1038/526187a.

[3] Brian A. Nosek, Charles R. Ebersole, Alexander C. DeHaven, and David T. Mellor. The preregistration revolution. *Proceedings of the National Academy of Sciences*, 115(11):2600–2606, 2018. doi: 10.1073/pnas.1708274114.

[4] Christopher D. Chambers. Registered reports: A new publishing initiative at cortex. *Cortex*, 49(3):609–610, 2013. doi: 10.1016/j.cortex.2012.12.016.

[5] Roger D. Peng. Reproducible research in computational science. *Science*, 334(6060):1226–1227, 2011. doi: 10.1126/science.1213847.

[6] Mark D. Wilkinson, Michel Dumontier, IJsbrand Jan Aalbersberg, et al. The fair guiding principles for scientific data management and stewardship. *Scientific Data*, 3:160018, 2016. doi: 10.1038/sdata.2016.18.

[7] David Elliman. The rigid ground truth framework: A protocol for catching ai confabulation in speculative research, 2026. URL https://doi.org/10.5281/zenodo.20468553.

[8] David Elliman. Relocating trust to the verifier: A truth-maintenance system for an ai-generated theory of everything, 2026. URL https://doi.org/10.5281/zenodo.20574473.

[9] Jonathan Barrett. Information processing in generalized probabilistic theories. *Physical Review A*, 75:032304, 2007. doi: 10.1103/PhysRevA.75.032304.

[10] Bob Coecke and Aleks Kissinger. *Picturing Quantum Processes: A First Course in Quantum Theory and Diagrammatic Reasoning*. Cambridge University Press, 2017. doi: 10.1017/9781316219317.

[11] Wojciech H. Zurek. Decoherence, einselection, and the quantum origins of the classical. *Reviews of Modern Physics*, 75:715–775, 2003. doi: 10.1103/RevModPhys.75.715.

[12] Wojciech H. Zurek. Quantum darwinism. *Nature Physics*, 5:181–188, 2009. doi: 10.1038/nphys1202.

[13] Joseph P. Simmons, Leif D. Nelson, and Uri Simonsohn. False-positive psychology: Undisclosed flexibility in data collection and analysis allows presenting anything as significant. *Psychological Science*, 22(11):1359–1366, 2011. doi: 10.1177/0956797611417632.

[14] Timothy Lebo, Satya Sahoo, and Deborah McGuinness. Prov-o: The prov ontology. W3c recommendation, World Wide Web Consortium, 2013. URL https://www.w3.org/TR/prov-o/.

[15] Lucien Hardy. Quantum theory from five reasonable axioms, 2001. URL https://arxiv.org/abs/quant-ph/0101012.

[16] Giulio Chiribella, Giacomo Mauro D'Ariano, and Paolo Perinotti. Probabilistic theories with purification. *Physical Review A*, 81:062348, 2010. doi: 10.1103/PhysRevA.81.062348.

[17] Lluís Masanes and Markus P. Müller. A derivation of quantum theory from physical requirements. *New Journal of Physics*, 13:063001, 2011. doi: 10.1088/1367-2630/13/6/063001.

[18] Samson Abramsky and Bob Coecke. Categorical quantum mechanics, 2008. URL https://arxiv.org/abs/0808.1023.

[19] Peter Selinger. A survey of graphical languages for monoidal categories. In Bob Coecke, editor, *New Structures for Physics*, volume 813 of *Lecture Notes in Physics*, pages 289–355. Springer, 2011. doi: 10.1007/978-3-642-12821-9_4.

[20] Alexander Wilce. Symmetry, self-duality and the jordan structure of quantum mechanics, 2011. URL https://arxiv.org/abs/1110.6607.

[21] Howard Barnum and Alexander Wilce. Local tomography and the jordan structure of quantum theory. *Foundations of Physics*, 44:192–212, 2014. doi: 10.1007/s10701-014-9777-1.